Ontology and Information Retrieval: The Case of the Fine Arts

Daniel Isemann

Thesis submitted for the Degree of Doctor of Philosophy

Department of Computer Science
University of Dublin
Trinity College

20 March 2012

To my late great-uncle Ludwig
and my late uncle Roland

Declaration

I declare that this thesis has not been submitted as an exercise for a degree at this or any other university and it is entirely my own work. Published or unpublished work of others wherever it is included, is duly acknowledged in the text.

I agree to deposit this thesis in the University's open access institutional repository or allow the library to do so on my behalf, subject to Irish Copyright Legislation and Trinity College Library conditions of use and acknowledgement.

Summary

This thesis explores methods of computational ontology for information retrieval in a knowledge rich domain. The example case we selected for our study is the domain of the fine arts. Various aspects were under examination, including integration of existing data into a system of ontology-informed knowledge representation, modelling details of such a representation and, in particular, user interaction with knowledge represented in this way through a graphical interface.

There is considerable literature documenting theoretical foundations and practical research efforts that bear relevance for an ontology of visual art and its application to aesthetic image retrieval. This includes work in computational ontology, research into informal conceptualisations of art and studies on user retrieval behaviour in this domain. There are a number of existing classification systems at the intersection of knowledge representation and art and heritage expertise; these systems, in particular the CIDOC Conceptual Reference Model, a top-level ontology for the domain of cultural heritage, have been used to explore ontology-informed representation and access to fine art in our implementation.

We employed methods of ontology development and knowledge representation technology, recently developed in the context of the Semantic Web, in the creation of a prototype application for representing and accessing information on works of fine art. The design of our system was informed by research into the retrieval behaviour and information needs of laypersons and experts in our domain. This enquiry was carried out using standard methods of qualitative research and statistical analysis.

We carried out a study among 21 art specialists and 27 laypersons on the choice of terminology for searching digital images of paintings. We found statistically significant differences between the two groups, with specialists focussing more on metadata and on abstract terms and laypersons more on concrete conceptual descriptors of the image subject. We worked closely with art experts at a major national gallery to learn what kind of information they record about works of art and to help us translate their retrieval needs into our system design.

Our implementation focuses on a number of points along a processing chain from a tabular legacy representation of information on visual art to a graphical user interface through which description logic queries may be formulated to access and retrieve such information. Different aspects of this chain were explored through the development of three systems built during the course of my study: *Ontology Populator*, *ART Ontology* and *Artfinder*. These systems were developed in the Java programming language and the Web Ontology Language OWL 2.

Ontology Populator and its pre-processing algorithms use a set of heuristics which is aimed at enriching a relational data set during its transformation into a knowledge base. This covers, in particular, the conversion of free text descriptions of dates into ISO 8601 compatible date expressions.

ART Ontology captures some of the conceptualisations and distinctions made by practitioners in

art and heritage curation. It forms a layer between the CIDOC CRM top-level model and an arbitrary set of data on visual art and is able to emulate temporal inferences for the purposes of querying the associated knowledge base. In our application of the CIDOC CRM we encountered difficulties and discovered incoherences in the model which were reported back to the CIDOC CRM developing community.

Artfinder is a visual query builder which allows users, unfamiliar with knowledge representation techniques or query languages, to access a knowledge base. The structure of the Artfinder user interface is automatically determined by details in ART ontology and the corresponding knowledge base.

An evaluation of *Artfinder* was conducted in which 10 art experts and 10 laypersons were asked to produce queries following a 14-point task protocol. While our evaluators described the use of *Artfinder* as intuitive and welcomed many of its options, we found that only a third of the expressions generated through the system corresponded to our predictions. Simple queries, however, could be produced consistently after a short exposure to *Artfinder*. There were statistically significant differences in expert and layperson usage of the system, with experts showing a tendency to overspecify and laypersons a tendency to underspecify in their queries.

Acknowledgements

I would like to acknowledge and thank a number of people who made this dissertation possible. First of all, I would like to thank my supervisor Prof Khurshid Ahmad, who was always available and attentive to my problems. My earlier interactions with my co-supervisor Dr Mícheal Mac Airchinnigh, in particular his eye for the bigger picture, are gratefully acknowledged. Moreover I would like to acknowledge the hospitality of Trinity College and the School of Computer Science and the Faireacháin project (IP/2009/0595) which provided me with a scholarship.

I would like to thank the experts at the National Gallery of Ireland for their interest in my work and the generous donation of their time which made my project possible. First and foremost in this context I would like to thank Dr Adriaan Waiboer, Curator of Northern European Art, for the many hours he invested in my project and for his advice which went well beyond the domain of fine arts. In addition, many thanks are due to Anne Hodge, Curator of Prints and Drawings and her colleague Niamh MacNally, Donal Maguire from the Centre for the Study of Irish Art and Raffaella Lanino from the Collection Management team, as well as Leah Benson and Andrea Lydon.

Furthermore, I would like to thank Dr Tim Fernando, Dr Arthur Hughes and Dr Eamon Mullins for their helpful advice regarding technical questions and also the many people in and around Trinity College who provided nuggets of wisdom when they were needed. A special thanks goes to the "lads" in the lab, Aaron, Barry, Nick and Peter.

On a personal note I would like to thank my family and this explicitly includes those who have become my family for their fantastic support through the years. I would also like to take the opportunity to thank my "Irish family", my house mates, and here in particular Shane and Robert, for moral and, at times, nutritional support. Heartfelt thanks go to Victoria for her marvellous encouragement during the final write-up and her active help with the transcripts.

Especially big thanks are also due to the many people who I have met in Trinity College, but who will hopefully stay with me "extra muros". These include in particular: Serena, who I would like to thank for the sheer endless amounts of positive energy she has to disperse; Anne for providing, in more than one sense, the measure of all things; and Jools, our occasional differences in "ontology" notwithstanding, a most extraordinary friend!

Finally, I have to name the triumvirate without whose friendship my time in Trinity would not have been the same: Brendan, Kevin and Gerard. Thanks for your support, thanks for an open ear, thanks for countless hours of side-splitting, slipper-kicking merriment.

Contents

Chapter 1

Introduction

The GOOGLE-powered ART PROJECT (2011) allows for the virtual exploration of 17 art galleries and other places of heritage interest.[1] The project is one of many, advocating and facilitating access to works of art to an ever widening public which years ago would have been available only to the select few. Art and heritage images are associated with a rich and traditional body of expert knowledge, and curators of these images make careful decisions on what to display in a museum and which collateral information to present with the selected works. The ambulatory mode of exploration provided by ART PROJECT creates a novel experience by virtually recreating the physical space in which art is displayed. This thesis is about mapping and providing access to the conceptual space in which the information about works of fine art resides.

Computational ontologies and other Semantic Web technologies form the context in which we define and motivate our research questions (Sections 1.1 and 1.2). We will then outline the content and structure of the thesis (Section 1.3) and summarise our key conclusions (Section 1.4). The introduction concludes with a commented overview of the publications that have resulted from the research conducted for this thesis (Section 1.5).

1.1 Context

The emergence of the World Wide Web has changed the way we deal with and think about information. In the context of the Web, we have become used to accessing apparently relevant data faster than ever and in novel ways. Different modes of data — audio, visual and textual — are easily combined and often neatly integrated to maximise the information content of the data retrieved. More and more data and information is ever more readily available from a great variety of sources — both authoritative and non-authoritative.

This environment has also given rise to new ways of organising data. Some of the earliest attempts to bring order to the Web such as the YAHOO directory[2] were modelled on a traditional cataloguing approach. But soon the rapidly growing search engine industry adopted a different paradigm. Crawlers and site ranking algorithms nowadays facilitate the retrieval of relevant documents from the haystack many million times a day. Text mining and information extraction technologies attempt to derive high quality information from unstructured text.

In 2001, a vision for the future development of the World Wide Web was proposed: the *Semantic Web* [12]. According to this vision, the Web of today will gradually emerge into a web, where docu-

[1] `http://www.googleartproject.com` (last accessed 27/10/2011).
[2] `http://dir.yahoo.com` (last accessed 14/08/2010).

ments and data are semantically annotated so that machines can interpret and process their meaning and interact sensibly according to this meaning.

The idea that a human readable presentation should not be the primary form in which data is stored and maintained was recognised independently of the Semantic Web. This is evidenced by the existence of database backed web services, for instance, weather forcasts or online phone books. In the context of the Semantic Web, however, this 'data-centric' view is even more emphasised and combined with the paradigm of a decentralised 'web'. This combination holds the promise to alleviate the burden of information integration on a far greater scale than has hitherto been achieved.

A number of technology layers that need to be realised in order to arrive at the full vision of the Semantic Web were postulated. The Semantic Web 'layer cake' comprises three sets of technologies: Conventional web technologies, which include text-level encoding (e.g. UNICODE) and unique addressing (URI/IRI), together with the *Extensible Markup Language* (XML) which provides a syntactic framework for documents and data, form the basis of the Semantic Web. This underlines that the Semantic Web is not envisaged as a separate entity competing with the World Wide Web, but rather as something which will coexist and complement the Web of today and into which this Web may gradually evolve. The second set consists of existing semantic technologies. These include schemata for richer annotation and interchange of data, such as the *Resource Description Framework* (RDF), for taxonomic structure among concepts and relationships (RDFS) and for more complex restrictions and axioms (OWL, SWRL). Query languages, similar to those known from database systems, can facilitate access to these semantic layers. The third set of layers is that of future web technologies. Current semantic technologies can reduce the burden of information integration and inference on the part of the retriever. However, the burden of the validation and reliable interpretation of the information conveyed through the semantic stack remains. The third section of the 'cake' deals with these issues: unifying logic provides the basis for proven facts, that together with digital signature and encryption may perhaps lead to 'trust'. User interfaces and applications, while not part of the core stack, may access and use any of the semantic layers, from 'data interchange' to 'trust' (Figure 1.1).

Some illustrations of the Semantic Web architecture do not include applications and some seem to suggest that interfaces and applications require all the layers of the semantic stack to be completed.[3] In our opinion, however, applications and intelligible user interfaces at any level of the layer cake will be crucial for the success of semantic technologies.

The layers of taxonomies and ontologies carry most of the burden of expressing semantics on the Semantic Web. Taxonomies have been used for centuries in biological classification and have later entered other domains. Ontologies in the sense of a more generalised form of taxonomy have been discussed in computer and information science for about 20 years, but have attracted special interest in connection with the Semantic Web. Their function in this context is to explicitly express a shared conceptualisation of a given domain and thus help to put individual pieces of data into proper relation. A description of ontologies, their roots in philosophy, their theoretical underpinnings and practical implications in computer science is contained in Chapter 2. We will point out that there are differences in opinion as to how ontologies should be constructed and what they should aim to achieve.

There are a number of early adopters of computational ontologies and work in biology and medicine is perhaps the most salient instance. The enormous amounts of data generated in genome-

[3]Compare the original Semantic Web layer cake at `http://www.w3.org/2000/Talks/1206-xml2k-tbl/slide10-0.html` (last accessed 15/09/2011) and [68], as well as `http://www.w3.org/2007/03/layerCake.png` (last accessed 17/09/2011).

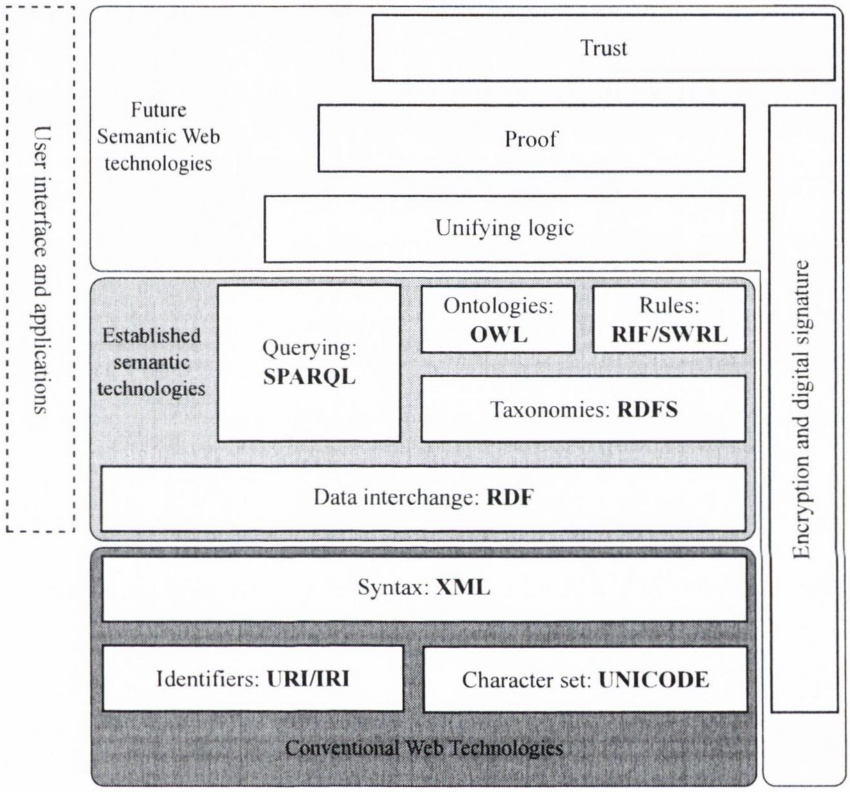

Figure 1.1 – The Semantic Web Layer Cake as an illustration of how different technologies build upon each other and where the current frontier of semantic technology lies. The layers up to and including the ontology layer have received implementations so far which are either an officially endorsed standard or else widely used by practitioners in the field. The acronyms for these standards are shown in bold letters in the diagram. The main difference between the stack shown in this figure and the layer cake advocated by the World Wide Web Consortium (`http://www.w3.org/2007/03/layerCake.png`, last accessed 17/09/2011) is that the "User interface and applications" box reaches down across all of the semantic levels and is not placed above "Trust" (compare [132, p. 264] for a similar layer cake).

related research requires scientists from different laboratories to frequently consult biomedical databanks for gene or protein sequences or bibliographic references on highly specialised problems. To integrate such information from distributed sources and facilitate access to it, a number of large scale ontologies were developed. Some of these ontologies may even be used to predict functional attributes of novel genes or proteins. However, problems related to data integration and logical inference persist (see [137] for an overview).

Scientific domains lend themselves to ontological description. The task of ontology in these domains is to terminologically disambiguate what is conceptually already unambiguous. A similar situation is found in the legal domain, where terminology is precisely defined, but has complex interdependencies. The need for 'question answering', as opposed to 'document retrieval' is particularly prominent for law professionals, who often have to take into account numerous exceptions to a certain rule. Legal ontologies and Semantic Web for the law are vibrant areas of research (see [7] for an overview).

One area that has received considerable attention from an early stage of Semantic Web development is the field of cultural heritage. Here, terminology is perhaps less rigourously defined and less fixed than in law or science. A call was made for a *Semantic Web for Culture* which, among other things, should allow for a representation of changes in meaning over time [159].

Cultural heritage as an application domain is also advocated in the "Use Cases and Requirements" document that forms part of the World Wide Web Consortium's original recommendation of the Web Ontology Language (OWL). As an example, the World Wide Web Consortium describes the use of an ontology of antique furniture including information on style, materials, creation date and cultural affiliation of artefacts for organising a collection of images depicting such furniture.[4]

Cultural heritage information often goes hand in hand with visual documentation in the form of images. Cultural heritage images, in particular visual art, have a rich associated body of background knowledge, which stems from centuries of art-historical research and curatorial efforts. In some cases, this background knowledge has been codified into classification systems and ontologies which are available in electronic format. We will discuss and compare some of these ontologies and resources for ontology construction in detail.

In the context of semantic annotation and retrieval of multimedia objects it is often argued that because of the sheer size of collections, any form of semantic markup would necessarily have to rely on a fully automated process for generating such markup. However, images of fine art are in comparison to photographic images much fewer and generally of longer lasting interest. This justifies a higher manual or machine-based effort in creating semantic markup for them. A comparable markup may not be feasible or worthwhile for everyday images that are continuously being published in large numbers. The FLICKR website, for example, hosts over six billion digital photos and receives thousands of new uploads every minute.[5] The digital 'volume' of cultural heritage images is increasing steadily, but at a much slower rate. The State Hermitage's digital collection[6] of cultural artefacts and works of fine art grew from 11,987 items in 2008 to 16,279 items in 2011. When broken down to subcategories, this growth averages about one painting, drawing and print every two days, and roughly four items per day in total. Similarly, the Tate online collection[7] has grown by just under two items per day over the last three years (Table 1.1).

---

[4] See Section 2.2 of `http://www.w3.org/TR/webont-req` (last accessed 1/09/2010).

[5] `http://www.flickr.com` (last accessed 17/09/2011).

[6] `http://www.hermitagemuseum.org` (last accessed 17/09/2011).

[7] `http://www.tate.org.uk` (last accessed 17/09/2011).

		Feb 2007	Sep 2008	July 2010	Sep 2011
State Hermitage	Paintings	2506	3092	3658	3760
	Drawings	1160	1522	1777	2205
	Prints		597	949	1124
	All artefacts		11987	15466	16279
Tate Online	All artefacts		68971	69834	71049

Table 1.1 – Development of the size of the online collections of two notable art museums: the *State Hermitage Museum* in Saint Petersburg and the *Tate* in the United Kingdom. The figures were established using the advanced search tools of the respective institutions' web pages.

1.2 Research Questions

Our principal research question is this:

> *What can well-structured background knowledge contribute to information retrieval in a knowledge rich domain?*

In order to explore this question, by using cultural heritage as an exemplar, we have developed an ontological model of important aspects of art related information and implemented and evaluated an ontology-based retrieval system for images of fine art. This process was guided by a number of secondary research questions, which contribute to different aspects of the principal research question:

1. *Which categories are used to describe and search for information on works of art?*

An answer to this question is a first step towards establishing a sound ontological description of our application domain. The structure of such an ontology of art is the subject of the second specific research question:

2. *Which ontology design principles are relevant for an ontology of fine art?*

Finally, in order to understand the contribution that ontologies like these can make to information retrieval, we look at usability aspects. Systems for querying the Semantic Web often require the user to input a query string in a complex and inaccessible query language. Even if an interface is provided for query construction, the use of such an interface typically expects at least a basic familiarity with knowledge representation languages or dedicated query languages [13], [33]. We are especially interested in a logically rich description of the domain. This is in contrast to many existing Semantic Web applications. Our last question therefore addresses the possibility of a successful communication between an untrained user (i.e. untrained in knowledge representation technologies) and a complex ontology:

3. *Can end-users successfully formulate complex ontology queries?*

A possibly important variable in this context, distinct from the intricacies incurred by formal knowledge representation, is the level of domain knowledge on part of the user. The annotations provided by experts are based on theoretical or experiential knowledge. Experts use a succinct style in their prose and their narratives are typically for their peers. The extent to which the rich annotations provided by experts can be retrieved by an arbitrary end-user is not clear.

Each of the questions is addressed by a mixture of literature review and methodological research. The second and third question were additionally tested in the context of an implemented system. The following section will give a brief overview of the work carried out for this thesis.

1.3 Thesis Outline

The research which is detailed in this thesis can be roughly divided into four parts. The first part concerns background research into how laypersons and experts categorise and conceptualise works of art. This comprises a number of qualitative interviews with heritage professionals (Section 4.1) and an online survey which was taken by a mixed group of art experts and laypersons (Section 4.2). In the second part, we developed an ontological framework for the description of works of art, which is based on the first part, on established ontology engineering methodologies and on the analysis and re-use of existing resources (Sections 3.1 through 3.4). In the third part, we implemented a user interface for querying a set of data described according to this ontological framework and a (straightforward) retrieval system for query answering (Sections 5.1 and 5.4). Finally, the fourth part consists of an evaluation of how effective the user interface is for translating a given user intention into the language of the underlying ontology. This evaluation was again carried out by a mixed group of domain experts and laypersons in a number of one-to-one evaluation sessions (Sections 5.5).

Broadly speaking, part one and part two represent work that was carried out in order to answer the first and the second research question respectively. Part four provided most of the contribution towards answering the third research question. Part three, on the other hand, was a necessary prerequisite for part four. It is worth pointing out that the work done for part three is independent of the chosen application area and applies equally to other knowledge domains if these can be modelled in a similar ontological framework.

1.4 Key Conclusions

In the course of our research project we interacted with art and heritage experts from the National Gallery of Ireland and from the Department of History of Art and Architecture in Trinity College. Our aim in this was twofold. On the one hand, we wanted to glean an insight into expert conceptualisations of art and make these conceptualisations fruitful for image search, perhaps also for the wider public. On the other hand, we wanted to understand experts' needs and expectations towards image retrieval systems, as they form a core user group of the technologies that we employed in their domain.

With respect to our first research question, we found that non-visual metadata is of key importance in searching for works of visual art, when compared with other imagery. Perceptual descriptors, on the other hand, such as terms referring to colour or texture, are rarely used. Moreover, we found indications for systematic differences in the terminology employed by specialists and laypersons when searching for fine art. While both groups used conceptual descriptors most frequently, specialists put greater emphasis on non-visual descriptors relative to laypersons.

The design and development of usable ontologies for fine art, our second research question, faces a number of conceptual and technological issues. Conceptually, much of the information traditionally recorded by art historians is less clearly defined than is perhaps the case in genome-research. The

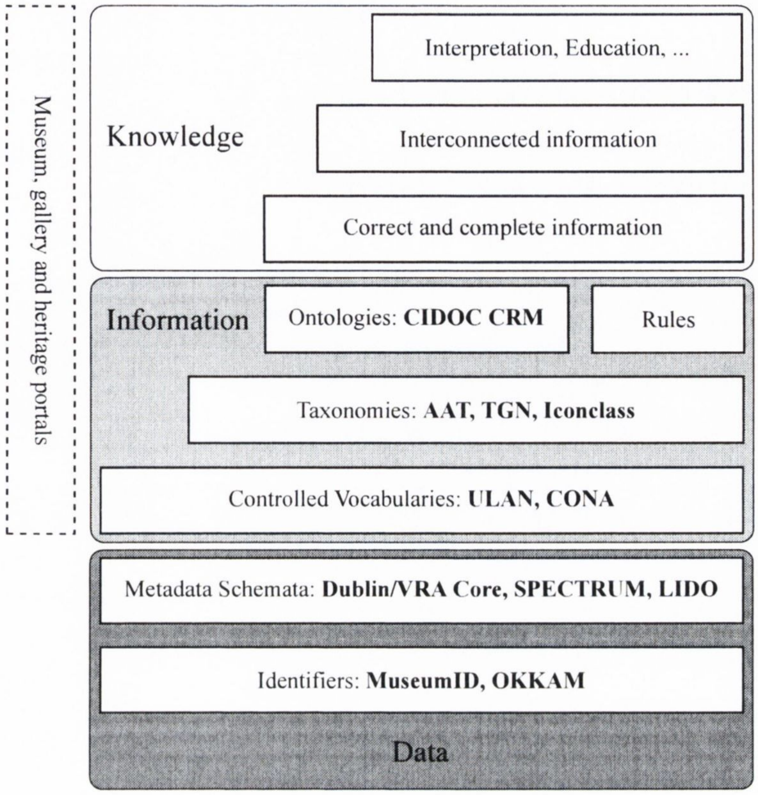

Figure 1.2 – A cultural heritage layer cake analogous to the Semantic Web Layer Cake. The acronyms printed in bold denote exemplary instances of the layer in question. In general these instances are more numerous and less standardised than the implementations of enabling technologies in the Semantic Web layer cake. The examples included in this figure make no claim to be complete.

historical outlook of the subject requires effective modelling of temporal and geographical information, but often temporal estimates are vague or implicitly tied to the dating of other events in current research. Event centred modelling of information was generally welcomed by experts and so was the inference of implicit information. With respect to the latter, however, the strict logical character of inference was not always appreciated. Some experts expressed the desire that records which "nearly" as opposed to strictly fulfilled certain criteria should be considered in retrieval scenarios as well.

Technologically, the implementation of ontologies for the fine arts is aided by, if not entirely dependent on, standards and implementations, which address the different Semantic Web layers in the chosen domain. In fact, the 'layer cake' we introduced above for the Semantic Web in general has an equivalent for the cultural heritage domain. At the bottom, tools for creating or reusing unique object identifiers and metadata schemata facilitate the collection and expression of relevant data. The core semantic layers have been filled by a range of terminology and ontology resources. This second stratum, if it is connected with the underlying data, may perhaps be called 'information'. Whether rules and inference engines can turn the sum of information into something more ambitious such as 'knowledge' remains an open question (Figure 1.2).

To test our third research question we implemented a prototype retrieval system comprising a representation of terminological (TBox) and assertional (ABox) knowledge about works of art. The assertional part was automatically generated by an ontology populator. A graphical query interface was designed and developed to generate description logic queries for accessing the knowledge base.

We evaluated the system with a mixed group of ten domain specialists and ten laypersons. Participants were asked to produce queries using a 14-point task protocol designed specifically to tease out the mapping between user conceptualisation and equivalent in the language of our ontology.

We found that most of the participants in our evaluation cohort described our interface as intuitive to use and well suited to their needs. In the objective part of the evaluation, we found that in general only a relative majority of participants reflected the expectations and premises on which the design of our system was based. Perhaps less than half of the queries submitted through our system were completely in line with our users' conceptualisations. As was the case with our image retrieval survey, we found indications for differences in behaviour between specialists and non-specialists, with the former tending to overspecify and the latter showing a tendency to underspecify their queries.

1.5 Publications and Presentations

In the course of this project a number of papers covering different aspects of the research were submitted to the scrutiny of professional audiences of both the computer science and the cultural heritage communities. Publications are listed in the order of publication date, together with a succinct summary.

Isemann, D., and Ahmad, K. (2009). Navigating cultural heritage in style. Sketching an ontological representation of metadata: the example of artistic period and style. *Museum Ireland*, 19, 149-155. Reviewed by editorial board. [85]

> *The paper gives a non-technical introduction to ontology, as it is used in computer science, and to the CIDOC CRM, an ISO standard top-level ontology for the cultural heritage domain. An event-centred modelling of information about cultural artefacts is explained and exemplified using the Mona Lisa as an example. The benefits of automatic reasoning in connection with such an information model are outlined.*

Isemann, D., Mac an Airchinnigh, M., and Ahmad, K. (2010). Navigating heritage on well-grounded terminology and robust ontology. In *Specialized Language in Global Communication – Proceedings of the XVIth European Symposium on Language for Special Purposes (LSP), Hamburg (Germany) 2007*, W. von Hahn and C. Vertan, Eds., volume 30 of the series 'Sprache in der Gesellschaft. Beiträge zur Sprach- und Medienwissenschaft', Peter Lang, Frankfurt am Main, Berlin, Bern, pp. 78-89. Fully peer-reviewed. [88]

> *The paper compares three terminological resources with a view to the vocabulary and the structural information they contain to describe the human figure in the context of aesthetic imagery. The three resources — Iconclass, WordNet and the Foundational Model of Anatomy — differ in their domains and intended applications. In order to evaluate the information contained in these resources from the angle of visual art they are compared against terms extracted from 14 text books on anatomy for artists. Differences in terminological coverage and the level of detail of part-whole relations are discussed. It is concluded that structurally neither Iconclass, nor WordNet could support precise inferences about the human figure. An experimental ontology of body parts is introduced, which supports such inferences and has sufficient domain coverage, but is considerably*

smaller than the extensive Foundational Model of Anatomy, which is designed for medical applications.

Isemann, D., and Ahmad, K. (2011a). Query terms for art images: A comparison of specialist and layperson terminology. In *Proceedings of the 25th BCS Conference on Human-Computer Interaction* (Newcastle, UK, 4th-8th July 2011). Fully peer-reviewed. [87]

The paper presents a study on the use of keywords for searching digital images of works of art. 48 participants were asked to supply query terms for the same set of paintings. Some of the participants were domain specialists as they had received a formal education in art history. Similarities and differences between the terminology employed by specialists and non-specialists and implications for art image retrieval systems are discussed.

Isemann, D., and Ahmad, K. (2011b). ART ontology: Experiences with CRM-based organization of heritage images. In *Proceedings of CIDOC 2011 – Knowledge management and museums* (Sibiu, Romania, 4th-9th September 2011). Reviewed and accepted based on abstract. [86]

The paper discusses two aspects of an ontology-based system for the retrieval of aesthetic images from a purpose-built knowledge base. The first part is concerned with knowledge modelling using the CIDOC CRM. The paper discusses the modelling of time-related information and a number of difficulties that arise in the application of the CIDOC CRM in its current form. The second part of the paper introduces 'Artfinder' a novel user interface for creating OWL DL queries and retrieving matching results. Key results from an evaluation of 'Artfinder' that involved 20 participants are presented.

Furthermore, a number of presentations without associated publications were made to various audiences. Excluding those that were given to members of the School of Computer Science at Trinity College in Dublin, the following two remain:

Isemann, D., Ahmad, K. and Mac an Airchinnigh, M. (2007). Annotating Fine Art Images. *Poster Presentation at 2nd Lewis Glucksman Memorial Symposium of the Trinity Long Room Hub* (Trinity College, Dublin, 13th June 2007). Invited participation.

The relationship between visual images and their contextual information is explored. The scope of the research project is outlined.

Isemann, D. (2008). Cyber Visits to a Gallery. *Presentation at the National Gallery of Ireland Research Day* (National Gallery of Ireland, Dublin, 6th March 2008). Reviewed and accepted based on abstract.

Different notions of image similarity are introduced and compared, ranging from physical to 'semantic'. The merits and limitations of the Iconclass system for facilitating semantic access to images and ways of overcoming its limitations in light of recent developments in computational ontology are discussed.

Chapter 2

Literature Review

Literature from various academic disciplines has informed this research project. Much of it is centred on the concept of ontology as a formal domain model that rests on the shoulders of a long philosophical tradition (Section 2.1). The question of how to categorise, describe and retrieve aesthetic and fine art images has been the subject of academic research ranging from art history to information and computer science (Section 2.2). In Section 2.3 we discuss explicit ontologies and other structured resources which codify knowledge about visual artefacts. Finally, Section 2.4 reviews research and existing implementations that apply ontologies and Semantic Web technologies to the visual arts and heritage domain. Section 2.5 summarises our review of literature and related work.

2.1 Notes on Ontology

The notion of ontology, as it is understood by information and computer scientists, was inspired by work in ancient and modern philosophy. We examine the philosophical notion of ontology first and then characterise its meaning and use in the context of computer science, together with a definition of the terminology regarding ontologies and their components that we will be using throughout this thesis. Finally, we briefly discuss different positions concerning the usefulness of a philosophical approach to ontology engineering.

2.1.1 Philosophical Origins

The term *ontology* has its roots in the Latin word *ontologia* which in turn derives from the Greek words ov, $ov\tau o\varsigma$ (being) and $\lambda o\gamma o\varsigma$ (science, theory). Based on these roots *ontology* is sometimes defined as "knowledge of being" [126, p. 401] or "science of being in general" [102, p. 670]. According to the philosophical literature the Latin form of the term originated in the specialised philosophical discourse of the 17th century [16], [101], [108]. It was thought to have been coined in the year 1613 [140]. More recent findings, however, suggest that it occurred as early as 1606 in Jacob Lorhard's "Ogdoas Scholastica".[1]

Philosophical ideas related to the study of ontology are much older than the term itself. They can be traced back to Aristotle, Plato and even Parmenides [16], [31], [45]. In the modern era Christian Wolff, Immanuel Kant and Martin Heidegger are widely credited with major contributions to the field [24], [105], [108].

[1] See `http://www.formalontology.it/jacob-lorhard.htm` (last accessed 7/03/2010).

Ontological question	Possible Answers	
Which entities are of primary importance?	things, objects (cf. Substantialists)	events, processes (cf. Fluxists)
How many types of entities should be accounted for?	entities of all levels of aggregation or granularity (cf. Adequatists)	only one or a few distinguished classes of entities (cf. Reductionists)
What does ontology make statements about?	an independently existing reality (cf. Realists)	theories, languages or belief systems ("Non-Realists"[2])

Table 2.1 – Dichotomies in philosophical ontology arising from the answers given to fundamental ontological questions by different schools of thought. The table summarises a discussion by Smith [140].

Smith gives a high-level overview of philosophical ontology and defines it as "the science of what is, of the kinds and structures of objects, properties, events, processes, and relations in every area of reality." [140, p. 155]. In the course of history different philosophers have responded to the challenges of this field with widely differing ontological conceptions. According to Smith, philosophical ontologists and their theories can be classified along a number of dichotomies: *Substantialists* vs. *Fluxists*, *Adequatists* vs. *Reductionists* and *Realists* vs. *Non-Realists*[2] (see Table 2.1).

The primary entities favoured by substantialists and fluxists are sometimes referred to as *continuants* and *occurrents* respectively. Smith points out that the dichotomy between fluxism and substantialism is less pronounced among adequatist philosophers since they tend to accept continuants and occurrents on an equal footing.

The relativistic approach of determining the ontology of a specific theory, rather than the ontology of an objectively existing reality, gives rise to talk of 'an ontology' or 'ontologies' instead of the traditionally uncountable term 'ontology'. The idea to extract ontological assumptions from a theory or system of belief is often referred to as making the *ontological commitment* of that theory or belief system explicit. It gained currency through the work of the philosopher Willard Van Orman Quine [24, p. 118] [140, p. 156]. Smith, however, points out that eliciting the ontological commitment of a theory as a method of philosophical ontology is not equivalent to denouncing ontological realism. In particular he emphasises that Quine himself (at least early Quine, who Smith is mainly interested in) was a realist philosopher who believed that the ontological commitment of various theories in the natural sciences could be the key to a general ontological understanding of all of existence. The key idea is that a successful theory that is widely considered to be true of reality should be a good starting point for a realist ontology.

2.1.2 Ontology in Computer Science

While it is not explicitly stated in the literature, it seems likely that the sense of 'ontology of a theory' mentioned in the previous section, has led to the adoption of the philosophical term 'ontology' by the knowledge engineering and Artificial Intelligence (AI) communities. John McCarthy was possibly

[2]Smith does not give a label to the position opposed to realism. He treats the distinction between realist and non-realist views under the heading of "Internal vs. External Metaphysics" and, especially the relativistic approach to ontology in the information sciences, in the section "Conceptualisations". Cf. [140, pp. 157f. and 161f.].

the first to use the term in this new context in the late 1970s.[3] However, it was not until the 1990s that the term gained currency. By then 'ontology' had come to describe system components that formalise terms, concepts and relations pertaining to a certain domain [56], [115]. Subsequently, other definitions, sometimes modifications of earlier ones, were put forward in the literature [18], [68], [147]. The definitions of ontology in the new millennium include one in the *Encyclopedia of Database Systems*:

> In the context of computer and information sciences, an ontology defines a set of representational primitives with which to model a domain of knowledge or discourse. The representational primitives are typically classes (or sets), attributes (or properties), and relationships (or relations among class members). The definitions of the representational primitives include information about their meaning and constraints on their logically consistent application. [58]

This characterisation lists the typical elements of computational ontologies — classes, attributes and relationships. Formally, an ontology can be seen as a logical theory, i.e. a set of statements in some logical language (cf. [57], [60]).

There have been attempts at formalising the structure typically exhibited by ontologies, usually involving a taxonomic hierarchy of classes. A simple, rigorous definition along these lines in a limited context is given by Merrill [110]. Formally defined concept lattices are also a variety of ontologies [29], [48].

A formal, set-theoretic definition of ontologies and related concepts is given by Cimiano [22]. The definition reflects the typical elements, that have come to be expected in applied ontology contexts as well as the relationships between these elements (e.g. the separation of conceptual structure and terminology layer). This formal account attempts to encompass a wide variety of ontologies, including all ontologies expressable in the Web Ontology Language (cf. Section 3.1.2). We will characterise the most important elements of ontologies according to this account and illustrate each element with examples from the domain of art history.

Concept identifiers, *concepts* or *classes* represent concepts in the domain of discourse that the ontology covers. Classes in an ontology are usually ordered into a *class hierarchy* or *taxonomy* by partial order. The class hierarchy has a unique top element with respect to this order and is required to contain a unique most specific parent class for any two classes in the ontology, where "most specific" means minimal according to the partial ordering. The partial ordering is commonly interpreted as *subclass-superclass* relation. The top element represents the most general class which subsumes all other classes in the ontology, often interpreted as "Thing" or "Entity". The unique most specific parent class of two classes is called the *least common subsumer* of these two classes. In practical ontology systems there is sometimes a distinction between *named* and *anonymous classes*. Named classes are representational primitives, whereas anonymous classes are complex class descriptions built up using Boolean operators or other mechanisms. In this context "class" is sometimes only applied to named classes, while anonymous classes go by other names like "restrictions" or the aforementioned "complex class descriptions". If Boolean class operators are present in an ontology formalism then the union of two classes is their least common subsumer. For example, in a visual arts ontology possible named classes could be: 'artist', 'work of art', 'painting' etc. In terms of the subsumption relation, 'painting' would be a subclass of 'work of art'. An example for an anonymous class would

[3] Compare Nicola Guarino's and Christopher Welty's comments on this issue [61, p. 61].

be 'portrait and painting' which could represent portrait paintings as opposed to portrait drawings or portrait sculptures.

Relation identifiers or *relations* for short represent relations on the classes in the ontology. Binary relations have two classes as their *domain* and *range* respectively. If n-ary relations are considered their *relation signature* is an n-tuple of classes in the ontology.[4] For example, in an art ontology there could be a relation 'is_creator_of' between an artist and a work of art. Relations in an ontology may be structured by a partial order called *relation hierarchy*. This order is such, that whenever two relations are comparable according to the order their relation signatures have the same arity and all ontology classes in the signature of the subrelation are in a subclass relation to the corresponding classes in the signature of the superrelation. For binary relations in particular if r_1 is a subrelation of r_2, then the domain of r_1 is a subclass of the domain of r_2 and the range of r_1 is a subclass of the range of r_2.

Attributes represent binary relations, which have a class in the ontology as a domain and a data type like string or integer as a range. For example, there could be an attribute 'has_width_in_cm' which assigns a floating-point number to a painting.

Instances or *individuals* represent specific objects in the domain. Usually instances are not considered to be part of an ontology proper but part of a *knowledge base* for an ontology [22, p. 15]. In more applied settings however this distinction is sometimes blurred [95], [118]. Instances are thought of as instantiating one or several classes in the ontology. They may be asserted as direct instances of a class or indirectly instantiate a class in virtue of instantiating one of its subclasses. They may also be inferred to instantiate a class as a consequence of the axioms associated with an ontology (see the paragraph on *axioms* below). For example, Mona Lisa or Leonardo da Vinci could be instances in an art ontology. Mona Lisa could instantiate the class 'painting', whereas Leonardo could instantiate the classes 'artist' and 'man'. Indirectly, Leonardo could also instantiate the class 'human being' if 'human being' was a superclass of 'man'.

In the same way that instances of classes represent individual objects that instantiate or belong to a class, *relation* or *attribute instances* are concrete instantiations of relations and attributes respectively. In the case of n-ary relations, these instances can be thought of as n-tuples of individuals. In the case of attributes, they can be thought of as ordered pairs of an individual and an instance of a data type. Like class instances, relation and attribute instances are part of the knowledge base of an ontology. For example, the fact that Leonardo has created Mona Lisa could be an instance of the 'is_creator_of' relation and the fact that 'Mona Lisa has_width_in_cm 53' could be an example of attribute instantiation.

Classes, relations and attributes can have *intensions* or *glosses* associated with them that explain their meaning. These can take the form of natural language definitions or more formalised definitions along the lines of formal concept analysis [48]. Either way they are generally interpreted as non-extensional definitions, i.e. the definitions should not consist of a list of all individual instances but rather of a description of characteristics that instances of a class, relation or attribute share. For example, in an art ontology the class 'painting' may have the following defining gloss associated with it:

A painting is an artefact which results from the practice of applying paint, pigment,

[4]The natural interpretation of the relation signature is that whenever we interpret a relation identifier as an n-ary relation r in the conventional mathematical sense with signature $< C_1, ..., C_n >$ then for every $< i_1, ..., i_n > \in r$ we have: i_k is a direct or indirect instance of C_k, for all $1 \leq k \leq n$. For the meaning of "direct or indirect instance" see the paragraph on "*Instances* or *individuals*".

> colour or another medium to a surface or support. The medium is commonly applied to the support with a brush but other objects may be used. Paintings may have for their support such surfaces as walls, paper, canvas, wood, glass, lacquer, clay or concrete, and may incorporate multiple other materials including sand, clay, paper, gold leaf as well as other objects.[5]

Note that the fact that the gloss contains lists of concepts does not make it an extensional definition. It would be extensional if it tried to list all instances of painting instead, which is not practicable.

A *lexicon* for an ontology and knowledge base contains signs for classes, relations, attributes and individuals. These signs can be of varying complexity, ranging from plain strings or labels associated with a language tag to complex linguistic structures. In an art ontology, for example, the labels "work of art", "artwork" and "piece of art" could all apply to the same concept 'work of art'. Similarly ("Mona Lisa", en) and ("Gioconda", it) could both be signs associated with the individual Mona Lisa.

Sometimes ontologies have *axioms* associated with them. They express general statements about the domain in some language of formal logic. For example, the following statement in first-order logic could be part of an axiom system for an art ontology:

$$\forall x\Big(self\_portrait(x) \quad \leftrightarrow \quad work\_of\_art(x) \wedge \exists y\big(depicts(x,y) \ \wedge \ is\_creator\_of(y,x)\big)\Big)$$

Thus, if Dürer was the creator of a painting that depicts himself, the painting could be classified as a self-portrait.

In order to apply ontologies to practical problems the elements which we characterised in this section need to be implemented in concrete knowledge representation formalisms or ontology languages (see Section 3.1.2).

2.1.3 Thesauri and Ontologies

Not all of the elements we discussed need to be present in every ontology. Sometimes, some of these elements may be missing or may only be present in a trivial sense. Lassila and McGuinness present an informal ten point scale of increasing ontology complexity [95]. We briefly characterise their ten point spectrum in terms of the terminology we introduced in the previous section (see Figure 2.1).

The weakest form of ontology they describe is a *controlled vocabulary*, i.e. a finite list of terms that unambiguously denote a concept, but that has otherwise no additional structure. Controlled vocabularies are equivalent to an ontology in which the concepts form a flat list (or a flat list beneath a single top element) and have one and only one lexical term associated each. Relations, attributes, instances, intensions and axioms would be absent from such an ontology.

A richer form of ontology are *glossaries* which assign a natural language gloss or definition to each term that clarifies the meaning of the term for a human reader.

Additionally to the information contained in controlled vocabularies and glossaries, *thesauri* provide synonyms and sometimes a hierarchical structure among their concepts. This structure however does not necessarily represent a strict subclass hierarchy but might instead be interpreted in the sense of 'broader concept/narrower concept'. Synonyms can be modelled in the form of several terms in the lexicon that are attached to the same concept. A non-strict class hierarchy differs from the class hierarchy we discussed in the preceeding section in its interpretation of instantiation. If an individual

[5]This is a modified version of the first paragraph of the article on "Painting" on the English language Wikipedia, http://en.wikipedia.org/wiki/Painting (last accessed 19/09/2011).

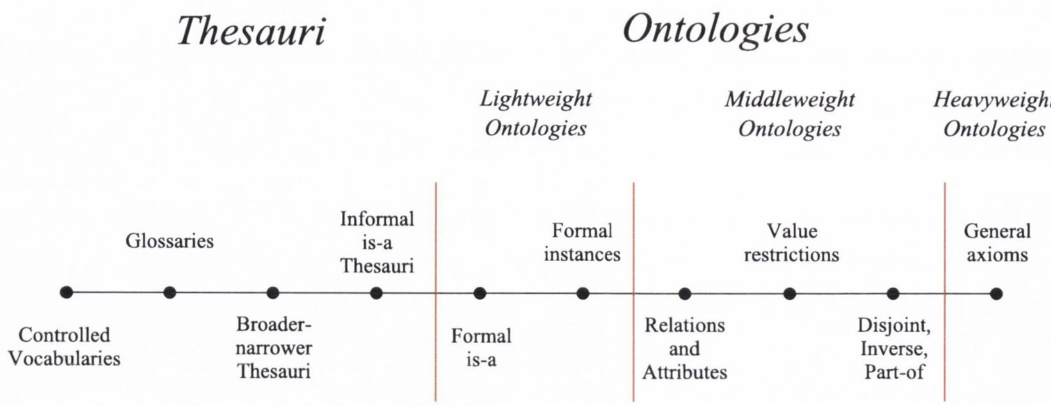

Figure 2.1 – A spectrum of semantic encodings based on a diagram by Lassila and McGuiness [95]. A similar spectrum of ontologies was published in [154] and a modification thereof in [60]. The terminology that applies to different sections is indicated above the spectrum.

instantiates a class directly then it instantiates all superclasses of this class indirectly. This does not hold for non-strict hierarchies. Note, however, that thesauri, as understood by Lassila and McGuinness, do not usually contain explicit instances. The authors further distinguish between thesauri which do either not contain a concept hierarchy at all or one that can only be interpreted as a loose 'broader/-narrower' association, and those thesauri with a class hierarchy that in all but very few cases could be interpreted as formal subsumption. The difference is quantitative in nature rather than qualitative and the authors justify separating the two categories pragmatically, as the latter applies to "many of the 'naturally occurring ontologies' on the web" [95]. They refer to this kind of hierarchy as an "informal is-a" hierarchy.

A strict class hierarchy in the sense of indirect individual instantiation without exceptions is what makes the difference between a thesaurus with informal is-a and an *ontology with formal is-a*. This corresponds to the ontology model we have outlined above except that no relations, attributes, instances or axioms are present.

An ontology in which formal instances or individuals are explicitly present forms the next step on the scale. The step after that is made up of ontologies which contain relations and/or attributes, but no non-trivial relation signatures. That means that the relation signature of an n-ary relation in the ontology is always $\top^n$ where $\top$ is the top element in the class hierarchy. This amounts to no restriction at all. Similarly, all attributes in the ontology would have $\top$ as a domain and their range would have no data type restriction associated with it.[6] The third last category of ontologies are those whose relations may have a non-trivial signature and whose attributes may have a domain and range restriction associated with them.

Axioms are the only element of our ontology model that none of the preceding types of ontology have exhibited. Among the ontologies that contain axioms Lassila and McGuinness informally distinguish two categories. One which only contains disjointness, inverse and part-of statements and one without any restrictions whatsoever. Disjointness of classes and the definition of inverse roles or

[6]Because the account by Lassila and McGuinness is given in frame-based language the authors do not distinguish between what we have called relations on the one hand and attributes on the other. In their set-up both relations and attributes are 'slots'. Some of the reconstructions of their ontology categories in our semi-formal model are therefore somewhat artificial.

properties are common mechanisms in many ontology formalisms. It is not clear what Lassila and McGuinness mean by "part-of", but presumably a weak axiomatisation of a mereological relation, including transitivity and possibly reflexivity (for possible axiomatisations of such a relation in the Web Ontology Language OWL see Section 3.3). The final category comprises ontologies with unrestricted, general axioms. Their expressivity is only limited by the logic that is used to formulate the ontology and its axioms. For some ontology modelling languages this logic can be very expressive. The CycL ontology language for example goes beyond first order logic in its expressivity.[7] These general axioms can also be expressed in a rule language or a rule system.

The presence of a strict class hierarchy is mentioned in [95] as a necessary criterion for a knowledge representation artefact to qualify as an ontology proper. This places the divide between ontologies in a wider sense (as we characterised them following Cimiano [22] on pp. 13-15 above) and ontologies in a narrower sense between the forth and fifth point on the scale (see Figure 2.1). Among explicit encodings of semantics as they are represented in this scale a distinction is sometimes made between *lightweight*, *middleweight* and *heavyweight semantics*.[8] Among ontologies proper some authors distinguish between *lightweight* and *heavyweight ontologies* [46] and the term *middleweight ontology* has also been suggested as an intermediate category. According to this terminology lightweight ontologies only contain concept lattices with formal is-a links, middleweight ontologies contain relations and basic restrictions (such as are expressible in OWL) and heavyweight ontologies contain arbitrary rule based axioms.[9]

Unless otherwise specified, we shall from now on use the term 'ontology' to denote what we called "ontologies in a narrower sense" above and what is labelled "Ontologies" in Figure 2.1. We will call the logically weaker forms of semantic encodings, that do not contain a strict subclass hierarchy, 'thesauri' (in the wider sense).

Apart from their degree of axiomatisation, ontologies have been classified according to other criteria. Six characteristic dimensions of an ontology project have been identified [70]:

(i) expressiveness (corresponds to the spectrum discussed above)

(ii) size of the relevant community

(iii) conceptual dynamics in the domain

(iv) number of conceptual elements in the domain

(v) degree of subjectivity in a conceptualisation of the domain

(vi) average size of the specification per element

Another dimension which is sometimes considered, is the generality of an ontology. Typical distinctions made here, in order of decreasing generality, are: general or common ontologies, top-level or upper-level ontologies, domain ontologies and application ontologies [50, pp. 29-34]. We will discuss an example of a top-level or upper-level ontology in Section 2.3.3.

[7] `http://www.cyc.com/cycdoc/ref/cycl-syntax.html` (last accessed 19/09/2011).

[8] Compare John Sowa's suggestion for a gradation of explicit semantics and the ensuing discussion on the ontolog forum, `http://ontolog.cim3.net/forum/ontolog-forum/2010-07/msg00000.html` (last accessed 15/11/2010).

[9] Personal communication with Frédéric Fürst, first author of [46] (4/10/2010 – 19/10/2010).

2.1.4 Discussion: The Contribution of Philosophy

While the connection between philosophical ontology and ontology in computer and information science is often mentioned, as far as the historical roots of the concept are concerned, there are widely differing opinions as to how big the actual contribution is that traditional philosophy can make to computational ontology engineering. These opinions range from a strong emphasis on a pragmatic approach which focuses on concrete implementations and problem solving [67][10] to a view which effectively construes knowledge representation and ontological engineering as a form of applied philosophy [50], [141]. Intermediate positions are also found. It has, for example, been argued that while logic, semantics, and philosophy of language offer a valuable contribution to ontology construction, an explicit endorsement of a metaphysical ideology (like realism or conceptualism) offers no methodological advantage in this respect [111]. Similarly, the conceptual analysis and techniques used by philosophers have been proposed as the basis for a methodology to build and validate information systems ontologies, independently of the "deep problems of existence, such as God, life and death" [62, p. 61] to which they were traditionally applied in philosophy (cf. also [61]).

2.2 The Informal "Ontology" of Visual Art

Concepts and relationships pertaining to the field of visual art, their meaning and mutual interaction, have been discussed from various points of view. Before we turn to resources that have attempted to capture conceptualisations of visual art with some degree of formal structure, we outline relevant scholarly discourses in philosophy, art history and information science and also the perspective of the retriever of visual art imagery.

2.2.1 The Philosophical Perspective

Traditional philosophical ontology (cf. 2.1.1) has concerned itself with the question of ontology of art. However, a look at recent introductions to this field suggests that there is little agreement among philosophical aestheticists and their competing designs, with the exception that 'ontology of art' should be interpreted as 'ontology of works of art' [128], [149]. Yet, even this has been disputed by some. Thomasson characterises the philosophical discussion as follows: "determining the ontological status of works of art is extremely difficult [...]. Indeed works of art [...] have been placed in just about every major ontological category — including those of mental entities, imaginary objects or activities, physical objects, and abstract kinds of various sorts." [149, p. 79]

It appears that many of the ontological models of art proposed by philosophers are somewhat removed from a common sense understanding of the topic. One contributor to the debate for instance admits that his view "entails that works of music cannot be heard, nor paintings seen" [149, p. 81].

Philosophical studies of visual art or two-dimensional visual art have addressed how these specific art forms can be defined, how our perception of such art can be characterised and what constitutes the specific aesthetic value of such art [100]. These deliberations "distinguish a painting's 'plastic form' from its 'illustrative content'" [100, p. 629] and "a painting's subject [...] from its content" [100, p. 627].

Perhaps the most important lesson that can be drawn from the philosophical investigation of art

[10] See also http://www-ksl.stanford.edu/kst/what-is-an-ontology.html (last accessed 28/09/2010).

is a heightened awareness of the distinction between the intellectual and creative achievement of the artist on the one hand, and the physical embodiments of the created work of art on the other. This has given rise to a distinction, in the philosophical debate, between forms of art that have only one concrete embodiment (painting, non-cast sculpture) and other forms that can have several equally valid embodiments (print, cast sculpture, literature). An additional degree of complexity is added in cases where a performance aspect is involved (drama, music).

Thomasson argues that philosophers should not disregard common sense conceptions of art and that the "beliefs — at least of those in the art world" [149, p. 87] should be taken seriously when constructing an ontology of art [ibid., Section 4.2]. It is to these beliefs, in particular the ones held by art historians, that we turn next.

2.2.2 The Art Historical Perspective

In his work "The Story of Art", Ernst Hans Josef Gombrich, one of the eminent art historians of the 20th century, does not attempt to give a definition of art: "There really is no such thing as Art. There are only artists." [51, p. 15]. He points out, that 'Art' has meant different things at different times and in different countries and that an objectivisation of the concept has done more damage to artists than good.

A number of authors have likened visual art to language [9], [30], [36]. Sometimes this analogy is taken further than a discussion of references or (hidden) symbolism in art. Berger, for example, considers it possible "to translate the language of oil painting into publicity [images]" [9, p. 140] and speaks of the "pictorial grammar of [an artistic] tradition" [10, p. 75]. According to an association dedicated to the study of visual literacy [36], however, agreement on a common definition of 'visual literacy' has yet to be achieved.[11] It remains unclear whether the correspondence between visual art and language carries beyond the level of metaphor. Approaches to automatic image annotation, that have employed machine translation techniques to 'translate' image features into keywords are perhaps of interest in this context [37], [90].

Erwin Panofsky has given an account of the successively more complex levels of image interpretation and image understanding in the visual arts as early as 1939 [121].[12] Panofsky sketches his iconographical method, "which concerns itself with the subject matter or meaning of works of art as opposed to their form" [ibid., p. 3]. By *form* Panofsky means the physical appearance, "the general pattern of colour, lines and volumes" [ibid.]. *Subject matter* or *meaning* is divided into three subcategories: *natural meaning*, *conventional meaning* and *intrinsic meaning*. An example for natural meaning is the recognition of objects or events depicted in a work of art such as human beings, plants, buildings, tools etc. Also the "mournful character of a pose or gesture, or the homelike and peaceful atmosphere of an interior" [ibid., p. 5] fall into this category. Natural meaning, according to Panofsky, is largely accessible by our practical experience as human beings and independent of our familiarity with cultural traditions. Conventional meaning, is determined by the cultural tradition the work of art is part of. An example for conventional meaning given by Panofsky is the identification of a male figure with a knife as St. Bartholomew. Understanding conventional meaning "presupposes a familiarity

[11]Compare the homepage of the *International Visual Literacy Association*: "each visual literacist has produced his/her own [definition of visual literacy]" (`http://www.ivla.org/org_what_vis_lit.htm#definition`, last accessed 20/09/2011).

[12]According to [73], Panofsky developed his theory "as early as 1962" [73, p. 602]. Originally however, Panofsky's book *Studies in Iconology* was published in 1939 by Oxford University Press and then reissued with minor modifications in 1962 and 1972.

with specific *themes* or *concepts* as transmitted through literary sources" [ibid., p. 11, emphasis in the original]. Intrinsic meaning "is apprehended by ascertaining those underlying principles which reveal the basic attitude of a nation, a period, a class, a religious or philosophical persuasion [...] condensed into one work." [ibid., p. 7]. According to Panofsky it can be grasped in virtue of what he terms "*synthetic intuition*" [ibid., p. 15]. However, he does not give clear-cut examples for the intrinsic meaning of a work of art.[13] For Panofsky, synthetic intuition, the primary source of discovering and interpreting this kind of meaning, is to some degree "subjective and irrational" [ibid.].

2.2.3 The Information Science Perspective

Information science, as an interdisciplinary field studying the organisation and management of information, has concerned itself with visual information contained in images. In search of a structured model of image content, some information scientists, including librarians and computer scientists, have taken up Panofsky's three-layered approach. Layne proposes three levels for the categorisation of image subjects, *Generic Of, Specific Of* and *About*, which are based on Panofsky's three distinctions of meaning. To these she adds four facets, *who, what, where* and *when* which apply equally to each level. The resulting 3×4 matrix for the classification of image descriptions is known as the Panofsky/Shatford model [134].

Literature from at least four different fields — art, cognitive psychology, information sciences and content-based retrieval — is used to create a framework for image indexing by Jaimes and Chang [89]. A global distinction between *visual content* and *non-visual content* is introduced. Non-visual content is information about an image, which cannot be directly seen in the image.[14] It is broken down into three categories: *Physical attributes* (e.g. location of the image, owner, size of the image, compression of a digital image), *biographical information* (e.g. author, date, title, material, technique)[15] and *associated information* (e.g. a caption, an article or a sound recording associated with the image). Visual content is divided into ten levels. Level 1 addresses the *type or technique* of an image (e.g. painting, black and white photograph, digital colour image). Levels 2 through 4 cover *global distribution* of colour or texture (as expressed in a histogram for example), *local structure* (basic shape or colour of certain regions in an image) and *global composition* (spatial arrangement of local structures) respectively. Together the first four levels refer to the purely perceptual properties or — in the terminology of the authors — the *syntax* of an image. It would seem that these levels correspond to Panofsky's *form*. Levels 5 through 10 in turn address the *semantics* of an image. They are derived from applying the generic of, specific of and about levels of the Panofsky/Shatford model to a) individual objects in the image or b) the depicted scene as a whole. Thus the six levels are: *generic objects, generic scene, specific objects, specific scene, abstract objects, abstract scene.*

Burford et al. suggest a taxonomy for image content, which comprises nine levels [17]. *Perceptual primitives* cover colour and simple texture descriptions, *geometric primitives* two- and three-dimensional non-representational forms. *Visual extension* refers to an interpretation of a scene from

[13] What comes closest to an example are two statements about the art of Michelangelo and Leonardo da Vinci. Panofsky says that trying to understand a fresco by Leonardo as a document of his "personality, or of the civilisation of the Italian High Renaissance" [121, p. 8] instead of just in terms of the objects or biblical stories depicted, was dealing with it at the level of intrinsic meaning. However, he does not illuminate what the meaning would be in this case. The same holds for the statement about Michelangelo.

[14] The term 'non-visual content' may perhaps be a misnomer as it does not refer to anything that an image contains, but rather to information otherwise associated with it.

[15] As the authors point out, there is an overlap between non-visual information and the type/technique level of visual content. Moreover, it would seem that information concerning the material of an image should be part of the physical attributes.

visual cues which requires some inference (e.g. depth, occlusion, perspective). *Semantic units* are general or specific names that refer to depicted content. Interpretations of a scene that rely on 'environmental knowledge' belong to the level of *contextual abstraction*. Such interpretations and associations that require culture specific knowledge form the level of *cultural abstraction*. *Technical abstraction* refers to associations that depend on detailed specialist knowledge. Emotional and affective associations, subjective to some extent, form the layer of *emotional abstraction*. The final level, *metadata*, does not refer to image content, but to other information about the image (e.g. image format, aspect ratio). Burford et al. point out that descriptors and search terms for images can be ambiguous between categories or even belong to multiple categories at the same time.

Burford and colleagues' model bears similarities to both Panofsky's distinctions and the ten levels of Jaimes and Chang. Perceptual and geometric primitives correspond to Panofsky's 'form'. Visual extension, contextual abstraction and the general names among the semantic units belong to Panofsky's 'natural meaning'. Cultural abstraction, technical abstraction and perhaps specific names represent Panofsky's 'conventional meaning'. Emotional abstraction, however, seems to cover only part of what Panofsky refers to as 'intrinsic meaning'. The metadata level has no equivalent in Panofsky, but corresponds to Jaimes and Chang's 'non-visual content'. The visual content of an image, however, is structured slightly differently than in Jaimes and Chang's proposal. In particular, the levels of Burford et al. make no systematic distinction between depicted objects and depicted scene, with the exception perhaps of the 'visual extension' level, which appears to refer to the scene.

For an analysis of descriptors of personal photographs, Lee and Neal recently combined the classification models of Burford et al. [17] and Jaimes and Chang [89] into a model for semantic photograph description [96].

An overview of the literature on image content models is the starting point of a synthesis of such models into a classification framework of image descriptions [73]. The framework that Hollink et al. propose synthesises a number of classifications, specifically a distinction between primitive, logical and abstract image features [38], the ten-levels of Jaimes and Chang and the Panofsky/Shatford model (see above). The resulting model maintains a degree of 'backward' compatibility with the original systems.

The proposed framework has three top levels: the *non-visual level*, the *perceptual level* and the *conceptual level*. Each of these levels corresponds to a general area of information about an image that can be expressed in image descriptions. Non-visual information about an image is defined by two characteristics. It is objective (i.e. independent of interpretation), and it cannot be directly derived from the image. Examples of non-visual information are *creator, creation date, culture, location, material, relation to other works, style/period* etc. Non-visual information is often referred to as *metadata*. Perceptual information about an image is "directly derived from the visual characteristics of the image" [73, p. 608]. It corresponds to what Panofsky called *form*. Examples of perceptual information about an image include descriptions of the colour, shape and texture of the image as a whole or certain elements (parts) of the image, as well as the absolute and relative spatial position of image elements to each other. Finally, conceptual information about an image addresses "the semantic content of the image" [73, p. 609] or, in Panofsky's terms, the subject matter or meaning. World knowledge is needed in order to arrive at conceptual information. Panofsky's classification of *meaning types* forms the basis of Hollink et al.'s threefold division of conceptual information into *general concepts, specific concepts* and *abstract concepts*. For the identification of general concepts in an image "only everyday knowledge of the world" [73, p. 609] or, as Panofsky called it, "practical experience" [121, p. 15] is required. Examples for general concepts are ape, banana, car, woman etc. Specific concepts require "domain-specific knowledge" [73, p. 610]. Examples for specific concepts are King Kong or Mercedes-Benz 300 SL. Note that specific concepts do not necessarily refer to unique individuals. They may still represent a class of individuals albeit one that often requires a certain familiarity with the domain. Finally, "knowledge used at [the level of abstract concepts] is interpretive and subjective" [ibid.]. Describing "the content of [an] image as a species threatened with extinction" [ibid.] is given as an example for abstract level information.

Even though the models of Hollink et al. [73] and Burford et al. [17] were developed independently, they show similarities as they both draw on similar literature. It seems that the categories of [17] can mostly be mapped into the model outlined in [73] and vice versa. An exception to this are perhaps Burford and colleagues' categories of 'cultural abstraction' and 'technical abstraction' which may present a refinement of the 'specific' level of Hollink et al.

In terms of differences between the models, Burford et al. draw a clearer distinction between the issue of generic and specific appellations on the one hand and the context dependence of descriptors on the other (i.e. whether general world knowledge or culture-specific knowledge is required for the interpretation of a descriptor). These two aspects are somewhat confounded in the model of Hollink et al. In contrast, Hollink et al. introduce specific subclassifications and tie their model together with a detailed schema like VRA Core.[16] This, perhaps, produces a less theoretical, more operational framework.

The classification frameworks proposed by information scientists are general and application independent. In part, however, these models are based on studies which have attempted to illuminate the needs and expectations of the users of particular image retrieval systems or the approaches that retrievers have adopted to solve a particular image retrieval task.

2.2.4 The User Perspective

It has been said that the client queries submitted to a particular photographic archive in the 1980's often "fell into a 'no-man's land of categories'" [39, p. 534]. Since then several studies on image queries and image descriptions have tried to chart this *no-man's land*, often as part of the iterative development cycle of image retrieval systems.

Jörgensen studied image descriptions supplied by 107 postgraduate students in library science for a set of six images [91]. Participants were given three different tasks: simple image description, image retrieval and description from memory. Using grounded theory [49], Jörgensen develops twelve 'attribute classes' from the descriptions and concludes that *objects*, *people*, *color*, *location* and narrative elements (*content/story*) are the most frequently used classes. No significant differences between the three tasks are reported.

Fidel analysed 100 image requests using a modified version of Jörgensen's attribute classes and describes image retrieval tasks as being positioned on a continuum between two extreme poles, the *Data Pole* and the *Objects Pole*: "At the Data Pole, images are used as sources of information, while at the Objects Pole, images are needed as objects" [44, p. 189]. An example of retrieval tasks near the Data Pole are maps of a certain geographical area or medical slides of certain aspects of human anatomy. Examples of retrieval at the Objects Pole are requests for "a very specific kind of image of a person or event, or for any image that represents a specific idea or object" [ibid.]. Intermediate forms

[16] `http://www.vraweb.org/projects/vracore3/index.html` (last accessed 26/09/2011)

of retrieval between the two poles may occur and Fidel suggests that art historians are a user group for which this is typically the case.

Smeulders et al. classify user aims in image retrieval into three broad categories: *target search*, *category search* and *search by association*. In target search the user wants to retrieve "a precise copy of the image in mind" [139, p. 1351], in category search the user is looking for an arbitrary image from a certain class of acceptable images and in search by association the user starts out with no specific aim in mind.

While some studies hypothesise about the role that visual arts and art historical terminology may play in image retrieval [44], [91], other researchers have targeted this domain empirically. Hastings studied intellectual access to art images in a computer-based setup [65]. She categorises queries into four different levels of complexity, ranging from queries for place, time and artist name to queries for image subject, meaning and 'why-queries'. It is concluded that some of the simplest queries can be answered without images, while some of the most complex require access to full-text secondary subject resources. Comparing these findings to a survey of web-based access to paintings, Hastings emphasises the need for manipulating (e.g. grouping and enlarging) high-resolution images for certain query types [66].

Chen works with a sample of art history students to study image queries and image retrieval [19], [20]. In his analysis he employs the frameworks developed by Jörgensen, Fidel and others. He reports that participants mainly used Enser and McGregor's unique category (i.e. they queried for specific, usually named objects or persons) without refining terms (cf. [40]). Like in Jörgensen's original study, the attribute classes of *location*, *object* and *people* were frequently used by the participants. Differences were found in the categories of *content/story*, which was of no significance in Chen's study and *art historical information*, which was used nearly twice as much [19]. Chen reports a positive correlation between participant's experience with Online Public Access Catalogues (OPACs) and self-reported levels of success in an image search task [20].

In a separate study, Chen has surveyed the image-using behaviours of museum professionals from six different museums in the United States [21]. Participants were working in curatorial, IT, archival and a number of other roles. The study attempts to put image retrieval and image needs into a wider context, including technological as well as social aspects. The findings illustrate a wide variety of different purposes in image retrieval between professional groups and sometimes among the same group and point to redundancies in the management of pictorial information in museums. Most of the participating museums were relying on in-house cataloguing and indexing rules, only two used standardised classification schemes.

Most image retrieval studies vary in their experimental set-up, for instance in the task set to participants, and also in the system they employ for classifying image descriptors. In cases where a limited comparison between different studies is possible some of the results in the literature seem to be contradictory (cf. [73, p. 624f.] and [39, p. 534]). In our opinion little research has been done which allows for a direct comparison of aesthetic image conceptualisations by persons with different levels of domain expertise. Hastings uses a sample containing students of library and information science and students of visual arts, but does not comment on differences between the two groups [66]. The participants of Chen's study were recruited partly from an introductory and partly from an advanced course in medieval art, but no systematic discussion of differences is provided beyond a comparison of the average number of search keywords or phrases submitted [19]. We have addressed this issue in a comparative survey involving art specialists and laypersons (Section 4.2).

2.3 Semantic Resources for Art Annotation

We will now turn to a review of explicit semantic resources and ontologies, that can inform the construction of an ontology about fine art images or can be used directly in such a construction. The word "explicit" is meant to emphasise that these resources display at least a minimum of formal structure, which allows them to be located at some point on the spectrum of semantic encodings that we have discussed before (Figure 2.1). This is in contrast to the preceding section, which covered informal or implicit conceptualisations of the domain of fine art.

We select six structured resources that cover the domain of cultural heritage and visual arts for closer inspection. Each of these has a different emphasis. The Getty vocabularies, four in total, are thesauri that were designed to cover the domain of artefact metadata. By contrast, the Iconclass hierarchy was developed for a description of the subject matter of an aesthetic image, not its metadata. Lastly, the Cidoc Conceptual Reference Model (CRM) is an upper ontology for the domain of cultural heritage information. Finally, we will also discuss GeoNames, a general purpose resource for structured geographical knowledge.

2.3.1 The Getty Vocabularies

The Getty Research Institute, a part of the J. Paul Getty Trust, that is "dedicated to furthering knowledge and advancing understanding of the visual arts",[17] maintains a program for creation and maintenance of structured vocabularies for categorizing, cataloguing and retrieving information about works of art.

To date, the Getty vocabulary program has produced four component vocabularies, that "contain structured terminology for art, architecture, decorative arts and other material culture, archival materials, visual surrogates, and bibliographic materials".[18] Apart from cataloguing and retrieval, the vocabularies are intended to serve as general research tools. The four component vocabularies are:

(i) the Art and Architecture Thesaurus (AAT)

(ii) the Union List of Artist Names (ULAN)

(iii) the Thesaurus of Geographic Names (TGN)

(iv) the Cultural Objects Name Authority (CONA)

All four resources are structured as thesauri "in compliance with ISO and NISO standards". The basic unit of information in each case is a 'record', mostly representing either a concept or an individual object, depending on the vocabulary. Records can have multiple names or terms attached to them, which represent synonymous expressions or names in different languages that point to the same underlying concept or individual. The relationship between variant names attached to the same record is referred to as "equivalence relationship". Relationships between records include a hierarchical relationship ("broader and narrower contexts") and a number of "associative relationships" (a collective term that covers "cross-references, relationships, or links between terms that are not hierarchical or equivalent"). The associative relationship effectively has different subtypes in different Getty vocabularies.

[17] `http://www.getty.edu/research/institute` (last accessed 30/06/2011).

[18] Any quotes or information related to the provenance, size or intended application of the Getty vocabularies in this section is taken from the Getty vocabularies homepage `http://www.getty.edu/research/tools/vocabularies/index.html` and its sub pages (last accessed 27/06/2011).

Each of the four component vocabularies covers a different aspect of heritage documentation, but they all have a focus on metadata in common. Perceptual qualities are covered to some degree in the AAT and the description of subject matter, while not central to the Getty vocabularies, is addressed in (the forthcoming) CONA.

AAT

The Art and Architecture Thesaurus (AAT) is the oldest member of the family of vocabularies and perhaps the one with the most general scope. Its development started in the late 1970s and was driven by the need for coherent and standardised terminology as computerized cataloguing of art collections began to emerge. The AAT was inspired by the Medical Subject Headings Thesaurus[19] and was first published in 1990 in both printed [123] and machine-readable format. The AAT is intended as "a general classification scheme for art and architecture" that "may be used to describe art, architecture, decorative arts, material culture, and archival materials". As such it contains a wide range of art historical and common sense concepts.

Top-level divisions (called 'facets') of the AAT are *Activities*, *Agents*, *Associated Concepts*, *Brand Names*, *Materials*, *Objects*, *Physical Attributes* and *Styles and Periods*. These top-level facets are further subdivided. The *Agents* facet, for example, contains *People*, *Organizations* and *Living Organisms*. Subcategories of the *Objects* facet, to give another example, include *Built Environment*, *Furnishings and Equipment* and *Object Groupings and Systems*, among several others.

Hierarchical relationships can have a number of possible interpretations in AAT. According to the online documentation they "link generic classes of objects, actions, or concepts to their members or species". A subconcept in the hierarchy can be a "type of, kind of, example of, or manifestation of" its superconcept. In the *Styles and Periods* facet, for example, *fin de siècle* subsumes both *belle époque* and *Decadent Movement*. In the *Activities* facet, *civil wars* are a type of *wars*, which are a type of *armed conflicts*, which in turn are a type of *events*.

For each record the AAT contains a preferred term and, possibly, "synonyms in the plural, singular, various forms of speech, variant spellings, terms in various languages, and synonyms of different etymological roots". Terms are given language designations from a controlled list of languages and several other flags, including whether the term is historical (i.e. not used any longer) or current. For each language there is only one preferred term. The overall preferred term for a record, called *descriptor*, is by default the preferred American English term, often in its plural form. Under the current editorial rules of the AAT "[a]ll terms in a record [...] are considered *equivalents* (i.e., *synonyms*)" (emphasis in the original). In particular, legacy data in which near-synonyms and sometimes hyper- or hyponyms were included in the term list of a record, is gradually being corrected.

ULAN

The Union List of Artist Names (ULAN) started out as a project to merge and coordinate different controlled vocabularies for use by J. Paul Getty Trust documentation projects. It was first published in 1994 in printed [14] and machine-readable form. The scope of ULAN is to record "proper names and associated information about artists", which are generally understood to be "creators involved in the conception or production of visual arts and architecture", but some performance artists are included

[19] `http://www.nlm.nih.gov/mesh` (last accessed 30/06/2011).

as well. While ULAN's main focus is on artists, it also contains records of persons, who are related to one or several artists, like donors, patrons and family members (spouses, siblings, etc.).

A record in ULAN either refers to an individual artist or a group of artists working together for one or more projects (these groups are referred to as *corporate bodies* in ULAN). Apart from the preferred name for an artist ("the indexing form of the name most often found in scholarly or authoritative publications"), ULAN lists variant names, pseudonyms, married names (for women) and artist names in other languages. Moreover, ULAN contains information about the gender, nationality and professional roles of an artist, as well as various relationships to other people or groups, including student-teacher relations.

For the famous Dutch painter, draftsman, and printmaker Rembrandt van Rijn, for example, ULAN records 40 different variants of his name, 25 persons that were in some way related to him (mostly his students), six different professional roles (artist, draftsman, printmaker, etcher, painter, teacher), and his nationality and gender, place of birth and death, as well as the time and place where he was active as an artist (Amsterdam, 1631-1669).

Originally, ULAN was constructed as an alphabetised list of clustered artist names. The current version contains hierarchical relationships, but its structure is much flatter than that of the AAT. Hierarchical relationships in ULAN have two different functions. They either link individual instances to their top-level 'facets', namely, *Corporate Bodies*, *Non-Artists*, *Persons*, *Artists* and *Unknown Artists*. Or else, and this is only relevant within the *Corporate Bodies* facet, they link parts of corporate bodies to their complex wholes. For example, the *Gobelins Dye Works* are part of the *Gobelins Tapestry Manufactory*, which are in turn part of the *Gobelins* factory in general.

TGN

The third vocabulary, the Thesaurus of Geographic Names (TGN), was created in response to requests for a controlled vocabulary of geographic names.[20] TGN was never published in print, instead it was made available in 1997 in machine-readable form.

It contains records on geographical entities of both political and physical nature: *Egypt (nation)* is an example of the former, *Nile River* an example of the latter kind. This is reflected by two separate hierarchies that both have *World* as its root. TGN contains records of entities that have ceased to exist, but are of historical interest: *Pharonic Egypt* is one of many such examples in the *political entities* hierarchy, the protocontinent *Pangaea* is an example of a *former physical feature*.

Records in TGN have several associated names, including vernacular and English names, variants in other languages, variant transliterations of non-Latin names, historical names, official names, abbreviations and nationally and internationally standardised letter codes for geographic places. Similar to AAT terms, TGN names have language labels associated with them and each language has at most one preferred name. The overall preferred name (*descriptor*) is by default the preferred name in the vernacular of the country or geographic region the entity belongs to, although for non-Latin alphabets a transliteration is usually chosen.[21]

[20]"The development of TGN was informed by an international study completed by the Thesaurus Artis Universalis (TAU), a working group of the Comité International d'Histoire de l'Art (CIHA), and by the consensus reached at a colloquium held in 1991, attended by the spectrum of potential users of geographic vocabulary in cataloguing and scholarship of art and architectural history and archaeology. The initial core of the TGN was compiled from thousands of geographic names in use by various Getty cataloguing and indexing projects, enlarged by information from U. S. government databases, and further enhanced by the manual entry of information from published hard-copy sources." (`http://www.getty.edu/research/tools/vocabularies/tgn/about.html`, last accessed 22/09/2011).

[21]Note, however, that in contrast to the default preferred name, the default setting for the *display* name in the online

Associative relationships between TGN records include *distinguished from* (expressly emphasising that two entities are different from each other), *ally of*, *successor of* (e.g. if an inhabited place has been rebuilt elsewhere due to natural disaster). Hierarchical relations in TGN typically represent (geographical) part/whole relationships. In particular, an entity in TGN is embedded under the lowest level entity that contains it in its entirety. This means that rivers or mountains spanning several countries appear as a direct child of the continent they are located on, not the individual countries they overlap with. Similarly, the historical region of the Abbasid Caliphate appears as an immediate child of *World*, because its extension spanned several continents.[22]

CONA

Finally, the Cultural Objects Name Authority (CONA) is the latest addition to the Getty vocabulary program. Its development started in 2004 and it is expected, that the vocabulary will be made available through a web interface in 2012. Records in CONA represent "cultural works, including architecture and *movable works* such as paintings, sculpture, prints, drawings, manuscripts, photographs, textiles, ceramics, furniture, other visual media such as frescoes and architectural sculpture, performance art, archaeological artifacts, and various [other] objects" (emphasis in the original). The focus is on "works cataloged in scholarly literature, museum collections, visual resources collections, archives, libraries, and indexing projects". As with the other three vocabularies, the initial core of entries was taken from various Getty projects.

CONA is for works of art what ULAN is for artists. It is interesting, however, that while the idea of a controlled list of artist names, perhaps inspired by artist encyclopaedias, has been pursued for well over 20 years, a similar effort for artefacts has only started recently.

We will conclude this section with a discussion of the representational depth and width, the strengths and, perhaps, weaknesses of the Getty vocabularies.

Notes on the Getty Vocabularies

An overview of the developments in the Getty vocabulary program shows, that the individual vocabularies are of considerable size and have been developed and tested for a long time before they are released to a wider public (Table 2.2). ULAN has, furthermore, experienced a considerable expansion of the number of artists it covers over the last three years. The size of the vocabularies is smaller than that of only the most comprehensive of resources. The reference work for artist biographies, for example, known as *Thieme-Becker-Vollmer*, contains over 250,000 biographical entries on artists to ULAN's 202,720.[23] Similarly, AAT and TGN have fewer entries then some general purpose resources of comparable nature.[24]

The strength of the Getty vocabularies, however, lies in their formal structure when compared to standard reference works, like *Thieme-Becker-Vollmer* and in their domain-specific design, when compared to other electronic resources. The *GeoNames* database (Section 2.3.4), for example, only has a most rudimentary coverage of historical places (especially of such places that used to exist

portal of the TGN, is *English* and not *Vernacular*.

[22]Note, that continents appear in the *political entities* hierarchy in TGN.

[23]Compare `http://rzblx10.uni-regensburg.de/dbinfo/detail.php?titel_id=8997&bib_id=gwlb` (last accessed 22/09/2011), for information about Thieme-Becker-Vollmer (German).

[24]*WordNet* [41], for example, or the *GeoNames* database (Section 2.3.4).

	Work		Size (May 2011)		Growth (2008-2011)	
Vocabulary	began	published	Records	Names	Records	Names
AAT	late '70s	1990	34,880	245,530	2.6 %	87.4 %
ULAN	1984	1994	202,720	638,900	68.9 %	118.1 %
TGN	1987	1997	992,310	1,711,110	8.8 %	55.6 %
CONA	2004	2012[25]	—	—	—	—

Table 2.2 – An overview of the development and growth of the different Getty vocabularies

but do not any more). The only entry in the GeoNames database, that represents the ancient city of *Carthage*, for example, is an entry of type 'ruin(s)', in other words, it refers to the present day remains of the city, rather than the ancient city that once was the capital of a Mediterranean empire. On the other hand, while WordNet [41] covers art historic styles and periods, one would not expect it to be as comprehensive or to carry the same authority in classifying these, than AAT, which was designed by art professionals. WordNet, for example, only lists 'High Renaissance' as part of 'Renaissance', but not 'Early Renaissance' and 'Late Renaissance', which are found in AAT.

Another strength of the Getty vocabularies is their coverage of multiple languages through the names and terms that are associated with records. All three of the already published vocabularies, show strong growth in the number of names they contain (cf. again Table 2.2). According to the online documentation this is largely due to recent efforts to add translations and synonyms in other languages. This is in line with the fact, that the scope of Getty vocabularies is only restricted by the material content that a given vocabulary is supposed to cover, but nearly unlimited in terms of geographical and temporal extension. The scope of AAT and ULAN "ranges from Antiquity to the present and [...] is global." Similarly, "the coverage of CONA is global, from prehistory through to the present" and TGN includes "current and historical places" and even a facet called *Extraterrestrial Places*.

2.3.2 Iconclass

The Iconclass classification system was explicitly constructed for the semantic description of image content [23]. It consists of a hierarchically ordered collection of 28,000 definitions and 14,000 keywords covering "objects, persons, events and abstract ideas that can serve as the subject of an image".[26] Iconclass pioneer van de Waal envisaged "a systematic overview of subjects, themes and motifs in Western art". After van de Waal's death the system was completed by a group of scholars. The results have been published as a substantial collection of 17 volumes and various electronic editions (1990-2001). Currently Iconclass is hosted by the Netherlands Institute for Art History.[27]

[25] Online browsing of CONA is scheduled to go live in early 2012, while licensable files may become available "within a few years of [2011]".

[26] Any quotes or information related to the provenance, size or application of the Iconclass classification system in this section is taken from the Iconclass homepage `http://www.iconclass.nl` and its sub pages (last accessed 29/09/2011).

[27] The complete hierarchy can be explored through the 'Iconclass 2100 browser' at `http://www.iconclass.org` (last accessed 29/09/2011).

The Iconclass hierarchy is structured by an alphanumeric code similar to the Dewey Decimal system for libraries.[28] There are 10 top-level categories in Iconclass which have a single digit number associated with them:

> ABSTRACT, NON-REPRESENTATIONAL ART (0), RELIGION AND MAGIC (1), NATURE (2), HUMAN BEING, MAN IN GENERAL (3), SOCIETY, CIVILIZATION, CULTURE (4), ABSTRACT IDEAS AND CONCEPTS (5), HISTORY (6), BIBLE (7), LITERATURE (8) and CLASSICAL MYTHOLOGY AND ANCIENT HISTORY (9).

The category "Human Being, Man in General" spawns six children, ranging from a biological description of man to the broader category of "man and environment".[29] "Society, Civilization, Culture" includes aspects like "family, descendance", "recreation, amusement" and "education, science and learning". "History" covers depictions of historical events and persons. The hierarchy under "Abstract Ideas and Concepts" comprises, among other things, "Causation", "Reasoning" and "Aesthetic Sense". The category "Abstract, Non-representational Art", which has no further subdivisions, was introduced later to the Iconclass system as a collective label for works to which traditional iconography does not apply.

If systematically used for annotation, Iconclass allows to establish connections between images of a similar subject. A painting which shows Hercules battling with the Hydra could be annotated with the Iconclass identifier "94L322 [...] the Hydra of Lerna is killed by Hercules". Another one which, in addition, depicts "Iolaus, who sears the roots of the severed heads with burning brands" could be tagged with the code "94L3221". An image of Hercules as an infant, strangling two serpents, may receive the Iconclass code "94L122". In virtue of the common prefix "94L" in their annotations, all of these depictions could then be identified as belonging to "(story of) Hercules". The killing of the Hydra that depicts Iolaus could be seen as refinement of the subject matter of an image which does not contain Iolaus, etc.

Because of the alphanumeric code structure, every category in Iconclass can only have a limited number of subcategories.[30] This appears to lead to some unnatural classifications in Iconclass. Thus, "the Twelve Labours of Hercules" are divided into a "first series" (labours 1-9) and a "second series" (labours 10-12), which subsume the individual labours. Similarly, the ten plagues of Egypt are divided into two parts.

The relationships between a category and its subcategories in Iconclass appear to be of a general associative nature. Iconclass does not distinguish between hyponyms, meronyms, instances or subcategories otherwise associated with their supercategory. As the categories are ordered into a tree structure without multiple parentage or cross-references, some semantic links between categories are not reflected in Iconclass. The entry for "Hydra of Lerna", for example, is classified among "monsters of the sea" and has no structural connection to the entry "Hydra of Lerna is killed by Hercules", other than that both labels appear under the top-level category of "Classical Mythology and Ancient History". In particular, "Hydra of Lerna" as an Iconclass subject can not be inferred from the presence of a subject tag "Hydra of Lerna is killed by Hercules".

---

[28] `http://www.oclc.org/dewey` (last accessed 29/09/2011).

[29] All quotes of Iconclass category names and hierarchical relationships are taken from the Iconclass 2100 browser at `http://www.iconclass.org` (last accessed 30/09/2011).

[30] Generally, a category can have at most nine or 26 subcategories, depending on whether a digit or a letter is expected at that point in the hierarchy. However, there are exceptions, e.g. when the systematic distinction between men and women is involved, which introduces an doubling of letters in the code.

2.3.3 CIDOC CRM

The *CIDOC Conceptual Reference Model*, sometimes also referred to as the *CIDOC object-oriented Conceptual Reference Model*,[31] or *CRM* for short, is a high-level ontology covering the domain of cultural heritage information [34]. The CRM was developed by the International Committee for Documentation (CIDOC) of the International Council of Museums (ICOM).[32] Development on the CRM as an object-oriented model began in 1996 on the basis of its predecessor, the *CIDOC Relational Data Model* [125]. The work was originally overseen by the ICOM-CIDOC Documentation Standards Working Group and in 2000 delegated to the CIDOC CRM Special Interest Group, which collaborated with the ISO working group ISO/TC46/SC4/WG9 with the aim of turning the CRM into an International Standard. In September 2006 the CRM was accepted as ISO standard 21127:2006: "Information and documentation – A reference ontology for the interchange of cultural heritage information".[33]

Successive versions of the CRM were published by CIDOC as English language documents made available through the CIDOC CRM homepage.[34] For some time version 5.0.2 of the CRM as defined in [25] was considered to be the final version of the model.[35] It was published in January 2010 and contains the definitions of 87 classes and 137 properties.[36]

The CRM, as maintained by CIDOC, does not include a machine-readable reference implementation. Other parties have, however, supplied such implementations and we shall discuss these in Section 5.3.5. Here we are concerned with the characteristics of the CRM as they are defined in [25].

The scope of the CRM is defined in two stages: the so-called intended scope, a general, high-level specification of what "the CIDOC CRM would ideally aim to cover, given sufficient time and resources";[37] and the practical scope, which is expressed through reference to the scopes of existing museum documentation standards which guided the development of the CRM.

The intended scope is "all information required for the exchange and integration of heterogeneous scientific documentation of museum collections" [25, p. ii]. "Scientific documentation" is meant to emphasise in this context, that information expressed or translated into the framework of the CRM can have the quality and level of detail expected by heritage professionals. This should not preclude the coverage of information intended for laypersons. Furthermore, "museum collections" is meant to refer to "all types of material collected and displayed by museums and related institutions [...]. This includes collections, sites and monuments relating to fields such as social history, ethnography, archaeology, fine and applied arts, natural history, history of sciences and technology." [Ibid.].

In contrast to the intended scope which is general and far-reaching, the practical scope can be

---

[31] `http://www.cidoc-crm.org/comprehensive_intro.html` (last accessed 13/04/2011).

[32] All information related to the provenance, scope, intended applications and characteristics of the Conceptual Reference Model in this section is taken from the current official CRM reference document [25] or the CIDOC CRM homepage `http://cidoc.ics.forth.gr/index.html` (last accessed 13/04/2011) unless otherwise stated.

[33] In November 2010 the status of ISO 21127:2006 was declared as "International Standard to be revised". According to the ISO webpage it is currently in the preparatory stages of revision. Compare `http://www.iso.org/iso/iso_catalogue/catalogue_ics/catalogue_detail_ics.htm?csnumber=57832` (last accessed 13/04/2011).

[34] `http://www.cidoc-crm.org/official_release_cidoc.html` (last accessed 13/04/2011). A French translation is part of ISO 21127:2006 and a number of other unofficial translations corresponding to various versions of the CRM are available as well.

[35] From January 2010 until October 2011 there were no new releases and it was stated on the CRM homepage that version 5.0.2 would be final. Since then two new versions 5.0.3 [26] and 5.0.4 [27] have been released.

[36] Note that classes E23, E43, E76 and properties P6, P18, P36, P47, P60, P61, P63, P64, P66, P77, P85 which were present in earlier versions of the CRM have been removed from later ones. Hence the overall number of classes and properties differs from the highest numbers used in class and property identifiers.

[37] `http://www.cidoc-crm.org/scope.html` (last accessed 14/04/2011).

seen as a description of what CRM de facto covers or at least aspires to cover. This may change over time and is currently defined as the union of the domains of 19 standards relevant to heritage documentation, including Dublin Core, MPEG7 and SPECTRUM [52].[38]

According to its definition, the CRM represents "a formal ontology intended to facilitate the integration, mediation and interchange of heterogeneous cultural heritage information" [25, p. i]. This intended application scenario is elaborated upon on the CRM web page: "The primary role of the CIDOC CRM is to serve as a basis for mediation of cultural heritage information and thereby provide the semantic 'glue' needed to transform [...] disparate, localised information sources into a coherent and valuable global resource".[39]

Though it was designed for the cultural heritage domain the CRM is sufficiently general to be considered as a top-level or upper ontology (compare Section 2.1.3). Even its most specific classes still represent concepts of great generality, such as *Information Carrier* or *Person*. Conceptualisations like 'work of art', 'painting', 'artist' or 'painter' are not provided by the CRM, but examples in the documentation indicate where their intended location in the class hierarchy would be. The CRM also provides a very general framework for modelling information about the temporal and spatial extension of events and objects. We will discuss this aspect of the CRM in Section 5.2.

2.3.4 GeoNames

GeoNames is a large geographical database that was compiled from approximately 50 different data-sources.[40] Sources included the National Geospatial-Intelligence Agency (U.S.), the U.S. Board on Geographic Names, the U.S. Geological Survey Geographic Names Information System, a range of geodetical and statistical services from around the world, like the Ordnance Survey (UK), and others. An additional source of input comes from users who can manually add and alter records through a wiki interface. The database is freely available under a creative commons attribution license.

Geographical 'features' (i.e. records) in the database are categorised into one of nine top-level classes and further subcategorised into one of 645 'feature codes', which describe the type of the feature in question. Examples of frequently occurring feature codes are 'populated place', 'stream', 'mountain' and 'administrative division'. The database contains over 7.5 million feature records and has over 10 million geographical names associated with these records. While general geographical information can be useful for many applications, there are a number of categories in GeoNames, which may be of particular interest in the context of cultural heritage (Table 2.3).

Relations between the entities in GeoNames can be explored by a number of web services. The 'children' service, for example, returns the administrative divisions and populated places lying within a given geographical entity. The 'hierarchy' service, on the other hand, returns the sequence of the successively more inclusive geographical units that a given geographical entity is contained within. Other services, return the closest feature or the closest populated place for a given set of coordinates or extract a short definition of a geographical entity from the corresponding English or German Wikipedia page.

In comparison with the Thesaurus of Geographic Names (see page 26) the GeoNames database

---

[38] See `http://www.cidoc-crm.org/scope.html`. For Dublin Core see `http://dublincore.org/documents/dces` and for MPEG7 `http://mpeg.chiariglione.org/standards/mpeg-7/mpeg-7.htm` (all pages last accessed 30/09/2011).

[39] `http://www.cidoc-crm.org/comprehensive_intro.html` (last accessed 30/09/2011).

[40] Any information related to the provenance, size or intended application of the GeoNames database in this section is taken from the GeoNames homepage `http://www.geonames.org` and its sub pages (last accessed 22/09/2011).

Feature Class	Feature Code	Feature Description	Entries
S.AMTH	amphitheater	an oval or circular structure with rising tiers of seats about a stage or open space	20
R.RDA	ancient road	the remains of a road used by ancient cultures	10
S.ANS	ancient site	a place where archeological remains, old structures, or cultural artifacts are located	2296
S.WALLA	ancient wall	the remains of a linear defensive stone structure	6
H.CNLA	aqueduct	a conduit used to carry water	106
S.ARCH	arch	a natural or man-made structure in the form of an arch	610
L.BTL	battlefield	a site of a land battle of historical importance	56
S.CSTL	castle	a large fortified building or set of buildings	3638
S.CH	church	a building for public Christian worship	227157
S.CVNT	convent	a building where a community of nuns lives in seclusion	70
S.HSTS	historical site	a place of historical importance	1213
S.MSTY	monastery	a building and grounds where a community of monks lives in seclusion	12567
S.MSQE	mosque	a building for public Islamic worship	4584
S.MUS	museum	a building where objects of permanent interest in one or more of the arts and sciences are preserved and exhibited	4563
S.OPRA	opera house	a theater designed chiefly for the performance of operas	37
S.PAL	palace	a large stately house, often a royal or presidential residence	250
A.PRSH	parish	an ecclesiastical district	45
S.RLG	religious site	an ancient site of significant religious importance	116
S.RUIN	ruin(s)	a destroyed or decayed structure which is no longer functional	12834
S.TMPL	temple(s)	an edifice dedicated to religious worship	4538
S.THTR	theater	a building, room, or outdoor structure for the presentation of plays, films, or other dramatic performances	250

Table 2.3 – Some 'feature codes' in the GeoNames geographical database that may be of interest for cultural heritage applications and the number of corresponding entries that are contained in the database (September 2011). Compare `http://www.geonames.org/statistics/total.html` (last accessed 22/09/2011).

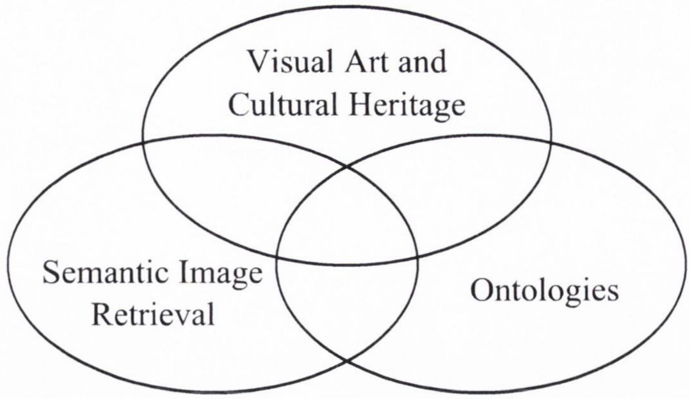

Figure 2.2 – Literature relevant to a semantic treatment of cultural heritage information from three different, but partially overlapping areas of research.

is larger, almost by an order of magnitude. This is reflected by its denser coverage, especially of smaller geographical features: when searching TGN for specific types of features, one finds 253 results for 'church', 35 for 'temple' and 16 for 'museum'.[41] The corresponding number of records in GeoNames is much larger (compare Table 2.3). On the other hand, GeoNames may be less well suited for a discussion of the historical aspect of geographical places. Some of the categories listed in Table 2.3 have an historical dimension and so do other categories, not included in the table, e.g. 'abandoned railroad' or 'destroyed populated place'. In each of these cases, the scope of GeoNames seems to dictate the documentation of the physical remains of a formerly functioning structure. This contrasts with TGN's records of historical political entities and its awareness of name changes over time. A further difference is, that TGN is curated and provides extensive references, while GeoNames can be expanded and modified by anyone. Finally, GeoNames can be downloaded for free, while TGN needs to be licensed.

2.4 Semantic Technology and the Field of Cultural Heritage

The resources we discussed in the preceding sections, give testimony to the fact that information recorded and required by heritage professionals is amenable to the structured modelling advocated by knowledge engineers. It is worth emphasising, that Iconclass and some of the Getty Vocabularies were originally published on paper and their electronic releases happened independently of recent developments in semantic technologies. What impact can these new technologies have on the field of cultural heritage?

To attempt an answer to this question, we will discuss exemplary research in a number of relevant fields. Because of the interdisciplinary nature of the research, relevant literature can be located at the intersection points of at least three thematic areas: Visual art and cultural heritage, semantic image retrieval and computational ontology (Figure 2.2).

In this section we will discuss research that is relevant for at least two out of these three areas. First we will review semantic image retrieval, especially the use of ontologies for image retrieval, but also some papers that address the semantics of fine art images without the use of ontologies. Then we will turn to more inclusive ontology-based heritage applications, in particular web portals and

[41] `http://www.getty.edu/research/tools/vocabularies/tgn/` (search carried out on 22/09/2011).

applications that store and handle background and contextual information on artefacts. Naturally, image retrieval often is an element in these systems, too. Finally, we will conclude with an outline and discussion of two conflicting positions with respect to the role that semantic technologies and ontologies should play in the heritage field.

2.4.1 Semantic Image Retrieval

In cultural heritage, in particular in the visual arts, image documentation plays an important role. Digital images of cultural heritage objects need to be retrieved from repositories and ontologies and semantic technologies may be utilised in this context.

Many authors have attempted to improve image access based on image subject or 'image content'. The term 'image content' is used in different senses by the image retrieval community. Sometimes it refers to 'semantic content', including higher level semantics or 'meaning' [17] (compare the discussion in Section 2.2.3). Often, however, and especially in connection with the phrase 'content-based image retrieval' (CBIR), it refers to lower level physical features of an image [98], [99], [139]. Ultimately, the two views represent aspects of the same problem. In each case the aim is to analyse and represent what an image shows in order to discover similar images or images that match a given query specification. The methods employed, however, and the specific problems tackled by the two communities are very different. Higher level 'semantic content' is usually addressed through terminology and text that is associated with an image, whereas CBIR relies on computer vision techniques such as colour histograms, feature vectors, segmentation and wavelet analysis.

Recent overviews of research into image retrieval have recognised the contribution that ontologies and semantic technologies can make to this field [39], [64]. We will first discuss the ontology-based approaches and then examples of computer vision based techniques as far as they have been applied to images of fine art. Finally, we will cover a number of approaches which combine CBIR with semantic technologies.

In an early example of ontology-based image retrieval, Hyvönen et al. describe the construction and use of a small, manually built ontology for the annotation and retrieval of photos of graduation ceremonies at the University of Helsinki [84]. They demonstrate that an ontology can not only be applied to annotation and retrieval, but also help a user in formulating an information need in a domain he or she might not necessarily be familiar with.

The Art and Architecture Thesaurus, WordNet [41], Iconclass and the Union List of Artist Names (see Section 2.3 above) were combined to allow for a semantic annotation of art images that covers both image content and metadata [72]. Metadata annotation follows the VRA Core Categories,[42] while image content annotations are structured according to a "sentence" template: *agent, action, object, recipient*. Users can annotate images, by filling each slot of the sentence template with an instance or class from one of the employed ontologies as a value. Semantic links between the four resources are included that allow for annotation suggestions (e.g. in the style category once the artist is known). A user interface helps in navigating ontology hierarchies and disambiguating terms. The annotation template is reused for searching the collection of annotated art images.

In an empirical study various query expansion strategies were compared with a view to improving precision and recall in image content search [74]. A number of paintings were annotated according

[42]At the time the most recent version was VRA Core 3.0 (`http://www.vraweb.org/projects/vracore3/index.html`, last accessed 26/09/2011). For the current version, VRA Core 4.0, see `http://www.vraweb.org/projects/vracore4` (last accessed 26/09/2011).

to the same guidelines. The authors used hyponym and other lexical relations from the WordNet database [41], to expand fifteen basic queries in three different ways: a conventional strategy using hyponymy, an uncontrolled combination of various lexical relationships of a certain depth and a controlled combination of such relationships which the authors proposed as an improved strategy. The three expansion strategies and the unexpanded query were evaluated with respect to precision and recall against a gold standard produced by two raters. It was found, that the proposed strategy performed better as measured by the F_1 score but not statistically significantly so.[43]

An ontology for annotating consumer photographs with respect to depicted objects, scene description, representation style and image quality was proposed and evaluated [117]. The authors found annotator agreement in the category of image quality, in particular aesthetic impression, to be low in comparison to annotations of depicted objects or scene descriptors. Constraints in the proposed ontology allow for limited inferences and consistency checking among annotations from different annotators.

Some contributions employing content-based techniques focus on aesthetic or heritage images. Low level semantics, such as the separation of paintings from photographs [28], the identification of drawing media (brush, chalk, graphite) [92], and the distinction between cityscapes and landscapes [156] were successfully addressed using computer vision approaches. Berezhnoy et al. were able to verify a long-standing claim by art historians about van Gogh's use of complementary colours with objective image analysis techniques [8]. An analysis of wavelet statistics was applied to problems of art authentication and determining the number of artists that may have contributed to a painting [104].

Higher level image semantics that are typically interesting in heritage — complex interactions of depicted people, named objects, named events or emotive content — are not accessible by current CBIR technology as pointed out by practitioners in the field (compare [99, p. 278]). This explains perhaps why heritage professionals have found CBIR solutions to be of limited usefulness [21]. The discrepancy between the physical properties of an image and the higher level semantics described by Panofsky for the case of visual art (Section 2.2.2) or the contextual information that may be of interest to image retrievers (Section 2.2.4) is referred to as the 'Semantic Gap'. In their attempt to characterise the semantic gap and review the solutions that have been made to overcome it, Hare et al. suggest to conceptually divide the gap into two smaller gaps [64]. The first one extends between low level physical image features and keywords or labels for the depicted objects. The second gap ranges from keyword annotation to 'full semantics', which involves relationships between the depicted objects, contextual information and other interpretations that a human viewer might make. To tackle the semantic gap, Hare et al. propose a strategy in which the first part of the gap is attacked 'from below' by automatic image annotation and the second part 'from above' by enriching image annotations with ontologies. They distinguish in this context between ontologies describing image content and ontologies capturing wider contextual knowledge (e.g. metadata).

Similar considerations have led some authors to a combined content-based and ontology-based approach to image retrieval. Lam and Singh combine a similarity analysis between image segments with a textual analysis of image annotations enhanced by WordNet [41] and conclude that their hybrid approach leads to more robust retrieval of semantically relevant images than either of the two approaches on its own [94].

Similarly, Wang et al. combine image analysis with textual analysis in an ontology-based ap-

[43]The F_1 score is defined as the harmonic mean of precision and recall and is intended to provide a balanced measure of retrieval accuracy [158].

proach [161]. The employed ontology is not restricted to the textual side, but links textual description to visual image features at the ontology level. In an evaluation in the domain of animal images they find that on average a text ontology-based approach outperforms keyword-based image search, and a multimodal ontology (combining text and image features) can result in a significant improvement over both keyword and text ontology search.

Vrochidis et al. have attempted to demonstrate the usefulness of a combined ontology and visual-based retrieval approach for heritage images [160]. In their hybrid retrieval architecture users can search images from a number of different heritage collections, including photographs of coins, inscriptions and paintings, either by supplying an example image or by specifying constraints on annotation ontology concepts. Images submitted for visual-based retrieval are segmented and MPEG-7 features extracted from the whole image and its segments. These features are compared against pre-computed MPEG-7 descriptors from the heritage image collections. Ontology-based search makes use of five custom-built ontologies, which represent the generic annotation structure and details specific to the individual collections. A graphical user interface facilitates ontology inspection, browsing and constraint formulation. In the hybrid retrieval mode a user can choose one of the two search options and the system will then recommend images which are similar to the result set according to the other search mode. If the user, for instance, performs a query by image example, the system will attempt to discover ontological similarities between the images that are found and in a second step supply conceptually, not necessarily visually, similar images. To make the system potentially interoperable with other ontologies in the cultural domain, the custom-built ontologies are mapped into the CIDOC CRM (see Section 2.3.3).

A combination of Semantic Web technologies and content-based retrieval methods is also an element in some cultural heritage portals (e.g. [1]), which we will address in the following section.

2.4.2 Semantic Portals for the Heritage Domain

Image retrieval is important in the heritage domain and is almost always a part of applications or prototypes in this field. Some applications, however, have a scope that goes beyond a narrow retrieval scenario and incorporates the transformation and integration of heritage related information in a wider context. An application of this type are *semantic portals* for the heritage domain. Semantic portals are a special case of information portals, which attempts to collect and publish data from distributed sources using Semantic Web standards [106], [127]. It is expected that much of the content of the Semantic Web will be made available through semantic portals.

Addis et al. present the SCULPTEUR project, a semantic portal and eLearning application that provides cultural heritage content from five major museums including the Uffizi in Florence and the National Gallery and the Victoria and Albert Museum in London [1]. Additionally to keyword search and controlled vocabulary search, both with logical operators, users are offered *search by content* and *search by concept*. Content search includes colour-based retrieval by selecting a value from the gamut and query by image example or, in the case of three-dimensional objects, query by 3D models expressed in a virtual reality modelling language. For concept-based search museum metadata is mapped into the CIDOC CRM (Section 2.3.3) and users can explore a simplified graphical visualisation of the resulting ontology. Different query components can be combined and are summarised in a 'query basket' for the convenience of the user. The SCULPTEUR project does not appear to be online anymore.[44]

[44] The location referenced in [1], `http://www.sculpteurweb.org` could not be accessed (last attempt 4/10/2011).

MuseumFinland, a semantic portal of Finnish museums, aims to achieve two central desiderata of the Semantic Web vision: intelligent applications and interoperability of distributed content [83]. Corresponding to these two goals respectively, MuseumFinland offers a semantic view-based search engine and a semantic recommender system, which attempts to discover semantic associations within the global collection data. The data was transformed from its literal form in museum databases to a semantic representation in a two step process, which involved manual disambiguation. The exploration of the data through the system is facilitated by nine facets or 'views', which are each connected to one of seven ontologies (the 'actor' and 'location' ontologies are used by two facets). When a facet value is selected, the results for each direct subcategory are computed and shown to the user. The faceted search is complemented by a traditional keyword search which can be used for both searching collection objects with matching values and searching category names for multi-facet search. Once a collection item is selected the system recommends items which are semantically similar in various ways. For example, they may have been related to similar events or been manufactured at the same time. Indirect semantic links can also be discovered through a set of rules. For example, a 'grandfather' relation can be inferred from two 'father' relations. MuseumFinland is available online.[45]

Hyvönen regards cultural heritage as a promising application domain for semantic information portals [82]. He argues that such portals are attractive to both end-users and content-providers for a number of reasons. For end-users these portals can provide "intelligent" semantic search options and a uniform view to heterogeneous content that is aggregated from distributed collections and different types of institutions. Content-providers can mutually enrich their organisations' collections and make them more attractive in their shared pursuit of promoting cultural knowledge. Once the content is aggregated it can be reused in other applications. From a pragmatic point of view, Hyvönen argues that semantic portals are more cost-effective and easier to maintain than traditional content management systems.

He identifies three major components in semantic heritage portals: a *content model*, a *content creation system* and a way to publish *semantic services*, either for humans through intelligent user interfaces or for other applications through web services [ibid.].

The creation of the content model is facilitated through the technologies in the Semantic Web layer cake applied to the heritage domain (compare Figures 1.1 and 1.2). However, there are obstacles to interoperability at the metadata level, like the harmonisation of date encodings (cf. Section 5.2.2) and at the ontology level, where syntactic interoperability between traditional thesauri can be achieved using SKOS, for example, but semantic interoperability poses a harder problem.[46] Hyvönen also discusses the use of logical rules in the heritage domain which has not yet produced accepted standards, but is nevertheless widespread in ontology-based heritage applications, for example in the form of semantic recommendations and explanations [ibid.].

Several kinds of content need to be created for semantic heritage portals. Hyvönen considers domain ontologies, knowledge bases (which he calls "instance-rich ontologies" [ibid., p. 765]) and terminology or labels. He also considers rules (in the sense of the Semantic Web layer cake, see Figure 1.1) to be a kind of content, but does not discuss the creation of rules in any detail. Automatic content creation from existing legacy data can be achieved, but requires human input in some cases: "the essential task is to move from term space into concept space by changing literal terms [...] into corresponding concept URIs [in] domain ontologies" [ibid., p. 766] (see Section 5.4).

---

[45] `http://www.museosuomi.fi` (Finnish, last accessed 4/10/2011).

[46] Hyvönen's 'ontology level' covers the levels of 'controlled vocabularies', 'taxonomies' and 'ontologies' in our 'Heritage Layer Cake' (Figure 1.2). For SKOS see `http://www.w3.org/2004/02/skos` (last accessed 3/10/2011).

Heritage portals publish information through semantic services. These can include faceted search, semantic auto completion, ontology browsing and semantic recommendations and other formats. The importance of visualisation techniques is emphasised, especially in the context of geographic (maps) and temporal (time-line) information [ibid.].

A different visualisation aspect, namely that of ontology visualisation or graphical support for formulating semantic queries is a component encountered in many semantic heritage portals. The visualisation efforts made in this connection can roughly be classed into three groups: hierarchy visualisations, graph-based visualisations and faceted search.

As a concept hierarchy is an integral part of many ontologies, it is natural to present it to the user as a visualisation aid, in the form of an indented list or as an expandable, interactive hierarchy. Hierarchy interfaces are used for viewing, browsing and specifying ontology constraints (e.g. by selecting a class) [160].

A graph structure of labelled nodes and edges corresponds well to the structure of most ontologies and is often used in generic ontology visualisation software (compare Figure 3.2 on page 49). Graph-based visualisations in semantic portals generalise hierarchy visualisations in two ways. They can represent relations other than the subsumption relation more easily and they expand two-dimensionally and thus present more information on a single screen [1].

Faceted search can take various visual forms which may involve hierarchies, tabs or windows presented next to each other [83], [107]. The distinguishing characteristic of faceted search is that, as the user navigates along various facets and selects values, he or she is only presented with those values or subfacets which contain items matching the current selection. In this way only available information is browsable and a situation where no hits are returned to a query cannot arise.

In Chapter 5 we will propose a graphical user interface for formulating queries against an ontology and knowledge base that combines a graph structure with the visual appearance of an advanced search interface.

2.4.3 Discussion: Two Approaches to a Semantic Web of Culture

More than half a decade has passed since the first call for a *Semantic Web for Culture* went out, sketching in broad strokes a vision of cultural knowledge representation that builds on philosophical principles and is able to chart 'meaning' as it changes over time [159]. Since then, various applications and implementations have made inroads towards a semantic web of culture, as we have seen. We will briefly discuss two of the more recent contributions that have attempted to answer, in general, what such a semantic web of culture may look like and how it should be pursued: Schreiber [131] and Doerr [35].

At the beginning of both outlines stands the observation that cultural heritage is an enormously diverse phenomenon and that integration of the equally diverse sources of knowledge and documentation practices is the primary challenge. The conclusions and recommendations that are draw from this shared diagnosis, however, differ.

Schreiber puts his treatment of digital heritage in the context of an emergent 'Web Science' [131]. In the future, such a science may study the impact that web technology has on human behaviour and society as a whole [136]. He suggests that web-based heritage experience should be seen as complementary, not necessarily as secondary, to real museum visits. He also advocates open access to heritage content and new license schemes in this context. In terms of the knowledge engineering aspect of digital heritage collections, he argues for a number of key points: because of the great variety of documentation practices, knowledge representation schemas for heritage should be based on the principle of minimal ontological commitment (i.e. minimal semantic constraints) and only partial alignments between knowledge organisation systems should be attempted, as unification, according to Schreiber, is unfeasible. Furthermore, he proposes additional research into knowledge extraction from natural language, as much of heritage related information is recorded in textual form. Lastly, he advocates alternative approaches to reasoning, because of the enormous size of heritage data on the web: "Given such amounts of data, the traditional notions of correctness and completeness make no sense." [131, p. 109]. It is in this context that graph based query expansion strategies such as the one presented in [74] come into play.

Doerr, like Schreiber, recommends that semantic technology, in particular ontologies, should be employed to integrate the great diversity of cultural heritage information, both with respect to different subdomains (e.g. history of arts, modern arts, archeology etc.) and with respect to different functions of information keeping in cultural heritage (e.g. collection management, presentation, description, interpretation) [35]. He criticises that most integration efforts "concentrate on finding minimal common description elements [...] as *finding aids* rather than documentation" [35, p. 467]. This, according to Doerr, fails to take the wider historical context of heritage objects, from which they derive much of their relevance, properly into account. As an alternative to this approach, he outlines the design principles of the CIDOC CRM (Section 2.3.3): the practice of assigning identifiers to heritage objects should be described by an ontology as part of the historical reality of the documentation process and not be resolved outside the ontology; similarly, terminological categories (i.e. types, classes) and their creation should be described together with the documented objects; these categories and other immaterial objects (e.g. messages, ideas of historical importance) should be explicitly modelled; lastly, and perhaps most importantly, historical information should not be centred on objects or persons, but on discretely modelled events. Thus designed, the CIDOC CRM, according to Doerr, serves a dual purpose: as a guide for good conceptual modelling practice in the heritage sector and as a global model for information integration and data migration. While current ontology languages (see Section 3.1.2) do not yet support the full range of constructs that the CIDOC CRM requires, especially not for the modelling of types, Doerr points out that faceted classification, which has emerged from within the heritage community, can be seen a precursor of description logic (see Section 3.1.1), which is still in its infancy as far as applications in cultural heritage are concerned.

Comparing the two visions of a semantic web of culture we find, that they give different, sometimes contrary answers to a number of key problems (Table 2.4). Schreiber focuses on feasibility aspects and stresses the importance of finding a small common denominator for the parties involved, because he considers agreement beyond that point impossible. His goal is a practical short to medium term solution that can handle the quantities of data typically found on the web. Doerr, on the other hand, advocates a carefully considered and negotiated modelling of curated knowledge, that tries to uncover conceptualisations that are shared across the heritage sector, even though conflicting documentation practices may sometimes conceal this consensus. It is of interest in this context that the CIDOC CRM was developed bottom up by reengineering and integrating many existing documentation schemas and has over the course of its development acquired considerable generality. Doerr emphasises that a sound solution that respects the practices of heritage professionals, while at the same time exploiting the possibilities of computational ontology requires time: "Whoever wants to deal with the subject [of cultural heritage conceptualisations] effectively must be prepared for a long knowledge modelling phase" [35, p. 481].

	Schreiber	Doerr
ontological commitment	"minimal"	sufficient to represent historical context
information integration	partial alignments only	full alignments across institutions and disciplines
role of societal consensus	restricted to naming conventions	extends to the conceptual level
automatic reasoning	scalable alternatives to traditional logic	description logic as an extension of "faceted classification"
main focus	scalability of applications	fidelity of information

Table 2.4 – Two visions of a "Web of Culture" and their main positions ranging from theoretical to practical aspects. To an extent, the two views appear to reflect the non-realist/realist dichotomy in philosophical ontology (cf. Table 2.1).

2.5 Summary

We have surveyed literature documenting theoretical foundations and practical research efforts that bear relevance for an ontology of visual art and its application to aesthetic image retrieval. The field of computational ontology, partially rooted in a philosophical tradition, sets the scene for the deployment of a new conceptual framework to an old problem: the description, organisation and retrieval of images of fine art. However, this framework by itself would be barren if it could not draw on a source, besides the principles that govern the formal arrangement of concepts. This source is the domain expertise in the field of visual art. Philosophical, art historical and application oriented aspects of such expertise were reviewed and their utility for art and heritage ontologies was discussed. The structured resources that are available in the field of art and heritage are at the intersection of ontological modelling (or its precursors) and domain knowledge. But although they were developed with broad applicability in mind their combination and integration faces technological and conceptual obstacles which will concern us in greater detail in the following chapter. At this point we have reviewed the state of the art in applications and portals that employ such resources and sketched open problems and differences in general approach which provide the context in which our research questions are situated.

Chapter 3

Methods

In this chapter we cover the theoretical foundations and practical techniques that are relevant for our project. We start by discussing knowledge representation and manipulation technologies and their formal semantic underpinnings. We then give an overview of ontology development methods and illustrate a detailed knowledge modelling problem at theoretical level using transitive propagation as our exemplar. We then describe practical modelling challenges, especially consistency and coherence checks in relation to the integration of a top-level ontology. We discuss also the challenging domain of image subject matter and, lastly, cover the qualitative research methods that we applied in our user study and our expert interviews.

3.1 Knowledge Representation and Formal Semantics

In order to represent objects and their attributes on a machine, one requires conventions of syntax which allow for the description of concepts and relationships between such objects. If such a representation is to be used to automatically infer new facts from ones that have been explicitly represented, it has to build on a sound theoretical foundation.

3.1.1 Logical Foundation

Formal logic provides a framework for such a representation. The family of *description logics* in particular has interesting properties for representing and reasoning with terminological knowledge.

Description logics (DLs) have been studied for several decades, originally under the names of *terminological systems* and *concept languages*. In the 1980's they were introduced into the field of knowledge representation to overcome certain deficiencies associated with frame-based systems and semantic networks. Frame specifications are tailored for representing knowledge about objects or classes, but are less flexible in representing knowledge about relationships. Moreover, both frame-based systems and semantic networks lack formal semantics. In description logic, however, the basic modelling unit is the axiom which can connect concepts and/or binary relations (called *roles*). A set of description logic axioms can then be interpreted and given model theoretic semantics just as is the case in other systems of formal logic. Finally, description logics provide increased expressive power compared to standard frame-based languages.

The case for description logic as the underpinning of ontology languages has often been made in the literature (e.g. [43], [77]; for a more comprehensive discussion see [5]). We will explore the correspondence between DLs and ontology languages in the next section. In the remainder of this section we will introduce a basic description logic formally and explain the language extensions and nomenclature of more expressive DLs in an informal way.

The principle language constructs in description logics are in line with basic knowledge modelling components (concepts, roles, individuals) as we have encountered them in our description of ontology structures in Section 2.1.2. These elements can be combined using precisely defined concept constructors to form complex concepts (e.g. the union of two concepts) and role constructors to form complex roles (e.g. the inverse of a role). Concepts and roles thus constructed can then become part of concept or role inclusion axioms, which together form the TBox (terminological "box").[1] Statements asserting that an individual is an instance of a concept or a pair of individuals an instance of a role make up the ABox (assertion "box").

A prototypical description logic is the *Attributive Language with Complements* ($\mathcal{ALC}$), sometimes also expanded to *Attributive Concept Language with Complements*. $\mathcal{ALC}$ was introduced by Schmidt-Schauß and Smolka [130]. It is representative of how description logics work, as it forms the basis for a range of more expressive description logics, which have been defined and studied in the literature [79], [80], [81]. Some of these extensions serve as the logical underpinning of ontology languages (cf. Section 3.1.2). Concepts, roles and the function of so-called concept constructors in $\mathcal{ALC}$ are defined as follows:[2]

Syntax of $\mathcal{ALC}$ (Definition). Let N_C and N_R be sets of *concept names* and *role names* respectively. Then the set of $\mathcal{ALC}$ concepts is the smallest set such that

1. $\top$ (*top*), $\bot$ (*bottom*) and every concept name $A \in N_C$ is an $\mathcal{ALC}$-concept,[3]

2. if C and D are concepts and $R \in N_R$ is a role name, then $C \sqcap D$ (the *intersection of two concepts*), $C \sqcup D$ (the *union of two concepts*), $\neg C$ (the *complement of a concept*), $\forall R.C$ (the *universal restriction of a concept by a role*) and $\exists R.C$ (the *existential restriction of a concept by a role*) are $\mathcal{ALC}$-concepts.

Statements in $\mathcal{ALC}$ are interpreted either as terminological or as assertional knowledge:

TBox (Definition). A *general concept inclusion* (*GCI*) is of the form $C \sqsubseteq D$, where C and D are $\mathcal{ALC}$-concepts. A finite set of GCIs is called a *TBox*.

ABox (Definition). Let N_I be a set of *individual names* disjoint from N_C and N_R. For individual names $a, b \in N_I$, a possibly complex concept C and a role name $R \in N_R$, expressions of the form $C(a)$ (*concept assertion*), $R(a,b)$ (*role assertion*) and $a \neq b$ (*different individuals assertion*) are called *assertional axioms*. A finite set of assertional axioms is called an *ABox*.

The semantics of $\mathcal{ALC}$-concepts, GCIs and assertional axioms is defined relative to an interpretation, which assigns a formal meaning to concepts, roles, individuals and axioms. Once the domain $\Delta^{\mathcal{I}}$ of an interpretation and the meaning of every element in N_C, N_R and N_I is fixed, the meaning of complex concepts and axioms is uniquely determined. Two interpretations which have the same domain and

[1] Some authors prefer to separate role inclusion axioms from concept inclusion axioms. In their terminology "TBox" refers only to concept inclusion axioms, while role inclusion axioms are collected in the so-called "RBox".

[2] The definitions given here mostly follow the exposition by Tobies [150].

[3] Concept names are sometimes also called *atomic concepts*. The top and bottom concepts ($\top$, $\bot$) need not be introduced as primitives, they can be defined using concept negation ($\neg$) and one of the other Boolean operators on concepts ($\sqcap$, $\sqcup$).

agree on each of the atomic syntax elements, will also agree in every other aspect and hence be identical.

Semantics of $\mathcal{ALC}$ (Definition). An *interpretation* $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$ is an ordered pair, where $\Delta^{\mathcal{I}}$ is a non-empty set, called the *domain of* $\mathcal{I}$, and $\cdot^{\mathcal{I}}$ is a function, called *interpretation function*, that maps every concept name to a subset of $\Delta^{\mathcal{I}}$, every role name to subset of $\Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$ and every individual name (if there are any) to an element of $\Delta^{\mathcal{I}}$. Additionally the interpretation function $\cdot^{\mathcal{I}}$ is required to map every complex concept to a subset of $\Delta^{\mathcal{I}}$ such that

$$
\begin{aligned}
&\top^{\mathcal{I}} = \Delta^{\mathcal{I}},\\
&\bot^{\mathcal{I}} = \emptyset,\\
&(C \sqcap D)^{\mathcal{I}} = C^{\mathcal{I}} \cap D^{\mathcal{I}},\\
&(C \sqcup D)^{\mathcal{I}} = C^{\mathcal{I}} \cup D^{\mathcal{I}},\\
&(\neg C)^{\mathcal{I}} = \Delta^{\mathcal{I}} \setminus C^{\mathcal{I}},\\
&(\forall R.C)^{\mathcal{I}} = \{x \in \Delta^{\mathcal{I}} \mid \textit{for all } y \in \Delta^{\mathcal{I}},\ \textit{if } (x,y) \in R^{\mathcal{I}},\ \textit{then } y \in C^{\mathcal{I}}\},\\
&(\exists R.C)^{\mathcal{I}} = \{x \in \Delta^{\mathcal{I}} \mid \textit{there is some } y \in \Delta^{\mathcal{I}} \textit{ such that } (x,y) \in R^{\mathcal{I}} \textit{ and } y \in C^{\mathcal{I}}\},
\end{aligned}
$$

where C and D are concepts and $R \in N_R$ is a role name. Similarly, in the following definitions, C and D are concepts, R is a role name and a and b are individual names.

An interpretation $\mathcal{I}$ *satisfies a GCI* $C \sqsubseteq D$ iff $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$. An interpretation $\mathcal{I}$ *satisfies a TBox* $\mathcal{T}$ iff it satisfies every GCI in $\mathcal{T}$. In this case, $\mathcal{T}$ is called satisfiable and $\mathcal{I}$ is called a model of $\mathcal{T}$.

An interpretation $\mathcal{I}$ *satisfies an assertional axiom* $C(a)$ iff $a^{\mathcal{I}} \in C^{\mathcal{I}}$, it satisfies an assertional axiom $R(a,b)$ iff $(a^{\mathcal{I}}, b^{\mathcal{I}}) \in R^{\mathcal{I}}$, and it satisfies an assertional axiom $a \not\doteq b$ iff $a^{\mathcal{I}} \neq b^{\mathcal{I}}$. An interpretation $\mathcal{I}$ *satisfies an ABox* $\mathcal{A}$ iff it satisfies every assertional axiom in $\mathcal{A}$. In this case, $\mathcal{A}$ is called satisfiable and $\mathcal{I}$ is called a model of $\mathcal{A}$.

This concludes our formal definition of the syntax and semantics of $\mathcal{ALC}$.

Extensions to $\mathcal{ALC}$ usually introduce a new concept constructor (e.g. $\leq 1R.\top$), a new role constructor (e.g. R^{-}) or new role axioms (i.e. $R_1 \sqsubseteq R_2$).[4] The stipulations for interpretations $\mathcal{I}$ are suitably extended to cover the additional constructors and perhaps restrict the concept of satisfiability. An important extension is the inclusion of transitive roles.[5] The resulting description logic is known as $\mathcal{S}$ [150, p. 93].

Constructs of description logic that go beyond $\mathcal{ALC}$ are encoded by single letters and the name of a logic that permits one or more of these additional constructs is formed by stringing the corresponding letters together and appending it to the name of the basic logic (technically, the order of the letters does not matter, but for most description logics a canonical order was established in the relevant literature). When a construct is stronger than another one (i.e. offers genuinely greater expressivity), only the letter representing the stronger construct appears in the name. An overview of the most important constructs used to extend $\mathcal{ALC}$ and their corresponding letters is contained in Table 3.1.

[4] Note that $\mathcal{ALC}$ itself does not have any role constructors or role axioms.

[5] Formally this can be done in different ways, either by introducing a second and separate set of *transitive role names* N_{R_+}, or by introducing a predicate Tr which applied to a role name R yields a TBox axiom $Tr(R)$ to the effect that R has to be mapped by an interpretation function $\cdot^{\mathcal{I}}$ to a transitive relation.

[6] In order to ensure decidability, certain restrictions have to be observed. Roughly speaking circular dependencies in complex role inclusion axioms need to be avoided. For a technically precise account of these restrictions see [79].

[7] Certain conditions must not be violated if cardinality restrictions are not to lead to undecidability. Roughly speaking,

		Explanation	DL syntax	Remarks
	$\mathcal{S}$	$\mathcal{ALC}$ with transitive roles	—	—
Role	$\mathcal{H}$	Role hierarchy	$R_1 \sqsubseteq R_2$	—
related	$\mathcal{R}$	Limited complex role inclusions[6]	$R_1 \circ R_2 \sqsubseteq R_3$	This subsumes $\mathcal{H}$.
	$\mathcal{I}$	Inverse roles	R^-	—
	$\mathcal{F}$	Functionality	$\leq 1R.\top$	—
Concept	$\mathcal{N}$	Cardinality restrictions[7]	$\leq nR.\top, \geq nR.\top$	This subsumes $\mathcal{F}$.
related	$\mathcal{Q}$	Qualified cardinality restrictions[7]	$\leq nR.C, \geq nR.C$	This subsumes $\mathcal{N}$.
	$\mathcal{O}$	Nominals	$\{a_1, ..., a_n\}$	—
	$(\mathcal{D})$	Use of data types[8]	—	A restricted form of concrete domains.

Table 3.1 – The nomenclature and meaning of description logic constructs commonly used to extend $\mathcal{ALC}$. Compare `http://www.cs.man.ac.uk/~ezolin/dl` (last accessed 28/09/2011).

3.1.2 Ontology Representation

The implementation of our image retrieval system is built on many officially endorsed or de-facto established technology standards which will briefly be introduced and reviewed in this section. Where appropriate, the standards we used will be compared and contrasted with similar standards and their choice over alternatives will be motivated.

Ontology Representation Languages

Description logics form the logical underpinning of ontology representation languages, thus providing them with a well-defined formal semantics. In terms of their syntax, however, a number of ontology languages built on standards that were developed by the Web community for a more limited form of semantic markup. In 1999 the World Wide Web Consortium (W3C) published a recommendation for the *Resource Description Framework* (*RDF*).[9] RDF was intended as a standard to record metadata about Web (or other) resources in a uniform way that would enable applications to process and exchange "machine-understandable" information about such resources. RDF descriptions consist of sets of object-attribute-value triples, that are called *statements*. Both objects and attributes, the latter are usually called *properties*, are resources; values can be resources or primitive data types. Structurally the RDF model closely resembles semantic networks or frame-based representations, depending on the point of view taken. Syntactically RDF is built on top of the *Extensible Markup Language* (*XML*).

Parallel to the development of RDF, the W3C also worked on *RDF Schema*[10] (*RDFS*), a suggested extension, which adds an extensible, object-oriented type system to the RDF Model. RDFS includes the ability to express a type (or class) hierarchy as well as a small range of constraints, e.g. basic cardinality constraints. It forms a simple ontology language.

roles that are either transitive or are defined by means of complex role inclusion axioms or have subroles that are either transitive or defined by such axioms cannot be used in cardinality restrictions. For a precise account see again [79].

[8]The exact meaning of $(\mathcal{D})$ seems to vary with respect to which data types and data type predicates (if any) are supported. It usually signifies a very restricted form of concrete domains [3], which can otherwise easily lead to undecidability [4], [103].

[9]`http://www.w3.org/TR/1999/REC-rdf-syntax-19990222` (last accessed 28/09/2011).

[10]`http://www.w3.org/TR/rdf-schema` (last accessed 17/05/2011)

DL	Basis for	Remarks
$\mathcal{SHIF}$	OWL 1 Lite	—
$\mathcal{SHOIN}$	OWL 1 DL	More expressive than $\mathcal{SHIF}$.
$\mathcal{SHIQ}$	OIL	More expressive than $\mathcal{SHIF}$.
$\mathcal{SHOIQ}$	DAML+OIL	More expressive than both $\mathcal{SHOIN}$ and $\mathcal{SHIQ}$.
$\mathcal{SROIQ}$	OWL 2	More expressive than $\mathcal{SHOIQ}$.

Table 3.2 – Description logics of the $\mathcal{SH}$ family and their correspondence to ontology languages.

One of the earliest ontology representation languages, which drew on all three influences just mentioned — description logics, frame-based systems and RDF/XML — was the *Ontology Inference Layer* (*OIL*) [43], sometimes also spelled out as *Ontology Interchange Language* [42, p. 71]. OIL's formal semantics and reasoning is based on the $\mathcal{SHIQ}$ description logic [78, Appendix C]. OIL also introduced the idea of a layered architecture to Semantic Web languages, providing subsets of different expressivity (*Core OIL*, *Standard OIL*, *Instance OIL* and *Heavy OIL*) for different applications. This idea was later inherited by OWL 1 and, in a modified form, by OWL 2 (see below).

The *DARPA Agent Markup Language* (*DAML*), a language for expressing semantic content on the Web, was initiated in 2000. DAML was developed with a view to creating a semantic markup language that would eventually become a Web standard: It was envisaged that the language would be delivered in two portions, DAML-ONT and DAML-Logic. The ontology language DAML-ONT was in many ways similar to OIL. Like OIL it was built on the already existing or proposed standards of XML, RDF and RDFS. Furthermore, like OIL, it drew inspiration from object-oriented and frame-based systems for its knowledge modelling primitives and was to be grounded in description logics. DAML-ONT was combined with OIL and the experiences from the Simple HTML Ontology Extensions (SHOE) project to form *DAML+OIL*[11] [77]. The DAML+OIL language is equivalent to the description logic $\mathcal{SHOIQ}(\mathcal{D})$.[12]

A revised version of DAML+OIL, together with the RDF Model and RDF Schema, became the basis for the *Web Ontology Language* (OWL), which was developed by the W3C Web Ontology Working Group as part of the World Wide Web Consortium's Semantic Web effort. Together with an updated version of RDF[13] and the finalised version of RDFS[14], the Web Ontology Language became a W3C recommendation in February 2004.[15]

OWL is defined as a collection of three increasingly more expressive sublanguages: *OWL Lite*, *OWL DL* and *OWL Full*. OWL Lite supports class hierarchies, unions and complements of classes and a few simple constraints.[16] OWL DL, in addition, supports the use of nominals and unqualified cardinality restrictions (compare Table 3.1). It corresponds to the description logic $\mathcal{SHOIN}(\mathcal{D})$.

[11] `http://www.w3.org/TR/daml+oil-reference` (last accessed 25/05/2011).

[12] Compare "DAML+OIL is equivalent to the $\mathcal{SHIQ}$ DL [...] with the addition of existentially defined classes (i.e., the oneOf constructor) and *datatypes*" [77, p. 7 (emphasis in the original)].

[13] The recommendation consists of six documents, the first of which can be found at `http://www.w3.org/TR/rdf-primer` (last accessed 12/05/2011). It includes references to the other five.

[14] `http://www.w3.org/TR/rdf-schema` (last accessed 12/05/2011).

[15] Similarly to RDF, the original recommendation for OWL consists of six documents, the first of which can be found at `http://www.w3.org/TR/owl-features` (last accessed 12/05/2011).

[16] The language definition of OWL Lite makes no explicit provision for unions and complements, but it could be shown that unions and complements can be reconstructed using other constructs and that OWL Lite corresponds to the description logic $\mathcal{SHIF}$.

OWL Full provides the same modelling primitives as OWL DL, but without restrictions attached to them. The expressivity of OWL Full goes beyond description logics. It allows the use of concept names in place of instance names in statements, which is contrary to the separation of ABox and TBox that is characteristic for description logics.

While OWL was widely adopted as an ontology representation language, following its recommendation by the W3C, theoretical inquiries and practical experiences with the language led to a revision of the OWL standard in 2009 [53].

OWL 2 is based on the description logic $\mathcal{SROIQ}(\mathcal{D})$. Unlike OWL, it no longer contains structural legacies from the frame paradigm, but retains syntactic compatibility with OWL ontologies that are expressed in RDF syntax.[17]

The OWL 2 standard offers a number of alternative syntaxes with various advantages and disadvantages. The normative RDF/XML serialisation,[18] is compatible with the OWL 1 RDF syntax. The OWL 2 Functional Syntax[19] corresponds directly to the structure of ontologies and is used to define OWL 2 ontologies in the abstract. An alternative XML based serialisation, called OWL/XML, allows for better support from existing XML tools, than RDF/XML does. Finally, the Manchester Syntax provides a readable serialisation of OWL 2 DL ontologies: using English language keywords in place of formal operators, Manchester Syntax approximates something like controlled English (assuming that class and property names are suitably chosen) [76].

We used OWL DL for the first versions of ART Ontology (cf. Section 5.1) and adopted OWL 2 as our implementation language after it became available. In particular, the part-whole modelling of geographical units and temporal events in ART Ontology knowledge bases, crucially relies on the feature of sub-property chains which was new to OWL 2 (cf. Sections 3.3 and 5.2).

Table 3.2 gives an overview of the correspondence between description logics and ontology languages based on these logics.

Program-controlled Ontology Manipulation

Encoding description logic in a machine-readable, e.g. XML based format, is only the first step in creating a semantic framework or applications resting on semantic models. If OWL ontologies are to be integrated into software applications, they need to be processed and manipulated by program code. A straightforward way of doing this is to directly read, process (i.e. interpret as part of the business logic of a program) and write the contents of an ontology to and from one of the serialisations of OWL. This, however, is an inefficient and error prone process, as every manipulation of an ontology has to take care not to destroy the integrity of the document or violate global constraints. In this approach, each ontology based application would have to reinvent the "semantic wheel" as it were, as structurally identical ontology operations have to be carried out, regardless of the content of the ontology or the purpose of the application.

---

[17] The OWL 2 specification consists of a series of documents. The first one is located at `http://www.w3.org/TR/owl2-overview` (last accessed 26/05/2011). The specification defines three language "profiles", which restrict the expressiveness of full OWL 2 in various ways for the benefit of different application scenarios. OWL 2 EL is designed such that efficient automatic classification of large TBoxes is possible. OWL 2 QL is optimised for conjunctive query answering. And OWL 2 RL is designed to allow for efficient rule based reasoning implementations. Unlike OWL Lite, these profiles are motivated primarily by their formal semantic and reasoning properties and not by syntactic restrictions that do not directly correspond to restrictions in expressivity. Because of the backward compatibility of OWL 2, however, OWL Lite and OWL DL can technically be seen as OWL 2 profiles as well. After the introduction of OWL 2, some have adopted the convention to refer to the original OWL language as "OWL 1".

[18] `http://www.w3.org/TR/2009/REC-owl2-mapping-to-rdf-20091027` (last accessed 26/05/2011).

[19] `http://www.w3.org/TR/2009/REC-owl2-syntax-20091027` (last accessed 26/05/2011).

To remedy this situation, several ontology processing libraries and APIs for different programming languages and different knowledge representation formalisms have been created. A number of such programming frameworks for accessing and manipulating semantic data algorithmically and the characteristics and desiderata of Semantic Web frameworks in general are discussed in [67]. Such frameworks need to be able to load and save semantic data from and to a variety of concrete representations of a Semantic Web language. This could be a file format, like the ontology language serialisations already mentioned, or a database or other persistence solution. The framework is required to provide an efficient storage solution for semantic data, as well as being able to execute basic modifications and providing access to more complex navigation, querying and reasoning services. The latter part is often realised by providing interfaces to other software (see Section 3.1.3 about reasoners). Furthermore, in an application context or when multiple users are involved, such a framework ideally also takes care of technical aspects like event handling and concurrent threads.

Examples of semantic frameworks are the Jena framework[20] or the API on which earlier versions of the Protégé ontology editor were based.[21] The OWL API is designed specifically for OWL 2, with the intention to be "the API of choice for anyone working with OWL 2 ontologies" [75]. For ontology applications that want to utilise OWL 2 (as opposed to RDF say) and remain unbiased with respect to the particular syntax (i.e. serialisation) in which ontologies are stored or the way in which ontology persistence is achieved (e.g. triplestore) the OWL API offers a flexible and powerful tool.

The reference implementation of the OWL API is written in Java, which makes it particularly easy to integrate it with well known GUI toolkits like Swing. More importantly the reference implementation provides an efficient in-memory storage solution for ontologies so that rapid application prototyping without an immediate need for another ontology persistence solutions is possible.

Our prototype implementation is based on the OWL API (cf. Section 5.3.4).

Ontology Editors

The flexibility of a semantic framework like the OWL API, which allows for powerful manipulation and easy integration of ontology encoded knowledge with other software elements, is of key importance for ontology-based applications. Sometimes, however, in the process of testing or quick prototyping or when ontology manipulations from domain experts unfamiliar with knowledge engineering techniques are required, it is necessary to manually manipulate an ontology or knowledge base. At the layer between the program-controlled access to an ontology and the manual manipulation of a serialised version of an ontology in text form, a number of ontology editors or knowledge representation environments have emerged, that use graphical user interfaces. Apart from the convenience of ontology manipulation, such editors usually also provide ways of visualising all or part of an ontology or knowledge base, for example in the form of an indented hierarchy.

A widely used example of such an editor is the *Protégé Ontology Editor and Knowledge Acquisition System*,[22] or *Protégé*, for short. Protégé was first developed at Stanford University and is currently maintained in a collaborative effort between Standford University and the University of Manchester.

Protégé began as a platform for frame-based knowledge representation and currently supports both frame-based ontologies and OWL ontologies. When the system was expanded to cover OWL

[20] `http://openjena.org` (last accessed 27/05/2011).

[21] `http://protege.stanford.edu/plugins/owl/api` (last accessed 28/09/2011).

[22] `http://protege.stanford.edu` (last accessed 10/05/2011).

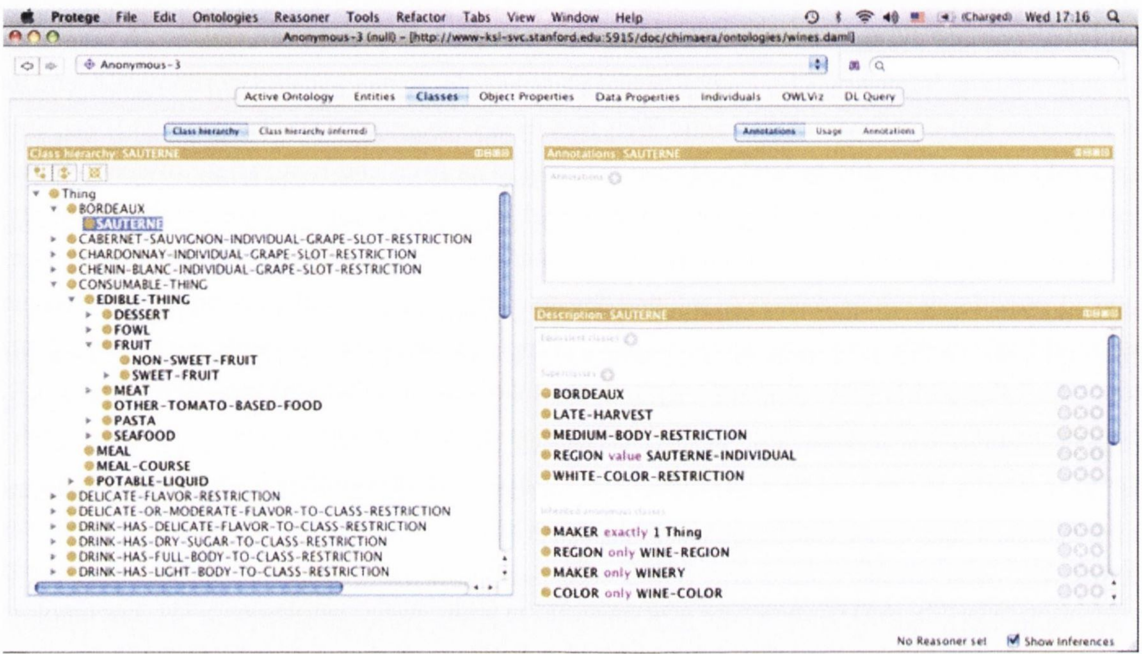

Figure 3.1 – A screenshot of the Protégé 4.1 system, showing the class hierarchy of the Wine ontology.

ontologies, initial versions used the existing frame-based machinery for the internal representation of such ontologies. This is the case for versions of Protégé up to the 3.x series. Releases of Protégé with version numbers 4.x on the other hand, which are maintained in parallel, are based on the OWL API and, starting from version 4.1, provide full support for OWL 2 (but no support for frames as of this writing).

The functionality provided by the OWL API enables Protégé 4.x to process OWL ontologies in a number of different serialisation formats, including RDF/XML, OWL/XML, OWL Functional Syntax, Manchester OWL Syntax and Turtle,[23] a subset of N3.[24] Once an ontology is read, its classes, object properties (roles), data properties (roles with concrete domain values) and individuals can be inspected in different tabs of the graphical user interface (Figure 3.1).

Statements about classes, properties and individuals can be introduced or modified, using the graphical interface only (e.g. subclass statements, data and object restrictions). More complex axioms can be formulated through the *Class expression editor*, which facilitates the creation of complex class descriptions in the relatively intuitive Manchester OWL Syntax.

Protégé is a free, open-source Java platform, built in an extensible plug-in architecture. There are plug-ins for visualising the relationships between classes that are well documented: OntoViz,[25] OWLViz,[26] and OntoGraf[27] (Figure 3.2).

Other plug-ins facilitate the importing of schema structure and data from relational databases into ontologies[28] or integrate natural language processing capabilities with ontology engineering.[29] DL reasoners provide plug-ins that integrate with the ontology development environment to provide

[23] `http://www.w3.org/TeamSubmission/turtle` (last accessed 20/03/2012).
[24] `http://www.w3.org/DesignIssues/Notation3` (last accessed 20/03/2012)
[25] `http://protegewiki.stanford.edu/wiki/OntoViz` (last accessed 12/05/2011).
[26] `http://protegewiki.stanford.edu/wiki/OWLViz` (last accessed 12/05/2011).
[27] `http://protegewiki.stanford.edu/wiki/OntoGraf` (last accessed 12/05/2011).
[28] `http://protegewiki.stanford.edu/wiki/DataMaster` (last accessed 12/05/2011).
[29] `http://protegewiki.stanford.edu/wiki/Natural_Language_Processing` (l. acc. 12/05/2011).

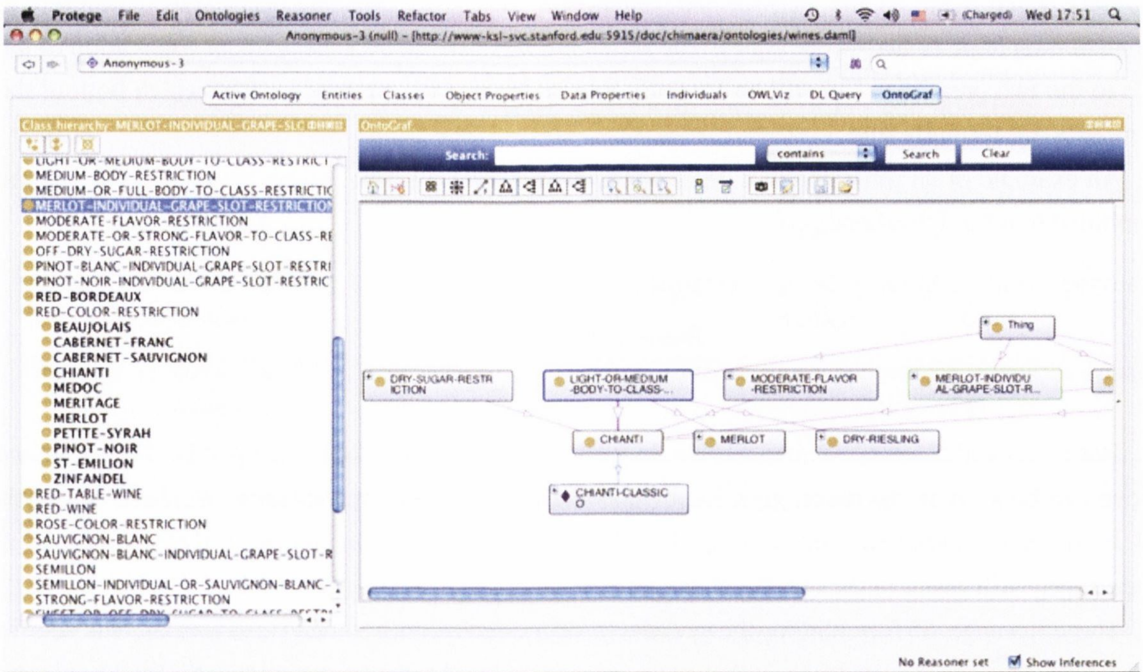

Figure 3.2 – A screenshot of a knowledge base visualisation, involving several classes and one individual, that was created with the OntoGraf plug-in.

services like subsumption testing and instance checking. Reasoning will be covered in the following section.

Different versions of the Protégé Ontology Editor were used in the course of our project for manipulating ontologies and knowledge bases, checking inferences and visualising parts of our knowledge representation.

3.1.3 Reasoning

A fundamental notion in logic is that of *inference*, the drawing of a conclusion from a set of premisses. The field of logic, in its most general sense, can be defined as the study of valid inference. The process of drawing inferences or conclusions is also known as *reasoning*.

In modern, formal logic the language in which premisses and conclusions are expressed and the notion of inference are explicitly, precisely and formally defined. It is customary that the notion of truth is formally defined as well [148]. The definition of inference, however, does not depend on the definition of truth. Inference is specified as a set of permissible syntactical transformations or deductions in some calculus. The validity of such deductions can be determined mechanically. The calculus is *sound* if and only if it fulfils the following condition: Whenever premisses $A_1, ..., A_n$ are formally true and conclusion B can be formally deduced in the calculus, then conclusion B is also formally true.

An example of a formal deduction is an application of the rule of *modus ponens*, which is part of many calculi for logical languages containing *material implication* ($\rightarrow$):

$$\begin{array}{ll} \text{Premiss 1:} & A \\ \text{Premiss 2:} & A \rightarrow B \\ \hline \text{Conclusion:} & B \end{array}$$

Here A and B can be any well-formed sentence in the logical language under consideration (e.g. propositional logic or first-order logic). The statement B is derived by a mechanical application of the rule to the two premisses and will be formally true under any interpretation under which both of the premisses are formally true.

An example of an inference in a description logic is the following, where two TBox axioms are combined to infer a third one:

$$\begin{array}{ll} \text{Premiss 1:} & A \sqsubseteq B \\ \text{Premiss 2:} & B \sqsubseteq C \\ \hline \text{Conclusion:} & A \sqsubseteq C \end{array}$$

Here A, B and C are not sentences, but concepts in a description logic (see p. 42). The inference above can be seen as the description logic equivalent of the well-known *Modus Barbara* of classical Aristotelian syllogisms (substituting the concepts *Greeks*, *Humans* and *Mortals* for A, B and C respectively will result in the canonical example of Modus Barbara).

The advantage of formulating the premisses of a conclusion not only in a formal, but also in a machine-processable language is that valid inferences can be drawn automatically by a program. In the field of computational ontology, such programs are called *reasoners* or *classifiers* and they form an important part of Semantic Web and ontology related technologies. Reasoners allow for automated consistency checking of ontologies and can infer facts implicitly contained in a given ontology knowledge base. Tableau calculi form the basis of the most commonly used reasoners. We used four widely used multi-purpose reasoners at different stages and for different purposes in our project:[30]

RacerPro. The RacerPro reasoner ('Renamed ABox and Concept Expression Reasoner') employs a tableau-based calculus for the description logic $\mathcal{SHIQ}$ [63]. Role axioms and reasoning with nominals ($\mathcal{SROIQ}$) are currently not integrated [63, p. 2f.]. RacerPro supports data types including integers, real numbers and strings. It is optimized for various DL related application scenarios, ranging from large TBoxes combined with many small ABoxes to small TBoxes with one large ABox.

HermiT. The HermiT Reasoner employs, what its creators call a "hypertableau" calculus [114], which aims at reducing nondeterministic behaviour (branching) in the execution of the algorithm [135]. HermiT also contains several other novel optimizations, including strategies to reduce the number of individuals that are generated in the process of building models for a given ontology and knowledge base (in order to demonstrate the consistency of a knowledge base, tableau algorithms try to construct a model for it or else show that no such model can exist). HermiT has been shown to perform well in ontology classification and has been able to classify several complex ontologies which no other reasoner could previously handle. However, the system is not optimized for reasoning with large ABoxes and can also display suboptimal performance when an ontology involves many complex relationship axioms (RBox).

Pellet. The Pellet Reasoner [138] is a tableau-based DL-reasoner for $\mathcal{SROIQ}(\mathcal{D})$ which grew out of "a proof of concept system to help meet the W3C's implementation experience requirements for the Web Ontology Language" [138, p. 51]. It supports data types, including numeric and date/time data types and user-derived types.

[30] There are various freely available and proprietary reasoners which have been optimized for different tasks and have different strengths and weaknesses. Some of them can only deal with a less expressive subset of the OWL 2 language. An overview of some of the reasoners available can be found on the W3C OWL Implementations page: `http://www.w3.org/2007/OWL/wiki/Implementations` (last accessed 4/05/2011).

Reasoner	DL calculus	Optimised for	Remarks
RacerPro	$\mathcal{SHIQ}(\mathcal{D})$, tableau	a large TBox with many small ABoxes; a small TBox with one large ABox	no support for complex role inclusions and nominals
HermiT	$\mathcal{SROIQ}(\mathcal{D})$, hypertableau	complex TBoxes without complex RBoxes	slow on complex RBoxes and large ABoxes
Pellet	$\mathcal{SROIQ}(\mathcal{D})$, tableau	nominal reasoning, incremental reasoning	good general purpose performance
FaCT++	$\mathcal{SROIQ}(\mathcal{D})$, tableau	taxonomic classification, inverse roles	late adopter of novelties (data types, OWL API)

Table 3.3 – A comparison of commonly used DL reasoners.

FaCT++. The FaCT++ reasoner ('Fast Classification of Terminologies') uses a tableau-based calculus for the description logic $\mathcal{SROIQ}$ and supports data types, including integers, doubles and strings [151].[31] Because of its optimisation for inverse roles FaCT++ performed well on our *ART Ontology* and was used in the final implementation of our *Artfinder* system (see Chapter 5).

Table 3.3 gives an overview of salient properties of these common general purpose reasoners.

3.2 Ontology Development Methods

The formal underpinnings of description logic and the tools and technologies we discussed in the previous section form the basis for building computational ontologies. In this section we will cover aspects of ontology development.

The principles of designing, building and maintaining ontologies and best practice patterns in ontological modelling have been discussed in the literature and many methods have been suggested (see [50] and [120] for overviews).

Some authors refer to their proposals in this respect as "methods", others as "methodologies". Gómez-Pérez and colleagues attempt to clarify this distinction by recourse to a, by now superseded, IEEE standard [50, p. 108]. We will refer to all proposals uniformly as "methods".

The proposed methods divide the ontology development process primarily into three areas: first, defining the purpose and scope of an ontology; second, building the ontology by creating concept and relationship terminology, definitions and axioms; and, third, maintaining the ontology and managing its life cycle. These three areas do not appear as strictly sequential and thus separated fields of activity, but are interlinked and mutually influence each other. Especially the last point, managing an ontology's life cycle, usually prescribes, more or less explicitly, a cycling back and forth between, for example, building a concept hierarchy and refining the scope of an ontology or evaluating and re-engineering certain aspects of an ontology.

[31] According to Tsarkov and Horrocks [151], FaCT++ is a $\mathcal{SHOIQ}$ reasoner, but as of December 2006, the reasoner supports the more expressive $\mathcal{SROIQ}$ description logic (compare `http://code.google.com/p/factplusplus`, last accessed 5/05/2011).

3.2.1 Defining purpose and/or scope

The main questions here are: Should the purpose of an ontology be explicitly considered in application scenarios or is it enough to define the scope of an ontology independently of applications? and: How is the scope of an ontology best determined and defined?

As far as the first question is concerned, it is important to consider, that there is a certain trade-off between application dependence and the potential for reuse in an ontology. Some methodological considerations, especially of an earlier date, are largely application independent or focus on the reuse aspect across applications [57], [97], [155]. Increasingly, however, more emphasis was put on tying a newly developed ontology to a core application or a set of core applications [118], [120], [143]. Delineating the application an ontology supports, helps in determining the ontology's scope and can resolve individual modelling problems.

Strategies that have been proposed for outlining the scope of an ontology (i.e. the knowledge domain that an ontology is supposed to cover), include brainstorming, "grouping" [155] and "trimming" [152] in order to arrive at a list of relevant terms. Some authors emphasise that at this stage terms, as linguistic entities, are sufficient and that semantic properties of terms (disambiguating, typing etc.) should be left for later [118], [155].

Several proposed methods, suggest defining the scope of an ontology through a list of competency questions, i.e. questions that the ontology should be able to answer [54], [118], [143]. Grüninger and Fox refine this into a two-stage process with informal competency questions first, and formal competency questions at a later stage [55]. Others have argued, however, that it may not always be necessary or even possible to specify precise competency questions at an early stage of ontology development and that such questions may play a more important role in evaluating an ontology than in determining its scope [155].

3.2.2 Building an ontology

The building phase is the core modelling activity where an ontology takes shape. Important differences between proposed methods concern the questions, whether and how existing ontologies should be reused, how terminology and definitions should be agreed upon and in which order concept and relation hierarchies should be created.

Ever since ontologies first entered the discourse on artificial intelligence and knowledge based systems, the idea of their reuse has been of central importance. It is therefore not surprising, that the majority of proposed ontology development methods expressly include instructions for the reuse of existing ontologies. The proposals differ, however, with respect to the development stage at which reuse should be considered. Staab and colleagues advocate looking for reusable ontologies early in the kick-off phase, around the same time when the scope of the ontology is made precise [143]. Similarly, Öhgren and Sandkuhl recommend that checking for candidate ontologies should happen "as soon as possible" [120, p. 374]. Noy and McGuinness make reusing existing ontologies the second step in their development method, after determining ontology scope [118]. In the METHONTOLOGY method [50, p. 125ff.] integration of other ontologies happens during the conceptualisation stage, but before the ontology implementation phase (coding), while according to Uschold and King [155] it could be part of either of those two phases, depending on the individual case. Some ontology development methods, however, do not contain explicit guidelines for integrating existing ontologies or similar resources into newly developed ones [55], [144], [152].

As we have indicated in the previous paragraph, some authors distinguish in the ontology building process, between a conceptualisation phase (design at the "knowledge level") and an implementation phase (coding in a knowledge representation language) [120], [155]. According to this view, the recommended outcome of the conceptualisation phase is often a natural language document that is sufficiently precise to serve as the basis for the implementation of the ontology in a concrete representation system. Reasons for keeping the two steps separated include, better and more traceable documentation, less rework on the ontology code if mistakes become apparent in the conceptualisation phase [50, p. 126] and no adverse expressivity or other limitations, due to early commitment to a particular representation language [155]. Grüninger and Fox suggest using first-order logic as generic formalism in the conceptualisation stage [55].

A concept hierarchy and, possibly, a relation hierarchy is at the core of most ontologies (compare Section 2.1.2). However, it is not immediately clear in which order one should proceed if such a hierarchy has to be constructed from scratch. Uschold and King [155] have addressed this point and mention three possibilities: top-down (from the most general concepts to the most specific), bottom-up (from the most specific concepts to the most general) and middle-out (from the most salient concepts to more general and more specific ones). These three approaches and the possible implications they may carry for the resulting ontologies are more fully illustrated in [50, p. 116-118]. While some authors have not addressed this question [97], [143] or left the choice between alternatives open [118], many have argued for a middle-out approach [50, p. 150], [152], [153], [155]. This is, if not directly supported, then at least motivated by research in psychology [93], [129]. At least one development method, however, favours a top-down approach [11]. The CIDOC CRM (cf. Section 2.3.3) is an example of an ontology which was, at least in the beginning, developed in a bottom-up fashion [35, p. 468].

Some ontology building methods differ from the strategies we have discussed so far in that they are tied to a particular core ontology which serves as a starting point for the development [146].[32] In these cases the line between methodological recommendations and a documentation of the core ontology can be blurred. Adopting such a method is conditional on the decision to use the core ontology. The intention here is that the core ontology facilitates interoperability between ontologies that have been created according to the same method and provides a guideline for modelling decisions that might otherwise be resolved in an ad-hoc fashion. Similar benefits may result from reusing a top-level ontology that is not tied to a development method. This is the approach that we have adopted (cf. p. 96).

Finally, it is worth mentioning in the context of ontology building, that the collective experience of semantic modellers has created a body of semantic design patterns, that form solutions to common modelling problems, analogous to object-oriented design patterns. This was aided by the increasing standardisation of representation languages and associated language constructions in recent years. Allemang et al. provide an overview of semantic modelling in RDFS and OWL and cover many recommended design patters and not recommended, but frequently encountered "antipatterns" [2].

3.2.3 Managing the ontology life cycle

Many method proposals consider the building of an ontology as part of a larger development context, that includes the identification and management of different ontology user groups (domain experts,

[32] See also the remarks on ontology design and best practices based on the Basic Formal Ontology at `http://www.ifomis.org/bfo/documents/manual.pdf` (last accessed 23/6/2011).

knowledge engineers, end users, etc.), control of the ontology building stages and the continued maintenance and evaluation of ontologies (e.g. [143]).

Different strategies have been suggested with respect to how changes and improvements to the ontology are controlled and agreed among stakeholders (e.g. evolving prototypes vs. incremental life cycle [50, p. 149]). For evaluation purposes some approaches suggest revisiting competency questions, if such questions have been defined at an earlier development stage [55], [118], [143]. A domain-independent evaluation method for ontological decisions, that is based on philosophical principles was proposed by Nicola Guarino and Christopher Welty [61], [62]. ONTOCLEAN, as the method is called, is more concerned with formal modelling correctness, than with application oriented characteristics, like domain coverage.

Considerations concerning the life cycle of an ontology are especially important in institutional and industrial settings and when the number of people involved in creating and using the ontology is large.

3.2.4 Discussion: Consensus in Ontology Development

The ontology development methods put forward in the literature are, aside from a review of other method proposals, mostly based on the experiences from a single project and in some cases only the development of a prototype. Generally, this is openly reflected by authors and accordingly none of the proposed methods pretends to hold all the answers to ontology development. Experimentation as a means of deciding which methodological approach would give the best results depends on too many variables and would have to remain restricted to toy problems. It was rightly pointed out that "it is not very likely that someone will pay twice for building the same complex ontology with different approaches" [50, p. 113].

Across the gamut of published field reports, however, a consensus on many of the individual questions under discussion appears to be emerging. Among the advice one encounters repeatedly are the following points: consider using competency questions for defining ontology scope, look for reusable ontologies sooner rather than later and construct concept hierarchies middle-out, rather than bottom-up or top-down. Like others before us, we attempted to translate what advice we gleaned from the literature into concrete decisions in our ontology development process.

3.3 Transitive Propagation

In the previous section we discussed general strategies for developing an ontology. As a rule, these development methods do not cover detailed aspects of ontology modelling. In this section we will discuss a pattern we have used for modelling temporal and geographical inclusion. We will first state the problem, then present three alternative solutions and finally justify which one we have adopted.

3.3.1 Problem Statement

A binary relation R on a set A (i.e. $R \subseteq A \times A$) is *transitive* if and only if whenever elements a and b and elements b and c are related by R then elements a and c are also related by R. In mathematical notation: $R\ \mathit{transitive} :\Leftrightarrow \forall a,b,c \in A : aRb \wedge bRc \rightarrow aRc$.

Listing 3.1 – Ontology O: Basic declarations and assertions in a transitive propagation example.

```
Declaration( Class( ex:Event ) )
Declaration( Class( ex:Place ) )

Declaration( ObjectProperty( ex:tookPlaceAt ) )
Declaration( ObjectProperty( ex:happenedIn ) )
Declaration( ObjectProperty( ex:fallsWithin ) )

Declaration( NamedIndividual( ex:StormingOfTheBastille ) )
Declaration( NamedIndividual( ex:Paris ) )
Declaration( NamedIndividual( ex:France ) )
Declaration( NamedIndividual( ex:Europe ) )

ClassAssertion( ex:Event  ex:StormingOfTheBastille )
ClassAssertion( ex:Place  ex:Paris )
ClassAssertion( ex:Place  ex:France )
ClassAssertion( ex:Place  ex:Europe )

ObjectPropertyDomain( ex:tookPlaceAt  ex:Event )
ObjectPropertyRange( ex:tookPlaceAt  ex:Place )

ObjectPropertyAssertion( ex:fallsWithin  ex:Paris  ex:France )
ObjectPropertyAssertion( ex:fallsWithin  ex:France  ex:Europe )

ObjectPropertyAssertion( ex:tookPlaceAt
                         ex:StormingOfTheBastille
                         ex:Paris )
```

Transitivity is an important property of equivalence and partial order relations. The latter are, in turn, a key ingredient in the formal axiomatisation of part-whole relationships (mereology). In the context of modelling structured objects in ontologies it is often desirable that a property propagates along a transitive relation. In mathematical terms, a binary relation $P \subseteq B \times A$ is said to propagate along a relation $R \subseteq A \times A$ if and only if $\forall a \in B$ and $\forall b,c \in A : aPb \wedge bRc \rightarrow aPc$. An example often mentioned in ontology literature is anatomical body parts and ailments that affect those body parts. The body parts are related by a transitive *part-of* relation and ailments are related to their body location by a relation *affects body part*. The *affects body part* relation can be seen as propagating along the *part-of* relation. This means if a certain disease affects the pulmonary valve, it will also affect the heart and the cardiovascular system as a whole.

An example from another domain are historic events that take place at certain geographical locations and, by virtue of transitive propagation, also in more inclusive geographical locations: if Napoleon's last battle took place at Waterloo we also know that it took place in Belgium and similarly in Europe once we know that Waterloo is in (is part of) Belgium and Belgium is in (is part of) Europe. We have used transitive propagation in our ontology to model the geographic location of events and also the 'temporal location', i.e. the link between an event and concrete time-spans.

Modelling transitive propagation in a logical framework is dependent on the expressivity of the underlying logic. In first order logic or stronger theories transitive propagation of one relation along another one is a straightforward axiom. In a knowledge representation language like RDF, on the other hand, no axiom or mechanism will permit the modelling of transitive propagation.

In between these extremes lie the possibilities provided by $\mathcal{SHOIN}$ and $\mathcal{SROIQ}$ [79] the two description logics underlying OWL 1 (i.e. OWL DL as specified in [122]) and its successor OWL 2 [113] respectively. In OWL 1 modelling transitive propagation is not supported in a straightforward way though a number of workarounds or extensions have been discussed in the literature [112], [124], [133]. OWL 2, on the other hand, was explicitly designed with transitive propagation in mind (among other things) [53, p. 310]. However, due to the restrictions that are placed on OWL 2 in order to keep the language decidable, complications can arise if everything that is semantically intended is explicitly asserted in a model.

For a detailed discussion of the modelling task, we will use the OWL 2 functional-style syntax as defined in [113]. Let our example ontology *O*, which we will subsequently extend, contain the declarations and assertions of Listing 3.1.[33] Ontology *O* contains a class called `Event` and a class called `Place`, three relations, `tookPlaceAt`, `happenedIn` and `fallsWithin` and four instances, three of which are specified to be places and one which is an event. The domain and range of `tookPlaceAt` are `Event` and `Place` respectively. Finally, our ontology contains three statements, namely that `Paris fallsWithin France`, `France fallsWithin Europe` and that the event `StormingOfTheBastille tookPlaceAt Paris`. The four named individuals and the three property assertions of our ontology are represented in the diagram in Figure 3.3. The individuals correspond to lozenges, the property assertions to solid arrows.

Our objective is to add axioms to ontology *O* that will allow us to infer the following three facts under the rules of OWL 2 DL inference:

I1:
```
ObjectPropertyAssertion( ex:happenedIn
                         ex:StormingOfTheBastille
                         ex:Paris )
ObjectPropertyAssertion( ex:happenedIn
                         ex:StormingOfTheBastille
                         ex:France )
ObjectPropertyAssertion( ex:happenedIn
                         ex:StormingOfTheBastille
                         ex:Europe )
```

Furthermore it may be of interest to infer the following geographical inclusion:

I2:
```
ObjectPropertyAssertion( ex:fallsWithin ex:Paris ex:Europe )
```

These two sets of inferences *I1* and *I2* correspond to the dotted arrows labelled (1) and (2) in Figure 3.3 respectively. In the context of this section *I1* will be our top priority, *I2* a second priority. Of course *I2* can be readily inferred by making the `fallsWithin` relation transitive:

T:
```
TransitiveObjectProperty( ex:fallsWithin )
```

Note that *I1* suggests that, while the basic fact about where the storming of the Bastille took place is stated using the property `tookPlaceAt`, it is the property `happenedIn` that we want to

[33] Listing 3.1 and each of the separately listed and labelled sets of ontology statements in this section, not only Listing 3.1, can be seen as a small ontology in its own right, sometimes containing only one statement. In these listings we will use the example namespace prefix "`ex`:" to comply with the syntax laid out in [113]. We have, however, omitted the merely technical statements, like prefix declarations and ontology IRI, because they are irrelevant to our discussion. Additionally, when ontology elements are mentioned in continuous text we have also omitted the namespace prefix "`ex:`" for easier readability.

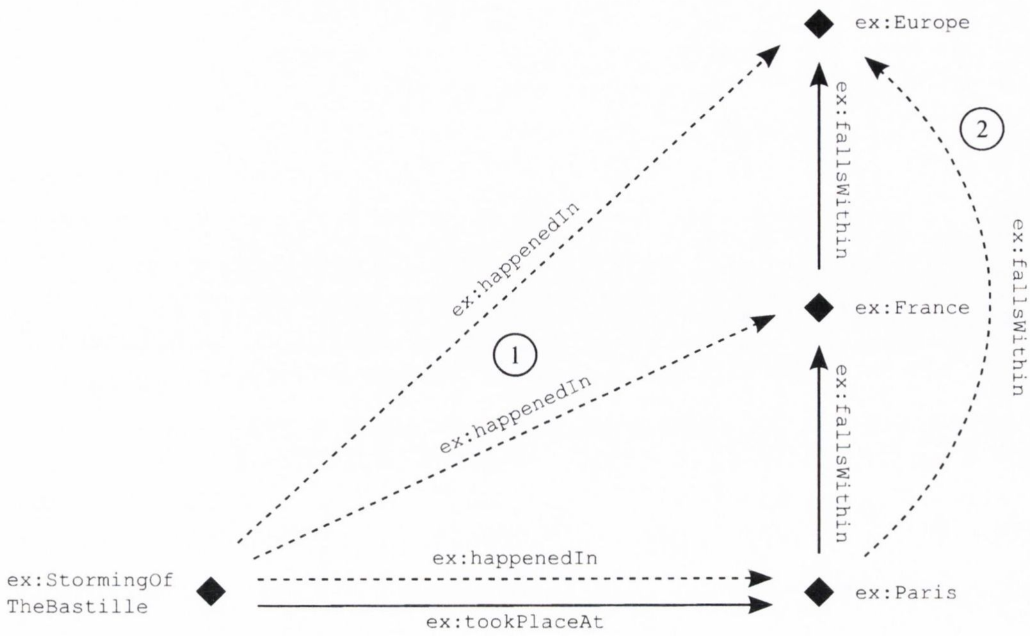

Figure 3.3 – An example of a transitive propagation problem. The solid arrows represent statements (i.e. relation instances) in the ontology O (Listing 3.1). The dotted arrows represent desirable inferences.

be propagating along the `fallsWithin` relation. There are several reasons for having two separate properties here. First, `tookPlaceAt` could be a property that is only used to describe 'atomic' cases of, in our example, the location of events. For example `tookPlaceAt` could express the location of an event only at a certain level of geographical granularity: 'took place at (city or town)', whereas `happenedIn` could be used to express the location of an event at all levels of geographical granularity. Second, `tookPlaceAt` could stem from a top-level or domain ontology like the CRM we have discussed earlier and it may be preferable to introduce a separate relation (i.e. `happenedIn`) that propagates in order to leave the axiomatisation of the original relation unaffected.[34] Note that in the presence of an axiom stating the equivalence of the properties `tookPlaceAt` and `happenedIn`, alternative 3.3.2 below will remain unchanged in its properties, while alternatives 3.3.3 and 3.3.4 will effectively become the same.

Modelling transitive propagation in OWL 2 is facilitated by the mechanism of sub-property chains. However, minor problems can arise if the property along which the propagation must happen is reflexive and if this reflexivity is explicitly stated in the model.

A binary relation $R \subseteq A \times A$ is said to be *reflexive on* A if and only if all elements $a \in A$ are related to itself by R, i.e. R *reflexive on* $A :\Leftrightarrow \forall a \in A : aRa$. In some mereological theories reflexivity is assumed as a basic property of part-of relations like `fallsWithin`. The intuition here is that everything is part of (or falls within) itself. If a part-of relation is explicitly stated to be irreflexive on A (i.e. $\forall a \in A : \neg aRa$) then it is interpreted as 'proper part of'.

If `fallsWithin` is interpreted as reflexive and transitive, one single propagation clause allows

[34]More precisely this means that we only introduce axioms that do not allow us to infer new instances of the original relation, i.e. the relation imported from another ontology.

us to infer the desired consequences *I1*. In mathematical notation, the propagation clause can be written like this:

$$\texttt{tookPlaceAt} \circ \texttt{fallsWithin} \subseteq \texttt{happenedIn}$$

Here the composition operator $\circ$ is defined as follows: If $R \subseteq A \times B$ and $S \subseteq B \times C$ are two binary relations, then their composition $R \circ S$ is the relation $R \circ S = \{(a,c) \in A \times C \mid \exists b \in B : aRb \wedge bSc\} \subseteq A \times C$.[35] The corresponding sub-property chain statement in functional style syntax is:

S1:
```
SubObjectPropertyOf( ObjectPropertyChain( ex:tookPlaceAt
                                          ex:fallsWithin )
                     ex:happenedIn )
```

In OWL 2 there are two ways of modelling the reflexivity of a relation. These two options correspond to the first two modelling alternatives we are going to present. The third alternative will involve no reflexivity statement. In the following sections we will illustrate these alternative approaches to modelling transitive propagation using the example ontology we introduced. We will argue that the third alternative is the preferable way of dealing with this modelling problem.

3.3.2 A Model with Global Reflexivity

The first option for modelling reflexivity, already available in OWL 1, is a strong global reflexivity statement for a given relation *R* which corresponds to the logical statement $\forall a : aRa$. Using the `fallsWithin` relation in the place of *R* the strong reflexivity statement *sR* becomes:

sR:
```
ReflexiveObjectProperty( ex:fallsWithin )
```

This means that in a model of our ontology, every element will be related to itself by the relation `fallsWithin`. Very often, however, this statement is too strong, especially when we specify a domain and range for the relation in question. In this case, a strong reflexivity statement like *sR* forces every element of every model of our ontology to be an instance of the domain and range of the reflexive relation. This is because the domain and range of the reflexive relation become equivalent to `owl:Thing` (the OWL/OWL 2 equivalent of *top*, see p. 42f.).

As a consequence if we would like to have the following set of 'desirable statements' *D* in our ontology, we will find that including them in the presence of *sR* renders our ontology inconsistent:

D:
```
ObjectPropertyDomain( ex:fallsWithin ex:Place )
ObjectPropertyRange( ex:fallsWithin ex:Place )
DisjointClasses( ex:Place ex:Event )
```

Using the ontology statements we have introduced so far we can now formulate the first of our three approaches to modelling the given transitive propagation problem. It consists of the basic ontology *O* merged with the statements *S*1, *sR*, and *T*. The resulting ontology is consistent and will

[35] Note that there are differing conventions concerning the composition operator. In some contexts what we have defined as $R \circ S$ is written in exactly opposite order as $S \circ R$. Under the latter convention the standard definition of functional composition becomes a special case of the composition of functional relations. However, since we are not concerned with functions, we follow the convention that is easier to read and that corresponds to the Manchester Syntax for OWL 2 sub-property chains (see [76]).

allow us to infer $I1$ and $I2$. Adding the statements D however will result in an inconsistency. The advantages and drawbacks of this solution are summed up by the following set of statements, where $\oplus$ is the ontology merge operator that combines the axioms of two ontologies, $\vdash$ is logical entailment in OWL 2 DL and $\bot$ represents inconsistency:[36]

$$O \oplus S1 \oplus sR \oplus T \nvdash \bot \tag{3.1}$$

$$O \oplus S1 \oplus sR \oplus T \vdash I1 \tag{3.2}$$

$$O \oplus T \vdash I2 \tag{3.3}$$

$$O \oplus sR \oplus D \vdash \bot \tag{3.4}$$

3.3.3 A Model with Local Reflexivity

The strong reflexivity statement *sR* together with *D* forces the class `Event` in our example to become equivalent to `owl:Nothing` (the OWL/OWL 2 equivalent of *bottom*, see p. 42f.). The presence of an instance of `Event` as stated in O then leads to the inconsistency 3.4. However, there is a second way of modelling the reflexivity of a relation in OWL 2. This second way corresponds to how we have defined reflexivity above, as reflexivity *on* a certain set A as opposed to global reflexivity on the whole universe of discourse. In this case A is the domain and range of the reflexive function and the corresponding reflexivity statement in mathematical notation is: $\forall a \in A : aRa$. A statement to this effect can be constructed in OWL 2 using self-restrictions. A self-restriction is a class that contains exactly those individuals that are related to themselves by a certain OWL object property. Hence, the following statement *wR* says that everything that is a `Place` `fallsWithin` itself or in other words, that `fallsWithin` is reflexive on the class `Place`:

wR: `SubClassOf( ex:Place ObjectHasSelf( ex:fallsWithin ) )`

This is a weaker reflexivity statement than *sR* and it is compatible with the set of axioms *D*. However, because of the global restrictions that the axioms of any OWL 2 DL ontology must satisfy, *wR* can not be stated together with T in the same ontology.[37] These restrictions, which are necessary to ensure decidability, require that axioms of the form `ObjectHasSelf( OPE )`, where `OPE` is an object property expression, do not contain transitive object properties. But in order to derive all of the desired conclusions $I1$ we need the transitive closure of `fallsWithin`. In our example we can achieve this by expressly adding $I2$ to the axioms of our ontology. This leads us to our second alternative.

Using the weaker reflexivity statement *wR* and the other ontology statements we have introduced, we can now spell out the second alternative. It consists of our basic ontology O amended by $S1$, *wR*, $I2$ and *D*. The resulting ontology is consistent and allows us to infer the set of statements $I1$. In the absence of the transitive closure $I2$ however, this inference does not hold. Adding the transitivity axiom T is incompatible with *wR*. The formal properties of this modelling approach are summed up

[36]Note that $\bot$ does not represent the *bottom* element of description logics here (cf. Section 3.1.1). Technically, for merging the ontologies the namespace related statements that we have omitted have to be added back in (see footnote 33). The entailments and non-entailments 3.1 to 3.11 in this section have been verified with the FaCT++ reasoner (version 1.5.1). They have also been checked with the HermiT reasoner (version 1.3.3) which found all of them to be correct, with the exception of 3.8. However, since the global restrictions on axioms in OWL 2 DL are unequivocal (compare `http://www.w3.org/TR/owl2-syntax/#Global_Restrictions_on_Axioms_in_OWL_2_DL`), this perhaps means that the implementation of the HermiT reasoner does not strictly follow the standard.

[37]See `http://www.w3.org/TR/owl2-syntax/#Global_Restrictions_on_Axioms_in_OWL_2_DL`.

by the following entailments and non-entailments:

$$O \oplus S1 \oplus wR \oplus I2 \oplus D \nvdash \bot \quad (3.5)$$

$$O \oplus S1 \oplus wR \oplus I2 \vdash I1 \quad (3.6)$$

$$O \oplus S1 \oplus wR \oplus D \nvdash I1 \quad (3.7)$$

$$wR \oplus T \text{ violates OWL 2 DL restrictions} \quad (3.8)$$

3.3.4 A Model without Reflexivity

Both of the solutions presented so far suffer drawbacks. While alternative 1 presents a consistent ontology that delivers the desired inferences $I1$ and $I2$, it effectively forces the domain and range of `fallsWithin` to be of type `owl:Thing` and thus restricts the possible interpretations of `fallsWithin` and the possible models of the resulting ontology. Alternative 2 on the other hand allows us to infer the desired set of statements $I1$, but only at the expense of explicitly adding the transitive closure of `fallsWithin` to our premises. In our minimal example, this is achieved by including the single statement $I2$. However, in order to provide a systematic solution for a larger ontology the transitive closure has to be added programmatically or by using rules. In either case this solution would involve a mechanism going beyond OWL 2 DL.

There is yet another approach to modelling the same facts that largely avoids the drawbacks mentioned above. Instead of using the sub-property chain $S1$ and making `fallsWithin` transitive and reflexive we can directly introduce `tookPlaceAt` as a subproperty of `happenedIn` (*SP*) and then add an alternative sub-property chain ($S2$) for `happenedIn`:

SP:
```
SubObjectPropertyOf( ex:tookPlaceAt ex:happenedIn )
```

S2:
```
SubObjectPropertyOf( ObjectPropertyChain( ex:happenedIn
                                          ex:fallsWithin )
                     ex:happenedIn )
```

The axiom $S2$ looks very strong, but it is safe as long as `happenedIn` is not a direct or indirect subproperty of `fallsWithin` and not used in any of the axiom types that can only contain so-called *simple* object properties, for instance self-restrictions or disjoint object property statements.[38]

We are now in a position to formulate the third alternative modelling. In its core, it consists of our basic ontology O together with *SP* and $S2$. In this case no reflexivity and not even transitivity of `fallsWithin` are needed to infer $I1$. Transitivity T, however, is compatible with the other axioms and including it in our ontology will allow us to infer $I2$. In the latter case, weak reflexivity wR cannot be added as stated by entailment 3.8 above. If reflexivity of `fallsWithin` is absolutely needed, the reflexive closure R can be added without conflict:

R:
```
ObjectPropertyAssertion( ex:fallsWithin ex:Paris ex:Paris )
ObjectPropertyAssertion( ex:fallsWithin ex:France ex:France )
ObjectPropertyAssertion( ex:fallsWithin ex:Europe ex:Europe )
```

[38]For an exact definition of *simple* object property expressions and the global restrictions imposed on OWL 2 DL ontologies in order to guarantee decidability see `http://www.w3.org/TR/owl2-syntax/#Global_Restrictions_on_Axioms_in_OWL_2_DL`.

The following set of entailments, together with entailment 3.3 ($O \oplus T \vdash I2$), sums up the most salient characteristics of alternative 3:

$$O \oplus SP \oplus S2 \oplus T \oplus R \oplus D \nvdash \perp \quad (3.9)$$

$$O \oplus SP \oplus S2 \oplus I2 \oplus wR \oplus D \nvdash \perp \quad (3.10)$$

$$O \oplus SP \oplus S2 \vdash I1 \quad (3.11)$$

As we have seen wR and T cannot be contained in the same ontology (entailment 3.8). However, we can consistently state the transitivity of `fallsWithin` and add its reflexive closure (3.9) or else state its reflexivity and add its transitive closure (3.10).

3.3.5 Discussion: A Critique of the Modelling Alternatives

Having presented three different alternatives and examined their formal properties, we can now ask which of these alternatives is preferable. If the relation along which the propagation must happen (i.e. `fallsWithin`) is assumed to be reflexive and transitive, then $S1$ suggests itself as a parsimonious solution. However, as we have seen, reflexivity and transitivity of a relation cannot be stated together without drawbacks (entailments 3.4 and 3.8). Yet, in practical applications this is unlikely to be a serious concern. Reflexivity of a relation does not appear to be interesting property in terms of querying an ontology: Looking for geographic entities that lie in Europe, most users will probably find Europe itself of little interest as an answer. It therefore seems preferable to drop reflexivity entirely and include the two statements SP and $S2$ and optionally T as suggested by the last alternative.

3.4 Practical Modelling Challenges

As we have seen in the previous section, some of the relations we use in our model can be naturally thought of as reflexive or transitive and transitivity in combination with other axioms can lead to interesting and welcome inferences once we employ a reasoner. However, we have equally seen that axiomatisations which are too rich, even if 'semantically justified', may cause problems (because of decidability restrictions, see p. 59f.).

This may be aggravated if a model is integrated with ontologies from other parties which carry their own restrictions. In these cases a methodological approach to semantic modelling involves a systematic inspection of the classes and properties that are imported from other ontologies and their axiomatisations. In this section we report on a number of incoherences in the CIDOC CRM that we have discovered in the course of such inspections, both at the abstract level (i.e. pertaining to the abstract definition of the CRM [25]) and at the level of particular CRM implementations.

At the end of this section we have included a discussion of a problem which is practically challenging for very different reasons: The modelling of image subject matter.

3.4.1 Some Issues with Relations in CRM

The abstract CRM definition does not contain separate records for the inverses of its relations [25]. Instead inverse relations are referenced indirectly in brackets, e.g. "P22 transferred title to (acquired title through)" [25, p. 43]. This circumstance has perhaps lead to a number of underspecifications in the axiomatisation of CRM properties.

At first this may seem surprising. After all the name of the inverse property is given in brackets, its intended meaning is determined by virtue of the meaning of its inverse and many of its formal properties are also determined by its inverse. If the original relation, for example, is transitive, its inverse will also be transitive; if the original relation is inverse functional, its inverse will be functional and so on. However, the absence of explicit inverse relations seems to limit the means of expression in the CRM document in certain cases. This concerns in particular subproperty axioms that involve inverse properties.

The property 'P9 consists of (forms part of)', for instance, links periods to their sub-periods, such that the "sub-periods into which the period is decomposed form a logical whole [...] and the sub-periods are constitutive of the general period." [25, p. 39]. The only example that is contained in the CRM documentation, "Cretan Bronze Age (E4) *consists of* Middle Minoan (E4)" [Ibid., emphasis in the original], seems to suggest that the property *P9* is intended to connect, among other things, art historical periods to their art historical sub-periods, e.g. 'Gothic' to 'Early Gothic', 'High Gothic' and 'Late Gothic' etc. On the other hand, the property 'P10 falls within (contains)' links, periods to other periods into which they happen to fall geographically and temporally without any further constraints. According to the CRM definition, the "difference with *P9 consists of (forms part of)* is subtle. Unlike *P9 consists of (forms part of)*, *P10 falls within (contains)* does not imply any logical connection between the two periods and it may refer to a period of a completely different type." [25, p. 39, emphasis in the original]. The example given for P10 is "the Great Plague (E4) *falls within* The Gothic period (E4)" [25, p. 40, emphasis in the original]. Based on the intended semantics, it would seem plausible, that the inverse relation of *P9* should be a subproperty of *P10*. The scope note that explains the difference between the two, essentially describes *P10* as an inverse of P9 minus certain characteristics. However, the CRM definition does not explicitly stipulate that the inverse of *P9* is a subproperty of *P10*.

A second point also involves the property 'P10 falls within'. As the property expresses a basic mereological 'part-of' relation between periods of time, one would expect it to be a subproperty of 'P132 overlaps with', which expresses a temporal overlap between two periods. This reasoning is purely based on mereological principles: if period *A falls within* period *B*, then this constitutes a special case of *A* and *B* overlapping.

The reason for these omissions may be that no explicit definitions of property inverses are contained in the CRM. Even though in some cases equivalent axiomatisations without inverses exist, examples can be constructed which demonstrate that explicit inverses may be desirable in a model such as the CRM.[39]

In other cases, the characterisations of relations in CRM appear to be logically too strong or at least misleading. We will illustrate this with an example. Symmetry and transitivity of properties in the CRM are not systematically stated in the individual property definitions, but derived instead from a statement in the preamble:

> Properties that have identical domain and range are either symmetric or transitive. Instantiating a symmetric property implies that the same relation holds for both the domain-

[39] Imagine, for example, a (not symmetric) relation *was influenced by* between artists, expressing direct or indirect artistic influence, and a relation *was teacher of*, relating artists to their artist students. Assume now that we would like to model that teachers always influenced their students, but students also always influenced their teachers. The most straightforward way to do this is to make both *was teacher of* and its inverse *was student of* a subproperty of *was influenced by*. Admittedly there is a way to achieve the same effect without an explicit inverse of *was teacher of*: introduce a (somewhat artificial) property *had student- or teacher-influence on* and make this property symmetric.

> to-range and the range-to-domain directions. An example of this is *E53 Place. P122 borders with: E53 Place*. The names of symmetric properties have no parenthetical form, because reading in the range-to-domain direction is the same as the domain-to-range reading. Transitive asymmetric properties, such as *E4 Period. P9 consist of (forms part of): E4 Period*, have a parenthetical form that relates to the meaning of the inverse direction. [25, p. xv, emphasis in the original]

As not all binary relations which have identical domain and range are either symmetric or transitive,[40] the statement seems to suggest, that all such relations (i.e. properties) in CRM happen to be of that nature. However, 'P9 consists of (forms part of)' is already an interesting point in case. According to the statement above it is transitive, therefore we may suspect, that at least in principle *P9* chains can exist, that allow to infer new instances of *P9* through transitivity, for example:

European Bronze Age	*P9 consists of*	*Cretan Bronze Age*
Cretan Bronze Age	*P9 consists of*	*Middle Minoan*
European Bronze Age	*P9 consists of*	*Middle Minoan*

At the same time the CRM definition characterises *P9* as "Quantification: one to many, (0,n:0,1)" [25, p. 39], i.e. as an inverse functional property. This would entail, continuing the above example, that *Cretan Bronze Age* and *European Bronze Age* could be inferred to be identical periods. In general, any *P9* chain of length three or greater would be collapsed into one of length two, setting all elements but the last equal to the first. Admittedly the CRM documentation does recommend, that functionality and inverse functionality should not be implemented:

> Quantifiers for properties are provided for the purpose of semantic clarification only, and should **not** be treated as implementation recommendations. The CRM has been designed to accommodate alternative opinions and incomplete information, and therefore **all** properties should be implemented as optional and repeatable for their domain and range ('many to many (0,n:0,n)'). [25, p. xiii, emphasis in the original]

But even so it would appear to be questionable, how much "semantic clarification" is provided by stipulations that can lead to undesirable results like the one outlined above. Especially as in this case, based on the informal semantics expressed in the scope note, the transitivity of *P9* seems to be less justified than its inverse functionality. From an implementation perspective it would be easier if the CRM defined characteristics, such as transitivity, symmetry and others, on a 'per property' basis as it is done in most implementation languages. This point will be illustrated in the following section where we discuss an implementation of the CRM provided by the University of Erlangen-Nürnberg.

3.4.2 Issues at the Implementation Level

Additionally to the difficulties in the abstract model, further or different restrictions may, perhaps accidentally, be present in implemented versions of the CRM. We illustrate our approach to testing the axiomatisations of properties with an example.

[40] Take the example of the relation *is parent of* on the set of parents.

An implementation of the CRM in OWL 1 DL, that we used at a time, contained the following statements about CRM relations:[41]

- Functional: *P22.transferred_title_to*, *P99.dissolved*
- Inverse functional: *P3.has_note*, *P57.has_number_of_parts*, *P81.ongoing_throughout*
- Symmetric: *P114.is_equal_in_time_to*, *P121.overlaps_with*, *P69.is_associated_with*
- Transitive: *P25.moved*, *P89.falls_within*

In order to see whether these stipulations make sense or not it is necessary to examine the scope of each relation and also its domain and range. In some cases the specifications appear to make sense. *P114.is_equal_in_time_to*, for example, is intended to express that two temporal entities have the same time span. Since it expresses an equality, the relation is obviously symmetric. *P121.overlaps_with* connects two 'places' that share some of their geometric extent. It is obviously symmetric as well. In the case of *P69.is_associated_with*, which connects two instances of 'E29 Design or Procedure', the intuition that it is symmetric is perhaps less clear. However, a look into the defining document of the CRM reveals that "this symmetric property describes the association of an E29 Design or Procedure with other Designs or Procedures. [...] The property is assumed to be entirely reciprocal."[25, p. 55]. It is therefore clear that modelling the property as (formally) symmetric is in line with the intended meaning. Finally, to give a last example of a convincing axiomatisation, *P89.falls_within* is a relation that connects two 'places', one of which is 'part of' the other in a mereological sense. As we have argued in the previous section for a similar `fallsWithin` relation, making this property transitive, is part of a good solution that enables one to reap the benefits of 'transitive propagation' when employing a reasoner.

In several of the above cases the rational behind the axiomatisation is less obvious, but can perhaps be justified with additional arguments. The relation *P22.transferred_title_to* connects an 'acquisition' to an 'actor', thereby expressing that the actor acquired whatever the acquisition instance is in turn linked to. By specifying that P22 is functional, any instance of the property can only point to one actor, i.e. any acquisition process can at most have one acquirer. The class 'actor' in CRM, however, also covers the case of 'groups' and therefore acquisitions which involve more than one buyer (e.g. a couple acquiring a work of art together) can be modelled by using a 'group' on the right hand side of this relation.

P99.dissolved is a relation that links a dissolution event to an instance of 'E74 Group', which is defined as a gathering or organization "of two or more people that act collectively or in a similar way due to any form of unifying relationship" [25, p. 29]. Making this relation functional entails, that any dissolution event can be the dissolution of at most one group. Plausible counter-examples may exist, however, in which a single event causes the dissolution of several groups: When a one-party system emerges, it is conceivable that the dominant political party dissolves all other parties by way of a single proclamation.[42]

---

[41] The implementation in question could previously be accessed at `http://www8.informatik.uni-erlangen.de/IMMD8/Services/cidoc-crm/erlangen-crm_090330_5_0_1.owl` (last accessed 14/07/2009). The closest version that is still accessible appears to be `http://www8.informatik.uni-erlangen.de/ecrm/erlangen-crm_090714_5_0_1_TQ.owl` (last accessed 1/08/2011). For more information on different CRM implementations see Section 5.3.5. Some of the aspects discussed in the following paragraphs were communicated to Bernhard Schiemann who was maintaining the CRM implementation at the University of Erlangen-Nürnberg at the time (see Appendix D).

[42] This example was given by Bernhard Schiemann in personal conversation (see Appendix D)

The property *P3.has_note* can be used to attach a note to any CRM entity. If it is made inverse functional, no two items can share the same note. However, notes in turn are connected through a data type property to a string that represents the text of the note and this data type property is not inverse functional, so that two distinct notes can have the same text (and even share the same string instance).

Finally, a third group of relations seem to have received an axiomatisation which is hard to justify and counter-intuitive. One may object to declaring *P25.moved* as transitive, purely based on technical considerations. According to the CRM definition, the domain of *P25* is 'E9 Move' and its range is 'E19 Physical Object'. However, 'E19 Physical Object' is declared to be an indirect subclass of 'E77 Persistent Item' and 'E9 Move' an indirect subclass of 'E2 Temporal Entity'. Furthermore, "E2 Temporal Entity is disjoint from E77 Persistent Item" [25, p. xvi]. Thus, 'E9 Move' and 'E19 Physical Object' are also disjoint classes. Then, however, the rule, that *P25* is transitive, will never apply, for there can be no individuals *a*, *b* and *c*, such that P25 connects *a* and *b*, as well as *b* and *c*, as this would entail that *b* is both an 'E19 Physical Object' and an 'E9 Move', contrary to the disjointness of these classes.

Another example is the property *P57.has_number_of_parts*, which is intended to express how many parts a multi-part physical object has. The example given in the CRM documentation is a specific chess set, that consists of 33 parts. Declaring *P57* to be inverse functional, entails that all chess sets (or other collections) that have 33 parts will be identical to each other. This is clearly undesirable.

P81.ongoing_throughout "allows a Time-Span's minimum temporal extent (i.e. its inner boundary) to be assigned an E61 Time Primitive value." [25, p. 58]. If this relation is made inverse functional, the consequence would be, that if two time-spans have the same inner boundary, they could be inferred to be the same time-span. This is counter-intuitive since the inner boundary specified by *P81* is subject to uncertainty about a time-spans true extent and should not determine its identity or non-identity with another time-span. Counterexamples can be constructed.[43]

The critical assessment of the relations discussed here and similar arguments for other relations were reported back to the University of Erlangen team, that was in charge of maintaining the CRM implementation in question. In response to this, about eight of the doubtful axioms were removed. In a later version, however, all relation axiomatisations of the above kind, even the ones that appeared to be justified were removed.

These examples, as well as the ones pertaining to the abstract definition of the CRM in the previous section illustrate the intricacies of correctly axiomatising relations or 'object properties', as they are called in OWL. They also illustrate the two extremes at either end of the axiomatisation spectrum: (i) too strong an axiomatisation, that leads to erroneous results and in the worst case inconsistencies, and (ii) an overly weak or minimalistic axiomatisation, that is adopted in order to avoid (i), but deprives the model of much of its inferential potential.

3.4.3 A Note on Modelling Image Subject Matter

Whether we talk about an image or look for an image online, the subject of the image is likely to be at the centre of our attention. Even in the metadata heavy field of the visual arts the subject, i.e. the conceptual content of an image, is still the most frequently employed category for searching.

[43] Let for instance, T_1 be a time-span that coincides exactly with a given day d. The inner (and outer) boundary of T_1 will then be d. Assume now, that T_2 is a time-span that includes T_1 and additionally, either the day before or the day after d, but (crucially) we do not know which of these two. In this case the best estimate for an inner boundary of T_2 that we can give is d, exactly the same as for T_1. However, we *know* that T_1 and T_2 are different, in fact we know that T_2 is twice as long as T_1.

This is illustrated by the results of our survey on art image retrieval (Section 4.2). Content-based image retrieval (cf. Section 2.4.1) attempts to climb the hill from low-level physical image features to image subject as humans would categorise it. For annotation-based approaches which do not rely on physical image content on the other hand, an effective modelling of image subject matter is perhaps the holy grail.

One of the problems with modelling image subject is that an image can depict an enormous range of objects, events, persons, emotional aspects and abstract ideas. A subject model with a degree of completeness therefore seems to amount to an encyclopaedic undertaking. This is evidenced by the size and effort that went into the compiling of Iconclass (compare Section 2.3.2).

As the most comprehensive classification system of image subjects in the visual arts, Iconclass is used by many heritage institutions.[44] But as a system of semantic image annotation it has its limitations. What appear to be conceptual asymmetries, for instance, that 'Bible' is a top-level category, 'Christian religion' located at the second level and 'Hinduism, Buddhism, Jainism' share a category at the third level, can sometimes be explained by Iconclass' focus on the Western tradition and could perhaps be dealt with. A more serious impediment from the point of view of semantic image content modelling is, that Iconclass does not attempt to capture semantic relations between depicted elements. The only relation the system knows is the hierarchical relation, which is not given a clear semantic interpretation as we have attempted to illustrate (Section 2.3.2). Iconclass categories encapsulate a complex meaning and the hierarchical system serves as a finding aid enabling a human user to discover a suitable category label (for both annotation and retrieval). Iconclass was not designed for, and cannot support, automatic inference on image subject matter.

In an attempt to model an area of a manageable size in detail and with attention to the relations a depicted object can enter in, we turned to the human figure as an example domain. The human figure has been the subject of depictions ever since the earliest rock art and many volumes have been published on the topic of human anatomy for artists. We collected terminology on (externally visible) human body parts from the table of contents of 14 textbooks on anatomy for artists or art students [88]. After the spelling and morphology of the terms were normalised and synonyms identified, we were left with 48 terms. These terms were entered as classes into an ontology of body parts and a concept of 'complete human body' was defined using these classes, a transitive part-of relation, and description logic restrictions.

Strict reasoning about visible elements in a two-dimensional image based on 'part-of' relations is a difficult task. This is due to the fact that three-dimensional objects appear as projections in which occlusion and especially self-occlusion can play a part: If a three dimensional object is visible in an image, this does not entail that all of its parts are visible as well. Still, our small model on human anatomy could possibly be used in a semantic recommender setup for image content: The query of a user who would be interested in a 'male torso' could be expanded by 'man' and 'navel', for example. For the following reasons, however, we did not pursue this line of research any further: First, the necessity of a detailed account of human body parts even in the case of images of the human figure may be doubtful. Only very few participants of our image retrieval survey submitted terms referring to the human body or its parts, although two of the images used in the survey showed a partially nude male figure. Second, a query expansion strategy, as opposed to precise inferences on image subject, can be successfully supported by a logically less rigorous approach (see e.g. [74]). The categories for

[44] An example which can be explored online is the photographic collection of art historical images *Bildarchiv Foto Marburg* at `http://www.bildindex.de` (last accessed 10/10/2011).

human body parts remained in our ontology, but played no part in the user evaluation of our system (Section 5.5). A detailed account of this aspect of our research has appeared as a publication [88].

Instead we have only addressed the following kinds of image subject descriptions in our system: the presence of historical persons in a depiction, the mode in which they are depicted and the ability to characterise aspects of their biography. These aspects are more accessible to ontological modelling and, as we learned from our interaction with experts, they are of great interest to art professionals. Perhaps such descriptions can also make a contribution to art historical interpretations that have in recent decades emphasised the social context and function of visual art as opposed to its iconographical content [9], [145].

3.5 Survey and Interview Methods

Having discussed the methods and approaches we applied in the development and modelling stage of our ontology, we will now describe the methodologies we have applied for our qualitative and quantitative surveys among domain experts and laypersons.

The methodological approach for the expert interviews (Section 4.1), the image retrieval survey (Section 4.2) and the feedback survey after system evaluation sessions (Section 5.5) is based on the module "Aspects of Survey Design" that was part of a Postgraduate Diploma in Statistics course in Trinity College, Dublin in the academic year 2007/2008. Unless stated otherwise, the following remarks on method are based on the course notes[45] and on the textbook by Dillman on the subject of survey design [32].

Two interviews were carried out with heritage professionals at the National Gallery of Ireland. The objective was to elucidate what is involved in curatorial work and to determine expert information needs with respect to image search, especially on the web. To cover the areas of interest we had identified, but at the same time leave enough room for the experts to give an account of their profession in their own terms, we decided on a semi-structured format for the interviews. This approach has the advantage of maximising the flow of information, while avoiding a completely unstructured approach, which places a heavier burden on the training and experience of the interviewer and has rarely been applied outside the field of sociology.

The goal of understanding curatorial work and a number of concrete research hypotheses (Section 4.1) informed the development of an interview guide. The guide was revised and sharpened after feedback from colleagues before the interviews were conducted (Appendix A).

With approval from the interviewees, the interview sessions were recorded and transcribed following accepted standards [109]. The interpretation of the interview results followed a mixture between bottom-up coding for the questions, which were concerned with general aspects of curatorial work and top-down coding for the questions which were derived from concrete research hypotheses. The second interview was less structured as it involved the demonstration of a content management system. In this case, different, sometimes dispersed, parts of the interview, had to be mapped to the questions of the interview guide they addressed, before an interpretation analogous to the first interview could be carried out.

For the two survey questionnaires research hypotheses and questions were formulated in accordance with methodological recommendations. These questions were translated into a set of survey questions. The first survey (Section 4.2), was designed as part of the Postgraduate Diploma course

[45]The course notes are no longer available online. An electronic copy of the notes can be provided.

mentioned above. The second survey was part of our system evaluation process (Section 5.5) which followed the pattern of a task-based evaluation involving questionnaires [116].

The ordering of survey questions in each case also followed recommendations on questionnaire methodology. In particular, questions concerning demographic information about the participants were asked at the end, while information on the expected duration of the survey was given at the beginning. These measures may have contributed to the high response rate in the first survey, which was anonymous and not researcher-administered.

For reasons of cost and practicality the sampling method for both surveys was non-random, combining quota sampling with elements of purposive and snowball sampling. This was deemed acceptable, because of the exploratory nature of our enquiry and the absence of a naturally given statistical population and sampling frame. Our approach is comparable to the methods employed in similar studies, e.g. [19], [66], [73].

For the coding of the results of the first survey a framework proposed and tested in the literature was used [73]. The framework is a synthesis of several other categorisations, thus maximising comparability to similar studies reported in the literature. The evaluation survey was more structured, comprising mostly of multiple choice questions. The few open-ended questions were coded in a bottom-up fashion, grouping similar responses together.

3.6 Summary

We surveyed and discussed methods and methodological approaches that we used to create concrete representations of knowledge and conceptualisations in our domain. Our discussion ranged from representation standards and enabling technologies to strategies and recipes for ensuring that the conceptual structure of such representations accurately reflects the domain for a given purpose.

It appears, that in the field of semantic representations, their creation and reuse, canonical methods, comparable to those in other science and engineering disciplines, have not yet been established. Due to differences in opinion, which we outlined in our literature review, both of a philosophical nature and with respect to heritage specific questions, it is possible that such canonical methods may never be established. On the other hand, the field of ontology modelling for the Semantic Web, and in particular for the 'Semantic Heritage Web' is still an emerging one. This is evidenced by the successive standards of ontology languages that have only recently begun to stabilise and by the shifting landscape of cultural heritage ontology implementations that are built on these standards. Textbooks on the use and best practices regarding modelling constructs new to the latest version of the Web Ontology Language have yet to appear.[46]

Like in other fields, the methodological approach in semantic heritage modelling, has an underlying trade-off between cost and effort on the one hand and effectiveness and accuracy on the other. We will close our discussion by contrasting what appears to be a practicable compromise, with what we consider an ideal method.

With respect to knowledge representation formalisms, the clearly defined semantics of formal logic offer an ideal reference point. With the move to description logics, implemented knowledge representation languages can now be seen as a syntactic variant of a theoretically well-understood logic. The decidability of description logics does, however, not guarantee that reasoning will effectively terminate. Performance issues may therefore have an influence on modelling decisions.

[46] Compare the section on transitive propagation.

Research into reasoning algorithms is ongoing and improvements can be expected from attempts to combine new reasoning techniques with large computing power.[47] In general, however, the knowledge representation tools that are available do not leave much to be desired.

In our discussion of ontology development methods we have already pointed out that a controlled comparison of different approaches in a real life scenario may be desirable, but is practically infeasible. Instead, it would seem that with the degree of standardisation at the technological level achieved, ontology development methods should expand from the treatment of conceptual and sometimes philosophical questions, to incorporate concrete modelling practices and ontology design patterns, analogous to ones in object oriented design. Such guidelines may lead to a discipline of 'ontology engineering' that lives up to its name. First steps in this direction have already been taken [67].

In the section on transitive propagation, we illustrated the deliberations that can inform modelling decisions. In practice, such deliberations are often cut short and a modelling solution is accepted once it fulfils the basic requirements of the context in which it arose. If unwanted side effects come to bear later in the development, this can lead to remodelling. Until best practices for many concrete problems are established in ontology engineering, a carefully designed top-level ontology may also provide some guidance in this context.

The CRM is perhaps the most ambitious top-level ontology for the field of cultural heritage. Developed bottom-up from a large number of database schemata and documentation structures, in a process of discussion and negotiation between domain experts and knowledge engineers, it combines the information needs of the sector with principled ontological modelling. In our extensions of the CRM we have attempted not to violate the *compatibility* constraints, as defined in the CRM reference document (cf. our knowledge modelling in Section 5.2). In particular, our extension is mostly a 'downward' extension of CRM concepts and properties and in cases where we introduced superproperties of CRM properties we observed the restriction of *query containment*, that is we utilised our properties only for querying the knowledge base, not for marking up data. Our experience, however, has shown that even a thoroughly developed model, such as the CRM, may need to be adapted in project contexts. As we hope to have demonstrated, the formal stipulations associated with a number of CRM elements appear to be barely consistent with the intended semantics. This is by no means a fatal flaw as the flexibility of ontology languages affords one the possibility to easily alter characteristics for local use.

In the long run, however, we hope that legitimate and well-argued cases for modifications from users of the CRM will eventually be considered in the official model, just like some of the suggestions we have made have found their way into implemented versions. An ideal approach to the curation of a top-level model, such as the CRM, may not only be independent from implementation languages, which the CRM is, but perhaps also separate precisely defined formal semantics on the one hand from an accessible documentation that lends the system its usability on the other. In such a model, the formal specification of the semantics would be normative and the documentation would have informative character. The advantages of a formally defined model are that its properties can be checked for inconsistencies and implausibilities with formal methods and that the intended interpretation of individual model elements can be illuminated by reference to sound semantics. Such a formal specification may also disregard implementation related considerations and include a maximal set of consistent information, such as a complete statement of disjoint classes and properties and in-

[47] `http://www.larkc.eu` (last accessed 17/10/2011).

verse properties as 'first class citizens', not mere labels. Parties who implement the CRM in various knowledge representation languages, could then effectively create a logically weaker version of the CRM and justify which aspects are included by reference to intended application scenarios or the particulars of the representation language. Such a top-level model would ideally be complemented by auxiliary domain ontologies. GeoNames is a good example and we hope that similar resources will be established for other domains. A knowledge base of historic events may be useful, which treats births, deaths, marriages, festivities and duels of historic persons as well as GeoNames treats geographical entities.

As far as image subject matter is concerned, an ontological approach that is not bound by economic constraints should perhaps start with a remodelling of Iconclass. The classification system is widely used and a mapping of Iconclass tags into an ontological framework would create a rich resource for querying and exploration. Such a remodelling should not stop at a recasting of Iconclass in a different syntactic framework, but amount to a full reengineering that structurally transforms Iconclass from what it is now, a finding aid, to a system of semantic image content descriptions that can be weakened to simpler descriptions or with additional information could be used to derive new and more complex descriptions of an image. It is beyond doubt, that such a transformation of Iconclass, or rather a mapping of Iconclass into an ontology of equal or greater complexity, would be a daunting enterprise, which would have to involve knowledge engineers as well as domain experts familiar with the Iconclass system.

Chapter 4

Conceptualisations of Art

An often quoted definition describes 'ontology' as "an explicit specification of a conceptualization" [56, p. 199]. A 'conceptualisation' in this context is an abstract and informal description of a certain domain. As a mental construct it may help people to conceptually organise and 'make sense of' a domain. In order to learn more about conceptualisations in the domain of visual arts, we have approached heritage specialists and academically trained art historians. We will first report on two interviews, which we conducted with heritage professionals, and then on a survey which contrasts the conceptualisations of art specialists with those of laypersons.

4.1 Expert Interviews

In order to improve our understanding of how domain experts approach works of art, what categorisations they apply and, to a lesser extent, what expectations they may have of computer-aided search and cataloguing techniques, we have conducted two interviews with employees at the National Gallery of Ireland. The interviews were concerned with the everyday tasks of curation and information management in a gallery, personal and professional preferences with respect to image search and the information typically recorded for heritage artefacts.

In accordance with the methodological considerations outlined in Section 3.5, we developed a guide for semi-structured interviews to address the following research questions about curatorial work:

(i) *Which information about visual art is recorded and for what purpose?*

(ii) *Which sources and systems are used to find information on works of art?*

(iii) *What strategies are employed for online image retrieval?*

(iv) *How can visual art information retrieval be improved?*

Additionally to these research questions, the research hypotheses for the survey, which was conducted shortly after the interviews, had an influence on our interview questions (Section 4.2.1).

Two expert interview partners were chosen, after consultation with a National Gallery curator who was already involved with our project, to represent on the one hand, the business of classical curation, and on the other, what one might call the emerging business of "digital curation". The version of the interview guide that was used for both interviews is included as Appendix A. Complete transcripts of the two interviews are provided in Appendix B.

At the time of the interview, the first interviewee (in the following referred to as 'expert A') was holding the position of "Assistant Prints and Drawings" in the National Gallery of Ireland. This position is part of the curatorial department and expert A was, in her own words, "acting as an assistant curator to [...] the prints and drawings collection" (Appendix B, p. 143). Previously she had held similar positions at museums in the USA and Italy. Her academic training comprised a first degree in art history and English, a masters degree in art history and an arts administration course.

The second interviewee ('expert B') was an assistant to the Centre for the Study of Irish Art, which is associated with the National Gallery. Previously, he had been employed in the education department of the same gallery. His academic background consisted of a joint degree in fine art and history of art and a masters degree in multimedia. His responsibilities included preparatory work for a new Web appearance of the National Gallery.

Both interviews were held on the 7th of May 2008 on the premisses of the National Gallery of Ireland. They differed slightly in nature: The first one, with expert A, closely followed the interview guide (Appendix A), while the second one, involved practical demonstrations of digital collection and information management systems by the interviewee and was more loosely based on the interview guide. The order of questions for expert B differed from our guidelines and some of the aspects were covered only indirectly by his explanations. Generally, however, both interviews covered the same topics. In the following, we will characterise the essence of the interviewees' statements with respect to a number of subjects. Literal quotes are taken from the interview transcripts in Appendix B and referenced by page number (in this thesis).

Aspects of curatorial work. Asked to describe the most important aspects of the profession of a curator, expert A names two different facets. First, the organisation of several exhibitions per year, which in turn need to fulfil two criteria: they need to be "art historically sound, under a certain theme or concept" (p. 144) and also "appealing to [...] the general public that comes to the gallery" (ibid.). The second facet comprises "increasing knowledge on the collection, [...] constantly updating our dossier files or [...] information on specific works of art" (ibid.). In this second process external researchers play an important part, for example, in the reattribution (i.e. reassignment of authorship) of works of art.

Expert B also touches on these two aspects of curatorial work, albeit not as explicitly. In his case, the two facets of curatorial activity were shifted to a digital context. Instead of preparing temporary, physical exhibitions, his responsibilities involved the preparation of a quasi-permanent 'online exhibition' as part of the newly designed web appearance of the gallery. This will enable the general public to access digital reproductions of works of art and archival materials which are hard to access otherwise. In this context he emphasises the importance of a proper contextualisation of works of art and their sometimes complex relationships to other works. As an example, he explains in detail one of the most complex items in the collection, a scrapbook from the 19th century which contains sketches and watercolours from several different artists and biographical notes and letters on one of the artists collected by an antiquarian. A custom made collection explorer system facilitates access to items within the scrapbook that are part of other items (e.g. sketches stuck onto a page) and allows all items to be viewed in their context, revealing, for example, content on the back of items, which may be by a different author, etc.

In terms of the second curatorial task — increasing knowledge on the collection — expert B mentions two examples, where the location that a work depicts, in one case a seascape painting and in the other case an old photograph, could be revealed using the Google search engine and in particular the Google Earth service: "This photograph of the collection has been in the gallery for about 20 years and nobody has ever identified where it was and [...] just 15 minutes sitting down with Google and we where able to find out where it was." (p. 168). In each of the two cases, new insights into an artist's biography were gained.

Information gathered by curators. The attribution of a work is, according to expert A, the most important piece of information that is recorded. Additional information recorded includes: "pigments used or the support the work is actually done on, [...] what type of paper it is and can we date the paper [...]. So all of that information about the media will be important to the file as well. And then any articles [...] that the work has been included in we would [...] insert that into the file, too. [...] Anytime it's referenced either in a scholarly article or whether it's appeared in an exhibition abroad or whether it's travelled." (p. 145). At a later point, expert A again emphasised the provenance of a work and also "any previous titles [...] the date of the work, which [...] can often change" (p. 148) and "whether it's a preparatory work for an oil within our collection [...] under [the] heading [...] related objects" (ibid.)

Expert B demonstrated the information typically kept by way of pointing out the different fields in the gallery's digital archival system: "things [...] that physically describe the item, its media, its dimensions [...] things like the collection that its part of, its provenance, where it came from [...] the access that's allowed to it [...] what collection it's in [...]. For some of the longer letters there is also transcriptions and there is a technical record, just so we have [...] a note of [...] what it was scanned with, how the digital version was created, when it was made, what the resolution was and [...] the collection its part of." (p. 154). An important emphasis for him is on "related items" for which the National Gallery had a custom feature added to the archival system, to hierarchically structure and virtually group items in the collection.

Information gathered on image subject. According to expert A, detailed information on the subject or content of an image was only kept in the past: "in the past you can see that they have written down absolutely in detail [...] what exactly is in the work. There were horses to the left, [...] the rider is to the right, there is a mountain in the distance and the overall colour is green [...]. So it was very, very detailed. We tend not to do that as much now." (p. 145)

The reasons given for this lack of content recording today include the availability of photographic records, greater economy in record keeping and lesser need for image content descriptions for the purpose of exhibitions: "what we try not to do now is basically detail exactly what's in the work cause we feel that people can see that themselves" (p. 146). Salient or unusual features would, however, be pointed out to the exhibition public and the older, detailed, hand-written descriptions, as far as they exists, are kept in the files.

Expert B addresses image content in the form of a subject hierarchy that constitutes one of four different ways in which the gallery's archival records can be searched in the future. The hierarchy includes subsumption relationships like 'objects' – 'furniture' – 'bed' and is expanded as the need arises (i.e. as new items are catalogued corresponding categories are added to the hierarchy). It was based on the subject hierarchy of the Tate web site.[1] Apart from subject based access, 'Art Search', the front end of the archive management system, allows users to access works by artist/author, by

[1] `http://www.tate.org.uk/servlet/SubjectSearch` (last accessed 16/10/2011).

keyword based search or through "browsing by collection".

Systems and sources of information currently in use. Additionally to the National Gallery's archive management system 'Art Search', expert B mentions 'Art Cat' a library management system for secondary sources (books and journal papers on art) and a number of online resources. Artnet[2] offers a subscription based online database of past and upcoming auction sales worldwide. It is, according to expert B, particularly useful to find lesser known and less published works by a certain artist. The Oxford Dictionary of National Biography[3] provides biographical records of "biographies of the men and women from around the world who shaped all aspects of Britain's past".[4] Finally, the National Irish Visual Arts Library[5] (NIVAL) has a collection that specialises in late 20th century and contemporary art and offers free access to an artist database on their Web page.

Expert A explains TMS, 'The Museum System', a commercial software that is used by the National Gallery for its collection management.[6] The system facilitates access to collection items for the curators in the gallery. It enables them to manage records on collection items, including textual information and allows them to create virtual collections of items in preparation for an exhibition. From a registrar's point of view, TMS functions as "a kind of a tracking system" (p. 147), which allows to locate works within the gallery and works loaned to other galleries and registers their insurance value and related information. In terms of web-based information retrieval, expert A mentions Google search, but also ULAN by the Getty Research Institute (cf. Section 2.3.1), Art Index[7] for sales and sales prices of works and an online database of the National Library of Ireland[8], which allows access to works held by other Irish institutions, often more efficiently than through the institutions' own homepages.

Vision of ideal art information retrieval. Asked about the shortcomings of current image retrieval processes, expert A laments the heterogeneity of available information sources and display structures. This necessitates exploring a multitude of web pages after an initial Google image search in order to find detailed information on each image of interest. However, she relativises her criticism, by saying that unsatisfactory image retrieval results may possibly be "a problem of us putting in very basic [...] query information [...] from the offset" (p. 152).

An ideal retrieval environment for images of fine art would, in her opinion, involve a "very unified search engine for these types of images" (ibid.), where results would "come up on one [...] image site" (ibid.) which presents "relevant information [...] as an adjunct to the thumbnail you're looking at" (ibid.). The examples that expert A uses to illustrate her vision of art image retrieval suggest an expectation of complete and correct result sets that are assembled from distributed sources: "In an ideal world I'd love to be able to [...] key in [...] 'Degas' drawings, search globally' and all of Degas' drawings [...] come up in addition to where exactly they are and the relevant information for each drawing." (p. 151). And similarly: "for instance Picasso has [...] images of [...] the 'Ballets Russes' [...] whether they be paintings, drawings, ceramics [...] and they're in the Hermitage, they're in

[2]`http://www.artnet.com` (last accessed 27/07/2011)

[3]`http://www.oxforddnb.com` (last accessed 27/07/2011), subscription-based.

[4]`http://www.oup.com:80/oxforddnb/info/quickguide` (last accessed 27/07/2011).

[5]`http://nival.ncad.ie` (last accessed 27/07/2011).

[6]Compare the TMS homepage at `http://www.gallerysystems.com/tms` (last accessed 16/10/2011).

[7]`http://www.artindex.com/home/index.asp` (last accessed 16/10/2011).

[8]It would appear that expert A referred to the prints and drawings catalogues and databases section of the National Library at `http://www.nli.ie/en/prints-and-drawings-catalogues-and-databases.aspx` (last accessed 16/10/2011).

Louvre, [...] they're in four museums in the States and you would be able to [...] have them all as a package." (p. 152).

Expert B, on the other hand, emphasises retrieval speed and quick access to information as key priorities. It appears that according to his opinion the basic information infrastructure is in place and that all that is required is faster responses to the questions a user may have. After further enquiries he clarifies that this does not primarily refer to faster connection times, but to a more efficient way of query answering which may involve improvements in subject access and keyword annotations.

4.2 Image Retrieval Survey

We discussed literature covering user conceptualisations and user terminology as it is employed in image search and in particular art image search as part of our literature review (Section 2.2.4). We saw that research in this field is difficult to compare and some results seem to point to contradictory tendencies. Moreover, although an influence of domain expertise on user behaviour has been postulated in general [73], little is know about possible differences between art specialists and laypersons with respect to the terminology they employ for image search.

We attempted to better understand the conceptualisations people have of aesthetic images through their choice of keywords during image retrieval. To that end we devised and conducted a survey which we will report on in this section. Part of the research presented here was published as a conference contribution [87].

In the following sections we will cover the survey design and set-up and give examples of typical answers and feedback we received, as well as an overview of the demographic profile of our participants (Sections 4.2.1 and 4.2.2). After that, we will explain our method for analysing the submitted query terms and report the survey results (Sections 4.2.3 and 4.2.4).

4.2.1 Survey Design

Our aim was to collect and analyse search terms submitted by participants in an image retrieval or simulated image retrieval setting for a set of images typical of our domain. Additionally, we wanted to probe the possible influence of specialist knowledge by collecting terms from participants of varying degrees of qualification in art history or related subjects.

Based on the methodological considerations outlined before (Section 3.5), we developed and conducted a survey. The following three research hypotheses guided the development of our survey:

(i) *Terminology applied by specialists for searching images of fine art is different from terminology applied by laypersons.*

(ii) *Target search (cf. [139]) is the most common image retrieval mode in the domain of fine art images.*

(iii) *Few people are aware of 'query by sketch' or 'query by example' approaches to image retrieval.*

The first hypothesis is central to our study, the other two are auxiliary. The difference in terminology mentioned in hypothesis (i) does not refer to a literal difference in the terms employed, but to a difference in the *type* or *kind* of terms employed with varying degrees of frequency. The coding system we used to classify terms is explained below (Section 4.2.3).

Page	Title/Topic	No. of Questions
1	Welcome	-
2	Experience with web search engines	3
3	Image retrieval goals for fine art images	1
4	Image retrieval - alternative methods	2
5	Querying for images of fine art	-
6	Image 1 (Pollaiuolo)	-
7	Image 1 - Questions	3
8	Image 2 (Schiele)	-
9	Image 2 - Questions	2
10	Image 3 (Monet)	-
11	Image 3 - Questions	2
12	Background in art history	2
13	Demographic questions	4
14	Thank you	-

Table 4.1 – Outline of the web-based survey: the page titles and number of questions presented on each page. The pages which contained the images and a number of pages with navigational function did not contain any questions.

The survey was implemented and conducted using an online survey service.[9] Apart from the ease of implementation and delivery to participants, an online survey appeared to be well suited for our context of keyword based image search.

Two kinds of invitations to take the survey were sent out by email:

- a general invitation to friends and colleagues
- a number of special invitations to individuals or small groups with a background in art history

The latter group contained an encouragement to send the invitation on to colleagues. After two days this was followed up by an email thanking those who had already taken the survey and encouraging further responses especially from male participants and from participants with an advanced background in art history, as these groups were underrepresented at that point. While the two types of invitations were intended to generate a response from potential laypersons and potential experts respectively, the judgement of participant status for our analysis later on was solely determined by self-reported expertise in art history (see Section 4.2.4).

The survey contained a total of 19 questions that were thematically grouped onto 14 pages (Table 4.1). Six of the survey questions addressed the topic of the image retrieval task (i.e. participants' experience and preferences in image retrieval), seven covered image queries and the final six covered demographic details including questions about the level of expertise in art history or related fields. Additionally to the questions that were directly motivated by the research hypotheses above, we included a set of questions about participants' experience with search engines and image search (survey page 2). Figure 4.1 shows one of the questions as it appeared in a web browser (the image is cropped to remove advertisements at the top and bottom that would have been visible to the participants while taking the survey). Appendix C contains a complete set of screenshots of the survey pages.

[9] `http://www.esurveyspro.com` (last accessed 20/07/2011).

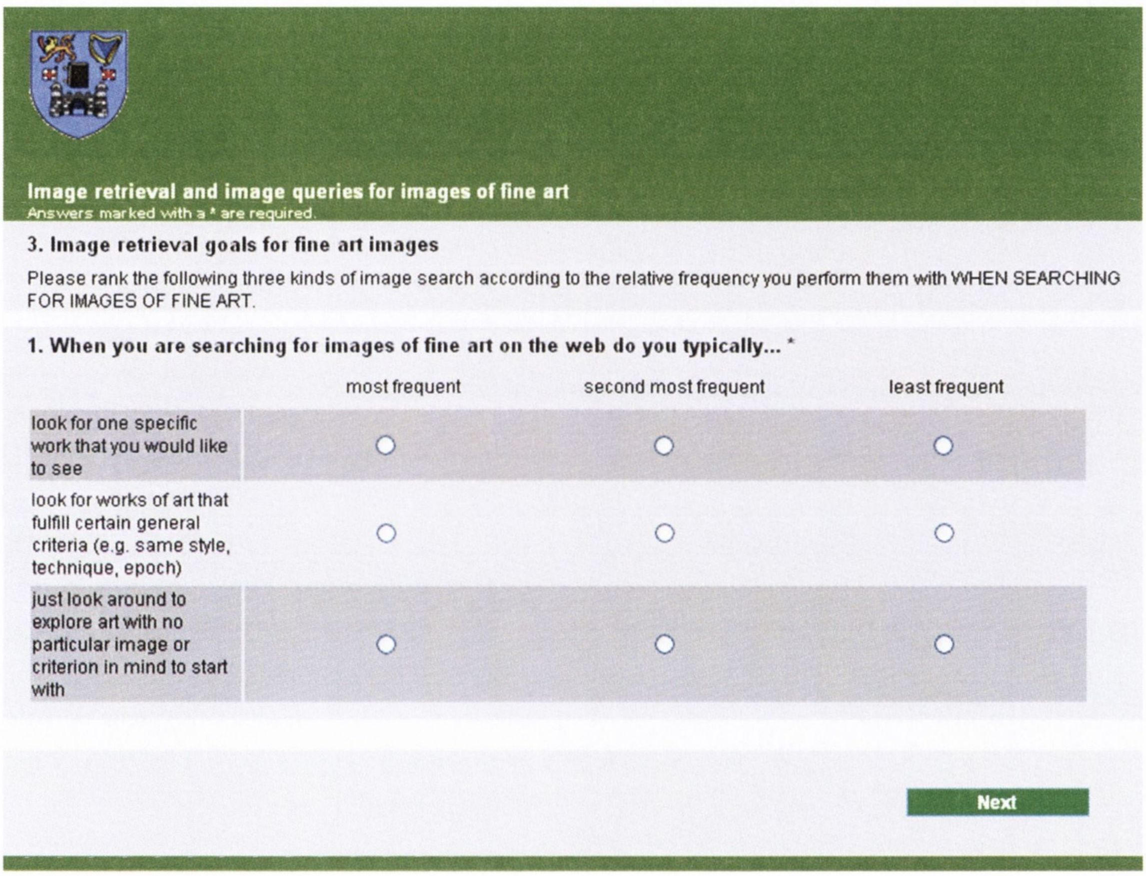
Image retrieval and image queries for images of fine art
Answers marked with a * are required.

3. Image retrieval goals for fine art images

Please rank the following three kinds of image search according to the relative frequency you perform them with WHEN SEARCHING FOR IMAGES OF FINE ART.

1. When you are searching for images of fine art on the web do you typically... *

	most frequent	second most frequent	least frequent
look for one specific work that you would like to see	○	○	○
look for works of art that fulfill certain general criteria (e.g. same style, technique, epoch)	○	○	○
just look around to explore art with no particular image or criterion in mind to start with	○	○	○

Next

Figure 4.1 – User survey: Question 3.1. Participants were asked to rank three different image retrieval goals according to their own experience with them. Appendix C contains a complete set of screenshots of the survey pages.

(a) Antonio Pollaiuolo, *Hercules and the Hydra* (c. 1475), Galleria degli Uffizi, Florence.

(b) Egon Schiele, *The Prophet* (1911), Marlborou.-Gerson Gallery, New York.

(c) Claude Monet, *Road at Louveciennes, Melting Snow, Sunset* (1870), Private Collection.

Figure 4.2 – Images of paintings by Pollaiuolo, Schiele and Monet used in the web-based survey. Participants were not given any textual information on the images.

The questions on experience with image search and image retrieval were placed at the beginning in order not to bias participants' answers to these questions by the retrieval tasks set later in the survey. This is particularly true for the question illustrated in Figure 4.1. All the retrieval tasks we set in the survey fall into the category of looking for one specific work (the first of the three choices presented).

In the section on image queries, participants were shown digital images of three paintings. The paintings chosen for the survey were *Hercules and the Hydra* by Antonio del Pollaiolo, *The Prophet* by Egon Schiele and *Road at Louveciennes* by Claude Monet (Figure 4.2). These images were selected to represent some of the diversity found in paintings of the Western tradition. They range from the 15th to the 20th centuries, from figurative to landscape art, exhibiting some of the characteristics of Renaissance, Expressionist and Impressionist art respectively.

Underneath each image a written instruction to participants read:

> *Please look at this image for a few seconds and then move on by clicking on the "Next" button.*

While participants were asked to dwell on the image only for a "few seconds" it was ultimately their own decision to move on to the next screen and some of them may have examined the images for longer. Participants were subsequently prompted to enter keywords into a free-form text box by the following question:

> *What keywords would you use if you were looking for this image online?*

The retrieval task we set to participants falls into Smeulders et al.'s category of *target search* [139]. Moreover, the retrieval task is located at Fidel's *Objects Pole* [44], because the image itself is the object of interest and not, for example, information contained within the image. In a real life scenario users may of course go on to extract information from such an image after they have retrieved it (e.g. information on the nature of Monet's brush strokes). This appears to be the reason, why Fidel names the retrieval patterns of art historians as an example of retrieval that is typically situated between the *Data* and the *Objects Poles*. However, she emphasises the object nature of the initial retrieval: "to make [an] inference, [the art historian] wants to retrieve all images, all objects, and each image must be viewed as a whole, as an object" [44, p. 190]. In our opinion, target search and retrieval at or near the Objects Pole is relatively common in the domain of fine art images.

Once participants had moved on from a page where an image was displayed (Appendix C, Figures C.6, C.8, C.10) to a page where they were asked to enter keywords into a free-form text box (Appendix C, Figures C.7, C.9, C.11), they were not able to go back to see the image again. This was done deliberately to prevent people from 'reading off' visually salient but perhaps trivial features of an image, as the image would presumably not be physically present in a real life retrieval situation. Hollink et al. describe this pattern as "bias due to visual cues in the image" [73, p. 616] and have opted against using test images in their study on user image queries. Instead their participants were provided with three sample texts and then asked to create a mental image that could be an illustration of this text. Image descriptions and search terms were provided with reference to these mental images. However, the authors report that 31% of the collected description or query fragments were "copied directly from the [provided sample] texts" [ibid.]. It would appear that the visual bias was traded off for a textual one.

In order to avoid this, we have chosen the approach described above. By not allowing test persons to look at the images when they were prompted for keywords we attempted to simulate a situation in which participants had to describe images from memory rather than read off visual detail. In the case of the second and the third image, where participants may have expected a prompt for keywords, we included 'blanking' questions such as "Have you seen this image before?". These questions were not used in our analysis, but served to briefly distract participants before they went on to enter search terms.

Following established principles for survey design (cf. Section 3.5), our survey concluded with a section in which participants were asked to provide demographic information about themselves.

4.2.2 Survey Response and Participant Demographics

We received a total of 61 responses to our survey through the web-based platform, of which 48 responses were complete (i.e. participants finished the entire survey, filling in all *required* answers). As only 52 emails were sent out initially, it appears that we had a good response rate. One reason for the high response may be that the invitation to take the survey "advertised well for it" as one participant put it. Additionally to the encouragement to take the survey in the email, the first page of the survey contained a statement that reassured participants that no "long winded descriptions" are expected of them (Appendix C, Figure C.1). Other participants commented positively on the survey's layout and ease of use.

Among the feedback from participants there were also some critical comments. One person said: "i tried to participate in the survey, but the consistency checks do not allow this: they do not permit a response for which two categories are equally least frequent by virtue of being 'never'". This comment referred to question 3.1, the matrix about image retrieval goals (Figure 4.1). Another participant answered the same question, in his own words, "incorrectly and then got stuck in a loop". These problems occurred, because a restriction was implemented that allowed participants to tick each column only once, effectively requiring participants to bring the specified retrieval goals into a strict order of preference. Apart from these two comments however, no other problems were reported.

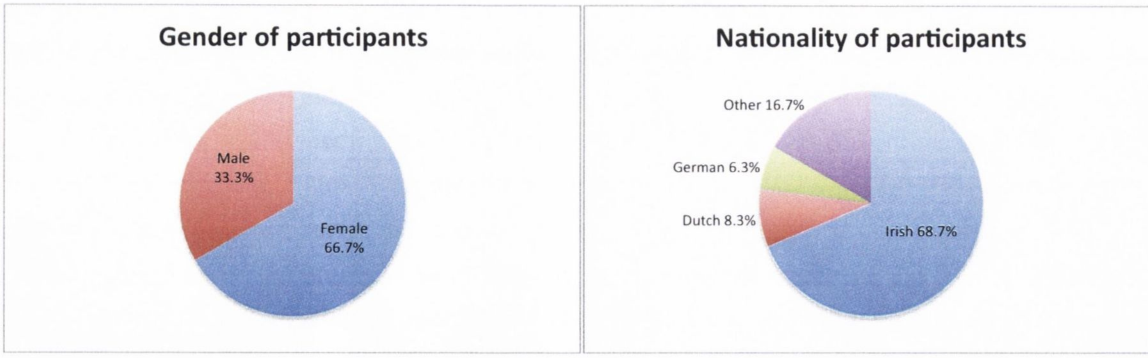

(a) (b)

Figure 4.3 – Participants' gender and nationality.

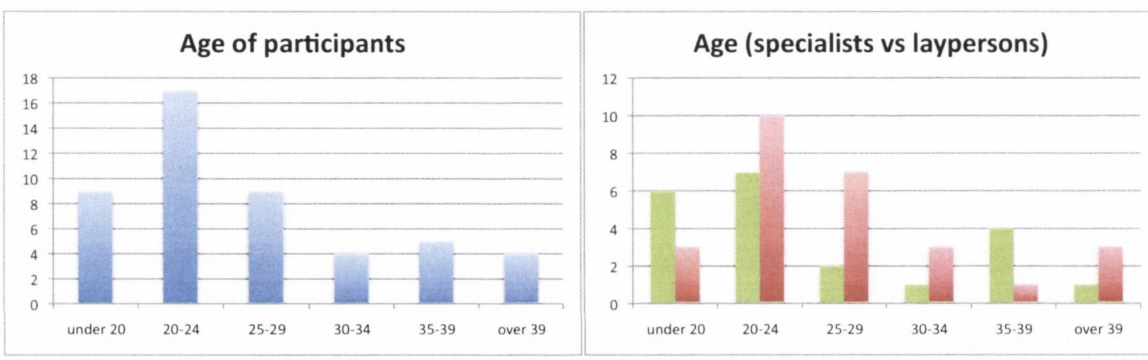

(a) (b)

Figure 4.4 – Age distribution of participants in general (a) and age distribution of specialists (green/left) compared to age distribution of laypersons (red/right) (b).

In at least one case, a participant did not complete the survey in the first attempt, but returned to it later to start over from the beginning and submit it as a separate response.[10]

For the analysis of the collected data only the 48 complete responses were used. Participants were given no incentive or encouragement to fill out the survey more than once and the set of demographic details (age, gender, nationality, level of education and personal interest in visual arts) were not identical for any pair of completed responses. We therefore assume that the answers were given by 48 distinct respondents and we will henceforth speak of participants rather than individual responses.

According to the demographic information provided, the survey was taken by twice as many female as male participants (Figure 4.3a). This was despite an explicit call for male participation in the second invitation email. The majority of participants were Irish nationals (68.7%), followed by Dutch, German and US American. In total, the survey was taken by participants of 10 different nationalities. Four participants were from a non-European background, but with one single exception all nationalities suggested a background in Western culture (Figure 4.3b).

The educational profile of the respondents showed that all participants had completed secondary

[10]The participant revealed this in personal conversation. There is no way of ascertaining that two responses were submitted by the same person or were not submitted by the same person based purely on the online tool.

(a) (b)

Figure 4.5 – Level of education among participants in general (a) and college level education in art history or related fields (b).

school, eleven were holding a third level degree, eleven a postgraduate degree and nine a Ph.D. (Figure 4.5a). At the time of the experiment, 14 participants had received university level education in art history for at least one academic year. Another seven participants had already completed such an education. Two in the latter group were holding a Masters degree and one a Ph.D. in art history. These 21 participants, who had either already acquired a certain expertise in art history or had chosen a specialisation that may eventually lead them to such expertise, are in our subsequent analysis referred to as 'specialists'. The remaining 27 participants either had never received a formal education in art history or a related subject (16 cases) or had only read it as a Leaving Certificate elective[11] or at an equivalent level (11 cases) (Figure 4.5b).

The age of participants ranged from 18 to 60 years, the average age was 27.2 years (median = 24, mode = 19, SD = 9.1)[12] (Figure 4.4a). At 25.9 years on average, our specialist group (median = 23, mode = 19, SD = 7.9) was slightly younger than our non-specialist group (mean = 28.2, median = 25, mode = 24, SD = 10.0) (Figure 4.4b).

The level of formal training in art history corresponded roughly to participants' self-reported interest in visual arts. 26 people suggested that they are more interested in visual art then most people but only five of those described themselves as "quite an expert" in this field. This option was intended to give laymen with no or little formal education in visual arts a chance to express that they have acquired a considerable knowledge by means of autodidactic study, for example. But all the people who rated themselves as "quite an expert" were in fact holding Bachelors, Masters or Ph.D. degrees in history of art or a related subject. By way of checking the plausibility of the answers we found that all college students or graduates in art history described themselves as more interested in visual art than average with only two exceptions. One college major in an art related subject stated that she had "no particular interest in visual art", which seems to be implausible. Either the person made a mistake or perhaps understood the question differently than laymen, with a different contrast set.

In terms of their experience with Web search, participants were asked how often they use search

[11]The Leaving Certificate Examinations are the final examinations in the Irish secondary school system. They are broadly comparable to A-Levels in England, Northern Ireland and Wales.

[12]While we have no reason to believe that our participants' ages come from a normal distribution, we report standard deviation as this appears to be common practice in similar studies (compare, for example, the guidelines of the American Psychological Association at `http://my.ilstu.edu/~jhkahn/apastats.html`, last accessed 13/10/2011).

engines in general and how often they use search engines for image and fine art image retrieval in particular. 79.2% said they are using web search engines like Google and Yahoo "nearly every day". Exactly two thirds of the participants are using image search with a frequency of between once a month and twice a week. 37.5% are searching for fine art images on the web with a frequency of once a week or more. 16.7%, on the other hand, said that they were never looking for images of fine art online. Since the questions proceeded from general (use of search engines) to specific (use of search engines for fine art image retrieval) one would expect no increases in the frequencies reported by each individual participant. Checking for inconsistencies of this sort we found such an increase in only one case. Due to the small number of inconsistent or implausible answers, however, none of the responses was excluded from the general analysis.

4.2.3 Analysis of Query Terms

The answers to questions 7.1, 9.2 and 11.2 of the survey,[13] in which participants were asked to submit search terms for each of the three images, were 'required'. Therefore each of the 48 persons who had completed the survey, had necessarily supplied three sets of keywords. In the case of the second painting (*The Prophet* by Egon Schiele), however, one participant submitted the string "I do not know". This was excluded from the analysis. Hence, a total of 143 keyword sets were recorded, ranging in length from one to 14 words.

Many participants structured the keywords they submitted with punctuation marks or line breaks, often to delineate multi-word terms. Examples of keyword sets submitted for each of the three paintings by specialists and non-specialists are shown in Table 4.2.

For the analysis of the search terms provided by participants we followed the methodology of Hollink et al. [73] (cf. Section 2.2.3). By using the unified framework developed there, our results become comparable to this study, but also to other studies, which use classification schemata that focus on a subset of the categories presented by Hollink and colleagues.

The query terms provided by survey participants were split into fragments following the same guidelines that were used by Hollink et al.[14] These guidelines cover both rules how descriptions or sets of keywords should be split into fragments and rules how those fragments should be classified into categories. Since our participants did not use free text descriptions, splitting the queries was mostly straightforward. Only in cases where participants had not separated their keywords by punctuation marks or line breaks, did we apply the guidelines.

The gist of the guidelines, as far as splitting keyword sets is concerned, can be summarised as follows:

- Separate words are separate fragments with some exceptions:
- Adjectives are classified together with their nouns unless the adjective indicates colour, shape, or texture.
- Words specifying an amount (e.g. numerals) are not classified separately.
- Prepositions that can be replaced by "and also" are not separate elements.

[13] That is the first question on page 7 and the second question of page 9 and 11 of the survey (cf. Appendix C, Figures C.7, C.9 and C.11).

[14] `http://www.cs.vu.nl/~laurah/guidelines.pdf` (Last accessed 23/07/2011).

Image	Keywords submitted	Participant information
Antonio Pollaiuolo *Hercules and the Hydra*	Man, battle, lion	Female, Norwegian, 22 years, third level educated
	Hercules, Hydra, Eurystheus twelve labours	Female, Slovak, 28, Ph.D.
	Pollaiuolo, Renaissance	Male, Irish, 35, M.A. in art history
	Hercules, slaying, dragon, club	Male, Dutch, 36, Ph.D. in art history
Egon Schiele *The Prophet*	Freud, dark, modern, male, nude, spine, torso	Male, Irish, 43, Ph.D.
	scary, death, klimt, dark, skeletal, men	Female, Irish, 34, Ph.D.
	cubist woman	Female, Irish, 18, College major in art history
	Egon Schiele, Austrian Expressionism, psychological angst in art, the nude in art, figure drawing	Female, Irish, 32, M.A. in art history
Claude Monet *Road at Louveciennes, Melting Snow, Sunset*	monet, impressionism, streetscape	Female, Irish, 20
	landscape, crepuscolare	Female, Italian, 34, Postgraduate degree
	light, illusion of space, good composition, airy, impressionist-captures a moment, plein air, sombre colours	Female, Irish, 20, College major in art history
	Landscape, street scene, Village, Sunset, Sunrise, Snow	Male, Irish, 29, B.A. in art history

Table 4.2 – Examples of keyword sets submitted for the images used in our survey, together with corresponding information on the participants. For each image the keywords above the line are examples submitted by laypersons and the keywords below examples submitted by specialists. Incisions made by participants in the form of punctuation marks or line breaks, are normalised to commas (,).

Only on rare occasions (six fragments out of 509) did we deviate from the literal application of the guidelines. The phrase "Eurystheus twelve labours" is an example of this. Following the guidelines strictly, "twelve labours" would be classified as one fragment (because "twelve" as a numeral is not separated) and "Eurystheus" would remain by itself. However, it seems likely, that the phrase was intended as one semantic unit (as in "Eurystheus' twelve labours") and that "Eurystheus" effectively acts like an adjective here.

The text fragments that resulted from the splitting process, were classified by the author, following the guidelines, into one of 52 categories: 12 *non-visual*, 10 *perceptual* and 30 *conceptual*.[15] The twelve non-visual categories are a subset of the VRA core 3.0 elements.[16] The perceptual categories cover information about colour, shape, texture, composition and technique pertaining to either the image as a whole or an individual element in the image ($5 \times 2 = 10$). The conceptual categories cover information about events, times, places, relations or general information referring to either the scene as a whole or an object in the scene ($5 \times 2 = 10$). Additionally, each fragment classified as conceptual had to be qualified as being either of general, specific or abstract nature ($10 \times 3 = 30$; compare the bullet points below). Some of these categories (e.g. ID number, rights and source of the information among the non-visual categories) are naturally not used for image querying and some of the combinations among the perceptual and conceptual subcategories are much rarer then others (e.g. as pointed out by Hollink et al. in their guidelines, time specifications are almost always about the entire scene). Only 30 categories of the 52 were necessary to classify the terms the participants of our study had submitted.

The most important rules for qualifying fragments as either general, specific or abstract are:

- Descriptions of concepts that require only everyday knowledge of the world are general.
- Descriptions of unique objects (e.g. proper names) or descriptions that require specialist knowledge (e.g. "Eurystheus twelve labours") are specific.
- Descriptions that someone from a *completely different* culture would not understand are specific.
- Descriptions are abstract if "the level of subjectivity is so high that differences in opinion about the interpretation are possible" [73, p. 618].

Classifying the fragments into categories sometimes involved a subjective judgement. Based on the third rule above, descriptors that were qualified by country adjectives (e.g. "Greek mythology", "French country town") were classified as specific because someone from a completely different cultural background might not properly understand their meaning. This is certainly a borderline case as other concepts like "lion" or "dragon" that not every culture might be familiar with were seen as universal enough to be treated as general.

In the study presented by Hollink and colleagues the initial classification of descriptors was validated by two additional reviewers and a concordance analysis. The authors report a correspondence of 70% and 75% between their main classification and those of the additional reviewers (compare [73, p. 618]). The category assignments we used for our analysis were made by a single human classifier without further validation. However, assignments to non-visual categories (e.g. name of the painter, regardless of whether it is correct or an educated guess) are usually straightforward. So are

[15] For the guidelines compare the previous footnote.

[16] See `http://www.vraweb.org/resources/datastandards/vracore3` (last accessed 30/10/2008).

Level	Count	%	(%)
Non-visual	127	25.0	(27.0)
Perceptual	48	9.4	(7.9)
Conceptual	334	65.6	(65.2)
Total	509	100.0	(100.0)

Table 4.3 – Distribution of image query terms over the three main category levels in absolute numbers, percentages and weighted percentages (in brackets).

assignments to the 'perceptual technique' category (mostly "painting"), general and specific object descriptions and general characterisations of the time (e.g. "winter") and place (e.g. "street") of the whole scene. These categories together received 402 assignments out of our total of 509. This corresponds to 78.9%. On the other hand, an important cause of rater disagreement mentioned in [73] are mix-ups between the 'perceptual composition' and the 'conceptual relation' categories. This is understandable due to the closely related nature of these classes. In our study, however, only five fragments in total were assigned to either of these two categories. For these reasons, we are confident that we could achieve similar rates of agreement to those reported in [73] given that additional reviewers were trained to apply the same guidelines.

Finally, in their analysis, Hollink et al. normalise the number of occurrences of a descriptor category against the overall length of the description it occurs in. This is done in order to account for differences in the length of descriptions between different participants (so that the preferences of more verbose participants would not be given greater weight) and to account for differences in length between the two tasks of their study. The first task involved longer free text descriptions and the second task shorter descriptions in the form of providing search terms, not unlike our task. All analyses in [73] are performed on these weighted numbers and the results given as weighted percentages.[17]

This kind of normalisation seemed to be less important to our study since we did not use free text descriptions and most queries were of similar length with little variation.[18] However, in order to make our results fully comparable to [73], we performed the same normalisation and, in the tables of the following section, give our results in unweighted percentages followed by weighted percentages in round brackets. Table 4.7, in which we compare our results to Hollink et al.'s, uses weighted percentages only.

4.2.4 Results

We present the results our survey has produced, first in terms of the general classification of keywords at the different category levels, then the differences between specialist and layperson terminology and a comparison of our findings to other studies. Finally, we briefly report the results regarding our participants' experience and preferences in image retrieval.

[17] Hollink et al. divide the number of occurrences of a descriptor category in a description by the total number of descriptors (description fragments) in that description. Based on this, a normalised count and based on that count a weighted percentage of categories is calculated. For the details of the calculation see [73, p. 617f.].

[18] The average number of fragments per description was 3.56 with a standard deviation of 1.91.

Scope	Object			Scene			Total		
Characteristic	Count	%	(%)	Count	%	(%)	Count	%	(%)
Colour	2	4.2	(2.5)	17	35.4	(34.4)	19	39.6	(36.8)
Shape	0	–	–	0	–	–	0	–	–
Texture	0	–	–	1	2.1	(4.4)	1	2.1	(4.4)
Composition	0	–	–	4	8.3	(5.6)	4	8.3	(5.6)
Type/technique	0	–	–	24	50.0	(53.2)	24	50.0	(53.2)
Total	2	4.2	(2.5)	46	95.8	(97.5)	48	100.0	(100.0)

Table 4.4 – Occurrences of perceptual categories in absolute numbers, percentages and weighted percentages (in brackets).

Two thirds of the search terms provided by participants in all three tasks were at the conceptual level, one quarter referred to non-visual properties and less than one in ten to perceptual properties of an image (Table 4.3).

Non-visual level. There are twelve types of non-visual descriptors in our classification framework, but the respondents' selection of such terms was dominated by only two categories: *Style/Period* queries make up over half (50.4%, 52.6% weighted); *Creator* accounts for one third (33.1%, 33.5% weighted).

Perceptual level. The distribution of query terms over the perceptual level shows that half of the perceptual descriptors were queries for *Type/Technique* and almost 40 per cent were colour related queries. Next to none of the query terms referred to individual objects or regions in the image (see Table 4.4).

Conceptual level. The distribution of query terms over the conceptual level suggests over 80 per cent of descriptors were of a general nature, just over 10 per cent specific and only very few abstract. More than half of the conceptual queries made reference to an individual object in the image, just under half to the scene as a whole (see Table 4.5).

Level		**General**			**Specific**			**Abstract**			Total		
Scope	**Characteristic**	Count	%	(%)	Count	%	(%)	Count	%	(%)	Count	%	(%)
Object	Event	7	2.1	(1.6)	0	–	–	0	–	–	7	2.1	(1.6)
	Place	9	2.7	(1.9)	0	–	–	0	–	–	9	2.7	(1.9)
	Time	0	–	–	0	–	–	0	–	–	0	–	–
	Relation	1	0.3	(0.2)	0	–	–	0	–	–	1	0.3	(0.2)
	Uncharacterized	134	40.1	(41.3)	28	8.4	(9.7)	2	0.6	(0.4)	164	49.1	(51.4)
Subtotal		151	45.2	(44.9)	28	8.4	(9.7)	2	0.6	(0.4)	181	54.2	(55.0)
Scene	Event	20	6.0	(6.1)	5	1.5	(2.1)	1	0.3	(0.4)	26	7.8	(8.6)
	Place	49	14.7	(15.2)	2	0.6	(0.5)	1	0.3	(0.1)	52	15.6	(15.9)
	Time	40	12.0	(10.7)	0	–	–	1	0.3	(0.1)	41	12.3	(10.9)
	Relation	0	–	–	0	–	–	0	–	–	0	–	–
	Uncharacterized	18	5.4	(5.8)	3	0.9	(1.0)	13	3.9	(2.8)	34	10.2	(9.6)
Subtotal		127	38.0	(37.8)	10	3.0	(3.7)	16	4.8	(3.4)	153	45.8	(45.0)
Total		278	83.2	(82.8)	38	11.4	(13.5)	18	5.4	(3.8)	334	100.0	(100.0)

Table 4.5 – Occurrences of conceptual categories in absolute numbers, percentages and weighted percentages (in brackets).

	All levels			**Perceptual**			**Conceptual**		
	Non-visual	Perceptual	Conceptual	Colour	Type	Other	General	Specific	Abstract
Specialists	**31.8**	9.1	**59.1**	31.8	63.6	4.5	79.7	10.5	**9.8**
Laypersons	**18.7**	9.7	**71.6**	46.2	38.5	15.4	85.9	12.0	**2.1**

Table 4.6 – Comparison of the occurrences of specialist and layperson descriptors in percent. Statistically significant results are printed in bold.

	All levels			**Perceptual**		**Conceptual**		
	Non-visual	Perceptual	Conceptual	Object	Scene	General	Specific	Abstract
Hollink et al.	0.9	11.9	87.2	45.6	54.4	74.4	16.4	9.2
Isemann et al.	27.0	7.9	65.2	2.5	97.5	82.8	13.5	3.8

Table 4.7 – Comparison of our findings to those of Hollink et al. (all numbers given are weighted percentages, cf. p. 85).

Specialists vs. laypersons. We have analysed the use of terms by specialists and laypersons. The striking difference between specialists and laypersons, in terms of the overall distribution over the three top-level categories, was that art historians tend to use more non-visual terms, laypersons by comparison more conceptual ones. The frequency of perceptual descriptions was nearly the same in both groups (Table 4.6).

The distribution within the non-visual level was very similar with both groups showing percentages for *Style/Period* (49.4% vs. 52%) and *Creator* (33.8% vs. 32%) very close to the overall values. There were differences in the *Material* category (more used by specialists) and the *Relation* category (meaning relation to other works) which was exclusively used by laypersons. Our observations at the perceptual level may suggest that laypersons talk more about colour and specialists query more for type or technique related information. Finally, conceptual descriptors are similarly distributed in terms of the object/scene distinction (51.8% object, 48.3% scene, for specialists vs. 56.0% and 44.0% for non-specialists). Perhaps the noteworthy difference at the conceptual level is that art historians seem to use abstract descriptions (9.8%) more often than laypersons do (2.1%) (Table 4.6).

We subjected the observed differences between specialists and laypersons to statistical significance testing. For the frequency distribution of specialist and layperson terms over different category and subcategory levels of our coding system, we tested the null hypothesis that observed differences between the two distributions are due to chance variation. Using Pearson's chi-squared test, we found that the distributions of specialist and layperson terms over the main category levels (non-visual, perceptual, conceptual) are significantly different ($\chi^2 = 11.77$, df $= 2$, $N = 509$, $p = 0.003$). Similarly, we found significant differences in the distributions over the general, specific and abstract subcategories of the conceptual level ($\chi^2 = 9.53$, df $= 2$, $N = 334$, $p = 0.009$). For distributions over the perceptual subcategories, on the other hand, we had to accept the null hypothesis. This may seem surprising, as some of the differences at the perceptual level appear to be stark, if viewed as percentages (31.8% vs. 46.2% colour, 4.5% vs. 15.4% other descriptors). In absolute numbers, however, only few perceptual terms were used by the participants of our study (48 in total).[19] The small sample size of perceptual terms limits the power of the statistical test.

Investigating the cases where we found significant differences further, we performed additional tests in an explanatory analysis to single out the categories which may plausibly be responsible for our findings. For the three main categories and the general, specific and abstract subcategories at the conceptual level, we singled out one category at a time and added the term count of the two others together. On the resulting 2×2 matrix we performed a chi-squared test to find out if differences in the category we singled out could possibly be significant. We found the differences at the non-visual level ($\chi^2 = 11.62$, df $= 1$, $N = 509$, $p < 0.001$) and at the conceptual level ($\chi^2 = 8.71$, df $= 1$, $N = 509$, $p = 0.003$) to be statistically significant. Within the conceptual level, we found the differences in the number of abstract descriptors to be statistically significant ($\chi^2 = 9.50$, df $= 1$, $N = 334$, $p = 0.002$). The chi-squared tests were carried out on the count data, but for illustrative purposes we have printed in bold the percentages that correspond to the statistically plausible differences between specialist and lay terminology in Table 4.6.

Comparison to other studies. Comparing our results to those reported in [73] we found some similarities but also some marked differences. The distribution of conceptual descriptors over the three abstraction levels — general, specific and abstract — is very similar, with slightly more general

[19]On average, this amounts to exactly one perceptual descriptor per participant over all three tasks.

descriptors in our study at the expense of abstract ones (Table 4.7). While the ratio of conceptual descriptors to perceptual descriptors is very nearly the same (8.3/ 7.3), their respective shares in the total number of descriptors differ. This is due to the fact that the third group, non-visual descriptors, were much more frequent in our study than in Hollink et al.'s (27% vs. 0.9%). As expected non-visual descriptors seem to play an important role in the domain of fine art images. A possible explanation for the low value of 0.9% given by Hollink and colleagues is their use of imaginary images. Another big difference is that perceptual descriptors in our study almost exclusively referred to the scene as a whole (97.5%) and hardly ever to individual objects (2.5%). The ratio is much more balanced in Hollink et al.'s study (54.4% and 45.6% respectively). Again the use of imaginary images could serve as a possible explanation here: In describing their mental images participants may tend to describe the perceptual properties of individual image elements more than they might otherwise do.

Our survey contained one question that addressed research hypothesis (ii) above: a possible preference for target search in art image retrieval. In question 3.1 participants were asked to rank three different descriptions of search behaviour as *most frequent*, *second most frequent* and *least frequent* with respect to their own practice when searching for images of fine art on the web (Figure 4.1). The three descriptions where chosen to resemble closely the three modes of image search postulated by Smeulders et al. [139]: *target search*, *category search* and *search by association*.

Of the 48 participants who have otherwise completed the survey, eight have skipped this question. Of the remaining participants, 72.5% said that "look[ing] for one specific work that [they] would like to see" was their most frequent way of searching among the alternatives presented. This option corresponds to *target search*. A clear majority of participants also said that "look[ing] for works of art that fulfil certain general criteria (e.g. same style, technique, epoch)" (*category search*) was second most frequent in their experience (70.0%). Similarly, 77.5% ranked "look[ing] around to explore art with no particular image or criterion in mind to start with" (*search by association*) as the least frequent of the given alternatives. Only 20% named category search and only 7.5% named search by association as their most frequent mode of searching for images of fine art. These results seem to support our hypothesis that target search is a dominant form of searching for images of works of art. This is perhaps in contrast to other domains in which target search is likely to play only a minor role.

Finally, our survey also contained two questions which addressed the content-based image retrieval techniques mentioned in research hypothesis (iii). Seven of our participants (14.6%) declared, that they were aware of non-keyword based approaches to image retrieval, such as 'query by sketch' or 'query by image example'. Given the choices of 'Unsatisfactory', 'Satisfactory', 'Quite good' and 'Excellent', three of these seven participants rated the performance of such methods as 'Unsatisfactory' and two as 'Satisfactory'.[20] The remaining two declined to give a performance rating.

4.3 Summary

We approached experts with a specialty in art history to learn how they conceptualise their field of expertise and we contrasted their use of terminology with that of laypersons for the same set of retrieval tasks. The high level results from our survey suggest, corroborated by statistical significance testing, that there may be systematic differences between a specialist's and a layperson's choice of search terms. While in absolute numbers the conceptual level is the most frequently used by both groups, it is, relatively speaking, less popular with specialists. This is in line with the picture we

[20]One person selected 'Excellent', but in the previous question said he had never heard of such methods.

received through our qualitative interviews. Asked what kind of information about a work is typically recorded, both of our expert interviewees stressed the physical description of media and materials and the provenance of a work, its travels through history, over the work's subject matter. It may therefore not be surprising, that the category of non-visual metadata was also more commonly used by specialist participants of our survey than by laypersons. The other alternative to conceptual descriptors, in the Panofsky inspired classification scheme of Hollink and colleagues, are terms referring to the perceptual qualities of an image. The use of these terms in our survey was sparse, both in the lay and in the expert camp. In light of the difficulties that content-based image retrieval (CBIR) has faced (compare [39], esp. p. 537) this is perhaps not surprising, for it is at this level that a content-based approach would be expected to make the biggest impact. This may also help to explain, why our interview partners and heritage experts surveyed by others [21], greeted CBIR approaches with scepticism.

Another reason for such scepticism, may of course be that content-based techniques are not widely known. This was one of our auxiliary research hypotheses. Nearly one in six of the participants of our survey reported that they had some familiarity with content-based methods. This appears to be a relatively high proportion for a technology which is still largely experimental. A possible explanation may be that a number of participants in our non-expert group have a background in computer science.

With respect to the second auxiliary research hypothesis, that target search is a dominant paradigm in art image retrieval, our findings give an indication that it may be valid. This was already hypothesised by Smeulders et al. [139], but in personal conversation we found that it is sometimes overlooked by image retrieval practitioners in non-art related domains.

While our survey did not directly address possible differences in search terminology employed for visual art and for other kinds of images, a comparison of our findings with published studies that apply a similar classification system, allow us to draw some tentative conclusions. Like Jörgensen [91] and Hollink et al. [73], we found that the conceptual level is by far the most frequently used. Unlike in these studies, however, non-visual descriptors played a much bigger part in our setup.

While we have taken care to design the survey according to established methodology and analyse our data in ways that would extract informative facets, our survey has shortcomings and limitations. The wording and the relative order of answer choices may sometimes have biased our participants towards a certain answer. An example of this is question 3.1 (Figure 4.1). In the third answer option, which describes browsing as a way of searching, the word "just" was intended to emphasise the initially targetless nature of this kind of image search. However, it also carries a derogatory connotation, making this option perhaps less attractive to some participants. In an ideal survey setup some of the more complex questions should perhaps have been rephrased or stated a second time for control or in negated form. However, we also had to keep the survey as short as possible as our participants were volunteers and a longer survey would invariable lead to a decrease in response rate.

A more fundamental consideration is, whether the format of a survey is appropriate for testing our research hypotheses. The auxiliary hypotheses may be best addressed by a survey, but in the case of our main hypothesis, at least, a case may be made for inquiring into it by way of an experiment. Such an experimental setup could involve observing participants, while they perform genuine retrieval tasks, as opposed to the mock retrieval tasks in our survey. This issue, however, touches on a problem which is inherently present in many image retrieval studies: How does one ensure, that the behaviour observed and search terms collected in a study, reflect genuine user needs and not merely a person's expectations of current search engine performance. In this context, we opted for a survey with mock retrieval tasks, which may trigger a more spontaneous response from participants, precisely because it is divorced from a real life retrieval setup.

Naturally, our survey also has limitations in how much its findings can be generalised to a larger population. Methodically, it lacks the basis for such a generalisation, as the sampling was non-random. In this, however, as well as in sample size and the number of tasks set to participants, our survey is broadly comparable to similar studies in the literature [19], [73], [91]. If we were to repeat our study or conduct a similar one, we would perhaps increase the number of search tasks and in doing so, take into account statistical considerations for determining sample size in comparative studies. To achieve a statistically significant result for the infrequently used perceptual descriptors, for example, our participants would have had to submit approximately twice as many keywords as they did.[21]

Future studies along similar lines may of course also test a number of hypotheses derived from the literature, that our survey has not addressed. One example is the location of art image retrieval tasks on Fidel's spectrum between the objects pole and the data pole [44]. At any rate we hope that future studies in this field will reuse a well-researched category system that is compatible with a wide range of other classifications (e.g. [17], [73]). Only studies which are comparable to each other can incrementally add to a fuller understanding of the picture.

Coupled with the review of relevant literature and methodological considerations, the results from our survey have influenced the design of our ontology. The relative importance of non-visual information, and here in particular the categories of style/period and creator, led us to a model which allowed users to contextualise a work of art in terms of its style and period, and an artist in terms of biographical details. In both cases aspects of temporal, geographical and event modelling played a part. These details of the model may in turn benefit an ontological description of image subject, as our survey showed that two out of five conceptual image descriptors were classified into one of the three categories: time, place and event. The development and description of our ontology-based prototype system is the subject of the next chapter.

[21] This holds under the assumption that the proportion of terms in the subcategories of the perceptual level would remain unchanged.

Chapter 5

System Development and Evaluation

The discussion in the previous chapter alluded to the problems facing an end-user who may not be familiar with the contents of an image database, especially that of an aesthetic image database, and also be unfamiliar with the logic underpinning a query system. Recall that in an aesthetic image database, each image has a number of attributes, and may have complex relationships with other images in the database. We have argued for the case of representing these sophisticated images in an ontology-informed knowledge base (Chapter 3) which will not only capture the salient attributes of an image but also the spatial and temporal relationships.

We created a graphical query builder, that structurally reflects a part of our ontology, but visually resembles the advanced search templates known to heritage professionals from collection management systems and to the general public from the web pages of first-rate museums and galleries.

Our goal was to develop an ontology that constitutes a structured description of works of visual art, especially paintings, drawings and prints created before the 20th century, with the objective of classifying and retrieving digital representations of such works from a repository of images. The information related to the works that is structured by the ontology should cover information on the artist, the style and the art historical period in which the work was created as well as the subject matter or content of the work. The intended user group of the ontology (or an application based on the ontology) are, primarily, art experts, such as curators in a museum or gallery.

In our evaluation involving experts and laypersons the query builder was able to synthesise queries from user input into a graphical interface and with the help of the ontology sometimes unified apparently different queries.

In this chapter we synthesise our findings from Chapters 3 and 4. We begin with a critique of ontology development methods elaborated by specific reference to the organisation, storage and retrieval of curated images (Section 5.1). Following this, we discuss the modelling of curated knowledge with specific reference to time and location records (Section 5.2). The outline of the architecture of our system is then presented. The system can build queries visually and can retrieve images from an ontology informed knowledge base (Section 5.3). The knowledge base is a subset of a generally available collection of over 40,000 art images (Section 5.4). It comprised 41 works by 31 artists for the system testing and evaluation which was undertaken with 10 experts and 10 laypersons (Section 5.5). In Section 5.6 the findings presented in this chapter are summarised.

5.1 Developing ART Ontology

Our literature review and the growing number of scholarly and technical publications, indicates that neither the ontology development methods nor the related implementation systems have reached the level of maturity to allow the creation of arbitrary applications straightforwardly. Software engineers involved in the development of systems like Protégé (see Section 3.1.2) have provided heuristics for building ontology-based information systems. Consider, for example, the seven-step method by Noy and McGuinness [118]:

(i) Determine the domain and scope of the ontology

(ii) Consider reusing existing ontologies

(iii) Enumerate important terms in the ontology

(iv) Define the classes and the class hierarchy

(v) Define the properties of classes — slots

(vi) Define the facets of the slots

(vii) Create instances

We will discuss our system development strategy with reference to these steps.

5.1.1 Domain and Scope of ART Ontology

The domain our ontology will cover relates to the retrieval of arbitrary aesthetic images from a well-curated library of images. Such is the case with art galleries. Curators annotate images with a variety of image external metadata and include image internal properties. Similar structures are found in forensic science image collections and others.

The people retrieving images from these well-curated collections often do not share the knowledge of the curators — there is a knowledge gap which might be filled-in, to an extent, by an ontology-informed knowledge base.

The domain of our ontology are aesthetic images (visual art), its scope is the representation of information on such images to describe, classify and retrieve them from a repository. As a paradigmatic case we decided to look at two-dimensional works like paintings, drawings and prints. Other forms of visual art like sculpture, ceramics and architecture will have much in common with these but also be different in various ways. Sculptures for example may not have a canonical point of view from which they are expected to be looked at. Thus, digital representations of the same sculpture may differ significantly.

Works of these kinds — paintings, drawings and prints — are commonly classified into various categories with respect to what they show, how they show it and the circumstances of their creation. Such categorisations are naturally of interest to our project. Since modern and contemporary art, however, often have the objective to defy categorisations and broad consensus on its classification has yet to be reached, we decided not to focus on such art forms.

The scope of our ontology was additionally sharpened by a survey on terminological preferences in art image retrieval (see Chapter 4). One of our findings in this survey was that the identity of the

1	Which paintings were done by Leonardo Da Vinci?
2	Which paintings originated from Leonardo's environment?
3	Which paintings were produced by Italian painters?
4	Which paintings exemplify Renaissance style?
5	Which paintings are from periods related to the Renaissance?
6	Which paintings did Leonardo have knowledge of?
7	Which paintings are portraits?
8	Which paintings are self-portraits?
9	Which paintings show Hercules and the Hydra?
10	Which paintings show a man fighting a monster?
11	What does the painting "Madonna of the Rocks (Louvre)" show?
12	Which paintings show a scene similar to "Madonna of the Rocks (Louvre)"?
13	Which paintings are located in Paris?

Table 5.1 – Initial set of competency questions

artist, the style and historical period of the work and the depicted subject matter were among the most frequently used categories for querying.

At an early development stage we formulated a list of competency questions which is reproduced in Table 5.1.

A number of these questions are concerned with the subject of an image, i.e. higher-level image content as opposed to physical image content, that is primarily addressed by content-based image retrieval (compare the discussion in Section 2.4.1). In the course of this project we explored one possible area of image subject in detail: the depiction of human personage and the human figure (cf. Section 3.4.3). Because of the difficulties of an objective representation of visual content, however, and in part due to the realisation that it is hard to predict which information is of interest to retrievers even in a limited domain, the focus shifted over the course of this project. The shift in scope was influenced by input from art experts we worked with.

As the focus changed, the set of competency questions that were in the range of our ontology changed as well. Eventually, only about half the questions in Table 5.1 could be answered by our ontology. Questions 1, 3, 4, 7 and 8 can be answered as they are stated. Questions 2 and 5 can be partially addressed and questions 6 and 9 through 13 cannot be answered by our ontology. On the other hand, the final version of our ontology can answer many questions that go beyond the first set of competency questions. A characterisation of the scope of our ontology at the time of our system evaluation, that is comparable to these competency questions, is the set of query tasks presented to the participants of our evaluation study (see p. 126, Table 5.3). Comparing these to the initial competency questions shows a shift of focus towards a richer characterisation of the biography of artists and depicted persons.

The evaluation queries in natural language have a correspondence in formal description logic statements which we will discuss in the context of our evaluation results (see p. 133, Table 5.4). It is these formal statements, rather than the natural language statements, that should be seen as answerable by the ontology. This is comparable to a proposal by Grüninger and Fox of a two-layered approach to competency questions [55]: After defining informal competency questions a number of formal competency questions are derived for which an answer can be inferred from the ontology (together with a suitable knowledge base).[1]

5.1.2 Reuse of Existing Ontologies

We discussed a number of candidate ontologies that could be reused in a project like ours (Section 2.3). In each case we had to make a judgement whether to reuse the resource or not.

A basic decision in each ontology development process is whether to map to a top-level ontology or not. As we mentioned, some authors consider this so fundamental, that they attempt to incorporate top-level ontology use into development methods for ontologies (see remarks on page 53). We have decided to use a top-level ontology primarily as a guide to cleaner modelling, but also to facilitate interoperability. Unless the application scenario is extremely narrow, the cost of using a top-level ontology seems to be justified. The CIDOC Conceptual Reference Model is a natural candidate, because it was developed as a core ontology for cultural heritage. Possible other candidates, such as BFO,[2] DOLCE [47], or SUMO[3] were developed for scientific domains or with no particular application domain in mind. Moreover, the CIDOC CRM has received numerous implementations in RDFS and OWL which are freely available (cf. Section 5.3.5).

In order to evaluate the conceptual structure of Iconclass, we examined some of the information the system contains on the topic of human anatomy and the human figure as the subject of an image (see Section 3.4.3). We found that relations in Iconclass lack a well-defined semantics and that major reengineering would be necessary in order to support meaningful inferences. This combined with a reorientation of our scope away from image subject led us not to use Iconclass or any of its parts in our system.

The Getty Vocabularies appear to be conceptually well designed for applications in the heritage domain (cf. p. 27f.). Like the CRM they have been thoroughly developed over a long time and grew out of a set of heritage records, from the bottom up. The vocabularies reflect the distinctions made by heritage professionals and have been widely used in the sector. With respect to their license status and their technical implementation, however, they cannot be readily included into an OWL or RDFS ontology. While the vocabularies can be freely explored through a web browser, a license is required for the use as XML data files or as relational tables. Even if licensed, however, the vocabularies still need to be converted into an ontology language, in order to be incorporated into our system. For these reasons we opted not to use the Getty Vocabularies. An exception to this is a small part of the AAT, which is concerned with art historical periods and which we recreated as part of an OWL file.

Our decision not to license the Getty Vocabularies, influenced our decision to reuse the GeoNames database for our geographical model. As we have pointed out (on p. 31ff.), the TGN may be better suited for historical geographical entities, which is appealing in heritage contexts. But GeoNames has a comprehensive coverage and offers free and direct access to the database through a number of APIs. This allowed for an uncomplicated extraction and conversion of the required data (cf. Section 5.3.5).

All of the resources discussed up to this point were found by using conventional search engines

[1] The details of Grüninger and Fox's proposal differ, of course. They suggest to use first-order logic as the language in which to formulate the formal competency questions. They also suggest to fix the formal competency questions in advance and to use them as a basis for specifying ontology axioms and evaluating an ontology in the form of completeness theorems [55].

[2] `http://www.ifomis.org/bfo/home` (last accessed 7/10/2011).

[3] `http://www.ontologyportal.org` (last accessed 7/10/2011).

or literature research. An alternative and more direct way of searching for ontologies are so-called ontology search engines. Examples of ontology search engines are *OntoSelect*[4] and *Swoogle*.[5] Such search engines have the advantage of providing access to ontologies that are intended to be reused (as they are publicly accessible on the web) and that are implemented in a concrete knowledge representation format. This eliminates or greatly reduces the technological and licensing obstacles to reuse. In an early phase of our project we used OntoSelect to search for relevant ontologies in our domain. However, we were unable to retrieve any resources of interest. The only ontologies we could find, for example, that included a class 'painter', were classifications of occupations, in which 'painter' was understood in the sense of 'house painter'. A difficulty with ontologies retrieved through ontology search engines is that there is usually little associated documentation and that the structured resources often have to speak 'for themselves' to reveal their intended meaning.

5.1.3 Enumerate Important Terms in the Ontology

The third step of the ontology development cycle is to collect a list of terms that appear to be important for the scope as it is defined. Terms collected in an initial brainstorming exercise were, for example, *work of art, painting, person/human being, painter, model, sitter, human body, human body part*. Eventually a more systematic approach for collecting terms was adopted. Some of the column names in our tabular test data[6] that we used to fill our knowledge base, naturally gave rise to class names in the ontology. Additionally, we extracted values from relevant columns in order of decreasing frequency. In this step multiword entries in database fields were broken up into their individual components. In the database, artists' geographical places of activity, for example, are often recorded as "Köln, Italien, Prag", "Niederlande, Italien und England",[7] i.e. as literal strings containing complex information. A frequency analysis on the individual components of these literals was performed and the most frequent ones were considered for inclusion in our term list. This applies not only to geographical terminology but to other categories as well. As the frequency distribution of these terms appeared to approximate a power law, we often found that a few terms in a given category covered the majority of entries in the database. The ten most frequent terms for art historical periods accounted for over 78% of period tokens in the database. Sometimes the frequency approach was combined with input from heritage relevant resources. The most frequently occurring art historical periods were expanded by super- and subperiods as modelled in the Getty AAT, to create a more comprehensive historical context. Advice from the art experts we worked with also resulted in the inclusion of new terms. The curator of the Prints and Drawings department of the National Gallery of Ireland, for example, provided terms denoting printmakers' professional specialisations, e.g. 'engraver', 'copper engraver', 'wood cutter' etc.

5.1.4 Classes and Class Hierarchies

The next step involves identifying, which of the collected terms should be modelled as classes in the ontology and which subclass relationships exist among classes. Many of the terms extracted from the database columns (as opposed to the column names) were naturally modelled as individuals. Among

---

[4]`http://olp.dfki.de/ontoselect?wicket:bookmarkablePage=wicket-0:de.dfki.ontoselect.Home` (last accessed 7/10/2011).

[5]`http://swoogle.umbc.edu` (last accessed 16/03/2012).

[6]Extracted from a German language digital art collection, see `http://www.zeno.org/nid/20003856844` (last accessed 14/08/2011). See also Section 5.4.

[7]The language of our data is German, cf. previous footnote.

the remaining terms we introduced those as named classes into our ontology that we expected to be used as query categories or in axioms and restrictions.

Once the classes are identified and their intended meaning is precisely defined, subclass relationships can be determined. The test whether one class is a subclass of another, rests on the semantics of the subclass relationship:

> A is a subclass of B if and only if every instance of A is also an instance of B.[8]

Here, the quantifier 'every' ranges over all conceivable instances, not just the ones that are included in our ontology or that we are interested in for a particular application. The intentional definition of a class provides a decision procedure for determining whether or not an instance is a member of that class:

> An instance x is a member of class A, if and only if x fulfils the definition of A.

Combining the two rules we arrive at a decision procedure for subclass relationships:

> A is a subclass of B if and only if every instance that fulfils the definition of A, also fulfils the definition of B.

If we recall the definition for 'painting' from page 14: "A painting is an artefact which results from the practice of applying paint, pigment, colour or another medium to a surface or support.", then we can see that 'mural', defined as "any piece of artwork painted or applied directly on a wall, ceiling or other large permanent surface."[9] is a subclass of painting. In particular, the perhaps prototypical property that a painting's support is mobile is not necessary according to the definition we have adopted.

Our approach in creating a class hierarchy could be described as a 'middle-out' (see p. 53). In particular, the first relationships that we often attempted to establish for a newly introduced class were subclass relationships in reference to the top-level ontology. In determining subclass relationships, the use of a top-level ontology can be helpful as it may reduce the number of systematic comparisons that are necessary for determining the position of a class in the hierarchy. Several high-level concepts in the CRM, for example, are declared to be disjoint. This entails that once a class is located 'below a certain branch' of the top-level ontology, it cannot consistently enter subclass relationships with classes from a disjoint branch. Some of our classes were identified with CRM concepts, e.g. 'human being' was identified with 'E21 Person' in the CRM. This may appear unusual, since the concept of person usually comprises more than just human beings, such as legal persons, but the definition for 'E21 Person', given in the CRM reference document: "This class comprises real persons who live or are assumed to have lived." [26, p. 11], suggests that its intended meaning is identical to the intension we associated with 'human being', regardless of the class name.

5.1.5 Properties and Restrictions

Steps five and six in Noy and McGuinness' ontology development method need to be 'translated' into the language of description logic ontologies (e.g. OWL ontologies). The frame-based paradigm in which Noy and McGuinness cast their methodological advice is focussed on the 'frame' as a basic

[8]Noy and McGuinness only state one direction of this equivalence, namely the one that is less relevant for determining the sub- or superclass status of two classes: "*If a class A is a superclass of class B, then every instance of B is also an instance of A*" [118, p. 8, emphasis in the original].

[9]`http://en.wikipedia.org/wiki/Mural` (last accessed 10/10/2011).

modelling unit. Frames correspond to both classes and instances in OWL ontologies. 'Slots', on the other hand, the equivalent of OWL properties, are always attached to frames and are not entities in their own right. As frame-based systems lack some of the expressivity of OWL ontologies with respect to properties, it is natural to interpret step five, "Define the properties of classes — slots" [118, p. 8], as including not only the definition of properties, but also that of a property hierarchy. Similarly, step six, "Define the facets of the slots" [ibid.], which includes determining domain, range and cardinality constraints for slots, may be generalised to include all property restrictions in OWL and property chain restrictions in OWL 2.

Because of the rich and detailed property structure in CRM our approach to defining properties and a property hierarchy, was more of a top-down approach than was the case with classes. Often we attempted to identify whether the CRM contains certain properties for the relations we wanted to express (as is the case, e.g., for linking a person to his or her birth date). In other cases we expanded the CRM property hierarchy downwards, refining the CRM property 'P62 depicts', for example, to represent different modes of depiction after consultations with National Gallery experts.

The domain and range constraints of our properties are often inherited from the CRM (i.e. domain and range of properties are indirectly restricted because the domain and range of their superproperties is restricted) and we did not specify these constraints explicitly for newly introduced properties. An exception of this are a number of data type properties which we introduced to represent the temporal extension of events (see below) and, for example, the property linking human persons to their biological human bodies for the anatomy model.

We already discussed the use of property chains, a feature which was introduced into the OWL 2 language as an improvement over OWL, in the section on transitive propagation. Property chain axioms played a part in our temporal and geographical modelling (Section 5.2).

5.1.6 Instance Creation

The final step of the method presented by Noy and McGuinness concerns the definition of instances. In our opinion two different kinds of instances need to be distinguished in this context. On the one hand, 'data instances' that represent some of the records we are interested in organising with the help of our ontology. These form the knowledge base of the ontology and we expect them to change depending on our current interest. On one occasion we may have a number of Renaissance paintings loaded into our knowledge base, on another one, self-portraits from the impressionist era. Distinguished from these instances, conceptually if not in terms of their implementation or formal properties, are instances that represent fundamental concepts in the visual arts and heritage field, but that are modelled as individuals because of their ontological properties. An example for this may be art historical periods, such as 'Baroque' or 'Post-Impressionism'. We have tried to make this distinction apparent by separating what is technically the 'ABox' of our description logic representation from what is conceptually the 'knowledge base proper' (cf. Section 5.3.1).

To get an idea of what the level of instance data in our system should look like, we manually modelled the properties and relations of the Mona Lisa by Leonardo da Vinci as an exemplary work of art. In this case we did not only rely on the information about the Mona Lisa that was contained in out data set, but also used additional sources like encyclopaedia entries on Mona Lisa. This singular model was intended as a target structure for the automatic population of our knowledge base from the data (cf. Section 5.4).

Throughout the development of the ontology, domain experts were consulted, over and beyond the expert interviews described in Section 4.1. In particular, the curator of Northern European art at the National Gallery of Ireland, Dr Adriaan Waiboer, was repeatedly approached on various aspects and details of the model.

5.2 Modelling Curated Knowledge

We will now look at more concrete aspects of our model, especially in connection with our top-level ontology, the CIDOC CRM. This includes a discussion of temporal modelling and of some difficulties that arise from explicit or implicit characterisations of concepts or properties in the CRM. Parts of this section are based on published work [85], [86].

5.2.1 Modelling Location

Philosophers and scientists may hold conflicting views on the fundamental properties of space and time. In pragmatic contexts, however, it appears that a sufficient agreement on the nature of space and time can be reached that allows us to 'locate' objects and events for the purposes of successful conversation. The record-keeping of cultural heritage information is such a context.

Spatial, especially geographical, information is important in a number of aspects in heritage documentation. In the data set we used, such information appeared in the context of the place of birth and death of an artist, the current location of a work and in others. Geographical information is usually recorded in the form of place names and the places denoted by these names can be related to each other in terms of spatial inclusion. The modelling of spatial information in the CIDOC CRM contains as basic elements: places which may include one another, physical objects or events which are located at certain places and labels and appellations which may be used to refer to places (cf. Figures 5.1).[10] Spatial inclusion between places was also illustrated by the example we have chosen for our discussion of transitive propagation in Section 3.3.

We recorded spatial information as instances extracted from the GeoNames database. We attempted to automatically resolve the place names we encountered in our database of works against the correct geographical features, selecting only geographical units from certain levels of granularity. The instances were given the type of 'E53 Place' and we used the CRM property 'P88 consists of' to order geographical entities into a mereological hierarchy. However, we did not make use of the mechanism of place appellations in CRM, but used `rdfs:label` and related annotation properties in our implementation instead.

5.2.2 Modelling Time

The representation of temporal information in the CRM is structurally similar to that of spatial information, but it introduces a number of additional complications. Events in the CRM are interpreted as spatio-temporal entities and their spatial and temporal aspects can be 'projected' separately by different properties into different classes. In the case of spatial information, the property is 'P7 took place at' and the class is 'E53 Place'. In the case of temporal information, the property is 'P4 has time-span' and the class is 'E52 Time-Span' (cf. Figure 5.2). Instances of both classes can be structured by a mereological (part-of) relation. Thus far, the two cases are analogous. Unlike places, however,

[10] Events are not part of Figure 5.1, but their relation to 'E53 Places' is documented in Figure 5.2, which covers temporal aspects. Both diagrams are taken from the CIDOC CRM defining document.

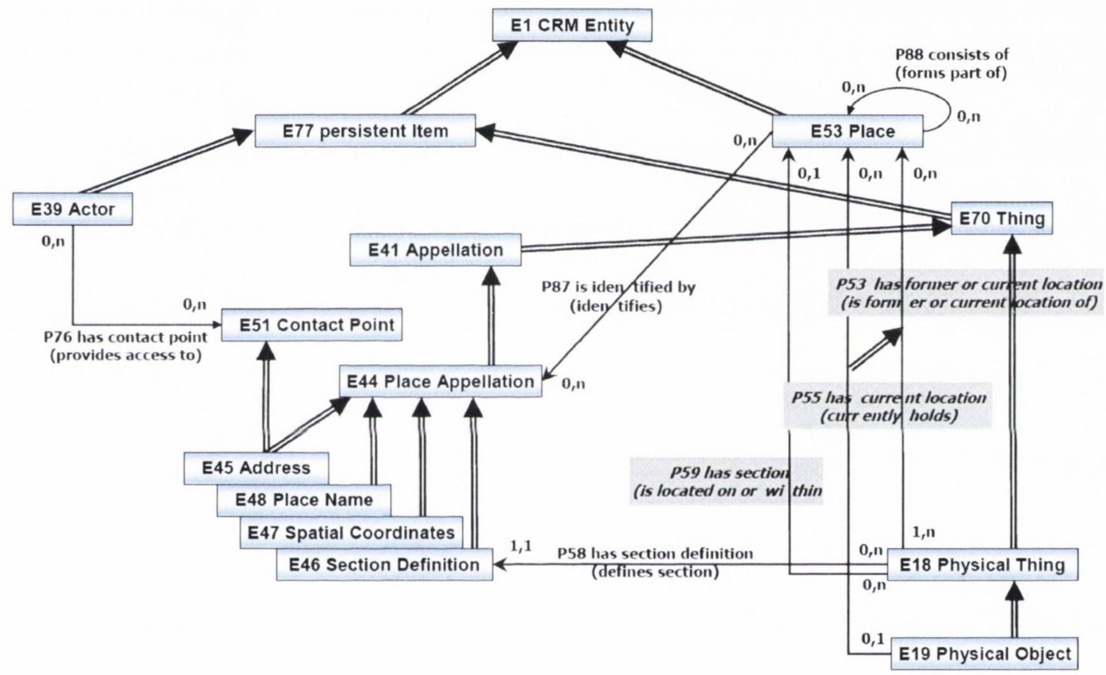

Figure 5.1 – Modelling spatial information in CRM. This figure appears as Figure 2 in the CRM reference document [25, p. xix]. Double arrows represent subsumption relationships (is-a) between CRM concepts and single arrows represent CRM properties, connecting a properties' domain and range. Note, that the double arrow connecting two single arrows to the right of the diagram signifies a subproperty relationship.

time spans in heritage documentation are usually not referred to by names, but by concrete, numerical dates. A date can be interpreted as a kind of name, but it has more than just referential qualities. Dates can be compared and calculations on dates can be performed. This introduces the need for an extra layer in the model, linking time-spans (which are ontology instances) to concrete data type representations on which comparisons or calculations can be performed. Two further aspects add to the complexity of time representations, especially in the heritage sector. First, the dating of events is often vague. Second, the literal strings used in heritage documentation to express time ranges are usually highly irregular and not standardised.

A number of considerations come into play when statements about the temporal position of historic events have to be made more precise. First of all, a large variety of different calendar systems exist or have existed, that make it difficult to compare and interpret recorded dates. This does not only hold for calendars from completely different cultures or based on different astronomical cycles (e.g. solar vs. lunar calendars), but also for the large variety of idiosyncratic, often regional or national calendars which resulted from a delayed, phased or abortive transition from the Julian to the Gregorian calendar. These circumstances make date references, such as "20th of March 1700", highly context dependent.

In the data set we used for our project (cf. Section 5.4) it was not clear, for example, whether artist birth dates prior to 1582 (i.e. the introduction of the Gregorian calendar) referred to the Julian calendar or the proleptic Gregorian calendar (i.e. the Gregorian calendar extended backwards to the time before its introduction) or whether this question was addressed at all in a uniform way by the data curators. However, this problem can, in principle, be solved by mapping date records from various calendar systems to a standardised calendar (or to a date range in a standardised calendar, if the original calendar system is uncertain). The proleptic Gregorian calendar can serve as such a calendar and the ISO standard 8601:2004,[11] which is based on it, as a standardised expression of date and time records.

Systematic representation of historical dates, poses another problem: That of the smallest unit of time that is expressible, in other words the granularity of the representation. Without the use of decimal fractions, for example, ISO 8601:2004 allows for the expression of combined date and time representations with a precision of one second, in the format `YYYY-MM-DDThh:mm:ss`.[12] With the use of decimal fractions, which can also be accommodated by the standard, ISO 8601 is not bounded by any granularity and can be used to express arbitrarily small time units. For most historical dates used in the arts and humanities, especially those prior to the 20th century, a granularity of one day would seem to be sufficient. For our project we have chosen the even coarser temporal granularity of one year. As we will detail below, this had mainly technical reasons, but at the same time it still appeared to be acceptable for the given domain.

Every event, at least in the domain of human affairs, has some duration or temporal extension, as opposed to happening in an instant. It is therefore appropriate to map historical events to time intervals rather than time points. A time interval can be expressed by two time points (its beginning and its end). If we use time and date representations in a calendar, however, we have to agree on a convention how to reference time points. This is necessary, because even an arbitrarily small unit in our calendar system does not unambiguously refer to a point in time. The expression '1599-06-06' does not naturally refer to a time point, but, arguably, signifies the time-span of one day in the 16th

[11] `http://www.iso.org/iso/catalogue_detail?csnumber=40874` (last accessed 20/10/2011).

[12] The standard also defines a number of variant formats; '`T`' is not a placeholder, but the actual letter '`T`' separating, by convention, the 'date' part from the 'time' part of the representation.

century. In the interest of precision, we need to agree on whether calendar expressions that denote the beginning and end of a time interval should be interpreted as being included in that time interval or not. If we adopt a granularity of one year for date expressions and the convention that date expressions signifying the interval end-points are included in the intervals, then the time-span of the year 1599 will be denoted by the interval '(1599, 1599)'. If we adopt the convention that dates are excluded, then the same time-span will be denoted by '(1598, 1600)'. And if we adopt the convention that they are excluded at the beginning, but included at the end of intervals, then that time-span will be denoted by '(1598, 1599)'. Any of these conventions or any other similar convention will serve the purpose, as long as it is always adhered to and well documented and communicated to the users of a system. The more fine grained a system of date expressions is, the less important the convention becomes in practice.[13] In our project we have adopted the convention that date expressions (in our case years) are included in intervals: the year 1599 would be expressed by the interval '(1599, 1599)'. How such an interval can in turn be represented in our knowledge base will be addressed in detail below (see p. 106).

Based on our considerations so far we seem to be conceptually set — technical implementation aside — for representing date appellations, like "1503", "1503-1519", "16th century", "first half of the 16th century", as they are encountered in heritage documentation. However, there are a number of other temporal descriptions, that are encountered with similar frequency. These include statements such as: "around 1503", "circa 1503", "circa 1503-1519", "1503 (?)", "after 1503", "before 1519".[14] These phrases are an expression of the uncertainty that surrounds much of the information on historical events.

There are several ways in which such uncertainty can be modelled. The VRA Core 4 metadata standard for the description of works of visual culture contains a 'DATE' element for recording temporal information.[15] This element has two subelements, 'earliestDate' and 'latestDate'. Each of these can in turn be qualified by a Boolean attribute 'circa' to express that the date is an approximation.

Another approach is to map historic time-spans to more than one or two dates. The *Simple Event Model* (*SEM*), a proposal for a representation of events on the Semantic Web, offers a total of seven time stamp properties: An unqualified time stamp, a time stamp for the beginning, a time stamp for the end, and an earliest beginning, latest beginning, earliest end and latest end time stamp [157]. The last four time stamps can be used to express uncertainty about the beginning and end of a period.

The CIDOC CRM contains two properties which effectively map time-spans to four dates: 'P81 ongoing throughout' and 'P82 at some time within' (cf. Figure 5.2). We will now describe the CRM's temporal model in detail and explain why we have chosen not to apply these properties but instead combined the CRM with properties similar to the time stamps of the Simple Event Model.

A number of classes and properties in the CIDOC CRM are involved in modelling temporal information (Figure 5.2). The model distinguishes between, what ontologists refer to as 'perdurants' (i.e. processes in time) and 'endurants' (i.e. entities that are fully present at any given moment).[16] The CRM class 'E2 Temporal Entity' represents perdurants, whereas the class 'E77 Persistent Item'

[13] If date-time representations have a precision of a microsecond, a user is unlikely to worry about whether the first and last microsecond are part of an interval, especially if that interval serves to date an archaeological artefact say.

[14] All of these cases were frequently encountered in the data set we used (Section 5.4).

[15] Compare the documentation of the standard at `http://www.loc.gov/standards/vracore/VRA_Core4_Element_Description.pdf` (last accessed 12/10/2011). The 'DATE' category can be used to assign dates to a range of events, such as artist birth and death dates, creation date of a work, etc.

[16] 'Perdurants' and 'endurants' are synonyms for 'occurrents' and 'continuants' respectively. The acceptance of both kinds of entities into one model makes the CRM an 'adequatist' ontology (for this terminology compare the discussion on p. 12).

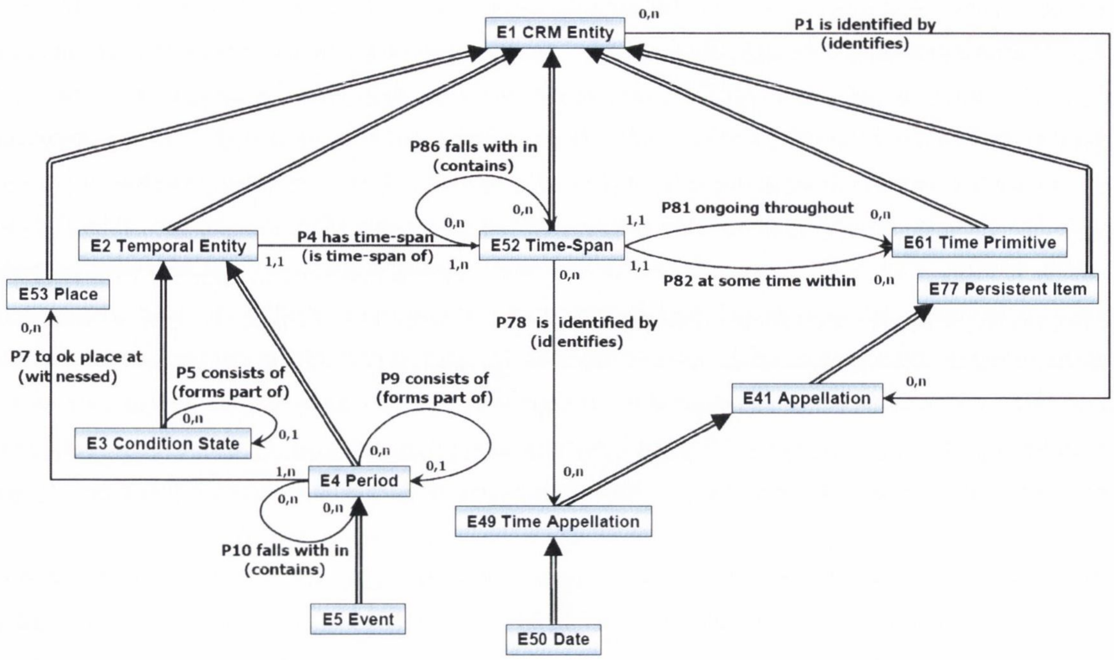

Figure 5.2 – Modelling temporal information in CRM. This figure appears as Figure 3 in the CRM reference document [25, p. xx]. Double arrows represent subsumption relationships (is-a) between CRM concepts and single arrows represent CRM properties, connecting a properties' domain and range.

represents endurants. These two classes are specified to be disjoint [25, p. xvi]. Cultural heritage artefacts, as suggested by the examples in the CRM definition, are naturally modelled as endurants, i.e. as (indirect) instances of 'E77 Persistent Item'. Temporal aspects about such artefacts, their creation, modification and destruction, are modelled as perdurants ('E2 Temporal Entity'). Examples of 'E2 Temporal Entities', given by the CRM documentation are: "Bronze Age", "the earthquake in Lisbon 1755" and "the Peterhof Palace near Saint Petersburg being in ruins from 1944 - 1946" [26, p. 2].

Information about 'E77 Persistent Items' is often recorded in the form of events in which these items participated. Events, which play an important role in the CRM, are modelled as a subclass of 'E2 Temporal Entities'. The relationship of an artist, for example, to his or her work is mediated through a production event. This has advantages over a simple link between artist and work: the temporal and geographical conditions of the production can be recorded and the role of other actors in the production process, e.g. a sitter for a portrait, can be expressed (Figure 5.3).

A number of properties are available in CRM to link 'E2 Temporal Entities' to 'E77 Persistent Items', e.g. 'P108 has produced', which connects production events to the physical man-made things they have created.[17]

'E2 Temporal Entities' have to be related by one and only one instance of the property 'P4 has time-span' to an instance of 'E52 Time-Span'.[18] The class 'E52 Time-Span' is characterised as follows in the CRM definition:

[17] Figure 5.3 illustrates the inverse property of P108, 'was produced by'.

[18] The restriction 'one and only one' can be read off the 'property quantifiers' on the arrows in Fig.5.2: '1,1' at the origin of the P4 arrow means, each instance of E2 takes part in at least one and at most one instance of the property P4.

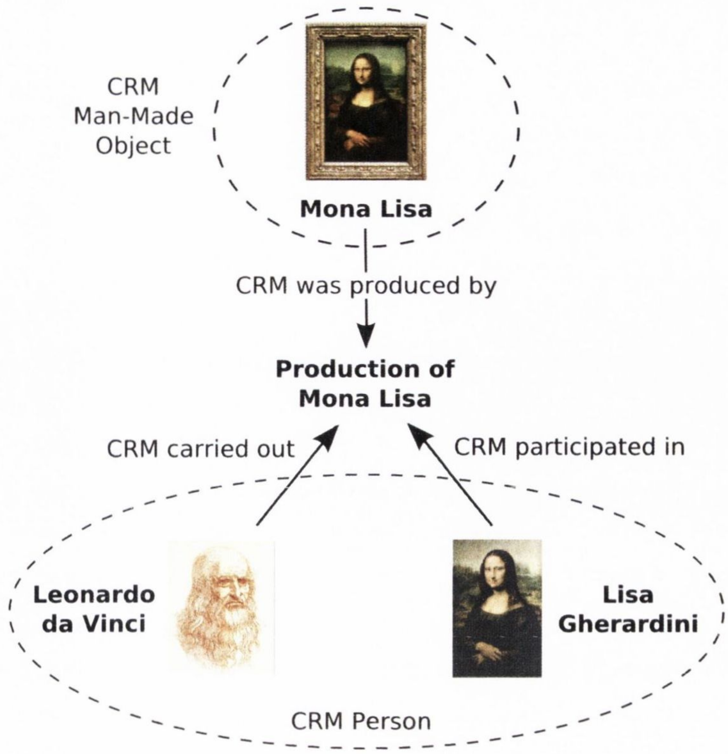

Figure 5.3 – Event-centred modelling of information in CRM, exemplified by the production event of Mona Lisa. This figure was published previously [85, p. 150].

> This class comprises abstract temporal extents, in the sense of Galilean physics, having a beginning, an end and a duration.
>
> [...]
>
> Since our knowledge of history is imperfect, instances of E52 Time-Span can best be considered as approximations of the actual Time-Spans of temporal entities. The properties of E52 Time-Span are intended to allow these approximations to be expressed precisely. An extreme case of approximation, might, for example, define an E52 Time-Span having unknown beginning, end and duration. Used as a common E52 Time-Span for two events, it would nevertheless define them as being simultaneous, even if nothing else was known. [25, p. 21]

The reference to Galilean physics is likely meant to express that time, as modelled in the CRM, should be seen as absolute, which is pragmatically adequate for all intended application contexts. Time-spans may be related to other time-spans through an instance of 'P86 falls within', which is intended as a purely mereological 'temporal part of' relation.

'E52 Time-Spans' are related to 'E61 Time Primitives' using the properties 'P81 ongoing throughout' and 'P82 at some time within' to capture uncertainties in temporal assignments. The intention behind the two properties appears to be to assign a core time to an event in which the event was definitely ongoing, and a wider time interval, which is perhaps too inclusive. P81 "describes the minimum period of time covered by an E52 Time-Span" [25, p. 58] and P82 "describes the maximum period of time within which an E52 Time-Span falls" [Ibid.]. Given the rather abstract nature of 'E52 Time-Spans' (compared to E49 and E61) it is surprising that they should be considered as "approximations of actual Time-Spans" [25, p. 21] and not representations thereof. After all the uncertainty surrounding the temporal extent of historical time-spans can be modelled at the level of 'E61 Time Primitive'. This is especially true, as the properties 'P81 ongoing throughout' and 'P82 at some time within' seem to be intended for exactly this purpose.

While we wanted to adopt a model of inner and outer time boundaries in principle and were encouraged to do so by our experts, we found that a concrete implementation of the CRM suggestions is not straightforward.

Any CRM implementation has to decide on a transition point between the abstract classes of the CRM and the concrete data types of the implementation platform. The CRM definition attempts to address this point with 'E61 Time Primitives'. But the details of the model are in conflict with the realities of OWL as an implementation language. The concrete data types most commonly associated with representing time information in OWL and RDF (xsd:date, xsd:dateTime)[19] can only represent 'atomic' time intervals (of the form day, hour, minute, second etc.). According to the semantics indicated in the CRM definition, however, 'E61 Time Primitives', include time intervals of arbitrary length (e.g. 27th of July 2012 to 12th of August 2012, the interval of the 2012 Summer Olympics).

One way to resolve this conflict, is to implement CRM properties with a range of 'E61 Time Primitive' (and similar 'primitive value' types) as data type properties with an untyped range. This was recently published as a recommendation [71]. At the time our prototype was created, however, this recommendation was not available and the OWL implementation of the CRM we originally intended to use[20] contained 'E61 Time Primitive' as an OWL class (i.e. not as a concrete data type) and used an especially introduced data type property to link 'E61 Time Primitive' individuals to concrete date and time representations (xsd:dateTime). Moreover, the implementation version suggested the use of two additionally introduced properties (i.e. properties not contained in the CRM definition) to express the beginning and the end of 'E52 Time-Spans'. Since the domain and range of the additional properties, however, were defined as time-spans (E52) and a superclass of 'E61 Time Primitive' respectively, these properties seemed to bypass CRM's own 'P81 ongoing throughout' and 'P82 at some time within', which have the same domain and a similar range.

Because of these difficulties we opted for an altogether different approach. 'E2 Temporal Entities' and 'E52 Time-Spans' were modelled as stipulated by the CRM. Time-spans, however, were linked directly to a concrete data type (for technical reasons we used xsd:int to represent calendar years, but xsd:date, xsd:dateTime or any other suitable data type could be used as well) through four data type properties: 'beganAfter', 'beganBefore', 'endedAfter' and 'endedBefore'. These four properties, as their names suggest, give an estimate of the beginning ('beganAfter', 'beganBefore') and the end ('endedAfter', 'endedBefore') of arbitrary time-spans. For contiguous time-spans these four values mark an outer boundary or most liberal estimate ('beganAfter', 'endedBefore') and an inner boundary or overly restrictive estimate ('beganBefore', 'endedAfter') of the temporal extent of a time-span. The following is an example of how the time-span of the Renaissance could be modelled:

'Renaissance' : 'P4 has time-span' : 'Time-Span of Renaissance'

'Time-Span of Renaissance' : 'beganAfter' : '1300'

'Time-Span of Renaissance' : 'beganBefore' : '1400'

[19] `http://www.w3.org/TR/xmlschema-2` (last accessed 14/08/2011).

[20] `http://www8.informatik.uni-erlangen.de/ecrm/erlangen-crm_090121_5_TQ.owl` (last accessed 14/08/2011).

'Time-Span of Renaissance' : *'endedAfter'* : *'1600'*

'Time-Span of Renaissance' : *'endedBefore'* : *'1700'*.

It is perhaps important to note that 'E52 Time-Spans' in this context are interpreted as the 'real' time-spans of events whose extent is possibly unknown. In particular the time-span of the Renaissance is not called 'Ca. 1350 to 1650' or similar, but only receives a numerical approximation through the four data type properties. In our project time-spans for any event 'X' were automatically generated and called 'Time-Span_of_X' (cf. Section 5.4). While this seems to be compatible with the CRM definition, it is not clear whether this is the only or even the dominant intended interpretation of 'E52 Time-Spans'. Many of the time-span examples given in the CRM definition are in fact numerical date ranges like "From 12-17-1993 to 12-8-1996" [25, p. 21] and therefore read interchangeably with examples for 'E61 Time Primitive', like "1994 - 1997" [25, p. 25]. It is by no means necessary that every time-span be linked to a primitive time representation through all four (or even any) of the data type properties. For some events it might, for example, not be possible or desirable to specify the inner estimates. The following is an example of how information about the temporal extent of the creation of Mona Lisa could be modelled:

'Production of Mona Lisa' : *'P108 has produced'* : *'Mona Lisa'*

'Production of Mona Lisa' : *'P4 has time-span'* : *'Time-Span of Production of Mona Lisa'*

'Time-Span of Production of Mona Lisa' : *'beganAfter'* : *'1503'*

'Time-Span of Production of Mona Lisa' : *'endedBefore'* : *'1519'*.

Of course there needs to be a conventional agreement as to whether or not the referenced dates (in our case years) are included or excluded in the time-spans in question. In order to work with dates as we would intuitively expect it, we need to make sure that our system deals with these statements in a certain way. According to the formal semantics of our knowledge representation language, the above statements just introduce a pointer to a numerical value. If we queried a knowledge base containing these statements for events that began after 1503, our answer set would duly contain the production of Mona Lisa. However, if we asked for events that began after 1502, we would not get the production of the Mona Lisa as an answer. This is despite the fact that we know that work on the Mona Lisa began after 1503 and therefore a fortiori it began after 1502. To address this we ensured that when users of our prototype system query for events that 'began after 1503', the formal request sent to our knowledge base retrieves all events that have a 'beganAfter' pointer to '1503' or a larger number. Similarly when users asked for events that 'ended before 1519', the formal query would include all events that had an 'endedBefore' entry of '1519' or a smaller number. In this setup there are combinations of 'beganAfter', 'beganBefore', 'endedAfter' and 'endedBefore' values that intuitively seem to be inconsistent (again 'intuitively' because they do not strictly cause an inconsistency according to the formal semantics of our knowledge representation language). For example, if we added the statement

'Time-Span of Production of Mona Lisa' : *'beganAfter'* : *'1520'*

this would, according to a natural interpretation of these statements, be in conflict with

'Time-Span of Production of Mona Lisa' : *'endedBefore'* : *'1519'*.

In general, the following three 'consistency conditions' need to be fulfilled for every event in our knowledge base:

(1) The maximum value of all 'beganAfter' relations needs to be smaller or equal to the minimum value of all 'beganBefore' relations associated with an event.

(2) The maximum value of all 'endedAfter' relations needs to be smaller or equal to the minimum value of all 'endedBefore' relations associated with an event.

(3) The maximum value of all 'beganAfter' relations needs to be smaller or equal to the minimum value of all 'endedBefore' relations associated with an event.

In terms of consistency checking it is not necessary to compare the values of 'beganBefore' and 'endedAfter' against each other although typically the latter will have higher values than the former for a given event. These consistency checks cannot be performed in the framework of our ontology language (OWL 2) alone, but require the application of additional rules or 'business logic' in the program code of an application working with such statements. In formulating the conditions above we have assumed that an event can be related to two or more values by an instance of 'beganAfter' or any of the other data type properties. This may or may not be desirable depending on the individual project. If several values are permitted, then 'beganAfter' can be declared a subproperty of 'endedAfter' and 'endedBefore' a subproperty of 'beganBefore' allowing for additional inferences in the case of partial information.

Using this setup, we exploited standard inference services that are available for ontologies to implement some basic temporal reasoning. Our aim was to let an 'E4 Period' that 'P10 falls within' another 'E4 Period' or an 'E52 Time-Span' that 'P86 falls within' another 'E52 Time-Span' inherit, in a sense, the outer temporal boundaries of the more inclusive period or time-span. A direct inference of this kind, however, is not possible. In our model 'E52 Time-Spans' receive their temporal estimates through four data type properties. In order to infer new data type values for time-spans directly, we would have had to combine an object property and a data type property in a property chain. Such a construction is not supported by our implementation language (OWL 2). In order to achieve the desired effect in querying the knowledge base, however, we introduced a property 'took_Place_During_TS' which connects temporal entities to time-spans. We declared this property a superproperty of 'P4 has time-span' and then used the feature of object property chains in OWL 2 to let it propagate along 'P10 falls within' and 'P86 falls within' (analogous to what we have discussed in Section 3.3). Time related queries submitted through our prototype user interface internally always make use of this property. In this way works of art that have no date related information associated with them, are still broadly accessible through temporal queries in virtue of being classified as a work of the 'Renaissance', for example, or any other of a number of frequently used art-historical classifications.

5.3 System Design and Implementation

The prototype system we have created, has two parts, a back end consisting of an ontology and knowledge base and components for processing these, and a front end consisting of a graphical user interface, that structurally corresponds to parts of the ontology. As an offline component we have created an ontology populator that converts information held in a database into an OWL knowledge base.

We will describe the overall architecture of our system (Section 5.3.1), the design of the graphical user interface (Section 5.3.2) and the input from domain experts we have received (Section 5.3.3). This is followed by details about the implementation of our system (Section 5.3.4) and the ontologies and resources we have reused (Section 5.3.5).

5.3.1 System Architecture

In this section we will describe the general architecture of our system, which is composed of

(A) an explicit knowledge representation which is comprised of

- (i) the CRM top-level ontology
- (ii) ART-T, the TBox element of ART ontology
- (iii) ART-A, an ABox element, comprising a representation of art historical periods
- (iv) a representation of a geographical knowledge base (extracted from GeoNames)
- (v) the "Knowledge Base Proper", an exchangeable and automatically generated ABox element holding information on works of art

(B) an application to access the represented knowledge (Artfinder) comprising three components

- (i) a graphical user interface, which acts as a visual query builder
- (ii) an off-the-shelf reasoner to execute the query
- (iii) a result manager to format the results for display

(C) The system is provided its knowledge base from a 'raw' data set on artists and works of art in an extract, transform and load process, which happens offline and involves the following components

- (i) a 'raw' data set (a digital art collection containing 41,543 works of art)
- (ii) Art Database, a MySQL database in which the entire digital collection is stored after some clean-up operations
- (iii) Ontology Populator, a program to load a subset of the collection (typically a few dozen to a few hundred works, 41 in our evaluation) into the "Knowledge Base Proper"
- (iv) Geo Populator, a program to extract geographic information from the GeoNames database

These components and their interdependencies are shown in Figure 5.4. The Figure reflects where external resources were integrated into our ontology. We are briefly going to elaborate on each of the three main areas (A, B, and C) before we discuss the design of the user interface and system implementation details in the following sections.

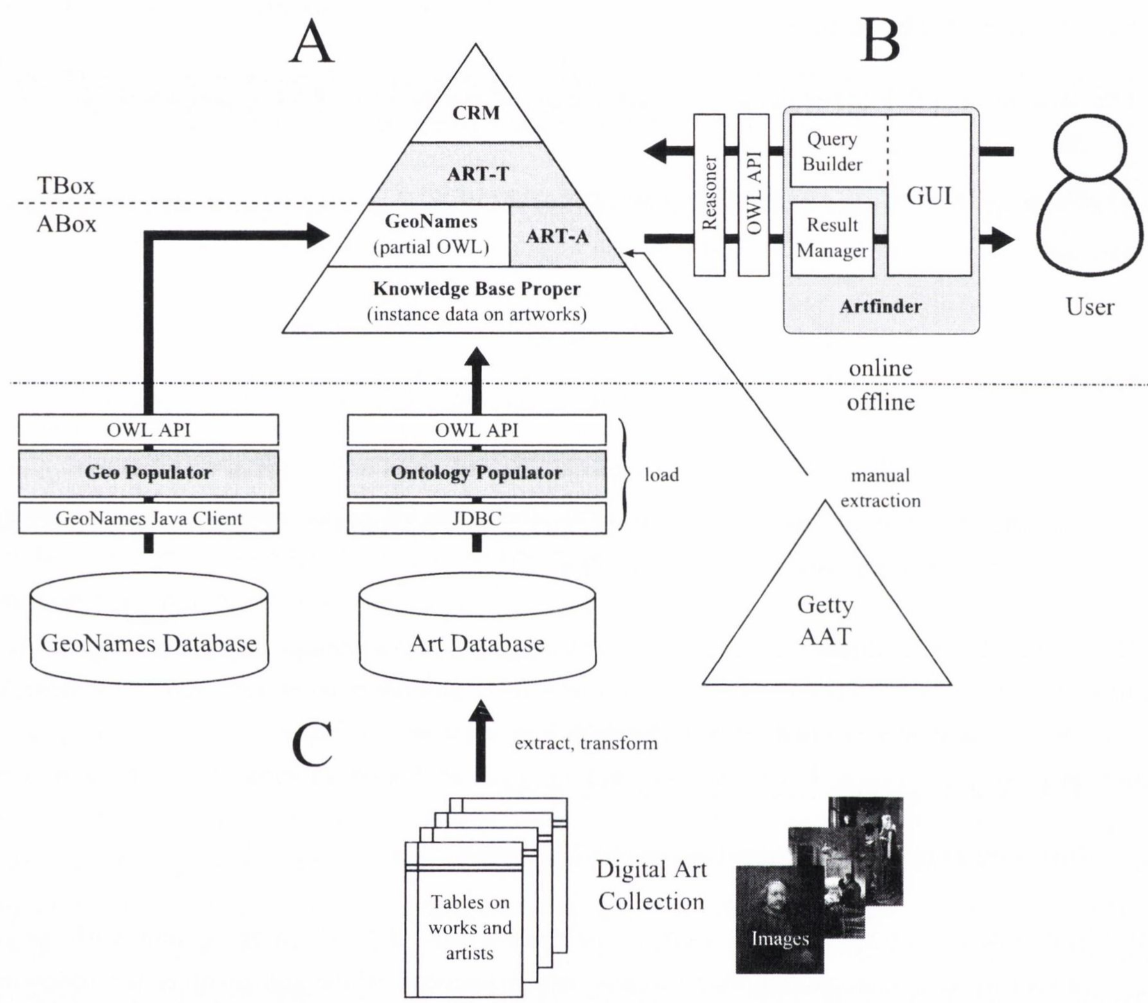

Figure 5.4 – An overview of our system components: A) *ART Ontology* (consisting of a TBox element ART-T, and an ABox element ART-A), together with the *CRM* top-level ontology and instance data on geographical locations and works of art; B) *Artfinder* comprising a graphical query builder, an off-the-shelf reasoner and a result manager for formatting the results; C) The offline pre-processing involved the storing of test data from a digital art collection in a MySQL database (*Art Database*) after some data clean-up; the population of a knowledge base of works of art (*Ontology Populator*); and accessing the GeoNames database through web services and extracting information relevant for our data set (*Geo Populator*). 26 concepts were manually included from the *Getty AAT*.

A) ART Ontology and Knowledge Base

The central back end component of our system is the ontology and its knowledge base. The distinction between ontology and knowledge base is usually drawn along the lines of TBox and ABox (cf. p. 42f.). If we regard the ontology, however, as the part that structures information about a domain and the knowledge base as the data belonging to that domain, then parts of our ABox should perhaps be considered as belonging to the ontology. This refers in particular to art-historical periods and geographical entities. In accordance with the modelling principles of the CRM and other, general ontological considerations art-historical periods are modelled as instances and are therefore part of the ABox. Similarly, nearly all the geographical information in our system is contained in the ABox. Concrete geographical entities, such as cities, countries and continents, are modelled as OWL individuals and only a shallow class hierarchy exists among their geographical types. Nonetheless, this kind of information has a more permanent character, than the assertions on works of art that can be swapped in and out and are, in some sense, the 'data proper'.

In terms of system development, the difference is, that the ontology in the wider sense (i.e. including 'art-historical' information in our ABox) is influenced by art expert input and the geographical information is taken from an authority database.[21] Assertions on works of art, on the other hand, are lifted from a digital art collection and can be exchanged for similarly structured data from the same or from another collection. These distinctions are reflected in Figure 5.4 and the concrete encodings corresponding to these different parts are maintained in different OWL files. The conceptual difference between the OWL fragments, however, should not detract from the fact that, in terms of their implementation, they are all made of the same material and are merged into one single ontology in our system at runtime.

B) Artfinder

The functionality of our front end can be divided into two stages. The start-up stage, in which the graphical query builder is created and the query stage, in which users enter information into the interface and the system returns matching results. During the start-up stage the graphical user interface is constructed, based on information contained in the ontology and knowledge base and a configuration file which specifies, for example, the category of objects that are to be retrieved and the ontology classes and relations that are important in a particular query set-up. The interface is constructed in such a way that elements of it, depending on the input received, contribute parts of the query in a certain predetermined way. Once the query is created, the mechanics of the retrieval are performed by standard reasoning software. The design of the interface is the subject of Section 5.3.2.

C) Ontology Populator and Further Pre-processing

The construction of our knowledge base started from the 'raw' data on works of art and artists in a commercially available collection of digital art.[22] The strategy we decided to use in order to create a sample knowledge base can be described as an extract, transform and load approach (ETL). The information contained in the art collection was extracted through the user interface that came with the collection (there was no export functionality similar to a database dump). The structure and quality of the data and idiosyncrasies of the publication format necessitated a number of transformations and

[21] `http://www.geonames.org` (last accessed 10/9/2011).

[22] For a detailed description of the collection, compare the section on "Test Data" on page 121.

clean-up steps. As a target structure for the transformation process we chose a database, because it is more robust than a spreadsheet application in performing large scale transformations and offers a more reliable and scalable persistence solution than an OWL file containing instance information. Our experience of working with large OWL files in the Protégé ontology editor has shown that such files can easily get corrupted. As the database forms an intermediate step between the original, unmodified data collection and the OWL knowledge base, two natural alternatives presented themselves for the structure of the database schema: it could be closely aligned with the schema of the original data, or it could be aligned with the OWL constructs used in our knowledge base. The second option is non-trivial if it is built from scratch and would therefore have to rely on an existing solution from another party. In the early stages of our project no database back end for ontologies was available that preserved the structure of OWL constructs. With the publication of such a 'native' persistence layer for OWL the second option became feasible [69]. Nevertheless, we decided to pursue the first approach. The reason for this was that the original data, once cleaned up, would not be subject to many changes, while the modelling details on the side of the ontology and knowledge base were in a flow during the project. It therefore seemed natural to preserve the result of the data clean-up as a source from which OWL knowledge bases of various sizes and characteristics can be created according to the current need. The creation of a knowledge base from the database is carried out offline (i.e. not as part of our system proper) by *Ontology Populator*. This program implements rules and heuristics for transforming part of the information contained in the database tables into OWL statements about individuals. A similar program, Geo Populator, accesses the Semantic Web service of the GeoNames database[23] and creates an OWL file of relevant geographical entities, modelled as individuals. A detailed description of the extract, transform and load process is given in Section 5.4.

5.3.2 Interface Design

The objective of our prototype system is to give users access to the information stored in a knowledge base. In order to achieve such access a user may:

(a) inspect a visual representation of the knowledge base, for example a tree diagram or a semantic network with labelled vertices and edges,

(b) browse a concept or relationship hierarchy, enhanced with instance information,

(c) interpret and manipulate subject-predicate-object triples in RDF based systems,

(d) utilise a query language, similar to SQL or other databases languages

While many of these methods are regularly used by the developer of a knowledge base, they have various drawbacks for an end-user who is unfamiliar with the conventions of knowledge representation. Visual inspection, even of interactive hierarchies or diagrams, does not usually reveal the full inferential potential of a knowledge base, but at best only basic inferences. Triples do not lend themselves to readable representations of complex (OWL) ontology constructs. And while query languages allow for querying the inferential closure of a given set of facts in a particular representation language and logic, and therefore grant access to information that is implicitly contained in a knowledge base, their statements can be quite far from the natural language expressions they are trying to model and their use requires a significant amount of training. In order to be in a better position to

[23] `http://www.geonames.org/export/ws-overview.html` (last accessed 29/08/2011).

answer our third research question, whether end-users can successfully formulate complex ontology queries (p. 5), we decided to adopt a different approach.

The ontology and knowledge base are the concrete, machine-interpretable representations of what our system "knows". Knowledge among humans, on the other hand, has no such concrete representation, at least none that can be universally agreed upon. When knowledge is passed on between humans this usually happens through the use of language, written or spoken. At the human-computer interface query languages are often used to retrieve information from structured data representations, especially databases.

Knowledge representation languages also have query languages associated with them. SPARQL is the W3C recommended query language for RDF.[24] Because OWL has RDF serialisations, SPARQL can be used to query OWL ontologies. However, it does not natively support OWL specific constructs and can not directly query entailments that stem from such constructs. SQWRL is a query language for OWL, that is derived from the SWRL rule language [119]. It is based directly on the DL semantics of OWL and not dependent on a particular serialisation of it. However, it has not received the same implementation support as SPARQL.

Finally, OWL itself can be regarded as a limited query language for OWL knowledge bases. Since class expressions can specify complex and nested conditions for membership, OWL classes can be interpreted as queries. Their instances would in this case be regarded as the result set to a query. Because of the absence of variables that OWL inherited from its description logic underpinnings, it naturally has limitations as a query language. It is not possible, for example, to formulate queries that return sets of tuples as results, like 'Return all student teacher pairs!'. Likewise, and more importantly, it is impossible to formulate a query in which two statements bind to the same variable, as is the case in a query for 'all artist couples who shared the same teacher'.

In our conversations with art and heritage experts, however, we found that their searches and information needs are often centred on the category of works, i.e. desired result sets typically contain works of art that fulfil certain criteria. OWL class expressions offer the possibility to chain restrictions together, as in 'Paintings that were created by someone, who was born in a place that lies in Italy'. They also offer the possibility to connect several branches of restriction chains through logical 'and' or logical 'or' links: 'Prints that were made by a painter or a draughtsman who was born in Italy and who died in France'.

In order to test whether complex and nested queries can successfully be formulated by users unfamiliar with OWL DL, a graphical user interface was designed and implemented. A balance had to be struck between an interface which is easy to use on the one hand, and which adequately reflects some of the structure of OWL DL queries, on the other. The guiding idea for our design of the graphical user interface (GUI) is that of the classical definition by 'genus proximum' and 'differentia specifica'. This type of definition goes back to Aristotle. It is characterised by defining a term by giving its nearest (i.e. most specific) type (its genus proximum) and supplementing that with information on what specifically distinguishes it from other elements of that type (the differentia specifica). Even though the term 'genus proximum' suggests only one genus, this part of the definition may contain several genera. While nowadays the structure of genus proximum and differentia specifica is no longer seen as necessary in mathematics and formal logic, it remains a dominating paradigm for natural language definitions and the way people construct concepts [162, p. 20].

[24] `http://www.w3.org/TR/rdf-sparql-query` (last accessed 20/08/2011).

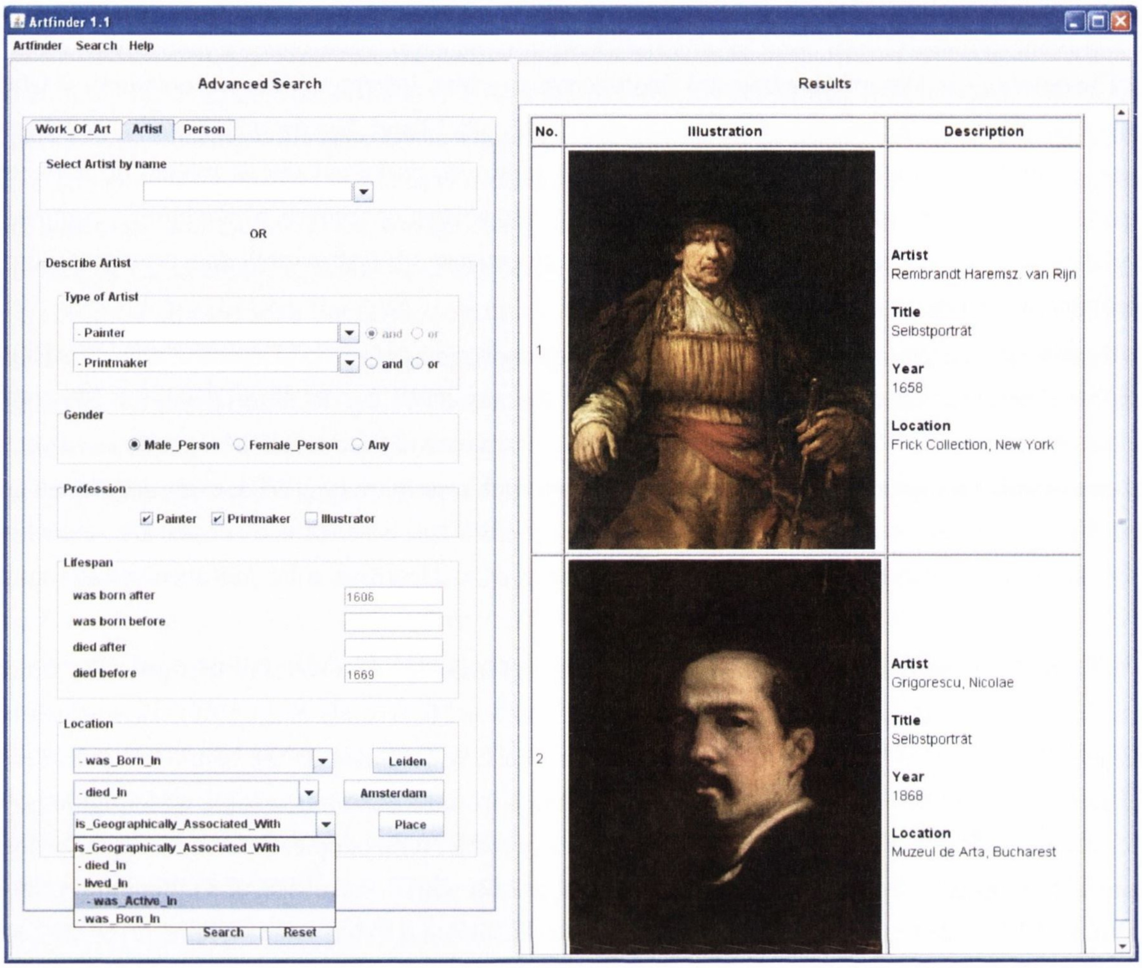

Figure 5.5 – A screenshot of the *Artfinder* graphical user interface.

Discussing our ontology with domain experts and preliminary tests suggested that a similar structure is found in many queries that are desirable to run against a knowledge base about works of art. In most queries a named OWL class (genus proximum) was refined by one or more OWL restrictions (differentia specifica). Naturally, this has to be seen more as a loose analogy, than a precise characterisation. A distinction that hinges on the difference between named OWL classes and OWL restrictions is arbitrary as it depends on the axioms in the knowledge model (a named OWL class and an OWL restriction may be logically equivalent in an ontology, as is the case with all 'defined classes'). Furthermore, it became clear from preliminary experimentation, that many natural language queries translated into OWL constructions that involved someValuesFrom-quantifiers, rather than allValuesFrom-quantifiers.

The interface we have created is divided into two sections. On the left, a user can form his or her query using several input fields and selection menus. On the right, a table displays the results of a query: Images and basic information about the works of art that match the specified criteria (Figure 5.5). In the following, we will describe the left hand side, where the query is constructed. We will first sketch the general idea and subsequently go into more detail.

In order to allow users to express the queries we have discussed above through a graphical interface, we wanted to structure the interface in the same way as the queries are structured. Certain segments of the interface should represent classes in our ontology and these segments would have to be linked by interface elements that correspond to object properties in our ontology. After experimenting with a number of different layouts, we decided to use tabs as the segments that represent classes and link them through labelled buttons, representing the properties, to switch between the tabs. The first tab is privileged in the sense that it represents the target category that queries are directed at. All other tabs generate restrictions on the class represented by the first tab. Which classes are represented by tabs and which properties are represented by the button elements are defined in a configuration file, which is consulted when the user interface is generated. Other details of the interface are derived from our ontology and the associated knowledge base.

The interface tabs correspond to a named class in our ontology. In the following, we will refer to this class as the 'eponymous class', since the name of each tab is derived from the name of this class. The main tab, in first position, corresponds to the type of object that is to be retrieved. In the standard case we are going to discuss now, the eponymous class of the main tab is the class 'Work of Art'. This means works of art were the target category that the participants in our evaluation set-up were looking for.

The interface elements on each tab serve to further restrict the class that the tab stands for. At the top of every tab, other than the first one, there is a combo box that allows the user to select an instance from the eponymous class of the tab. The combo box is populated by all the instances of the tab's eponymous class contained in our knowledge base. As soon as an instance is selected from this combo box, all other user interface elements on the same tab are disabled. This happens, because once a particular instance has been selected there is no need for further restrictions on the tab.

Additionally to the instance combo box, in case it is present, every tab in our interface contains two further groups of GUI elements. The first one serves to specify what we have called the 'genus proximum' above, the second one to specify the 'differentia specifica'. The 'genus proximum' elements let the user restrict the class that a tab stands for by choosing a named class from our ontology that is narrower than the eponymous class that the tab represents by default. Two different varieties of interface elements are offered to the user. The first is another combo box, this time not populated with ontology individuals, but with all the subclasses of the tab's eponymous class known to our ontology. This includes both direct and indirect subclasses. A possible hierarchy among the subclasses is visualised as an indented list in the drop-down menu of the combo box. The second variety of input elements consists of grouped radio buttons or checkboxes, that allow the user to select classes from our ontology. Groups of radio buttons represent mutually exclusive choices, while checkboxes allow for multiple selections. These elements are intended to give users direct access to salient and frequently used classes. Each checkbox or button is linked to a class in the ontology. The configuration file, that defines the search set-up, specifies, which classes are grouped together and whether they represent mutually exclusive choices or not.

Below the elements that specify the 'genus proximum' another group of interface elements addresses the 'differentia specifica'. All of these elements are divided into two parts, of which the first represents a (data type or object) property in the ontology and the second an input method for the user to enter a value for the property. In the case of data type properties, the property is represented by a simple label and the entry method is a text field into which the user can type a literal value. In the case of object properties, the property and its subproperty hierarchy, as defined in the ontology, is populated into a drop-down list. The user can select any property from this list, either the top property, which is selected by default, or one of its specialisations. His or her input will later be evaluated accordingly. The input element for object properties is a button, which, when clicked, takes the user

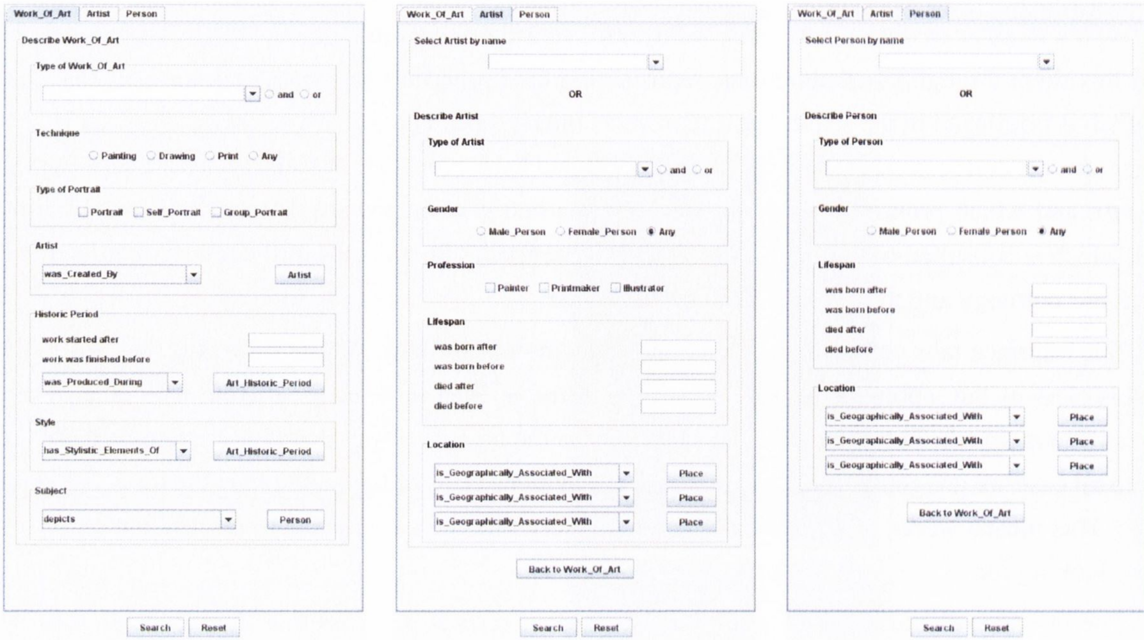

(**a**) The *Work of Art* tab. (**b**) The *Artist* tab. (**c**) The *Person* (sitter) tab.

Figure 5.6 – The three main tabs of the user interface in the configuration we used for our evaluation sessions.

to a tab that represents the range of the object property in the ontology. This tab in turn represents a class in our ontology in the same way that we have just described. This pattern can be iterated arbitrarily.

Not all of the elements or sets of elements we mentioned need to be present in every tab. Our main tab, 'Work of Art', for example, does not contain a combo box to select individual works directly. Choosing an instance from such a box on the first tab would be equivalent to querying for a work of art by selecting from a controlled list of titles. While this is perfectly meaningful, we were not interested in minimalist queries of this kind for the purposes of our evaluation.

Another kind of tab, that does not contain the full range of GUI elements, is a 'terminal' tab: The process of specifying class restrictions through tabs that are linked via object properties to other tabs has to terminate at some stage. At the end of a chain of object property links there needs to be a tab that contains only an instance selection box and perhaps input fields for data type property values. In a case where a terminal tab contained just an instance selection box, it was not actually implemented as a tab in our interface, but as a small, separate window, which pops up when needed and closes as soon as the user has made a selection. This simplifies the layout of our query input elements. In our evaluation set-up the 'Artist' and 'Person' tabs pointed to several 'Place' windows which contained a combo box populated with geographical place names. Figure 5.6 shows the three main tabs of our interface set-up, excluding the pop-up windows for place names.

5.3.3 A Note on Expert Input to System Design

In order to ensure that the functionality provided by the graphical interface can be interpreted correctly by potential users and better meets their expectations, early versions of Artfinder were tested by expert users in open-ended, semi-structured feedback sessions. Three heritage professionals from the National Gallery of Ireland were approached to comment on the interface and ontology design: a collection management expert familiar with the database systems used in the gallery, a curator in the department of prints and drawings and an assistant in the Centre for the Study of Irish Art at the gallery, who was involved in the creation of a visual archive search system.

The sessions were conducted on an individual basis and structured by an interview guide. The structure changed from session to session, but the material covered was roughly the same. The first two sessions differed only minimally and consisted of the following parts:

- Personal information on the participant
- Introduction to the digital art collection
- Introduction to ART ontology
- Example queries in natural language and description logic
- Discussion of modelling aspects in the ontology
- Introduction to the query interface
- Test queries
- Overall feedback and suggestions for improvement

The order in which these topics were covered changed for the third session. This session was structured as follows:

I Front end aspects
- Introduction to the digital art collection
- Introduction to the query interface
- Introduction to ART ontology
- Test queries
- Feedback on the interface

II Back end aspects
- Discussion of modelling aspects in the ontology
- Example queries in natural language and description logic
- Examples of inferred knowledge
- Feedback on modelling aspects

III Personal information on the participant

The rationale behind the change was to see if someone unfamiliar with the structure of the underlying knowledge base would still be able to successfully formulate queries using the graphical user interface.

Aspects of the ontology modelling that were discussed with the experts found their way into ART ontology. For instance, a detailed hierarchy of printmakers or suggestions on how possible relations between an artist and a geographical location should be ordered. Recall, that the design of the ontology was informed by expert input independently of these feedback sessions. Here, Dr Adriaan Waiboer, curator for Northern European Art in the National Gallery of Ireland was a chief influence.

After positive feedback on the user interface in the first two sessions and a demand for more options and fields, the complexity of the interface was increased from one session to the next until a stage was reached where we felt that additional features may be perceived as confusing.

5.3.4 System and Interface Implementation

Almost all of the code for our project was written in the Java programming language. There were several reasons for choosing Java over other languages in our context. The implementation of our interface benefited particularly from Java's object oriented outlook, as our interface elements are closely aligned with ontology elements, which resemble object oriented structures. Moreover, the reference implementation of the OWL API which we used for program-controlled ontology manipulation (see p. 46) is written in Java. The OWL API forms an interaction layer between the ontology and our application and pre-processing code. Querying the ontology is facilitated by reasoners to which the OWL API provides general purpose interfaces. Similarly, the part of our code that did not 'face' the ontology was also well supported by Java APIs: the Java DataBase Connectivity (JDBC) API for interacting with our database and a Java client library provided by the GeoNames team for using their web services. Lastly, the Java programming language is a convenient choice for creating graphical user interfaces.

The Artfinder application we created for accessing our knowledge base follows the "thick data, thin applications" paradigm of the Semantic Web. According to this paradigm rich information is contained in the data, in our case the ontology and knowledge base, and the application processing the data has comparatively little complexity in its business logic. The functionality that Artfinder provides over and beyond the GUI is 'thin' in this sense and closely tied in with the GUI elements.

We used Swing, the primary GUI widget toolkit of the Java programming language, to create the user interface for the Artfinder application. The ontology handling capabilities of Artfinder are based on the OWL API and ontology inference is delegated to a third-party reasoner (cf. Section 3.1.3).

In the object oriented design of Artfinder a key Java interface[25] is `SemanticGUIElement`. All the user interface elements that correspond to ontology elements and contribute to the creation of the query, have to implement this Java interface. Its main feature is that it ensures GUI elements define a method `evaluateUserInput()`, which returns an `OWLClassExpression` (this class is imported from the OWL API). The intended function of `evaluateUserInput()` is that it should determine a GUI element's contribution towards the query as a function of the states of the element's individual widgets which in turn depend on user input. Typically this is done by calling the `evaluateUserInput()` method of these widgets in a recursive fashion. Examples of GUI elements that implement the interface are drop-down menus for selecting individuals or classes, the grouped radio button panels mentioned in Section 5.3.2 and also the tabs containing these and other elements shown in Figure 5.6.

After starting up, Artfinder instantiates a semantic handler which bundles reasoning and ontology interaction services and a factory class for creating 'semantic GUI elements' (i.e. instances of classes which implement the `SemanticGUIElement` interface). The factory class is used for creating the

[25]We use the term "Java interface" to distinguish this meaning of "interface" from the sense of graphical user interface.

layout and links of the tabs that define a specific query configuration (see Figure 5.6 for an example of such a configuration). This layout is defined in a configuration file which contains the identifiers of OWL classes and properties that need to be represented in the GUI. The configuration determines, for example, in the shape of an OWL class, what kind of records are to be retrieved, works of art, artists etc. The factory class, which has access to the semantic handler and thus to a reasoner, completes the information provided in the configuration file by filling in subclass and subproperty hierarchies and instance information as found in the ontology.

Once a particular interface configuration is created, a user can specify a query filling in or manipulating the GUI elements. As soon as a query is submitted to the system, the evaluation function of the main panel (representing the category of objects that are to be retrieved) is called and (recursively) determines the OWL class expression corresponding to the query. This expression is then used by a reasoner to identify all instances of this class expression in our knowledge base. An illustrating image and basic information on these instances is provided to the user as the result of the query.

5.3.5 Implementation of Reused Ontologies

We have already discussed the reuse of existing ontologies (Section 5.1.2). From an implementation point of view, we can express the technological and licensing aspects in this context, as a list of questions: Does a machine-processable implementation of the ontology exist? Is it freely available or can it be licensed? Which knowledge representation format is it encoded in? Are modifications to the ontology necessary?

As we decided to use the CIDOC CRM as a top-level ontology, its implementation in a knowledge representation format compatible with our ontology (i.e. OWL or RDFS) was of key importance. A standardised and well-documented top-level model, the CIDOC CRM has received a number of implementations (Table 5.2). As the tabular overview shows, implementing an ontology that is specified in natural language in a concrete representation formalism is by no means straightforward, even if the specification is reasonably precise, as is the case with the CRM. Some of the implementations listed in Table 5.2 only implement parts of the CRM model. Proposals differ in how they implement various aspects of the model. Some include elements that are not part of the CRM definition. Naturally, implementations also reflect the evolution of the CRM across several versions. At the beginning of our project we used the implementation provided by Gangemi (5.2, number 5). For the most part, however, we worked with a series of implementations that is maintained by the University of Erlangen-Nürnberg [59]. The version we used for the evaluation was number 17, at the time the only implementation of CRM version 5.0.2 available. Reinhardt's implementation (number 11) deserves to be mentioned as the first one, to our knowledge, to use OWL 2 as an implementation language. Although our ontology uses OWL 2 constructs, we decided against this implementation, as some of its modelling details did not agree with the needs of our project. Since Reinhardt's implementation was developed with a view to linked data projects, the inverses of CRM properties are not modelled as properties in their own right, but instead are assigned RDFS labels that denote the meaning of their inverses. The University of Erlangen-Nürnberg team has recently begun to work on an OWL 2 implementation of their own and has published an OWL 1 DL implementation of version 5.0.3 of the CRM [26] (see Table 5.2). Even more recently an implementation of CRM 5.0.4 [27] was made available which is not yet reflected in Table 5.2.[26]

[26] It can be found at `http://erlangen-crm.org/111201/` (last accessed 20/03/2012).

CRM version	No.	Created by	Format	Release date	Notes	URL (shortened redirect)
3.2	**1**	FORTH-ICS	RDFS	October, 2001	partial definition	`http://tinyurl.com/444tvya`
	2	FORTH-ICS	DTD	February 5, 2002		`http://tinyurl.com/3dvzgcm`
3.4.9	**3**	FORTH-ICS	RDFS	November 20, 2002	partial definition	`http://tinyurl.com/3ernfan`
4.2	**4**	FORTH-ICS	RDFS	June 27, 2005	partial definition	`http://tinyurl.com/3z5rgwe`
	5	A. Gangemi	OWL	June, 2007	partial definition	`http://tinyurl.com/44lu5e5`
4.2.4	**6**	Erlangen	OWL 1 DL	September 1, 2008		`http://tinyurl.com/4ymqhdx`
4.2.5a	**7**	Erlangen	OWL 1 DL	December 16, 2008		`http://tinyurl.com/3jhxod7`
5.0	**8**	Erlangen	OWL 1 DL	January 20, 2009		`http://tinyurl.com/3dykpd8`
5.0.1	**9**	FORTH ICS	RDFS	March 31, 2009	partial definition	`http://tinyurl.com/3htwvnc`
	10	Erlangen	OWL 1 DL	July 14, 2009		`http://tinyurl.com/3kckfhs`
	11	S. Reinhardt	OWL 2	September 3, 2009		`http://tinyurl.com/3bant25`
	12	Erlangen	OWL 1 DL	September 22, 2009		`http://tinyurl.com/42lo4cg`
	13	FORTH-ICS	RDFS	November 23, 2009		`http://tinyurl.com/3unmgba`
	14	Erlangen	OWL 1 DL	November 25, 2009	March Edition	`http://tinyurl.com/3wc8h47`
	15	Erlangen	OWL 1 DL	December 17, 2009	November Edition	`http://tinyurl.com/3m8s28c`
	16	Erlangen	OWL 1 DL	March 2, 2010	November Edition	`http://tinyurl.com/4ymb6u9`
5.0.2	**17**	Erlangen	OWL 1 DL	July 7, 2010		`http://tinyurl.com/3jfo27d`
	18	FORTH-ICS	RDFS	August 11, 2010		`http://tinyurl.com/4ypnetj`
	19	Erlangen	OWL 1 DL	October 1, 2010		`http://tinyurl.com/3emgpdm`
	20	FORTH-ICS	RDFS	December 20, 2010		`http://tinyurl.com/3hmqrrp`
	21	Erlangen	OWL 1 DL	February 24, 2011		`http://tinyurl.com/3sccl43`
	22	Erlangen	OWL 1 DL	March 17, 2011		`http://tinyurl.com/3d5ffez`
	23	Erlangen	OWL 1 DL	April 4, 2011		`http://tinyurl.com/3mj4v9r`
	24	Erlangen	OWL 2	April/May, 2011	preliminary draft	`http://tinyurl.com/6a44zzj`
5.0.3	**25**	Erlangen	OWL 1 DL	November 22, 2011		`http://tinyurl.com/6v949sj`

Table 5.2 – A partial list of implementations of the CIDOC CRM in the last ten years. We mostly used the series of implementations issued by the University of Erlangen-Nürnberg (shortened to "Erlangen" in the table). The URLs provided in the last column redirect to the ontology files.

Geographical information in our knowledge base was taken from the GeoNames database (cf. Section 2.3.4). GeoNames offers a range of web services for accessing its data. These include services for searching the database with keywords ('search') and for finding spatially more inclusive geographical entities for a given geographical record ('hierarchy'). Results are provided in XML format or a number of other formats. We extracted the geographical entities relevant to our data collection based on the analysis of the columns in our database mentioned in Section 5.1.3. Geographical names occurring in our data were extracted from different columns and merged into a single list without duplicates. We created a 'geo populator' program to process this list and generate an OWL knowledge base with geographical information. For each location name on our list, the populator uses the 'search' service to obtain the most likely geographical record corresponding to this name. As our data is in German, this step benefited from the multilingual labels in GeoNames and provided us with a free translation of location names into English. The 'hierarchy' service is then used to obtain a sequence of successively more inclusive geographical regions (e.g. Florence, Municipality of Florence, District of Florence, Region of Tuscany, Republic of Italy etc.). During this step some of the lower level administrative regions, corresponding to boroughs, districts and parishes were discarded. Based on the feature codes in GeoNames (cf. Table 2.3 on page 32) the remaining toponyms were assigned to one of five ontology classes: populated place, first-level administrative division, country, continent or, if none of these four applied a general 'place' class ('E53 Place' from CRM). Links between geographical children and parents delivered by the 'hierarchy' service, were recorded in the knowledge base using a property to express spatial inclusion ('P89 falls within' from CRM).

Finally, a small portion of the Getty AAT (cf. Section 2.3.1), related to art historical periods, was used in our ontology. This was recreated manually and the AAT terminology and ID numbers were recorded using `rdfs:label` and other annotation properties.

5.4 Ontology Population

By ontology population we understand the entire chain of pre-processing steps that led from the original, unmodified data collection on art to the OWL file that represented our knowledge base. We will first describe the test data, then the data clean-up and finally the algorithm that created the OWL file from the cleaned up data.

5.4.1 Test Data

To apply our ontology to a knowledge base of works of art a set of test data was required. The image set and associated information that was used for testing our approach to ontological markup and querying was taken from a commercially available German language digital collection of art: "40000 Meisterwerke", published by The Yorck Project, Berlin.[27] The collection contains over 40,000 works that are classified into one of three broad categories: painting, drawing and print.[28] There are nearly 20,000 paintings, roughly 9,200 drawings and just over 12,500 prints. The time span covered by the collection ranges from antiquity up to the beginning of classic modernism. For copyright reasons the collection only includes works of artists that died before 1937. The main focus is on European art, but artists and works from the American, Arabic, Chinese and Japanese traditions are also reflected.

[27] The Yorck Project was since renamed to Zenodot Verlagsgesellschaft mbH.

[28] The corresponding category names in German are "Malerei", "Zeichnung" and "Grafik".

The edition is part of the series "Digitale Bibliothek" by Directmedia publishing, a sister enterprise of The York Project, and is accessible using a graphical user interface that is bundled with the publication. The information is contained on two DVDs, one for paintings and the other for drawings and prints. The records on works of art and artists are organised in tabular form similar to a database. As there are four separate tables in total: one for paintings, one for drawings and prints, one for the artists of the paintings and one for the artists of the drawings and prints. There is no overlap between the painting and drawing tables, but there is a significant overlap between the two artist tables (i.e. there is a considerable number of artists which produced paintings as well as drawings or prints and were hence listed on both tables). The image files illustrating the works of art were stored in three different sizes on the DVDs.

5.4.2 Data Transformation

To make it easier to access and manipulate the information we decided to feed it into a MySQL database. Unlike the images, which were available as JPEG files, the tabular information associated with the images was not readily accessible, for example as a text file. The tables could only be viewed through a proprietary graphical user interface. As the GUI application (an e-book reader program) only allowed a maximum of 256 table rows at a time to be copied and pasted into another application, the entirety of the 45,848 entries (41,543 works of art, 4,305 records on artists) in the four original tables was transferred into four separate sheets of a spreadsheet application[29] in a bit by bit process.

Before the tables were loaded into a database a number of clean-up operations were performed. The four tables were merged into two tables, for artists and works, removing the separation of works into paintings and prints and drawings due to the distribution format of the collection. As we mentioned there were a number of artists that appeared on both of the original artist tables, as they had produced paintings as well as prints or drawings that were included in the collection. We attempted to single out duplicate entries using comparisons of artist name and additional fields, such as birth date or birth place. This was done to ensure that two different artists sharing the same name were not accidentally identified with each other. In six cases the information contained in duplicate artist entries, although identical in every respect, was so sparse and so generic that we did not consider it save to assume the entries referred to the same individual. In some of these cases the 'artist name' consisted of a description of an anonymous artist, such as "French Master around 1650", and contained little information beyond that. Apart from these exceptions duplicates were unified into one record, reducing the total number of 4,305 artist entries from the two original tables to 3,846 in the merged one.

The digital art collection we used is structured as an electronic book with each artist and each work of art residing on a virtual page. A page on an artist is followed by the pages of all of his or her works until the page of the next artist in the alphabet appears. This interwoven sequence of artist fact sheets and artwork records helped us to re-establish the link between artists and their works in our database as the e-book page numbers were recorded in the tables we extracted. Modifying the original table structure we assigned a range of page numbers to each artist, that included all the pages of his or her works. If an entry in our artist table resulted from the merging of two duplicates, we recorded two page ranges, one for the artist's paintings and one for his or her prints and drawings. Finally, we

[29]We used OpenOffice.org Calc, as it afforded greater flexibility in exporting the data to comma or tab separated value files than comparable software. See `http://www.openoffice.org/product/calc.html` (last accessed 19/10/2011).

converted the 'category' fields for works of art, which signify whether a work is a painting, a drawing or a print from a textual to a numeric representation.

This data was exported from the spreadsheet application as a tab separated value file and imported into two tables of a MySQL database, 'Artists' and 'Works'. In this process, each table was expanded by an additional column containing an automatically incremented ID number as a primary key. A third table, 'Categories', was created to which the numerically encoded category information (i.e. painting, drawing or print) in the table of works pointed as a foreign key.

The full range of entries in most columns of the 'Artists' and 'Works' tables where extracted as a list, inspected and analysed.[30] We attempted to develop an idea of the nature and variation of the data typically recorded in each column. Data from columns of the same or similar 'type' was merged (e.g. Location, Country, Place of Birth, Place of Death, Place of Activity). Complex strings like "Paris, Arles, St.-Rémy, Auvers-sur-Oise" or "Gotik, Vorrenaissance" (Gothic, 'Pre-Renaissance') were broken up and the individual components compiled into lists of the 'basic vocabulary' used in each column. Perhaps the most complex and irregular entries are those referring to date related information. This led us to the development of an algorithm for converting text strings describing dates in the database to ISO 8601 compatible date representations.[31] Some of these strings, for example "29.09.1815", are readily parsed into ISO 8601. Others, like "1659/60", have to be mapped to a time span, and still others require more complex heuristics for their interpretation. Examples of the latter category include "um 1510" (around 1510), "vor 1480" (before 1480), "Januar 1604" (January 1604), "21. August 1872", "1. Hälfte 16. Jh." (1st half of the 16th century), to name but a few of the many different types of date strings. We used a Perl script with regular expressions to implement the various heuristics. Approximate specifications of dates, for example, are mapped to longer intervals if they are farther from the present time or if their years consist of round numbers. The entry "um 1600" (around 1600), for example, is mapped to the ISO 8601 interval "+1590/+1610", whereas "um 1611" is mapped to the shorter interval "+1610/+1612".

In summary, we can say that the data we used as the starting point of our knowledge base population appeared to fit Hyvönen's description of typical heritage legacy data quite accurately: "Content in memory organisations is usually available as relational legacy databases, whose annotations are literal terms and free text descriptions. Such annotations are often intended for human usage, use various syntactic conventions, are often semantically ambiguous, and may contain syntactic typing errors." [82, p. 766]

In sessions preliminary to our main evaluation we showed the database tables to three experts form the National Gallery of Ireland and they confirmed that our data is very close to the kind of information they would typically keep about works of art. The only noteworthy difference is that they keep more detailed records on the provenance of works, whereas our data only contains information on the current location of a work (cf. Section 5.3.3).

5.4.3 Knowledge Base Creation

We created a Java program for transforming part of the information held in the MySQL database into an OWL file containing instance information. As an OWL file encompassing all the information in the database would be too large to be held in the memory of a typical desktop computer, our *Ontology*

[30] Some of this work was carried out on the basis of the tab separated value files using Perl scripts.

[31] `http://www.iso.org/iso/catalogue_detail?csnumber=40874` (last accessed 20/10/2011). Compare also our previous remarks on ISO 8601 on page 102.

Populator program accesses only a subset of the works of art contained in the database. This is either done by filtering entries on certain criteria or by randomly selecting a given number of rows from the 'works of art' database table. For each selected work, the values of all attributes in the 'works of art' table and the values of all attributes for the artist that is the work's creator are extracted and used to create OWL instance information in the following way.

First, *Ontology Populator* creates an OWL individual representing the artist, replacing accented letters in the artist name and substituting underscores for whitespaces in order not to violate OWL naming constraints. The unmodified artist names, possibly containing diacritics, are preserved as RDFS labels. The artist's professions or professional specialisations (i.e. painter, printmaker, etc.) are then parsed (i.e. broken down into individual bits, in case the string contains more than one term) and checked against a mapping of 18 common entries for professions to classes in our ontology. The mapping is based on frequency considerations. In case of a match the artist individual is assigned to the corresponding OWL class as an instance. If no professions are given or if the given professions are not covered by the mapping, the artist individual is assigned to the more general class 'Artist'.

The field on an artist's profession is also used to determine gender. In German, the language of our database, expressions for professions have gender specific endings. If an entry for an artist's profession ends in "-in" (the equivalent of English "-ess" in certain contexts), then the artist individual is assigned to the class "Female_Person". If a profession is given, that does not end in "-in", the artist individual is assigned to "Male_Person" and if profession information is missing, the gender of the artist remains undetermined.

Ontology Populator creates ontology individuals representing an artist's birth event and an artist's birth time span and introduces the relationship instances between the artist, the birth event and the time span. This happens in accordance with the CRM time model (see p. 104, Figure 5.2). Artist birth dates are given as strings in the database. To convert these human-readable strings into a machine-processable date representation we reused the mapping of these strings to ISO 8601 dates (or date intervals) we had created at an earlier stage of our project (compare the remarks on "Data Transformation" above). *Ontology Populator* makes use of the mapping by way of a lookup table. The integer endpoints of the ISO interval were linked to the time span using the 'beganAfter' and 'endedBefore' data type properties (see p. 106). In the same way date related ontology assertions including event, time span and integer values, were generated for the death of the artist and the production (event) of the work of art.

The birth and death places (strings) recorded in the database are identified, if possible, with OWL individuals representing geographical locations. This is done in a lookup analogous to the one for professions. Birth and death events of an artist are linked to their respective locations through the CRM property 'P7_took_place_at'. The artist individual is linked through an object property 'was_Active_In', which is part of our ontology, to places of activity as recorded in the corresponding database attribute.

Finally, after the information on the artist has been created, an OWL individual representing the work of art is created and linked to the artist via the production event and the appropriate relations ('P108I was produced by' and 'was_Carried_Out_In_Leading_Role_By'). In order to avoid name clashes and make the ontology easier to inspect, both the title of the artwork and the name of the artist are used to generate the individual name for the work. As in the case of the artist name, accented letters and whitespaces are substituted and the unmodified title of the work is recorded as an RDFS label. If the title string of the work of art contains a German equivalent for the term "portrait" or one of the three German equivalents for the term "self-portrait" ("Selbstporträt", "Selbstportrait" or "Selbstbildnis"), then the OWL individual for that work is assigned to the corresponding class 'Portrait' or 'Self_Portrait' in our ontology. In the case of self-portraits, it is additionally stated that the work 'P62_depicts' the artist individual.

After these steps have been carried out, the next work (of the works that were initially selected from rows in the database) is processed in the same way. Note, that if two works by the same artist have been selected, there is no need to modify the procedure, because the artist individuals that are introduced will have identical names and thus be automatically identified with each other without collisions or violation of constraints. When this process has been completed, typically for a few hundred works of art, the 'ontology populator' saves the resulting ontology to a file, which then becomes the knowledge base in the narrower sense that we have described above (compare Section 5.3.1 and Figure 5.4). In our evaluation we mainly used a knowledge base of self-portraits as this offered a straightforward way to reuse the details given for the artist as biographical details for the sitter.

5.5 System Evaluation

An evaluation of the graphical user interface was carried out that involved art experts as well as laypersons. As we have seen, description logic can be used to model and express facts of considerable complexity about heritage artefacts. The *Artfinder* prototype was implemented in order to visually represent some of the relations and hierarchies in *ART Ontology* and to provide a simple interface to construct such statements as a query to a knowledge base. The objectives of our evaluation of the *Artfinder* prototype were to determine:

(a) How intuitive the generation of logical queries with Artfinder is.

(b) How accurate user generated expressions are.

(c) How queries generated for the same task compare to each other in terms of their logical properties.

We have attempted to address these objectives by an experimental setup in which participants were asked to translate a number of natural language queries into description logic expressions using *Artfinder*. A report on the evaluation experiment and its results has been published [86].

5.5.1 Evaluation Design

The improved graphical user interface was subjected to an evaluation by a mixed group of experts and non-experts. The participants individually completed a standardised evaluation session consisting of three parts. The first part consisted of an introduction to the project and some of the modelling aspects in the ontology in the form of a short presentation. The second part was a user experiment which involved the creation of 14 queries using the graphical user interface. Lastly the third part, required participants to complete a web-based questionnaire to collect feedback on their experience.

Experiment Sessions

All of the 20 experimental sessions were carried out on the premises of Trinity College Dublin, one participant at a time. Apart from the participant only the experimenter was present in the room for

1	A Van Gogh self-portrait.
2	A painting done by someone who also worked as a printmaker.
3	Drawings by people born before World War II.
4	A print from after 1815 by an artist that lived in France at one point.
5	A Baroque painting by an artist that worked in Germany.
6	Art that shows French 19th century painters.
7	Italian drawings showing people who died during World War I.
8	A post-1920 painting depicting the artist.
9	A portrait by a Dublin-born painter, showing a female sitter.
10	A portrait from between 1648 and 1850.
11	Prints by Dutch printmakers.
12	Group-portraits by Spanish painters.
13	Renaissance paintings that contain stylistic elements of Late Gothic.
14	A painting showing either a female painter or a male sculptor.

Table 5.3 – The 14 query tasks presented to the participants of our system evaluation.

most of the duration of each experiment. At particular points, however, the experimenter usually left the room for a short while (see below).

The participant was accommodated next to the experimenter at a workstation. Some of the sessions were run on a desktop computer and some on a laptop computer, but the setup for the participant was always "desktop-like", i.e. involving a 17" flat screen, a wired mouse and a keyboard.

The first part of each session consisted of an introductory presentation given by the experimenter to the participant. The presentation contained 13 slides, which can be found in Appendix E. Seven slides were holding content and six had navigational function for the presentation and the evaluation session as a whole. The introductory presentation served two purposes: to introduce the participant to the larger context of the experiment and to familiarise her/him with some of the modelling choices in the ontology so that elements in the user interface representing classes or relationships in ART ontology and their mutual interrelationships could be more easily recognised by the participant (see Figure 5.7).

Post-experiment Questionnaire

Each evaluation session was concluded by an electronic questionnaire, which is reproduced as Appendix F. The questionnaire started with questions to evaluate how much the query tasks corresponded to participants' information needs and wishes. In the next part we asked participants to comment on how intuitive the system was and suggest improvements. After that, participants were shown screenshots of sections of the graphical user interface they had just used, with certain selections made and highlighted. We asked them to choose which of a number of verbalizations (ranging from two to four choices) best described the intended meaning of the selection they were looking at. Alternatively, participants could formulate a free text description if they felt that none of the options matched their interpretation. The goal of these questions was to check whether certain logical interpretations which our system made implicitly were intuitive to users or not. This referred in particular to logical "and/or" connectives and to universal versus existential quantification (see Figure 5.8). Agreement with certain subsumptions (Are all self-portraits portraits?) and inferences (Does every animal scene show at least one animal?), some of which are incorporated into the knowledge

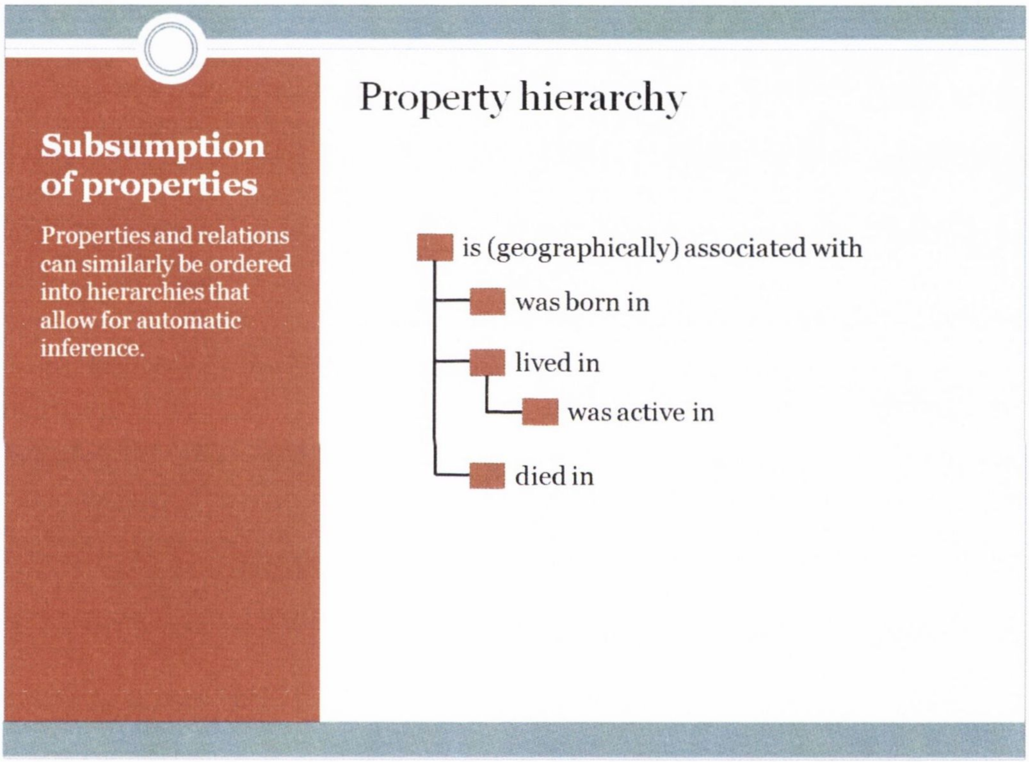

Figure 5.7 – A slide familiarising participants of our evaluation session with the idea of a property hierarchy.

model of our system was also checked.

5.5.2 Demographics of the Participants

A total of 20 people took part in our evaluation. All participants completed the same evaluation procedure. In particular the expert or non-expert status was not determined in advance, but based on participants' self-reported level of expertise in art history or similar subjects. However, people from different institutions were approached according to their expected expertise in this field. The 'expected experts' were taken from National Gallery of Ireland staff, staff of contemporary art galleries in Dublin and postgraduate students of art history in Trinity College. The 'expected non-experts' were mostly recruited among postgraduate students and staff in the School of Computer Science in Trinity College.

The questions relating to participants' demographic information — gender, nationality, age and educational background — were worded exactly like the corresponding questions in the image retrieval survey of Section 4.2. Eleven of the participants were female (Figure 5.9a). The majority of the participants (13) were Irish nationals. Other nationalities represented were Chinese (2), Dutch, French, German, Indian and Italian (one each) (Figure 5.9b). Participants' age ranged from 22 to 40 years, with four participants choosing not to report age related information. Among those participants who did reveal their age the average was 28.25 years with 26 years being both the most frequent answer (mode) and the statistical median (Figure 5.10).

In terms of their educational background 9 participants had completed an undergraduate degree, 10 were holding a post-graduate degree and one a Ph.D. (Figure 5.11a). Half of the participants have not received any university level education in art history or a related field. Among the other half, one person had read art history at an undergraduate level but had not completed the degree, three were holding a bachelors degree, five a masters degree and one a Ph.D. in the field of art history (Figure 5.11b).

5.5.3 Evaluation Results

The queries created by each participant were analysed in terms of their syntactic and logical structure. First of all the queries were assigned to the corresponding English language sentence (1 to 14 in Table 5.3) set as a task to the participant. In this process it became apparent that several participants had submitted several queries corresponding to the same natural language sentence and that on one occasion a participant had accidentally omitted sentence number 6 altogether. In the following we will use "query" to refer to the logical formula created by *Artfinder* and "sentence" when we speak of the English language sentences that were given to participants to prompt them for query generation. In cases in which several queries were submitted by the same user in response to the same sentence the queries were compared to each other. This was usually caused by participants hitting the "Search" button in the user interface several times instead of only once. In total there were 13 such cases. In all but one case however the submitted queries were identical to each other. In one exceptional case the participant had accidentally submitted his query to soon and re-submitted it after entering some more details. This was documented in the observation notes taken during the session and the incomplete query was then dropped from the analysis.

As no art experts were involved in phrasing the natural language sentences that were used to trigger the query generation, we included a question to assess how realistic our example sentences are.

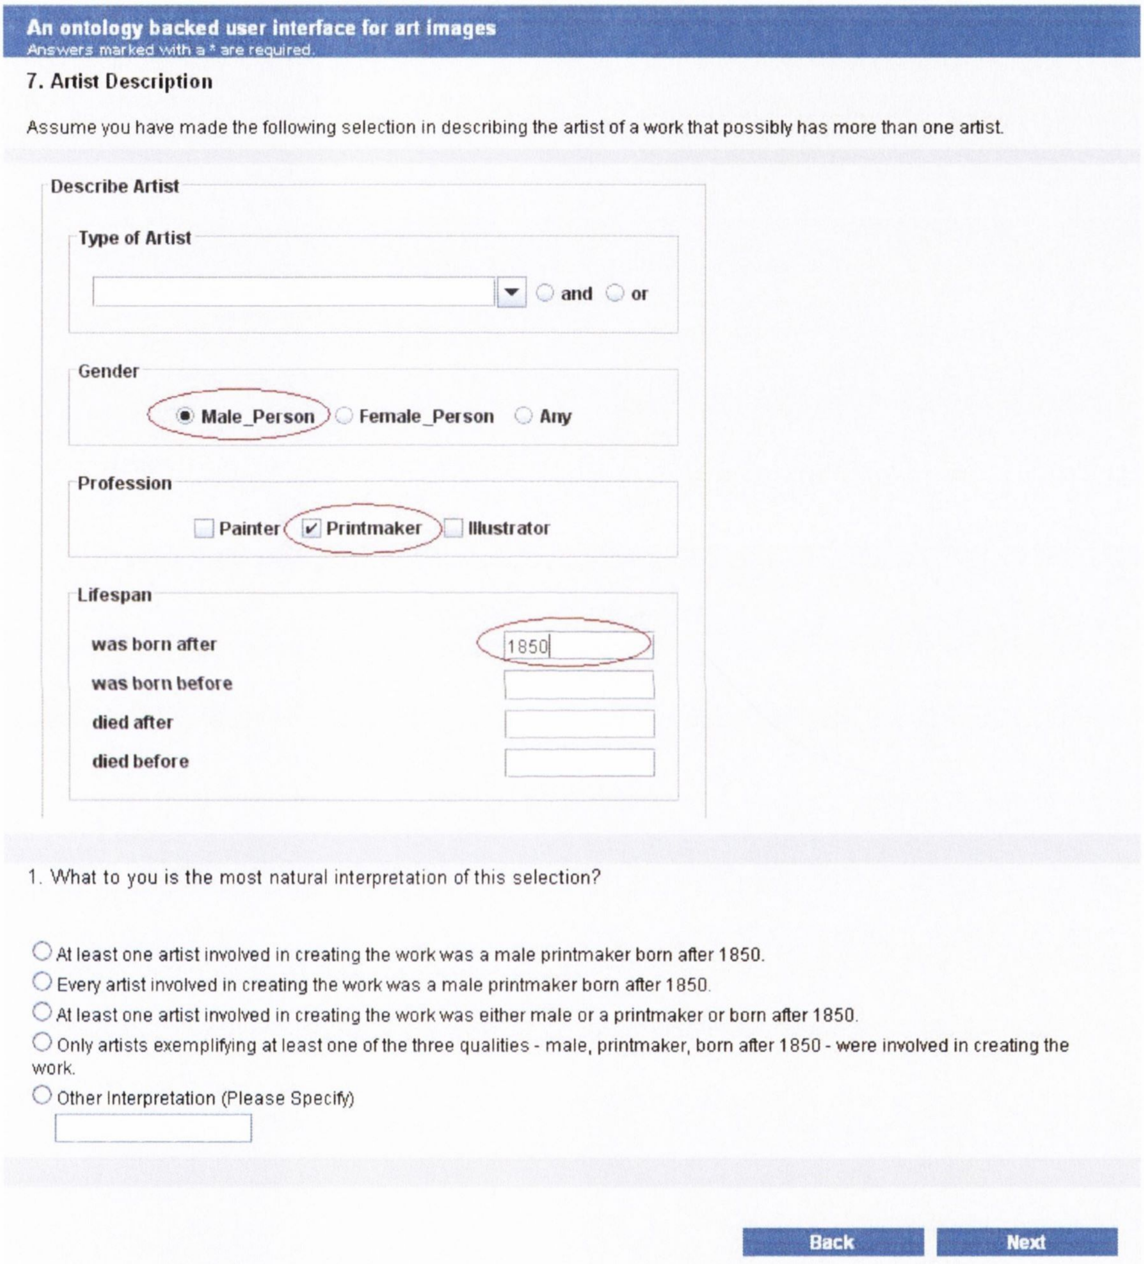

Figure 5.8 – A question from the post-evaluation questionnaire.

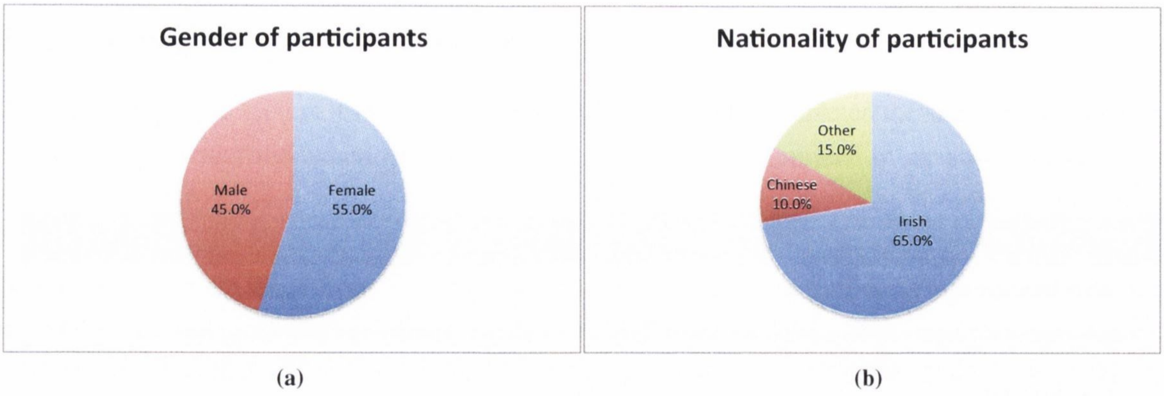

Figure 5.9 – Participants' gender and nationality.

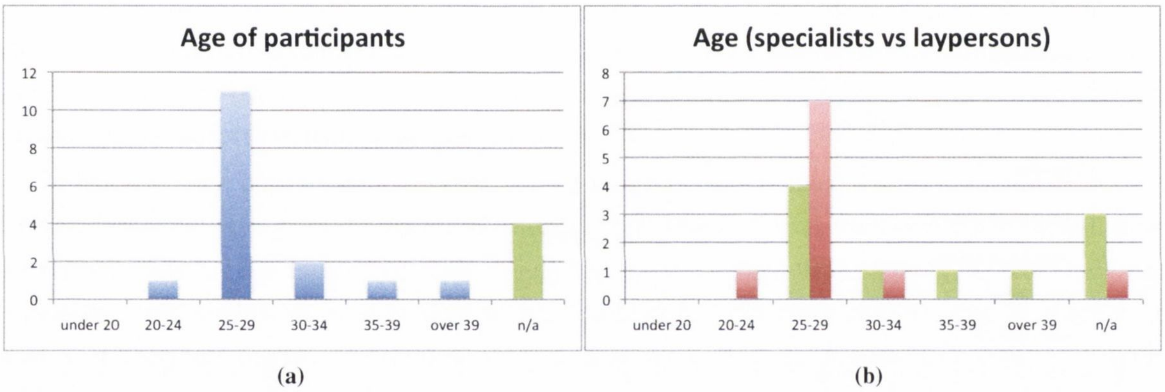

Figure 5.10 – Age distribution of participants in general (a) and age distribution of specialists (green/left) compared to age distribution of laypersons (red/right) (b).

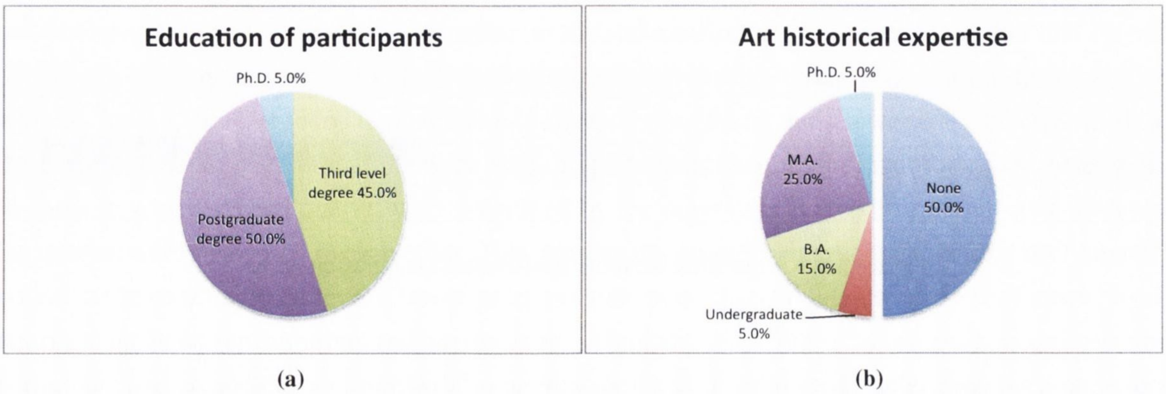

Figure 5.11 – Level of education among participants in general (a) and college level education in art history or related fields (b).

When asked, whether the sentences resembled what they would like to ask when looking for works of art, 25% of participants said "Absolutely", 40% opted for "Mostly" and the remainder selected "Somewhat". None of the participants chose the options "Little" or "Not really". We therefore assume that the sentences given to participants are not too dissimilar from genuine information needs.

We will discuss our evaluation results in response to the three objectives (a), (b) and (c) (cf. p. 125). First, we discuss immediate feedback on the usability of *Artfinder*, then the 'correctness' of queries by two different standards and finally the logical relations between queries generated for the same task.

Feedback on Query Generation through Artfinder

In reply to how intuitive they found the interface to use on a scale from 1 (not intuitive at all) to 7 (very easy to use, very intuitive), 90% of our participants chose a mark of 5 or higher. No participant selected 3 or less (M = 5.60, SD = 0.88). Each participant was encouraged to give an honest opinion and openly criticise or suggest improvements to the system. Among the suggestions for improvement were requests for better navigation between query tabs and more feedback and information for the user. The relative complexity of the user interface triggered different responses ranging from 'still too little' to 'too much':

- *The 'stylistic period' choices were somewhat limited – more detail in this area would be useful. It also does not provide query options for non-human elements in the paintings, for example an architectural element*
- *there were enough options to allow for a comprehensive search, its just more labour intensive than a keyword search*
- *I wonder if you want to get the system up and running for [...] a broader scope of people you might somehow make it easier – and easier means less options*
- *there are too many options, unless you have quite an advanced user this could be confusing – it currently requires quite a high level of competancy (sic) to use and the researchers we often deal with do not always have this*

All of the above commentaries were made by participants in our expert category.

Accuracy of Queries

Apart from the subjective opinion of our participants on the usability of the system, we also probed whether their intended meaning of their query expressions objectively agreed with the interpretation in our system. This was done by asking questions such as the one shown in Figure 5.8. Some of the ways in which *Artfinder* interprets user input are not made explicit and hence they might not be understood by users in the way they are intended. Multiple selections on the same tab, for example, are always interpreted conjunctively (logical 'and'), whereas restrictions along a relation are existentially and not universally quantified ('this painting depicts a person that...' as opposed to 'this painting only depicts persons that... if any'). In three questions about the meaning of certain selections a majority of individuals always selected the interpretation that was built into the system. Agreement levels ranged from 65% to 85% with an average of 73.3%. Across all three questions however only 45% of participants agreed with the system's interpretation in every single case. This leaves a majority of participants who expressed a different intention than the one encoded in the system in at least one of the three cases.

A different way of looking at the accuracy of the submitted description logic queries is to compare them against a 'correct' reference query. Of course, the term 'correct' has to be used with care here, because every rendering of a natural language expression in a formal language involves an act of interpretation. However, many of the terms occurring in our natural language sentences (Table 5.3) have homonymous equivalents in OWL classes or properties of ART ontology. In fact, each of the 14 query tasks set to our participants was inspired by an OWL class expressible in the vocabulary of ART ontology through *Artfinder* as it was configured for our experiment. Table 5.4 contains the 14 OWL class expressions (in Manchester OWL Syntax) corresponding to the query tasks.

1	Self_Portrait AND (was_Created_By VALUE Gogh_Vincent_Willem_van)
2	Painting AND (was_Created_By SOME Printmaker)
3	Drawing AND (was_Created_By SOME (was_Born_During_TS SOME (endedBefore SOME int [<= 1939])))
4	Print AND (was_Created_By SOME (lived_In VALUE France)) AND (was_Produced_During_TS SOME (beganAfter SOME int [>= 1815]))
5	Painting AND (was_Created_By SOME (was_Active_In VALUE Germany)) AND (was_Produced_During VALUE Baroque)
6	P62_depicts SOME ((died_During_TS SOME (beganAfter SOME int[>= 1801])) AND (was_Born_During_TS SOME (endedBefore SOME int[<= 1900])) AND (was_Born_In VALUE France))
7	Drawing AND (P62_depicts SOME ((died_During_TS SOME (beganAfter SOME int[>= 1914])) AND (died_During_TS SOME (endedBefore SOME int[<= 1918])))) AND (was_Created_By SOME (was_Born_In value Italy))
8	Painting AND Self_Portrait AND (was_Produced_During_TS SOME (beganAfter SOME int [>= 1920]))
9	Portrait AND (P62_depicts SOME Female_Person) AND (was_Created_By SOME (Painter AND (was_Born_In VALUE Dublin)))
10	Portrait AND ((was_Produced_During_TS SOME (beganAfter SOME int [>= 1648])) AND (was_Produced_During_TS SOME (endedBefore SOME int [<= 1850])))
11	Print AND (was_Created_By SOME (Printmaker AND (was_Born_In VALUE Netherlands)))
12	Group_Portrait AND (was_Created_By SOME (Painter AND (was_Born_In VALUE Spain)))
13	Painting AND (has_Stylistic_Elements_Of VALUE Late_Gothic) AND (was_Produced_During VALUE Renaissance)
14	Painting AND (P62_depicts SOME ((Female_Person AND Painter) OR (Male_Person AND Sculptor)))

Table 5.4 – The 'intended' description logic formula corresponding to each of the English language query tasks of Table 5.3 above.

Task	1	2	3	4	5	6	7	8	9	10	11	12	13	14	All tasks
Experts	40	50	10	50	40	0	0	30	20	80	30	20	70	0	31.4
Laypersons	60	40	40	30	30	0	0	40	10	100	60	30	70	20	38.1
All	50	45	25	40	35	0	0	35	15	90	45	25	70	10	34.8

Table 5.5 – Level of agreement between the submitted queries and the reference expressions (Table 5.4) for all 14 query tasks in percent.

Task	1	2	3	4	5	6	7	8	9	10	11	12	13	14
No. of equivalence classes of queries	7	6	14	8	9	20	18	11	14	3	6	11	7	12

Table 5.6 – Number of equivalence classes among the submitted queries with respect to equivalence in ART ontology.

In all but four cases the query submitted most frequently by our participants, modulo logical equivalence, was identical to this reference expression. In other words, in most cases a plurality of participants produced a query logically equivalent to the reference expression. However, the average rate of agreement with the reference query was only 34.8% (Table 5.5). Over all 14 tasks there was no significant difference between experts and laypersons in terms of producing the reference query (paired t-test: $t = 1.50$, $df = 13$, $p = 0.156$). A possible explanation why laypersons showed a tendency to do at least as well as experts in producing the target query may be that most of the participants in our lay cohort were computer scientists and thus very familiar with user interfaces and query specification. However, in some cases, domain expertise did make a difference. One of our lay participants, for instance, did not know what 'Baroque' was and therefore could not locate the term in our user interface.

Logical Relations between Queries

Instead of comparing the queries produced by our participants to a reference expression with respect to logical equivalence, we may also compare them to each other. One aspect we looked at in this context was how many equivalence classes of logically equivalent queries do exists for each of the 14 tasks. The dispersion of logically different answers given in response to a natural language sentence may perhaps be seen as measure of the ambiguity or complexity of the sentence. Since the axioms of ART ontology are normally available to the reasoner used by our system for retrieval, we formed the equivalence classes with respect to logical equivalence in ART ontology. This can in principle result in fewer equivalence classes than would be obtained with respect to description logic (OWL 2) equivalence without additional axioms.[32] The higher the number of equivalence classes, the more logically different queries were submitted for a task. In task 6, for instance, no two participants generated the same query (Table 5.6).

Comparing the dispersion of the submitted queries (Table 5.6) to the agreement levels with our

[32] In our experiment the two notions of equivalence only made a difference in one case: task number 12. In this case the number of equivalence classes was reduced from 12 to 11 by the ontology.

	weaker	equivalent	stronger	incomparable
Experts	10.7	31.4	17.9	40.0
Laypersons	26.6	38.1	5.0	30.2
All	18.6	34.8	11.5	35.1

Table 5.7 – Percentage of queries submitted by experts and laypersons, that were weaker, equivalent, stronger or incomparable with respect to the reference expressions (aggregated over all 14 tasks).

	weaker	equivalent	stronger	incomparable
with ART Ontology	18.6	34.8	11.5	35.1
without ART Ontology	11.8	34.8	10.4	43.0

Table 5.8 – Percentage of submitted queries that were weaker, equivalent, stronger or incomparable with respect to the reference expressions with and without taking ART Ontology axioms into account (aggregated over all 14 tasks).

reference expression (Table 5.5) we find a very similar pattern if we order the tasks by increasing dispersion and decreasing agreement. In both cases tasks 1, 10 and 13 are in the top bracket (high agreement, little dispersion) and tasks 6, 7 and 9 are in the bottom bracket (low agreement, high dispersion). In particular, task 10, a rather straightforward task, which occurred when participants had already familiarised themselves with the interface, was perceived as easy to execute according to feedback received during the evaluation sessions.

So far we have only considered equivalence between query expressions, i.e. mutual entailment. In formal logical systems, however, expressions may also be compared with respect to one-way entailment. If an expression or formula entails another one, it is said to be stronger than that other expression. If it is entailed by another expression, it is said to be weaker. If no entailment relation persists between two expressions, these may be said to be incomparable with respect to entailment.

We used our reference expressions again (Table 5.4) and grouped the queries submitted by our participants in comparison to these expressions into four categories: logically equivalent, weaker, stronger and incomparable. The notion of entailment in this context is understood to include the axioms available in ART ontology (as was the case with the equivalence classes of Table 5.6). In comparison to the reference expression, logically weaker queries will retrieve a superset of the results and logically stronger queries a subset. We found that overall 18.6% of participant queries were logically weaker and 11.5% logically stronger than the reference query (Table 5.7).

Categorising weaker, stronger, equivalent and incomparable queries separately for specialist and layperson participants we found statistically significant differences between the two groups ($\chi^2 = 22.26$, df $= 3$, $N = 279$, $p < 0.001$). Additional chi-squared tests that were carried out on the same data, suggested a tendency by experts to submit logically stronger and a tendency by laypersons to submit logically weaker queries in relation to each other. This confirms our impression during the evaluation sessions, that experts sometimes specified more details than appeared to be necessary and laypersons tended to leave out detail when they were in doubt.

Entailment is influenced by the axioms that are available as premisses for reasoning. If no axioms are available the entailment is purely logical, i.e. in the case of our system, entailment in OWL 2. If additional axioms are available these may result in more entailments, which are then said to follow from these axioms. One way of looking at the influence that ART Ontology has on the interpretation of queries, is to subtract its axioms from the reasoning process (i.e. execute the queries against the knowledge base without the class and property subsumptions defined in ART Ontology). Table 5.8 shows the changes in entailments, aggregated over all 14 tasks of our evaluation, that would occur if ART Ontology were deactivated: The result would be fewer queries that are stronger (10.4% vs. 11.5%) and, in particular, fewer queries that are weaker (11.8% vs. 18.6%) than the reference query.

5.6 Summary

In our implementation we have addressed a number of aspects along a processing chain from a tabular legacy representation of information on visual art to a graphical user interface through which description logic queries may be formulated to access and retrieve such information. We have focussed on a number of points along this chain which appeared to pose interesting challenges.

The conversion of heritage data in relational tables into the graph-like format of a Semantic Web language is a current topic in academic research and in the practice of heritage institutions (cf. [82]). With *Ontology Populator* and its pre-processing algorithms, we have developed a set of heuristics which is aimed at enriching the initial data set during its transformation. This covers, in particular, a range of free text descriptions of dates which are systematically converted into ISO 8601 compatible date expressions. Some of the aspects of these heuristics are specific to German (the language of our test data), but the principles will translate to other languages as well.

With ART ontology we have developed a prototype application ontology which captures some of the distinctions made by practitioners in the domain of art and heritage curation. ART ontology forms a layer between the CIDOC CRM top-level model and an arbitrary set of data on visual art. In particular, we have utilised the feature of subproperty chains in the Web Ontology Language (OWL 2) to emulate temporal inferences for the purposes of querying the ART ontology knowledge base.

Building on these components, we have designed and implemented *Artfinder*, a visual query builder the structure of which is determined to some extent by the representational details in ART ontology and the associated knowledge base created with *Ontology Populator*. Changes in the underlying knowledge base will be reflected in *Artfinder* and its graphical interface elements are arranged in a way that attempts to make query generation intuitive.

Finally, we have conducted an evaluation of *Artfinder* with 10 art and heritage experts and 10 laypersons. We found that the use of *Artfinder* was described as intuitive and many of its interface options were welcomed by participants. We also found that perhaps only a third of the queries generated with *Artfinder* by users unfamiliar with the system, may accurately reflect the user's information need. There were indications, however, that simple queries at least could be accurately and consistently produced after only a few minutes of exposure to *Artfinder*. We found statistically significant differences in expert and layperson usage of the system, with experts showing a tendency to overspecify and laypersons a tendency to underspecify with respect to our predictions.

Chapter 6

Conclusions

We began this thesis by asking the question whether or not well structured background knowledge could facilitate information retrieval in a knowledge rich domain. We chose the multifaceted domain of cultural heritage as an exemplar case to investigate this question. One of the facets of this domain is purely physical, the material and visual properties of works of art. The second facet is about meta-level information about the artist, sitters and the provenance of a work. A third facet is what the artist intended to signify in his work and a forth facet what the viewer thinks a work signifies. We would like to add in parentheses that the artist facet cannot be explored generally in a heritage sense as most of the artists are deceased or not physically present and therefore one has to rely on experts and other collateral information on a work, such as preliminary sketches, notes and letters. This multiplicity of facets gives us an index of the richness of the given domain of knowledge. Each facet relates to the philosophical question of what there is. We have addressed this question with the methods of computational ontology.

Researchers in computational ontology have focussed on questions of the syntax and semantics of how to represent an object and its attributes onto a computational system through psychologically motivated schemas like frame-based systems and logically motivated schemas like description logic. Here we use the logically motivated representation of objects and their attributes through assertional and terminological knowledge.

Up until very recently researchers in ontology tried to handcraft domain knowledge either by talking to experts or by using common sense both at the level of concepts and at the level of instances. However, one can find a large collection of domain specific thesauri, glossaries and dictionaries which are organised at varying levels of conceptual descriptions but may not be easily mapped onto a representation schema like OWL.

The conceptual description notwithstanding these pre-ontological collections of knowledge are systematically organised, e.g. as slot and filler templates with names for slots and constraints on fillers. The template is then mounted on relational tables and the tables onto a database. It is possible, as we showed to export the contents of a domain specific knowledge collection onto an ontology system.

We show in this thesis that one can build an information system based on a rich domain knowledge which can deal with a range of questions posed by experts or by informed laypersons. Throughout our project we kept in close contact with our domain experts at a major national gallery. Our cohort of evaluators included ten experts and ten laypersons. A software program was built which represents art related knowledge and maps a graphically generated user query onto a description logic expression.

6.1 Research Questions Revisited

Our principal research question was this:

> *What can well-structured background knowledge contribute to information retrieval in a knowledge rich domain?*

We have chosen fine art images as an example domain and specifically addressed this research question via three corollary questions which we will now revisit in turn.

1. *What categories are used to describe and search for information on works of art?*

Conventional wisdom has it, that an image says more than a thousand words. By this we mean, as a general rule, that what an image shows says more than a thousand words. In the case of ephemeral everyday photographs one may even find it hard to think about what else could be of interest about an image apart from what it depicts, for instance, where it was taken or who the photographer was. Given our focus on visual messages, it is perhaps not surprising that descriptions of the image subject feature prominently in searches for images of fine art as well. It may seem surprising, however, that as many as a third of the terms employed for searching visual art may not refer to visual aspects after all.

In order to illuminate the first corollary question, over and beyond what we found in scholarly publications on this subject, we conducted interviews with two members of the curatorial staff of the National Gallery of Ireland and we surveyed a sample comprising 21 art specialists and 27 laypersons on their image retrieval behaviour. We found significant differences between the two groups suggesting that, as the level of domain expertise rises, the importance of the image subject recedes and the significance of the provenance and other meta information about a work gains weight.

2. *Which ontology design principles are relevant for an ontology of fine art?*

We implemented *ART Ontology* as a layer between the CIDOC CRM, a top-level model for structuring and interchanging cultural heritage information, and a set of data arbitrarily selected from an existing digital collection. As the basis for our retrieval application, *ART Ontology* attempts to capture some of the distinctions and categorisations that are made by heritage practitioners and that are possibly of interest to the wider public for retrieving information on fine art.

In modelling the data of a digital collection of over 40,000 works of fine art, of which a subset, ranging from several dozen to several hundred, were connected to *ART Ontology* at any given time, we attempted to represent the information not only with a view to our particular application, but true to the intended semantics of the concepts and relationships in our model. A possible exception of this is the representation of numerical date expressions in which theoretical limitations of the Web Ontology Language forced us to adopt a model which emulates temporal reasoning for the purposes of querying our knowledge base rather than executing the reasoning true to the intended semantics of the relationships involved.

Based on this ontological principle, which is also a foundation of the CIDOC CRM, we examined the CRM and found a number of misleading constraints in the model. We adjusted these aspects in our implementation and we hope that our findings may influence future versions of the top-level model.

3. *Can end-users successfully formulate complex ontology queries?*

ART Ontology and a knowledge base automatically created from relational data by an ontology population program, were connected to a novel graphical user interface, *Artfinder*, the layout of which was influenced by the structure and content of the underlying represented knowledge. The interface can be used to generate description logic queries, which can then be executed by an off-the-shelf reasoner against the ontology and knowledge base. The functionality of *Artfinder* was incrementally expanded and improved in three individual sessions with domain experts and finally subjected to an evaluation by ten different experts and ten lay users.

Based on the results of our evaluation we may not be able to give a conclusive answer to the last corollary question. While our evaluators described *Artfinder* as intuitive to use and even compared the logic of its interface favourably to existing collection management software, only one third of the query expressions generated through *Artfinder* agreed with what we would have regarded as the accurate query in the language of our ontology. The dispersion ranged from 18 queries in agreement with our prediction, to 20 mutually different queries of which none was equivalent to the query we would have expected for the given task. To some extent, however, the knowledge encoded in the ontology may help to reduce this dispersion as it increased the number of queries that were, if not logically equivalent, then at least logically stronger or logically weaker than our target query. It also has to be emphasised that participants had no prior exposure to our system and were only given a brief introduction to the interface and the structure of the underlying knowledge model. We found indications in the course of the evaluation sessions, that even a brief familiarity with the system may improve the understanding of the interface and the ability to use it correctly.

While we found no difference between experts and laypersons with respect to the frequency with which they generated the target query, we discovered, as a parenthetical result, that experts tended to create queries which were logically stronger than the target and laypersons, in comparison, queries which were logically weaker.

In summary, an ontology-informed approach can be recommended for a domain with complex relationships like heritage, if a number of criteria are met: (a) ready access to domain experts, (b) an investment into a knowledge modelling phase which reflects domain conceptualisations correctly, and (c) a mechanism for querying a complex knowledge representation system that is both accurate and intuitive to use. A key feature of our system, and something that any ontology-based system has to achieve, is that queries that are off target can be put in a broader or narrower relation to a target query.

6.2 Future Work

Different lines of future work arise from the research presented in this thesis. Despite the large number of studies that have investigated user image retrieval behaviour and in particular the retrieval of aesthetic images, many questions in this area still remain open and could be investigated by follow-up studies to our survey.

For our ontological modelling, we examined the CIDOC CRM in detail and found, if not inconsistencies then at least incoherences and implausibilities in the model. A general report of these findings was given to the developer communities of both the abstract CRM model and its implementation. We have made arrangements to give a more detailed account of individual disputable aspects of the model in the near future and discuss what changes could, in our opinion, be made. We hope that this will have an impact on the CRM community and encourage them to explore and address the points under discussion.

Our evaluation of the *Artfinder* interface and the feedback received from users of the system has revealed potential for improvement. A different layout and additional functions may make the system both easier to understand and more expressive in the queries it can create. More importantly, the architecture of *Artfinder* is general enough to map a query interface not just to *ART Ontology*, but to other application ontologies and knowledge bases as well, as long as these are encoded in the Web Ontology Language. Exploring the application of our system to independently developed ontologies would be a useful step towards demonstrating the broader applicability of the system presented in this thesis.

The one key implementation we would wish to undertake, however, is a detailed modelling of the historical context of a set of strongly interrelated works of art that elaborates on our use of events in *ART Ontology*. Events have been a key theme in our project in several different ways: About ten percent of the image descriptors collected in our survey referred directly to events. Together with time and place descriptors, often indirectly associated with events, these descriptors accounted for over 40 percent of the total of the submitted terms. The event-centred approach of the CRM model provided the substrate for the art historical periods, the birth, death and creation events represented in our knowledge base. In individual feedback several participants of the *Artfinder* evaluation, including domain experts, expressly recommended the modelling of further events such as historical battles, revolutions, the reign of certain dynasties and so forth. In comparison to laypersons, we have seen that for art specialists the description of visual elements in an image becomes less and the historical context of a work more prominent, due to the different requirements of their practice and their different conceptualisations. In this matter event modelling is likely to make an impact. A knowledge base containing systematic records of everyday events, such as meetings between artists, travels, the signing of legal documents and the sale of artworks together with a robust documentation of the time and place of these incidents, as far as they are known, may help to establish connections which are otherwise much harder to discover. In such an undertaking, it is critical — and that is one of the principal lessons drawn from our research project — that systems developers have significant and ready access to art experts.

Expert practitioners usually act on the premise that scholarly disputes in their particular domain can, in principle, be resolved by employing the methods of their field, and that their findings can be integrated into a body of knowledge. If this idea is translated to a larger perspective which reaches across disciplines it may be identified with what Vanevar Bush, an early visionary of the World Wide Web, has called "the great record". In his vision, presented in an influential article "As We May Think" in 1945, he speculates that analogue processes such as "dry photography" and the "thermionic tube" may be the enabling technologies of the knowledge-based systems of the future.

Perhaps the current suite of Semantic Web technologies is the present day embodiment of Bush's "dry photography" and "thermionic tube" in the pursuit of the "great record". We believe that the paradigm of consistent knowledge integration within and across disciplines will continue to be tested until it is approximated more and more completely or its principled limitations emerge replacing it with a new one.

It is the goal of our study to be a modest contribution towards this end.

Appendix A

Interview Guide

This appendix contains a scaled down reproduction of the interview guide used for the expert interviews covered in Section 4.1.

Questions for National Gallery of Ireland curators

About you

What is your position in the NGI and how long have you worked in this position?

Have you worked in a similar position/function at other institutions?

You studied Art/ History of Art/ History/.... ?

Curatorial work

Could you describe the most important aspects of the job of a curator?

What information about an image is gathered and recorded as part of the curatorial process?

What information about *the content* of an image is gathered and recorded?

How do you use computer and web based resources in your *curatorial* work?

Do you use art databases or knowledge systems? If yes, which ones?

Exploring art on the web

What do you think about online image galleries? What is the biggest difference to conventional ones?

How often do you search for information about artworks on the web (privately or as part of the job)?

Do you typically look for textual information or for images? What is the ratio between these two?

If you look for images which of the three following statements best describes your goal:

- I browse the web in search for interesting art works without a particular target in mind
- I look for artworks of a certain category that share similar characteristics
- I look for one particular work of art

What resources do you use (Google, Yahoo, special art portals, special software)?

How do you go about keyword-based search? Do you have a method?

Have you ever heard of or used query by sketch or query by image example techniques?

What in your opinion are the shortcomings of current image retrieval systems?

How would exploring art on the web ideally look like in your opinion (be imaginative)?

How would you envision an art expert system (be imaginative/demanding)?

Appendix B

Interview Transcripts

This appendix contains the transcripts of the two expert interviews that we conducted in the course of our study. Both interviews are discussed in Section 4.1.

B.1 Expert A

Participant:	Niamh MacNally
Participant ID:	Expert A
Participant Position:	Assistant Prints and Drawings at the National Gallery of Ireland
Location:	National Gallery of Ireland
Date of Interview:	7/05/2008
Interviewer:	Daniel Isemann
Interviewer ID:	Interviewer
Transcriber:	Daniel Isemann

Expert A: ...quite informal really.

Interviewer: Oh, absolutely. It's nothing but an informal interview. Bla bla bla bla. Ok, it does something there, so we should be fine.

Expert A: Grand.

Interviewer: So, there are three blocks of questions. First very briefly about you, then about your curatorial work or the work you do in the National Gallery here and then, like, exploring art online and like your personal opinions or experience.

Expert A: Ok.

Interviewer: What is your position in the National Gallery and for how long have you been here?

Expert A: Ehm... this is my sixth year now and my title would be assistant prints and drawings. But basically I'm part of the curatorial department so it would be like acting as an assistent curator.

Interviewer: Curator.

Expert A: Yeah to specifically the prints and drawings collection.

Interviewer: Yes.

Expert A: Even though we would be very aware of the other schools within the collection our focus is on works on paper.

Interviewer: Ya, like Adriaan Northern Art or...

Expert A: Exactly.

Interviewer: Ah ya, ok. Have you worked in a similar position or function at other institutions before that?

Expert A: Similar, yes, at a museum in the States and in Italy.

Interviewer: All right and that would have been prints and drawings as well or just...

Expert A: No it wouldn't, no more general.

Interviewer: Ok, and you studied like sort of your background or your...

Expert A: I would have studied Art History and English in UCD and continued on a Masters in Art History and then followed on with... its called an Arts Administration course... also in UCD.

Interviewer: Ah well, the university doesn't really matter. And Arts Administration, ok. Just a rough idea so. Now concerning the curatorial work, I'm obviously a layperson and I have maybe vague ideas what it looks like, but if you just outline the most important aspects of... what is involved in the job?

Expert A: Ehm... Well the job that we do here would be organising up to three exhibitions per year in the print gallery. So our focus would be developing those exhibitions or shows that would be art historically sound, under a certain theme or concept but that would also be appealing to the public, the general public that comes to the gallery. And we would vary that not only across our schools here but also across different media and in addition to different themes, as I said, or concepts. And I suppose the second facet of our work here would be increasing knowledge on the collection, so constantly updating our dossier files or our... as I said information on specific works of art. So we would have a great number of researchers in here on a monthly basis. And a lot of them would be researching either a book on a specific topic but some could be, you know, experts on a specific artist and so they can tell us well actually that's the brother of the artist you say it is. Or that work is... they would re-attribute it. So if they are the current expert on a specific artist we would generally take their advice and add that information to the file.

Interviewer: Ok. And...

Expert A: So it's constantly kind of developing knowledge on the collection and publishing on that as well. So if there is any new acquisitions come into the collection we would hope to publish within a couple of years something on that specific work.

Interviewer: Ok, ok, that would be like a monograph then?

Expert A: Ehm... yeah it would be... yeah it would be like a monograph, but it would also... we would like to display it as well with this new information.

Interviewer: Ok, so that would be the piece of information you get, like not just the small one but say a table of...

Expert A: Well, this is it: In the exhibition space we would always display just a label text which would be about 100 to 150 words, but obviously the best way to publish would be within an art historical journal.

Interviewer: Ok. And now these experts you said you get them in on the occasion of an exhibition specifically? Or just sort of...

Expert A: No, on the off chance so we're constantly updating that information in our dossier files. And then if we come to do a show if we tend to include a work that has been re-attributed or whatever we will explain the reasons why or say...

Interviewer: Say it has been believed that it...

Expert A: It has been...

Interviewer: ...but now we think that...

Expert A: Exactly!

Interviewer: So, what information about an image is gathered and recorded in your files? So what do you keep in your records roughly?

Expert A: The most important thing would be the attribution of a work. So updating any attributions or any... If we have any question marks about a work it's always good to search around in other galleries and see what the believe, you know, the specific aspects about a work that's maybe similar to the work here.

Interviewer: Attributions means authorship, basically.

Expert A: Yes, yeah, who the artist is. So that would be the key most important information about a specific artwork we have here. But in addition to that I suppose we work in collaboration with the paper conservators. So they are constantly, if a new work comes into the collection, looking and examining the pigments used or the support the work is actually done on, whether it is board or paper or what type of paper it is and can we date the paper and is there a specific watermark that's typical of a certain artist or workshop or something like that. So all of that information about the media will be important to the file as well. And then any articles that have been... that the work has been included in we would, you know, insert that into the file, too. So anywhere it's appeared, whether it be in a...

Interviewer: Oh, so references to scholarly articles...

Expert A: Yeah, anytime its referenced. Anytime it's referenced either in a scholarly article or whether it's appeared in an exhibition abroad or whether it's travelled.

Interviewer: All right, ok, ok. All right, I see. So there would be the provenance and also the scholarly body of literature?

Expert A: Yes.

Interviewer: As far as the content of the image is concerned, so what the image actually shows, the scene and so on. Do you record any information about that? Is there sort of written description about what the piece shows?

Expert A: Ehm, well actually in our files, ehm... in the past you can see that they have written down absolutely in detail, you know handwritten, what exactly is in the work. There were horses to the left, you know, the rider is to the right, there is a mountain in the distance and the overall colour is green or something like that. So it was very, very detailed. We tend not to do that as much now.

Interviewer: I suppose now from what you say that was done in the past, it was to uniquely identify the... the... the... Like, is it for sort of cataloging reasons or ...

Expert A: Well that was more than likely or old system of cataloging.

Interviewer: Yeah. Whereas now you just keep a photographic...

Expert A: Ehm, we keep a photograph and we just have... you know, we just title the work, we don't necessarily detail everything that's in the work. Ehm... and I suppose people will gather that type of information if it's been written about in a certain book, do you know what I mean?

Interviewer: Yeah.

Expert A: But we don't go into that kind of very very detailed information.

Interviewer: Ah, I see, yeah. But you said you would have... you would still have these records?

Expert A: Yes, yeah. Oh, no they would always be kept but it's just I suppose that kind of time consuming system we wouldn't use any more.

Interviewer: And say for the purposes of the exhibition you'd then maybe only in the exhibition text highlight a few bits that are in the image and say like this is... or is that am I roughly correct there?

[a member of the caretaking staff enters the room briefly: there is a short interruption]

Expert A: Say that last question again.

Interviewer: The... oh yeah like you know... for the exhibition description text sort of like sometimes you would highlight to laypersons who look at the painting that a certain figure is Bacchus or something like that.

Expert A: Yes.

Interviewer: But that is then done for the...

Expert A: The public.

Interviewer: For the public and when the exhibition comes up. There is no sort of systematic record kept about who shows up in your...

Expert A: Well, in... certainly for label text we also, we title the work, we give the medium, we give... we don't often give the dimensions either, because you know they're seeing it. And what we try not to do now is basically detail exactly what's in the work cause we feel that people can see that themselves. But if there is something that needs to be highlighted the fact that ehm... Turner very much concentrated on the sun when he was doing his images of Petworth (?) Park or something like that, we would point that out even though they can see that the sun is a very important feature of the work or as you said if Bacchus is in it or something like that or ehm... Yeah, we would detail that kind of information.

Interviewer: Yeah. Ehm...

Expert A: Am I explaining that properly?

Interviewer: Oh, yeah. Absolutely, absolutely, absolutely. No, no, no. It's just... it's very good for me to get like... because I only have vague idea what sort of curators do and so it's just very good to detail the process a bit so I get an idea what your job is and... because obviously I'm coming from a slightly different... like say the main question I would be concerned with is...

Expert A: Yes, of course.

Interviewer: ...like say if I look for images online, how can I find them or something like that you know?

Expert A: Yes.

Interviewer: So in that sense it's interesting what can I see in this image, because I might not know the title, I might not know the artist or I might not even care about the title or the artist. I just wanna see images that contain a certain... So that's why I was asking about the image content.

Expert A: But say for instance now in the past we would have ehm... eh... let me see. Just to give you an example we have a work... this our kind of older, well this is our catalogue, you know, our essential guide and we have noted here that this is by Claude Vignon and it's a work showing Porcia.

Interviewer: Ok.

Expert A: But I suppose they have tried for the purposes you know our catalogues etc. to shorten things, whereas in... you can see here she's titeled "Porcia walking with her suicide by swallowing hot coals in the background. An illustration for these ?('Femme Forte')?" and the date of these illustrations. So none of that is included in the blurb or the information we provide there.

Interviewer: Ah right. Because that's of course... that's an important information. Like I wouldn't have... First of all it's hard to see what's going on in the background here and secondly it's basically her again. So it's like...

Expert A: It's an alternate scene. So that's just an example of how I can see that there are certainly shortening titles within our collection. And there is often times quite a lot of information like something like that that's just cut, you know? And certainly for our TMS system, out museum system that we're using, the database that we use, the primary title is now "Porcia".

Interviewer: Yeah.

Expert A: Whereas we have the secondary title you can go into... kind of you know... you can just click on the kind of ehm... the secondary title and it'll give all the extra blurb.

Interviewer: Ok.

Expert A: But I suppose just for people... you know, for easy access for people when it goes online maybe that's what they're aiming at, to keep it as simple as possible.

Interviewer: All right. So and that... What would... So that would be the book like the essential guide and this would be a more comprehensive reference?

Expert A: This would be our illustrated summary catalogue. So this would be the kind of bible we would use still today about you know the works in the collection. And all of this information is on our database but as I said maybe as secondary information you know that... the more detailed stuff is secondary. But you still can access it.

Interviewer: But you still can access it and even in an electronic form, sort of like.

Expert A: Mhm. Yeah you can.

Interviewer: because mentioning databases, this is actually my next point. So what kind of databases or so. You mentioned this TM...

Expert A: It's called TMS, it's "The Museum System".

Interviewer: Alright, ok.

Expert A: And it's a system basically for... ehm... curators to use... ehm... to access the collection. Ehm... it hasn't gone online yet, but we're organising that at the moment and ehm... so we use it for basically ehm drafting up lists for exhibition ehm, you know you can... you can drag and drop works...

Interviewer: Alright.

Expert A: ...if you... your prefered works for a show and create a kind of an exhibition.

Interviewer: A virtual exhibition.

Expert A: A virtual kind of exhibition within that. Ehm... and then curators are all the time as well adding information to that, ehm... on specific works, so the most popular works in the collection. Say for instance all this written information will be on TMS now because anything that's in the Essential Guide was seen as, as crucial to get on the database, eh... first. Ehm, in addition to that we have a kind of a, ehm, eh, what would you call it? A registrar's section so they can see basically exactly where the works are within the... both within the gallery and maybe if there are works out on loan, where exactly they are. So it's a kind of a tracking system for them as well. And then you things like insurance values and all of that type of thing would be on their system as well.

Interviewer: Alright.

Expert A: Then we have an exhibitions area, ehm... where, you know, the head of exhibitions basically has her, her own kind of site. So she's basically monitoring what exhibitions are happening within the gallery and what exhibitions maybe of ours have travelled abroad. So there's that kind of section, so there's quite a number of sections that we don't use at all. As curators we just, we just use the...

Interviewer: But it would all be part of that system? Or no there's different...

Expert A: It's all part of the one system yeah, it's all part of the one system. And... ehm... in addition to that... I suppose the reason why we use it mostly... still, is for locating stuff. So ehm... or seeing, basically... you know, at the moment I'm changing any attributions. So any attributions that we have noted as, as, as to be amended, I'm amending them on the database. So ehm... yeah so... although still we use it in a minimal way, just basically to say finding one object and we'll key in our NGI number here and then it'll come up where it is. But it'll come up a full front card you know stating...

Interviewer: Yeah.

Expert A: Do you want me to kind of?

Interviewer: No, no that's ok. So ehm... and the full front card contains also, like you have a digital image then as well of the... ?

Expert A: Yeah we have a digital image and we would have, you know, all of this information, ehm... the provenance, the eh... the media, the dimensions, the title, any, any previous titles... ehm... the date of the work, which is very important, can often change, ehm... whether it's inscribed or signed or ehm... you know, any, any... kind of adjunct information, whether it's a preparatory work for an oil within our collection that will be kind of ehm... under a heading, I think it's like eh... you know, related objects, so...

Interviewer: Yeah, yeah. So you can reference...

Expert A: We can reference, yeah.

Interviewer: Do you also use any web-based... or would you also use web-based ehm... tools, like say you want to know something about a certain artist, would you go out and you that like, say, maybe you have three paintings here or something like that, but there are obviously more, so would you go out on the web and like check to get a quick reference...

Expert A: Yeah, yeah, we would. Yeah we would, definitely, yeah.

Interviewer: And, and, what do you, sort of, what do you use there like? Is it just Google image search, which I would use, or is it a more specialised system?

Expert A: I suppose we would start of like that first... ehm... and... eh... what would we use? I suppose the, the databases we use mostly are the Getty index of names, say, have you heard of that one?

Interviewer: Is it ULAN, like Union List of Artist Names?

Expert A: ULAN! Yeah, yeah, yeah... we would use that a lot. We would use Art Index a lot for sales prices, eh checking up what has, you know... in descending order basically what works were sold where and what prices they would have reached. So that's very good for us... for, for kind of gauging eh... the likes of insurance values if we have works traveling out. So...

Interviewer: Oh and the ULAN, you just use the online browsable version? So you haven't that licensed here, say, in the gallery or like you, you just go to the Web page and...

Expert A: Ehm, no I think, like, we all have it on our desktops, ULAN, so... I don't know I'm sure we've all, we've subscribed for it.

Interviewer: Because you can use it for free online. Like you can go to the Web page and use it, but if you sort of want to use it locally or something like that I think you have to, you have to license it... So you...

Expert A: I think maybe...

Interviewer: ...can browse the names and you can browse the hierarchy, like if you want to use the whole thing to...

Expert A: To be honest, the people that we go through for ehm... databases or the likes of ehm, kind of, researching a specific artist or that type of thing, ehm, it's the library set up all of those databases for us eh, and gives our user name and passwords and that type of thing. So it might be actually interesting for you to speak to eh... Catherine Sheridan in the library. She's doing, she's doing the web page.

Interviewer: I met her at one of the meetings, like, Catherine Sheridan, I'm pretty sure.

Expert A: I think she'd be very important for you to talk to, cause not only is she aware of the ehm... web-based information, but she... you know she's very aware of the different types of Web sites that the curators and other people in the gallery are using as well...

Interviewer: Alright.

Expert A: ...cause she would generally set them up.

Interviewer: Ok. Oh yeah, it's good to know. Ehm, Catherine Sheridan, I definitely remember her, remember her name. I met her at a, at a small workshop we had over there in, in Trinity.

Expert A: Ok.

Interviewer: It was a good while ago, but like...

Expert A: I believe they have the, they do have the actual, they have chosen somebody out of the tendering process to do our web site, so...

Interviewer: Alright. So like for the, for the remainder this is sort of as much your professional persona as your private, like, persona as well. Ehm... oh no, first of all, what do you think of online image galleries?

Expert A: Well, you know we would use online image galleries in quite a basic to be honest ehm... I would generally just search for images for use on Power Point presentations here for the likes of lectures. Ehm... and that's what I would use image galleries for... that's about it.

Interviewer: And what do you think is the biggest difference to conventional galleries? Like... in terms of...

Expert A: I would say, possibly the quality of actual galleries eh... you know images would be better you know? I would think so. So we would tend to maybe try and steer for, if, if... the National Gallery in London have a great image of a Velázquez painting I definitely would use that one over one that Google has.

Interviewer: Oh... eh, no eh, when I said conventional galleries I meant actually physical galleries sort of like... basically what is the difference between exploring art in an online gallery as opposed to, say, like going to the National Gal...

Expert A: Oh eh going and seeing things in the flesh?

Interviewer: Yeah.

Expert A: Well I would say go and see the actual gallery, of course.

Interviewer: Yeah, yeah.

Expert A: But as a curator I think it's much more of a, a, enlightening experience to actually go and see something in the flesh, no?

Interviewer: Yeah. No, I, I, I would agree so... yeah, yeah, no.

Expert A: I think you get a much more rounded kind of impression, no?

Interviewer: I totally agree.

Expert A: And you're also getting added information from the curators that maybe doesn't exist online.

Interviewer: Good point, yeah, yeah. No I mean obviously some things you just can't...

Expert A: And you are also seeing things in a kind of a contextual way as well ehm... within a kind of a group show, so you're getting more of a... an all rounded as I said experience, rather then a... clicking on something on the... I just find ehm... it's more enriching, no?

Interviewer: Yeah.

Expert A: Maybe you totally disagree...

Interviewer: No, no, I don't disagree at all, like it's just... no. Ehm... so and how often would you search for information about artworks on the web now, say as opposed to your local database, like, basically just go out online and... you know things that wouldn't be on your database... Would that happen, like, a lot?

Expert A: Quite often yeah, I would, I would search through Google to be honest and then anything that comes up that, that is of interest, I would ehm... I would tend to use, I wouldn't be fearful of using it. You know if a couple of people are saying the same thing, I would generally consider that to be... and I would maybe double check it on, say, the Tate's Web site if it's, if it's connected with Tate or... you know, but I wouldn't, you know... a lot of people are very sensitive to using material online, but I, I have to say it can be very helpful.

Interviewer: Yeah. If you look for information about these artworks, say, using Google or anything, like, do you typically like... are you typically interested in a reproduction of the image or in textual information about an artwork, like?

Expert A: Ehm... both.

Interviewer: Both. And how would you rate the ratio roughly, like if you... shooting from the hip.

Expert A: I would... yeah I would say ehm... well to be honest we would use ehm... say if there is a work that we are putting on display that we haven't shown in a long, long time and I actually happen to not really know much about that specific artist, you know, because we have over 10,000 works in the collection, so ehm... you know, I'd do a bit of research on that artist and I would not only read around the books in our library but I would also search online and again see does... see does anything come up as to whether there are other works within Dublin institutions or... You know I often find that, that the online way is a great way of kind of accessing other institutions basically, do you know, at a much quicker speed. You know so...

Interviewer: Than ringing them up or...

Expert A: Then ringing them up or even seeing what they have online about a specific work, you know. We would do that quite a bit with the National Library, because their database is excellent. So yeah...

Interviewer: Ok. Alright. Ehm... what have we got left? Oh yeah there we have like, if you look for an image, which of the three following statements describes your goal, or if you look for images? Ehm, first of all: "I browse the Web and search for interesting works of art without a particular target in mind". So you just look around.

Expert A: Mhm.

Interviewer: Secondly: "You look for artworks that fall into a certain category and have something in common". And thirdly: "You look for one specific artwork, which you need to find".

Expert A: I would say I would do the last two.

Interviewer: The last two?

Expert A: I wouldn't generally do the first at all.

Interviewer: Yeah, ok. And which one of those... if... about evenly?

Expert A: Well, say, if we're doing an exhibition, it might be interesting to see what other, you know, museums have done... similar types of exhibitions and what maybe images they've used or artworks so seeing them as a kind of a group thing would be of interest yeah. And then ehm... I'm always looking for particular works of art, so I'd say maybe the last one most.

Interviewer: The most frequent probably.

Expert A: Mhm. Yeah.

Interviewer: Ok that was already done. Ehm... if you, if you, if you use Google, say, and you go and you do keyword based search, ehm... do you have a special method or is it just common sense.

Expert A: I think common sense, because I wouldn't say that as curators in the gallery we would be very technically savvy. So I'd say it's again quite a basic level...

Interviewer: So what...

Expert A: You know we would put in an artist and their dates and ehm... you know whether it's a printmaker or a pastellist or a miniaturist or something like that and... see what comes up that way.

Interviewer: Ok.

Expert A: It wouldn't be in depth.

Interviewer: Ok. Ok. Ehm... have you ever heard of or used so called query by user sketch or query by image example techniques?

Expert A: Query by sketch, what does that mean now?

Interviewer: It means, basically, you, you know, say, the rough composition of the painting and you do, like, a digital sketch of it and say: "Go off and find me paintings that match this general pattern!"

Expert A: You make a sketch?

Interviewer: You make a sketch, yeah.

Expert A: No, I've never done that before.

Interviewer: And, and the other one would be, query by image example would be, you have a digital representation, maybe a very small one, or maybe an image which is kind of similar and you feed that into the search process.

Expert A: Never done that either.

Interviewer: Ok.

Expert A: No. And I have never heard of any of us doing that.

Interviewer: Ok. Ehm... it's, I think it's early days like, so like, only like five or ten years ago people have started looking into that.

Expert A: Oh really? [laughs] Ten years is a long time.

Interviewer: But, like, I mean... When I say early days I mean there is actually no good online portal where you can... actually so...

Expert A: Right.

Interviewer: I was just wondering maybe you would, you know, sort of... more than... because I'm pretty sure... I never even... I hardly ever used it.

Expert A: Right, ok.

Interviewer: I tried to play around with it a bit. Ehm... ok, the last two questions are like, so if you could make a wish or you could like sort of be demanding and imaginative, what would exploring artworks on the Web, like, sort of regardless of technical restrictions or what you think technical restrictions would impose, what would that be like in, sort of, an ideal world.

Expert A: In an ideal world, I'd love to be able to kind of ehm... key in, you know, "Degas' drawings, search globally" and all of Degas' drawings come, come up, in addition to where exactly they are, ehm... and the, the relevant information for each drawing.

Interviewer: Yeah.

Expert A: That would be marvelous. That would be a, you know, a... curatorial kind of dream.

Interviewer: And...

Expert A: But that doesn't happen generally. You get a certain amount of images and then you have to kind of go off on other tacks and, you know, I just don't think there is a kind of ehm... as such a kind of very unified search engine for these types of images. Like I know that you can go into ehm... specific like art sites and eh... they'll bring up a certain amount for you, but... that would, that would just be, that would be great.

Interviewer: So actually I just realised, I skipped a question here which is related to what you said, so, so, like, what would be the biggest shortcomings of current image retrieval systems, like, ehm... is it that...

Expert A: I think you can get to, if you put in, kind of, again I think it's possibly a problem of us putting in very basic information, ehm, query information in the be... ehm from the offset. But I think you, you can, they can drag up or retrieve maybe three or four from you and you go, oh that's interesting, I wonder where that is and you have to then, click on that and then you have to go into a whole rigmarole of, say, the British Museums Web site in order to, you know, find out more information about that, rather then having it literally at your fingertips, like that, that type of relevant information ehm... literally like as an adjunct to the thumbnail you're looking at, but that would be marvelous if you, if you could click on something and instead of going, you know, down in, say, Google, going into several different sites... that it would all just come up on one, kind of... image site, you know?

Interviewer: Yeah. Ok.

Expert A: That's maybe just very far fetched.

Interviewer: No that's like, but that was the point of the question. That's why I said, be ehm, sort of, be imagin... what would be the ideal thing without, sort of, what...

Expert A: Like I think, to be honest, it's very possible.

Interviewer: It's very possible.

Expert A: But eh I, you know, say... It would be marvelous to be able to, instead of having to dredge through like tons of books in different libraries etcetera to be able to say, oh well, say for instance Picasso has, you know, images, you know the Ballets Russes or whatever, whether they be paintings, drawings, ceramics whatever, ehm... and they're in the Hermitage, they're in the Louvre, they're in the...

Interviewer: Yeah.

Expert A: They're in four museums in the States and you would be able to, kind of, have them all as a package...

Interviewer: Yeah. And the Ballet Rouge is a ballet now...

Expert A: Eh the Ballets Russes, it was this Russian ballet, that he ehm... basically, kind of, was very connected with ehm... He did a lot of drawings of... Let's see do we hav one here? Yeah. So ehm... yeah he was, he was, kind of, quite friendly with this man called Sergei Diaghilev, who was the director of the Ballets Russes and he did a lot of drawings of the likes of the ballerinas.

Interviewer: Alright, ok.

Expert A: Do you know what I mean, so it would be great to search under something like that, even "Picasso" and "Ballets Russes" and know exactly where in the world those types of work exist, but that's not to be, I think, just for the moment. You can go so far, but then you realise, oh, having read another couple of books, oh well, that didn't come up on my search or... you know? You're not getting the whole picture.

Interviewer: Ok.

Expert A: I don't know I'm sure, I'm sure that could be possible in the future.

Interviewer: I'm absolutely sure it's possible, it's probably just a question of, sort of, manpower and you know...

Expert A: Manpower. But then like, say for instance, all these Google search engines, there's incredible manpower, I can imagine, at something like that, so...

Interviewer: Also I'd say like, if the proper framework would be established then like if you say you go off on your search and like after a while you find in this other book that one didn't come up, then you as a sort of an expert on the field or having discovered that now, could key that in somehow...

Expert A: Exactly.

Interviewer: So it would come up in the future if someone ever again looks for it...

Expert A: Exactly, so this kind of collaborative thing between curators maybe, or further afield... but even today we were talking about Wikipedia, that you can, I didn't realise that people can, you know, amend or add information to it all the time, so it could be something like that. But say for instance I've noticed... certainly for our exhibitions, but in terms of an exhibition where there is a lot of loans, the likes of the Impressionist Interiors at the moment, that exhibition is pretty much made up of loans from abroad from all over the world and I'm sure that the curator there, of that particular exhibition spent months trying to locate these specific works that would be relevant to her theme of impressionist interiors ehm... whereas wouldn't it be great if we could just, kind of, go in under kind of themes like that and... for these things to pop up...

Interviewer: Yeah, yeah. I think it would.

Expert A: Yeah it would be great.

Interviewer: Man, shame we don't have it yet. So you basically, on the go you answered the last question already, "How would you envision an expert system on art?", you know, where all the knowledge, not just like the retrieval of the images, but where all the knowledge and the connections are gathered...

Expert A: Exactly, where there, there... not just under a specific artist but maybe how they, they appear under certain themes, like, ballerinas or whatever in art or... You know what I mean? So it, it...

Interviewer: Yeah.

Expert A: ...can come up that way as well, ehm... Yeah, I think that would be just marvelous. It would make everybody's live so much easier, but...

B.2 Expert B

Participant:	Donal Maguire
Participant ID:	Expert B
Participant Position:	Assistant in the Centre for the Study of Irish Art
Location:	National Gallery of Ireland
Date of Interview:	7/05/2008
Interviewer:	Daniel Isemann
Interviewer ID:	Interviewer
Transcriber:	Daniel Isemann, Victoria Vize

Interviewer: Off we go. So now if we talk, yeah there something happens, ok. That should be it. Sorry, for the...

Expert B: So these are the fields then. So they are things from the physical kind of description that physically describe the item, its media, its dimensions to things like the collection that its part of, its provenance where it came from eh... the access that's allowed to it by [incomprehensible], what collection its in, where it came from, things like that. For some of the longer letters there is also transcriptions and there is a technical record, just so we have, we remember, or we have a note of how, what is was scanned with, how the digital version was created, when it was made, what the resolution was and then that's the collection its part of.

Interviewer: Ok.

Expert B: So this is items subjects then. Every item in the database has subjects that are assigned to it and it basically allows the user on the front end then as a way of accessing material. So its not keyword searching or browsing. You can search by subject. So say in this case it's... they're painting a bed and so bed is subject as part of furniture and ink is the media. So you can search... so say if you type in "bed" everything, every item in the database with "bed" will come up. So its a way, it's just a keyword way of searching.

Interviewer: Sweet. I can see you have you have a little... sort of a tiny ontology there. Like objects, furniture, bed.

Expert B: Yeah. I can you all that in the front end. And this is something they developed really for us. Ehm, it's called "related items". What it does it allows you to create relationships between items. So in this case this letter, eh, it says is in the hierarchy, the hierarchy relationship it says it is part of this item GRA. GRA is, ehm... this item here, which is, eh the Gray's [?] collection item. So it's basically all of these letters are underneath this item in the hierarchy. So this item it says contains these items.

Interviewer: Ok.

Expert B: This item 197 is part of GRA.

Interviewer: All right, ok.

Expert B: So it's just relating two items together.

Interviewer: So you have like, kind of have like virtual groupings of collections?

Expert B: Yeah. So I just show you the front end, how it all works. So this is Art Search then, this is our... this will be like an online interface that people can use when it eventually...

Interviewer: Yeah.

Expert B: So there are four basic ways of searching. You can search by artist/author, subjects, collections, browsing and keyword searching. So just an example collections. Eh, these are the collections that are [incomprehensible] the National Gallery of Ireland, Archives, the CSIA Collections here, the Library Special Collections and this is from the Prints and Drawings Collection. So this is the GRA item, the Orpen item. You can select that item. Ehm, just gives you a little blurb about the collection and then you can view all of the items that are in that collection. So there is 31 pages, which is the 365 letters.

Interviewer: Ok.

Expert B: Then you can eh... see where is the bed. In fact I'll show you this one here. This is a letter then. So you can chose the letter, eh... It provides you with the image. This is the data record you saw in [incomprehensible]. Ehm, there's other information you can put in, like information about the artist. So on the basic page it tells, it shows you the image and these are the relationships here. It's saying this is part of the Orpen collection. Ehm... and the subjects for it are these items here. So it's things about love, emotions, his daughter, William Orpen's daughter, baby, pillow, eh... Mrs Saint George, the media, ink, [incomprehensible].

Interviewer: His mistress...

Expert B: So if I select "baby" it will bring up all the other items in the database that are...

Interviewer: related...

Expert B: ...also have baby, just that have babies in it. It doesn't have to be anything got to do with this item. It's just the subject "baby" and it will bring up other items with babies in them.

Interviewer: And can you use the little hierarchy you had there before like furniture, you know, objects, furniture, bed, like is that available as well on the...

Expert B: Yeah. So that's browsing by collection, you're just looking at the items. Subjects, you can search by subjects, this is the hierarchy here.

Interviewer: All right.

Expert B: So something like ehm... you can pick anything, emotions... sorry... it's actually it's optimised for Firefox. So... it should open it now.

Interviewer: Yeah, I suppose most people...

Expert B: Here, so subjects, ehm so... emotions... and it's just, it's a hierarchy of keywords. Ehm, so you can go to any of these kind of emotions, so something like happiness and it'll bring up an item which has got something to do with happiness.

Interviewer: All right. Now I'm particularly interested in that as a matter of fact, because that is sort of like the... the topic of my PhD research.

Expert B: Ok.

Interviewer: Like these ontology or taxonomy things. So if you could just...

Expert B: Well this is actually... this is something a lot of eh... a lot of databases use. We, I think we based ours on the Tate website.

Interviewer: All right, I was just about to ask where did this come from.

Expert B: Have you seen the Tate, their webs... the had their archive online as well and they are a very similar layout as this. And eh... you can search by keywords as well. So it's based on theirs really.

Interviewer: I have to check that out. I think I've been on the Tate page, but maybe I missed that one.

Expert B: And you can also search for subjects. So say if you type in "horse" it searches and brings up that there is two subjects "horse" and "horse racing". Select "horse" and it brings up items got to do with horse.

Interviewer: Yeah.

Expert B: So in this case there is actually from another collection. This is eh... this is Michelangelo Hayes [?] Collection. It's from that collection. And he made these... they're called fenestitiscal [?]...

Interviewer: Fenectististical [?]...

Expert B: And they are hand-painted wheels, which they spin around and you basically hold a mirror up to them and you look in the mirror...

Interviewer: you see...

Expert B: ...it just sections off one part and you see, it appears as if... it's like animation kind of. But we can actually import video onto this, so we're able to make a video.

Interviewer: Nice. Wow.

Expert B: So you can imagine if you hold a mirror up to it and look in the mirror you only see that there.

Interviewer: Yeah. Cool, it's really nice. Yeah. Good stuff yeah. I like it that you have the video straight away there, you know.

Expert B: Yeah.

Interviewer: Just do you have the hierarchy or the taxonomy you have there, does it have information about, say, human anatomy as well and body parts that sort of thing. That was a topic we had a look into... a while ago. Medical or maybe people and society? Actions, adults, children. Human anatomy, ah yeah there we go. Beard, eye, face, foot...

Expert B: See we build this as we make things, so it's only the words that are here really most of them relate to something that exists in the collection. So as we put new items in more words, keywords will emerge and we'll put them into the hierarchy. So I mean something like "beard" if there is a number of photographs with people with beards.

Interviewer: No, no, but that's exactly the thing we... we basically... we wanted to build up this little ontology. That's... was part of my... on the research day here as well. And... and we... we just picked human anatomy as a field, because we thought most images will contain people and you know like if you want to have a really detailed information including parts information and stuff like that. So we thought it would be a good point to start. So beard, face, yeah, yeah.

Expert B: So eh... I don't know where beard comes into it, but someone obviously identified it was a beard in this and eh... maybe if you're interested in 19th century beards you might...

Interviewer: Yeah, yeah, no.

Expert B: But eh... so the images are quite... they're scanned in very high resolution. So eh...

Interviewer: The guy here seems to be the... the... Impressive specimen of a beard.

Expert B: These are actually in the south of England. They are the artist Walter Osborne's photographs.

Interviewer: Ok... very good.

Expert B: So eh... that's collections for browsing, then subjects, then you can also search by artist and author. So it's just an alphabetical index and say eh Osborne, eh that's Orpen there also, Walter Frederick Osborne. This is his photographic album. So they are just pages from his album. Photographs and collected. And some of his, actually there is some of his letters there as well. Ehm... so I just show you quickly some of the more complex items.

Interviewer: Yeah.

Expert B: Eh one of the items that we... eh so this is... in the CSIA Collection there is smaller collections of artist material. So there is a... these are all individual artists' collections... and one of the collections is the Crofton Crocker Archive Collection. Crofton Crocker was an antiqarian and he, he had this... it's basically a scrapbook that he made that tells the story of the last days really of the life of the artist Thomas Foster. And Thomas Foster was a early nineteenth century artist. He was a real romantic and he at a very young age, I think it was only when he was about 29 or 28, he eh he was living in London and was kind of successful as a painter, but ehm for... nobody is really sure why... one night in his hotel he shot himself in the head with two pistols.

Interviewer: Jesus!

Expert B: And eh he eh so Crofton Crocker basically was alive, it was in London at the time and interviewed a few people around the time, see what was happening and basically detailed the story in this in this scrapbook.

Interviewer: He reconstructed...

Expert B: Yeah. And he also... sorry. He also, in the book he also eh stuck in different sketches and notes and stuff like that, by Foster, by other artists, by friends of Foster and anything he found about his life just put it into the book. But at the start of the book it was a scrapbook that he had that he had actually picked up in eh an antiques shop. So there is these other pictures by Briggs at the beginning. So there is a lot of different artists involved, people, individuals involved in the book. The book itself the artist author is Crofton Crocker because he mostly collected it and put it together. Then these are the pages of the book here. And at the beginning these are watercolour studies by the artist Briggs.

Interviewer: Briggs.

Expert B: So you can actually go in and view those if you want to.

Interviewer: Jesus, it's really kind of tricky, like you know the question of authorship and stuff like that...

Expert B: Yeah, exactly yeah.

Interviewer: ... it's not as easy any any longer.

Expert B: So they are the first pages of the book. Then around here it starts talking about Thomas Foster. These are works and eh... these are portraits of Thomas Foster by different artists. There's different... and then there is drawings by Thomas Foster stuck in and there is things on the back of things. One of the most complicated pages was this one here. So this is a page of the album.

Interviewer: Yeah.

Expert B: With eh... this handwriting is on the album page and it's by Crofton Crocker and he's writing about Thomas Fosters life, ok? So the artist author of this page is Crofton Crocker but it also tells you that eh there are items pasted to this page.

Interviewer: Yeah.

Expert B: And these are these three items here. Eh... so there is a number of relationships here. This item contains these items, these three items.

Interviewer: All right.

Expert B: It is part of the album. Eh... you can also we had a relationship that gave you other views of the item. So this is basically the item, this is this page with this page folded over.

Interviewer: Flipped in.

Expert B: Yeah, so you can see what it looks like as well. Ehm... sorry back to the mask here.

Interviewer: Jesus, yeah!

Expert B: And then there is also the thing that they developed called "hotspots" which eh... which allows you to highlight what is actually stuck on. So you can... it's basically these items down here.

Interviewer: Yeah.

Expert B: Say this is a jar ehm its an individual item in itself, it's an artwork by Thomas Foster and this is the record for this. But it's stuck onto this page by Crofton Crocker. Ehm... this one here is also by Thomas Foster, it's a little sketch as well that's on the page. And this item here was the most complicated because it's actually this is one page, ok? It's stuck onto this page but this is also a separate drawing, small little ink drawing stuck on the back of that. So... sorry it's actually a letter, it's a letter it's a letter from... to Thomas Crofton Crocker that he stuck in so...

Interviewer: So the letter is on the front and on the back he glued the...

Expert B: So this page is part of this page and... but this page itself has this item on the back of it. And that also has a separate artist author because it is a different item itself. So you can see they are very you can see even though this is a simple item because of where it's positioned in the album makes it very complex. And eh... the reason for doing all this is... it seems incredibly complicated, but the reason we wanted to do this was because if you type in, say someone comes in and they do a search for eh "jars" if they are interested in china jars and it brings up this item, ok? We didn't want to just be that it was a single item with a record floating around in the ether, that when you select it you know exactly where it comes from.

Interviewer: Yeah, yeah.

Expert B: It's stuck on this page here, it's linked to this page. So you can still go back to the page and see what [incomprehensible]

Interviewer: [incomprensible]

Expert B: ...so the context is always there. That's build through these relationships.

Interviewer: Yeah.

Expert B: So does that make sense?

Interviewer: Yeah. Absolutely. It's very fascinating I have to say.

Expert B: So that's eh... that was the most complex item really. But there are other items as well that we were working on, but eh... does that give you a good kind of idea of what's that?

Interviewer: Yeah, definitely. And that, by the way I'm very glad I switched the recording thing on because that... you answered like quite a few of the questions I have like on, on the go. So I'll just bring my little...

Expert B: If you want to eh... if you need any kind of screenshots or anything like that as well you can contact us.

Interviewer: Ehm... at a later stage maybe, because that would be... that would be interesting to... I mean when is this thing going to go online.

Expert B: Eh well it all depends on the website. Have you spoken to anyone in here about the website?

Interviewer: Ah no, no.

Expert B: It'll... it might be another half a year or so before it goes... the website is up and running. And only then can we start planning for this to go on so it could be a year away so hopefully not too long but could be a year.

Interviewer: No I might contact you about screenshots or anything like that indeed. That might be interesting.

Expert B: But eh if you're looking for more kind of information on eh how these subjects, keywords and things like that... the person to contact is in the library, is eh Catherine Sheridan.

Interviewer: Ah yeah, yeah, yeah.

Expert B: Have you spoken to her yet?

Interviewer: No, but ehm Niamh just pointed... pointed me to Catherine Sheridan.

Expert B: And eh... she's away. She's actually on her honeymoon at the moment. But she's definitely the person to talk to, because she has a background in library studies and she's also very good with databases and things like that.

Interviewer: Yeah. Ehm I met her as well ehm a good while ago she attended, she attended one of the workshops ehm we had together with Professor Ahmad in Trinity. So I... I was introduced to her like that would be like half a year ago or something like that. Yeah I might actually contact her at some point... So ehm my questions let's see. So the first section is basically just about you and your background. So what is your position here and how long have you been in the gallery?

Expert B: Ehm I've worked... this is the Centre for the Study of Irish Art.

Interviewer: Ok.

Expert B: So I've worked here for four years now and basically my position is assistant in the centre. You met Brendan. He is the administrator.

Interviewer: All right. Ok.

Expert B: Eh so it's really... there is only two members of staff here in the centre it's Brendan and myself.

Interviewer: Ok.

Expert B: And... ehm before that I worked as a I worked in the education department as a guide. And eh... then before that I was in university.

Interviewer: Yeah. And and what is your study background?

Expert B: Eh in university I studied... eh... my undergraduate was in eh eh NCAD, it's the National College for Art and Design and I did... they do kind of a joined course up there which is 50 percent history of art and 50 percent of the course is fine art, so painting and history of art.

Interviewer: Ok, ok. Brilliant. And you did... Niamh mentioned that you did... was that you? That you did a Masters in multimedia?

Expert B: Oh yeah and then I just finished a Masters eh last eh November in Multimedia up in DIT, that was in eh digital media and eh so yeah I did that.

Interviewer: Hence also that.

Expert B: Yes so, it was, actually we started working on this before I did the Masters, I did the Masters because so much of kind of research and the arts are involved now with digital media so I felt it would be kinda useful to do it and to have a knowledge of that kind of side of it and just to even be familiar with how things work, even if I wasn't to go into designing websites but just to know how they actually work, its kinda useful you know.

Interviewer: Yeah very good. So let's see, ehm ok you did that already, that has also been covered. Ehm yeah, eh do you use any other like data bases or knowledge bases apart from the purpose built one here, ehm..

Expert B: Yeah.

Interviewer: Yeah eh for your work, like?

Expert B: Well, in the gallery here... sorry, eh so this is the archive management system basically: that's Art Search, Art Cat is the library management system, so this is for cataloguing eh books.

Interviewer: Alright ok.

Expert B: And eh, this has, this manages the collections of ehm the library and the CSA so the easiest way to explain it is Art Search is for primary material. It's for actual archive material, where Art Cat is for secondary material: so it's everything that is like books, exhibition catalogues, eh general pages, chapters everything..

Interviewer: General pages.. yeah yeah

Expert B: And eh..so. One is archive and one is secondary material, or printing material.

Interviewer: Ok yeah.

Expert B: Ok so this is the library cataloguing system and it is used with libraries throughout Ireland and Britain, the system. I insist, seriously it's called [incomprehensible] and then there's lots of online web resource systems we use as well.

Interviewer: That would be the next question as well, especially web based or online resources, so give me a rough idea of what..

Expert B: Well, Artnet is eh very useful. This ehm, it's a catalogue of eh, or it catalogues all of the auction sales throughout the world really and eh including Irish auction houses like eh Adams on St Stephens Green and Whytes and people like that. And what you can do is you can search for eh any sale of any works that have come up over the last ten years or so, and eh so it's very useful for researchers because if you are researching any, if you are doing any work on an artist eh quite often the most difficult thing is sourcing images because books tend to print only the most well known works and any kind of, eh especially for more or less well known artists, and eh a lot of the Irish survey books will print only a handful of images by one artist so ehm and a lot of them. But a lot of them, all of the works would be going through these auction houses all the time, so say even someone like Orpen.

Expert B: Sorry...[incomprehensible]

Interviewer: Oh you need, you need to have a user account for this Artnet basically or..

Expert B: Yeah you can look at it if you wanna see it [both laugh]. I trust you. Yeah you have to subscribe to it... Eh so, they are just all ways of searching.

Interviewer: Yeah, yeah

Expert B: And really it's, as well as art historians, people involved like the art market is a huge business, and eh so people involved in the art market will use it to as well to see, to track prices of art works and things like that. So ehm...

Interviewer: There we have them yeah.

Expert B: There are a number of different Orpens, Sir William Orpen is the Irish painter.

Interviewer: Disambiguity of like country and year or date of.

Expert B: And this would bring up eh all of the results. So you can see there are 498 in total and they are all of the works that have gone through the auction houses in recent years eh so a lot of oil paintings. This was in eh just very recently in Christies on Thursday May 8th... or sorry that's not... That's coming up. The estimate is this and eh...

Interviewer: Alright yeah. [Expert B's phone rings.]

Interviewer: If you have to take this or something...

Expert B: No, whoever it was hung up, so.. Wrong number... [Expert B accepts a phone call, the interview ceases for about a minute]

Expert B: Sorry... So eh yeah these are works then, eh some of these are in upcoming sales actually so they have already got them in... eh... actually all of these, they have already got them in, ok sold for, so these are ones that were recently sold. This was in on, in February in Whytes, eh which is on eh Molesworth Street, and it gives you the estimate, it gives you the price it was sold for, eh where it was, and eh any other details there. So this and eh I mean you can see if anyone is using this or if a researcher is using this, they type in the word Orpen and they get 498 works by the artist rather than going through a book and getting maybe ten or twenty works by the artist.

Interviewer: Or Google images.

Expert B: Exactly yeah. And the image quality is good, I mean it's not very high resolution but you get a pretty decent image, because they're taken directly from the eh sales catalogues. That's quite good. Ehm so that's Art Net.

Interviewer: Ok, ok. And this sort of searching for images would be part and parcel of your daily job pretty much or like... [incomprehensible]

Expert B: Well if people come in and someone comes in with an enquiry I mean we get all sorts of users coming in here, people who are interested, who are doing research, postgraduate research, postgraduates and are looking for works by artists or.. [D's phone rings and he accepts the call, interview stops for a few minutes: 26.23- 29.41] [Interview resumes at 29.41]

Expert B: Sorry.

Interviewer: No, no worries.

Expert B: Ehm so another one we use is eh, this is the Oxford Dictionary of National Biography. Have you used this before?

Interviewer: Ok no, no. You have to subscribe to this as well?

Expert B: Well this was a printed dictionary for a longtime and it went online recently and ehm I'm not sure how ehm,,, There are literally tens of thousands of entries of eh biographies of ehm important individuals from Britain and eh Ireland as well.

Interviewer: So national means British and Irish, does it?

Expert B: Well Ireland anyway up til 1920 or so. Maybe I haven't, maybe there are people after 1920 as well, but ehm...So just again, an example: Orpen.

Interviewer: Yeah.

Expert B: And this is eh the biography of William Orpen.

Interviewer: Yeah, yeah.

Expert B: And its em, they usually get people who are specialists in the subject as well, so it's a very reliable source as well for any kind of biographical information. Bruce Arnold wrote eh William Orpen's biography.

Interviewer: Yeah ok, oh very good. It's very authoritative.

Expert B: Yeah, yeah. Ehm another one is eh Nivel, is the National Irish Visual Arts Library. And this is based up in NCAD College, in the college. And eh they have a similar collection to us here except they specialize really in eh contemporary art or late 20th century art, so it's from the 70s, 80s and 90s ok.

Interviewer: Oh ok.

Expert B: And ehm they have an artist data base that you can search by artist and ehm someone new like eh, there's an artist Dorothy Cross. Do you know her?

Interviewer: No.

Expert B: Eh or even just eh do you know Willie Doherty, he has a video piece in the Douglas Hyde gallery at the moment? Have you seen this?

Interviewer: Ah, no, no

Expert B: If you drop into the Douglas Hyde you will see it.

Interviewer: I can do that on my way back.

Expert B: He actually, he exhibited in the Venice Biennale this year as well.

Interviewer: Ok.

Expert B: Ehm that's strange, oh sorry. I put in O'Doherty. So this thing.. eh. And it's not an archive; it's not an image database. It's basically just eh it's like a library catalogue its ehm, it tells you information about the artist and then they have folders up there with things like newspaper clippings and things like that and magazine clippings and catalogues, and all that kind of ephemeral material. But it tells you a bit in the file being up there.

Interviewer: So that would be more sort of like, that would be more like textual information about the..?

Expert B: Yeah, yeah, it's a library rather than an image.

Interviewer: If you would, I mean, it's hard to do this in general but like if you would rate the ratio between looking for like visual resources and textual resources about art like ehm is that about on par?

Expert B: 100%?

Interviewer: Or yeah 100%, or..as far as your work is concerned?

Expert B: Ehm I suppose with the internet, I suppose visual really because ehm.. the textuals I suppose people still rely more on books, especially in the arts it's not like sciences or computers, I think a lot of journals are published online now, aren't they?

Interviewer: Yeah, yeah that's true yeah

Expert B: In sciences.

Expert B: Whereas in the arts it's still very much hard copy that people rely on. And all the big journals still they don't really have, they just, they are always a bit slower I think to move forward than people in the sciences or computers.

Interviewer: Yeah.

Expert B: So things like ehm even the journals that we have here in Ireland like CIRCA and the Irish Arts Review eh they don't have, I mean you don't, can't really rely on their websites for the texts. I mean it's beginning to.. CIRCA are putting all their back catalogue, or all their texts online now, but eh still people use, rely on the hardcopies for everything I think so for online resources I think maybe images still are more, are probably more used than eh..

Interviewer: More sought after...[incomprehensible] yeah, yeah.

Expert B: Because eh you can't get images basically in the books or anywhere else, or you can't get enough of them anyway.

Interviewer: Enough of them yeah, yeah. If you would like say like from ehm your personal experience, or also from the potential user group of this ehm eh collection ehm you know browsing facility you showed to me. What is the name of this system again? Like ehm...

Expert B: Art Search, the archive one.

Interviewer: Yeah the archive one, like where you can search by artists.

Expert B: Yeah, yeah.

Interviewer: Art Search. So ehm if you like compare the three different sort of styles for looking for an image like first I just browse around without a particular sort of, you know, I just wanna explore basically, ehm secondly I look for images which are in one category and share certain characteristics.

Expert B: Yeah.

Interviewer: Or in thirdly I look for one particular work of art?

Expert B: As way of searching?

Interviewer: So there is one defined target. If you compare these three how would you rate the mutual importance or the eh.. ?

Expert B: Ok I know what you mean yeah, well basically say if in the database there is an item somewhere that somebody wants to get to, the ways, there are basically four different ways to retrieve that item from the database and that's to, if you know nothing about the item and you just browse and go through the collections and find that you come across it and eh or eh that's one way of finding it. The other way is if you know something about it of if you are looking for something like the media or an item if you have a key word you can eh do a subject search so if it's like, if there's something in it like a baby, or a horse, or if eh you're looking by media you may find it. Then if you are looking for items eh by the artist or author you can ehm, so if you know the artist of an item you will come across it by their artist and then finally I suppose the most specific way eh of finding it is by keyword searching that if you have individual words that you know or associate it with the item, that you can..

Interviewer: For example in the title of the item, yeah, yeah.

Expert B: Now eh which people find most useful I suppose ehm when I use a database I would probably tend to go to ehm artist author first because I might have an artist that I would be interested in, if I was interested in Orpen, if I was to approach this database I would say, I want to find letter, I want to find material by Orpen, I'd either put Orpen in here, or I would do an artist author search.

Interviewer: Artist author, right yeah, yeah.

Expert B: If you are interested in the way the collection here is built, and what collections are in it and how its managed and what the layout of the collection is you go to collections cos then you can see the hierarchy of the collection, how it fits together and how items work and how/where items are actually positioned. And what's... an overview of what's in the collection.

Interviewer: Yeah.

Expert B: At the moment it's easy to see cos it's quite small. But as it gets bigger it could be an endless way of searching and then subjects is just another way of searching so I suppose the most useful.. for most people I would imagine Ð artist author, search by artist, the person who created the item is eh the most invaluable... [incomprehensible]

Interviewer: Yeah, yeah. And that would probably be with like a particular sort of, looking for a particular work of art and then going through artist author or basically only a group of [incomprehensible].

Expert B: Well I suppose the thing with archive items are they are not something that eh people generally know about it's not a famous painting, it's quite often a letter or something it will be people who want to eh see what the artist was writing about so you probably won't be looking for specific items you will just be looking for material by the artist instead. Like if I was writing a book on Orpen I'd want to consult as many of his letters and diaries as possible and eh I would have to just look and see what is here by Orpen and then browse through them and see what's relevant to me.

Interviewer: Yeah. Have you ever heard of or used ehm query by image sketch or query by image example techniques ehm as opposed to keyword?

Expert B: Say that.. keyword?

Interviewer: Like have you ever heard of or used query by image sketch or ehm query by image example techniques?

Expert B: What does that mean?

Interviewer: It means like you...say you look for a particular image and you only remember say the general layout of the composition

Expert B: Oh yeah yeah.

Interviewer: And then you give it a sketch as good as you possibly can and say like it looks something like that, and go off and find it.

Expert B: And give that to the computer?

Interviewer: Yeah, yeah.

Expert B: No, does that work?

Interviewer: Not really.

Expert B: Ok [laughs].

Interviewer: People have tried! That's true!

Expert B: People are trying to do it?

Interviewer: Yeah I think actually the Hermitage web page has a feature like that but I tried to.. but like I played around a bit like it didn't seem to work very well I have to say but..

Expert B: It would be very useful I'm sure because eh especially with a large collection, people..

Interviewer: Yeah you mean in a large collection where people couldn't ... ?

Expert B: They might just have in their head what something looks like.

Interviewer: It would be a prime case would be like you know I have forgotten everything about this thing except how it looks you know like I don't know who the artist is.

Expert B: But still I can only imagine If you had an idea of what it looked like, surely I would imagine that if you remember that there's a sheep in it or if there's a cow in it or if there's a boat in it, typing in boat, sea, island, those, the combination of words would be more successful way of finding something than trying to draw it because if you say like ehm I mean it's probably very similar I mean if you tried to draw a boat. I mean cos people's memory for composition I don't think... unless you have a very good visual memory eh the composition or something is harder to remember rather than the fact that there's a boat in it. I mean if you have a seascape with a boat in it eh I think to keyword searching you would just put in: sea, boat, island where to actually draw it you would have to draw it you would have to remember where the boat was in the picture, where the horizon was and things like that and whether people... I don't know if peoples' memories, visual memory is as good as that too.

Interviewer: It's not only, even if you have that vision then you have to be able to, you have to be quite a draughtsman as well to kinda to get this across like eh. It probably has its limitations like, or severe limitations probably I dunno like ehm..

Expert B: I dunno.

Interviewer: I mean, the other thing was about the like query by image example or.. What were you going to say there?

Expert B: I was just gonna say eh I suppose it's.. it would more like a fun kinda way or a toy more than anything, it's more of a gimmicky thing. I don't know how useful it would be to anyone trying to carry out any kind serious research that I mean to actually sit down well I don't know, I mean if you can draw it surely you have to have some pretty good idea of knowing what's in it and if you know what's in it you can just keyword search it but eh yeah.. I dunno how to, I mean it would take a lot of.. it would be relying a lot on the persons' memory and their ability to remember what a picture actually looks like.

Interviewer: Yeah. No, I'd say so too. I mean the other thing like query by image example could be interesting say if you have only a small of a thumbnail or something like that and you would like to find the exact same in bigger or something like and like.. you know for some reason like you know pretty much nothing about it or, then you could maybe go off and like retrieve the same thing in bigger, Or even just if you have an image, say a landscape, 17th century landscape blah blah blah and just go like I wanna to see similar ones, and feed that into a search engine and it comes back with.

Expert B: Yeah, yeah, yeah.

Interviewer: What do you reckon?

Expert B: I mean yeah eh, it would seem to work, I suppose it would have to. I suppose the difficulty would be proving how/ whether it is more effective than eh keyword searching then; I mean paintings are so like... I mean.. paintings can be I suppose for databases can be quite easily described in words just keywords and eh it is eh, I mean I suppose the whole basis of the database, like I mean it's there to provide eh people who are writing, people who are talking about images as a way of finding out about images that it's all text based, it's all the textual side of imagery really, databases, and research and writing interpretation. And ehm whereas I'd just wonder how I mean I'd find it difficult really to think of a scenario how that could be more effective than keyword searching, yeah. No? Do you know what I mean?

Interviewer: Yeah, yeah.

Expert B: Maybe there is ways maybe there will, but I can't imagine... How if someone came in... I know you can say maybe you could do this but surely there would be an easier way to do it by keyword searching. Is it something you are interested in?

Expert B: Eh not really no its no just I'm more of your school of thought that's why I look into ontologies and these like taxonomy and or sort of like you know you know kind of sort of key word mark up of images or in some way takes shape...likes maybe a bit clever or in a bit of a different way than just you know plain tags or something like that but like along these lines. I would imagine that these things: query by image example and query by sketch and so on have only like a minority application but there are some people who... like you know go like this is you know gonna be really... [incomprehensible]

Expert B: I was just thinking there, you probably could take a case study, like that's quite a complex image and ehm, I mean, do you know this painting?

Interviewer: No, no...

Expert B: Sorry this is a good example. Maybe I'm thinking about it because if you have, if you know paintings... I dunno, like if you saw that painting now.

Interviewer: Yeah, yeah.

Expert B: And if you look at it and now if eh which would you find it easier to describe, in words or in eh drawing?

Interviewer: It's weird; I mean the words that come to me are fairly generic, ehm like people sitting on a table ehm...

Expert B: So you could put "table" in.

Interviewer: Yeah I could put table in.

Expert B: Table, cos "furniture" would be something.

Interviewer: Yeah furniture, table ehm then like men sitting around the table, lady, portrait of a lady or something like that.

Expert B: I mean yeah you would be able to browse the keywords and look at like there's probably like conversation piece or something like that. There would certainly be table, furniture, chairs; eh interior would be an obvious one as well. Interior scene, so that would be one way of coming from it. Otherwise do you think you would be able to draw it?

Interviewer: Ehm probably not like I mean I could give it a... The problem is that you wouldn't get the position right for example you see I didn't remember there was four guys now I knew that there were two here you know.

Expert B: So what would you be expected, so you could draw the square?

Interviewer: I could draw the square here; I could maybe do a rough sketch of like the two people sitting here.

Expert B: Yeah but maybe this is the only painting in the world that is this format with a square in the middle as well, I dunno but...

Interviewer: The problem is however like if you draw it like say you don't remember the exact aspect ratio you know and this might be slightly out of place like the people are not exactly in the position where they are... and your off kinda... you know. I mean I played around with the Hermitage one a bit and I didn't find it very convincing or like I didn't have a feeling of you know, they also give you the opportunity you have like geometrical, you can insert geometrical objects like circles and stuff like that so if you, you know I tried to reconstruct some...

Expert B: You can put a square in here? Yeah you could put a square in here. You can put like a pinkish dot for a head but like it didn't really pick up on it very well because like because but like if you look at it, it's very dark here and its very light the face so there's a huge contrast but like semantically looking at things there's one head so I'd probably draw one block there you know because that's my recollection, my main perception of it.

Expert B: I think like you if you were able to provide a good enough drawing of it you would probably be able to find it fairly easy with keyword searching, if you knew that much information about the picture you could have done key word search.

Interviewer: So yeah, that's a very good point yeah

Expert B: Yeah, yeah.

Interviewer: Yeah that's a good point. So I think we are pretty much through I just have ehm. Oh yeah like the last three questions are more like ehm eh what would you consider the biggest shortcomings of current image retrieval images like ehm..? Say Google search or whatever you use?

Expert B: Well Google images? ehm...

Interviewer: Not necessarily whatever you use, just in general.

Expert B: Well Google images ehm basically I suppose the quality for a start with Google you never know, I mean it's pretty useless for eh getting quality images you can't, and that's understandable, mean the fact that their images are up there in print its probably infringing on copyright anyway but eh I mean they are great for... Google is great for doing research ehm especially for eh identifying locations in paintings and things like that. I mean quite often we'll have landscape paintings that will come in and someone will say: do you know where this is? And eh all you do... And eh there was a painting that came in recently and we had absolutely no idea who, where it was, we knew the artist, we didn't know when he painted it or eh where it was and eh by just doing a bit of work on the finding about on the artist we found out that he was in different locations at different times. He was in the north of France, he was in the south of England and eh he was in Ireland also and then sometimes by looking at the painting you can take a guess whether its north of France, south of England or Ireland. If you think its north of France you can do things like there was one we had, eh it was a lighthouse scene, it was a seascape, lighthouse and some boats and we just started googling lighthouses on the north of France and it brought up ehm thousands of images after going through a couple of them there was one or two that looked like it could be the location and then ehm so then eventually you come across the actual lighthouse eh when the lighthouse was there, eh you can go back to the information you know about the artist, when he was there so then you can tell people that the painting was painted by this artist when he was in the north of France during this period, and almost entirely based on what you've managed to find on Google and eh it's not using information that Google provides, it's using just the random bits of information that people, or pictures that people put up on Google or on their websites ehm. So that's great I mean you basically you have images, unlimited amount of images of pretty much everywhere in the world at your finger tips.

Interviewer: Yeah it's pretty cool.

Expert B: It's something that I don't know how anyone would have done even 15 years ago; you just wouldn't have that kind of resource to be have to be able to instantly look at some lighthouse on the north of France, that some German tourist probably took somewhere when they were travelling across Europe.

Interviewer: So you could identify like that very same lighthouse, I mean architecturally you could see that?

Expert B: Yeah when you are looking at them you begin to realize how many different designs there are for lighthouses and eh very few of them seemed to be the same. Well this is just one example now but this happens all the time and eh yeah there was one of the Osborne photographs, some of the Osborne photographs, eh we didn't know where they were ehm..I can show you one here.

Interviewer: That's brilliant, like the lighthouse story, I like that. Jesus I took up way more of your time than I said I would.

Expert B: No, no it's grand. Ok this was a picture ehm eh that Osborne had in his ehm photographic album. It doesn't say anywhere on the photograph where it is all we knew, all we could see it was a bridge and a town on the other side but we knew absolutely nothing of where it was so ehm we checked out Osborne's biography and eh during his early twenties he had gone to Antwerp to the school of art in Antwerp and while he was there he travelled to Germany. So we basically did a Google search of medieval bridges and bridges in Belgium and in Germany and went through a few on eh the internet and came up with ones like this, that kinda looked familiar, or that looked like the photograph. Eventually you come across one that looks exactly like it and we did and it turned out to be in this town called Koblenz do you know it?

Interviewer: Koblenz oh yeah, yeah.

Expert B: Ehm this is the bridge. I'll show you now. Do you know Koblenz?

Interviewer: No no, I have never been there but I like I know of the name. Wait a sec is Koblenz, is that in the south of Germany eh?

Expert B: Oh, do you know how I found it actually, it was actually on Google earth eventually, we found something that looked like bridges in Koblenz and then I checked it on Google earth and there was better images of it. But ehm..

Interviewer: You see the twin towers?

Expert B: These are the two towers of the church.

Interviewer: You see them there in the background of the image right? Brilliant.

Expert B: And you see them here as well? And eh so yeah it was just, now we know that this photo was taken he probably bought this photograph while he was travelling across Germany, he was in Koblenz. We never knew that he was in Koblenz but we knew that he had been in Germany but now we know that Walter Osborne was in Koblenz because of the photograph,

Interviewer: Brilliant.

Expert B: See before we would have waited for someone from Germany to come along and identify it or it probably would never have been identified and eh now it's in the eh database as being the bridge in Koblenz and Walter Osborne was obviously there and two years ago this would have never been known.

Interviewer: Yeah because even if someone from Germany comes along you wouldn't even know that this was Germany say.

Expert B: This photograph of the collection has been in the gallery for about 20 years and nobody has ever identified where it was and it was just 15 minutes sitting down with Google and we where able to find out where it was.

Interviewer: Brilliant. Deadly.

Expert B: I mean it can be very useful as well things like Google and just doing searches like that.

Interviewer: So my last question would be: if sort of like if the fairy came along and you had like, you could make a wish for the perfect image retrieval and art, fine art expert system like without you know technical limitations or anything, like what would it look like?

Expert B: Ehm well I think I mean I know we are getting faster, that the internet is getting faster all the time but it's still quite slow, you can't do everything immediately. It takes time and eh I suppose that really comes down to searching and eh I don't know if there is a better way of doing it but ehm I suppose the only thing is speed really if you can do things quicker that would be the ideal thing. Because the quality is there, now people are able to scan quite high resolutions but it's basically about making, about putting them online. Because I mean this is only available in the gallery at the moment which isn't really ideal; we want it to be available to everyone and on their home PCs. That's the first thing it doesn't really make, I suppose you have to think of these things from the point of view of the user not the people who are developing them so it's not really what I would want but what the user wants and the best thing is if they can go on in their PC at home and sit in their living room and browse these type of things and secondly that they can do it very quickly as well. Find out what they want eh instantly and that they don't have to spend ages eh typing in key words. I mean one of the ones I used most recently is, have you used the Irish Times Online National Archive?

Interviewer: Eh no.

Expert B: Have you seen it?

Interviewer: No, no I don't think so.

Expert B: They have an online, The Irish Times have their entire archive online now back to eh the 1850s and all the pages of the newspapers are scanned and they have used text recognition to eh to scan.

Interviewer: Ah OCR yeah, yeah.

Expert B: So you can type in say for example Orpen and it would bring up every article that has Orpen mentioned in it, and you can narrow it down by saying "William Orpen" but it just takes time you have to narrow down your searches, your years. Even at that William Orpen may have been mentioned in a period of thirty years and you have to browse through thirty years of eh references with everything with William Orpen in it and every reference with William Orpen in it so there I think speed is the real thing even though it's much faster than actually going to the National Library and having to go and ask for old newspapers and to have to scan through them physically. It's incredibly fast in comparison to that I mean there's no comparison at all. But I suppose people eh will always expect /want things to be faster. I suppose that's the main thing speed really.

Interviewer: But just to deliver those things faster to them like wouldn't let's say now in the case you made like with Orpen and you like say I had super super fast internet and it would be there in an instant and stuff like that I'd still have to sort of browse through [incomprehensible] it's not the internet it's the...

Expert B: It's not the internet, I think basically how the database is built and as you said how the keywords kinda function and people can put in the right words and immediately retrieve what they want and eh that's really how, I don't know how I mean I suppose that's based on subjects and keywords and things like that rather than I mean it doesn't matter how fast the internet is if people don't, if it's not searchable it's not searchable I mean it's like Google that you put in a word and the last thing you want is 2 million entries its kinda meaningless you are never going to see 2 million entries and eh but I mean there's nothing you can do about the internet because people are never going to eh tag things correctly or put the right keywords on things. But with databases I suppose that's what it's all about, making items that are as retrievable as quickly possible and that they, that it can be done very accurately as well.

Interviewer: Very good. Yeah let's leave it at that ehm...

Appendix C

Survey Screenshots

This appendix contains screenshots of the 14 pages of the Web-based survey discussed in Section 4.2. The screenshots are cropped to show only the main content of each page, with the exception of the first page, which contains advertisements placed by the *eSurveyspro*[1] platform which was used to conduct the survey.

[1] `http://www.esurveyspro.com` (last accessed 08/08/2011).

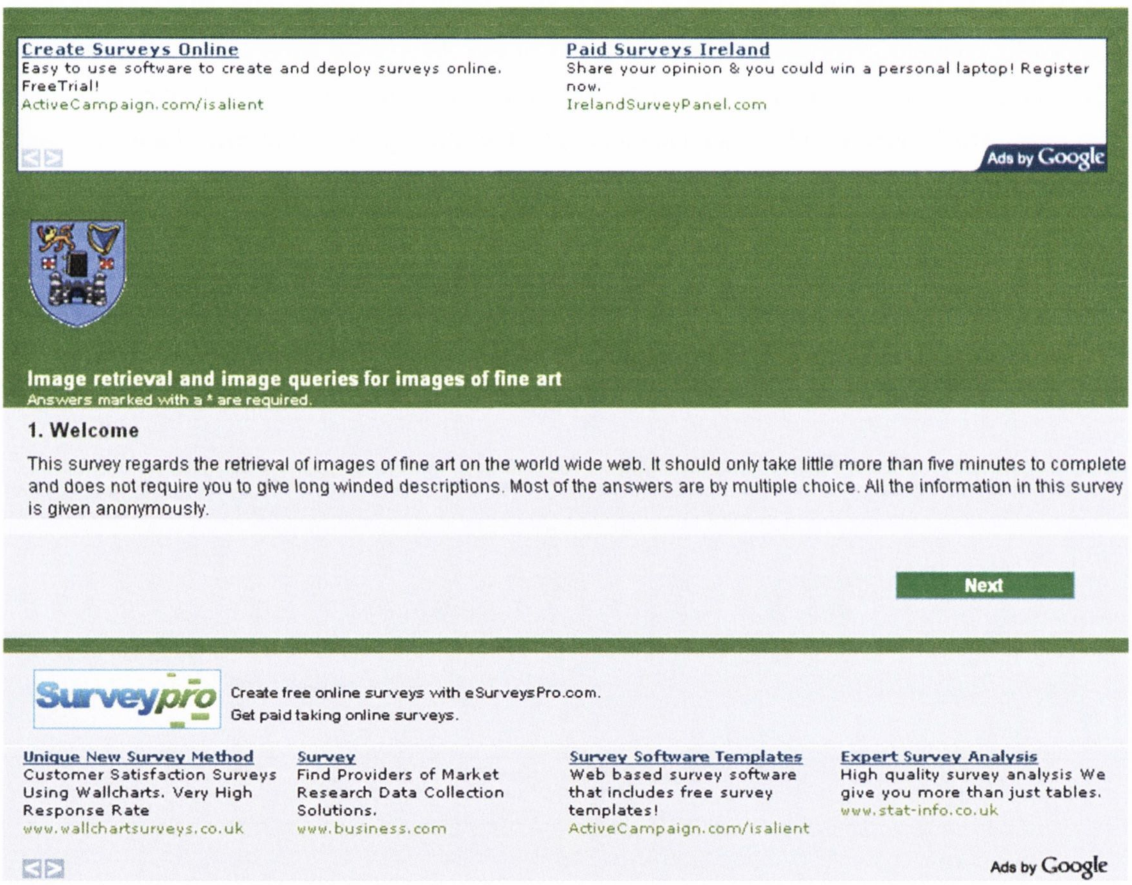

Figure C.1 – Page 1: Welcome

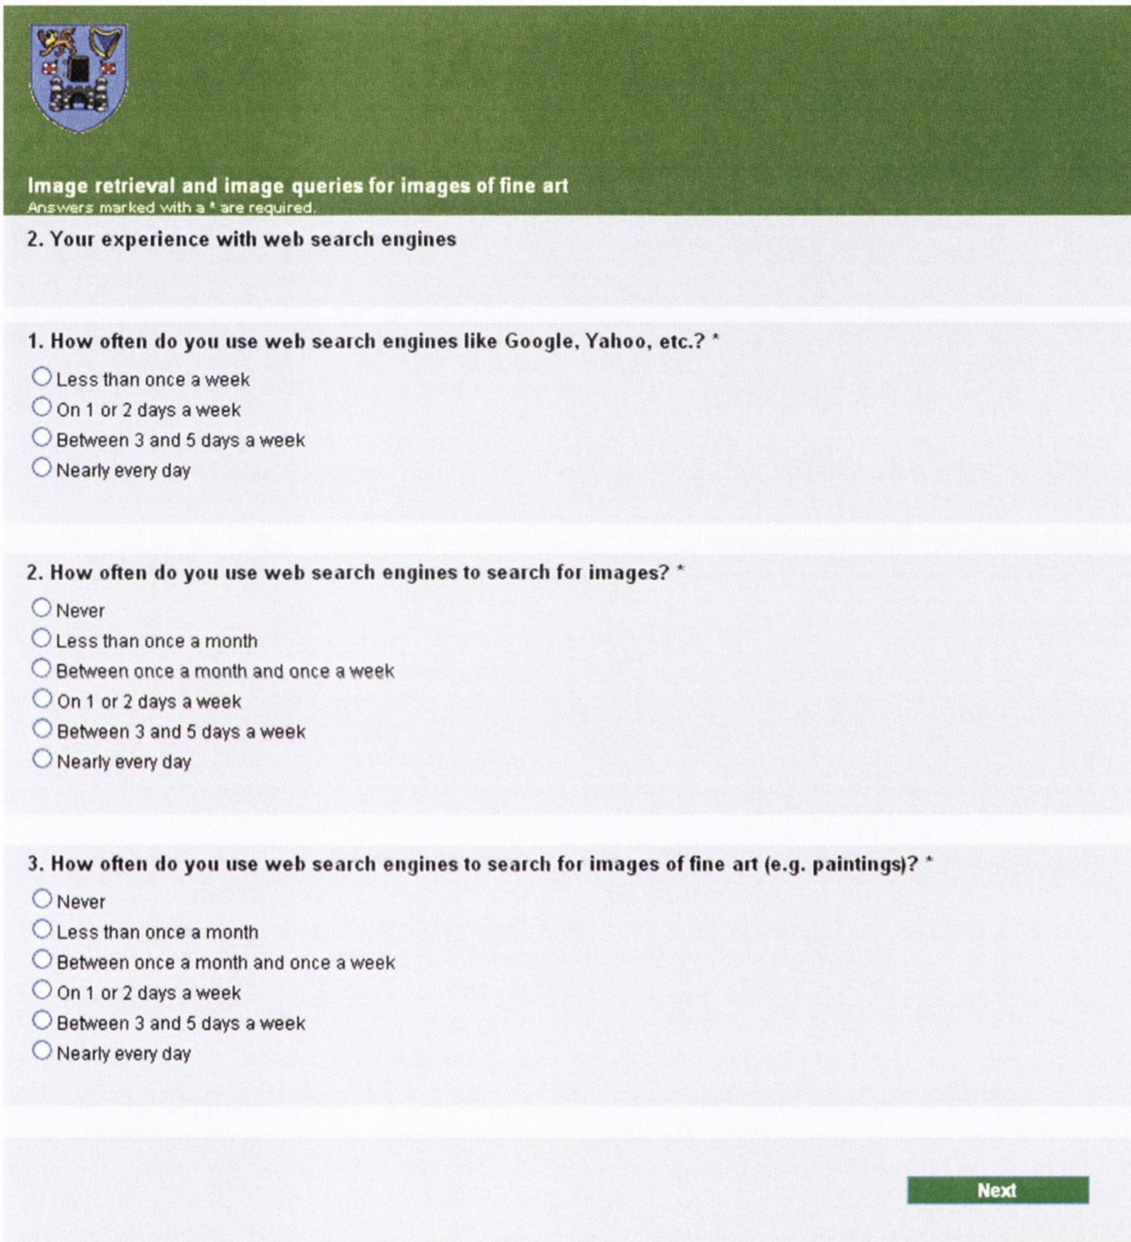

Image retrieval and image queries for images of fine art
Answers marked with a * are required.

2. Your experience with web search engines

1. How often do you use web search engines like Google, Yahoo, etc.? *

- ○ Less than once a week
- ○ On 1 or 2 days a week
- ○ Between 3 and 5 days a week
- ○ Nearly every day

2. How often do you use web search engines to search for images? *

- ○ Never
- ○ Less than once a month
- ○ Between once a month and once a week
- ○ On 1 or 2 days a week
- ○ Between 3 and 5 days a week
- ○ Nearly every day

3. How often do you use web search engines to search for images of fine art (e.g. paintings)? *

- ○ Never
- ○ Less than once a month
- ○ Between once a month and once a week
- ○ On 1 or 2 days a week
- ○ Between 3 and 5 days a week
- ○ Nearly every day

Next

Figure C.2 – Page 2: Your experience with web search engines

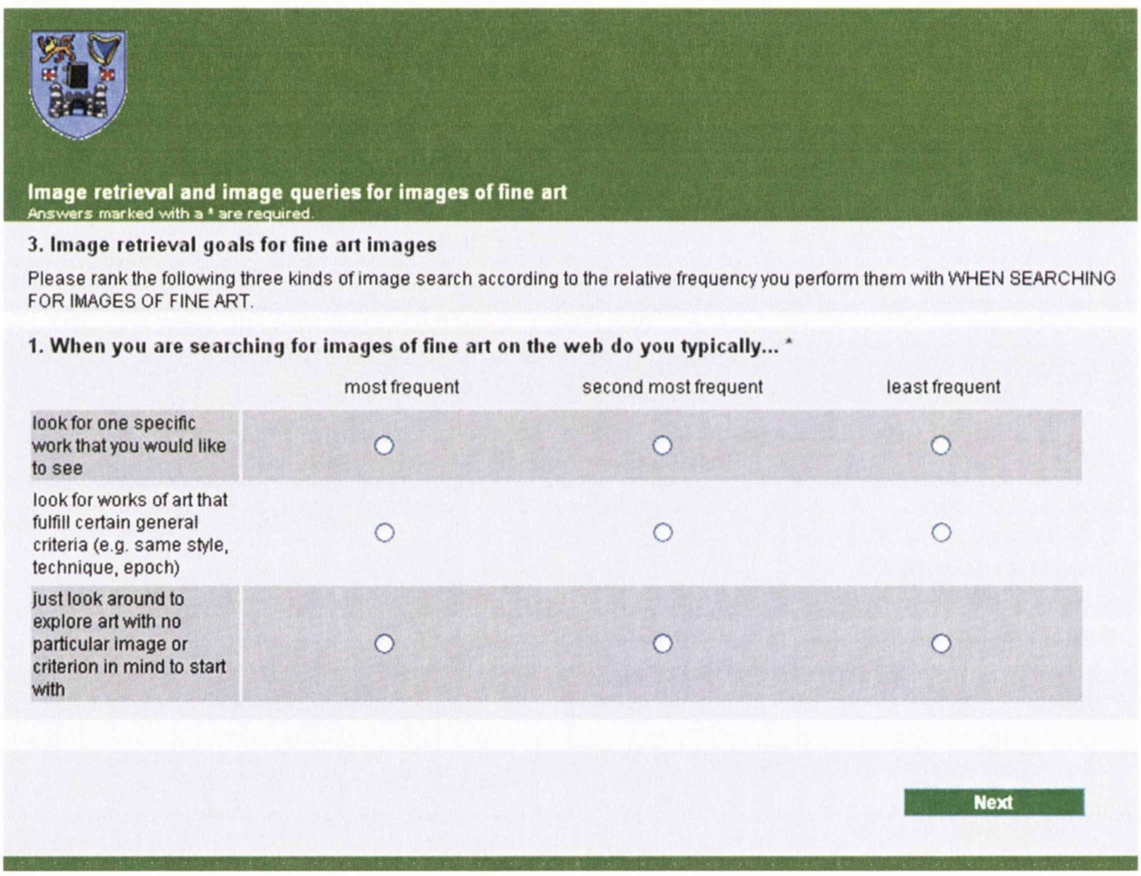

Figure C.3 – Page 3: Image retrieval goals for fine art images

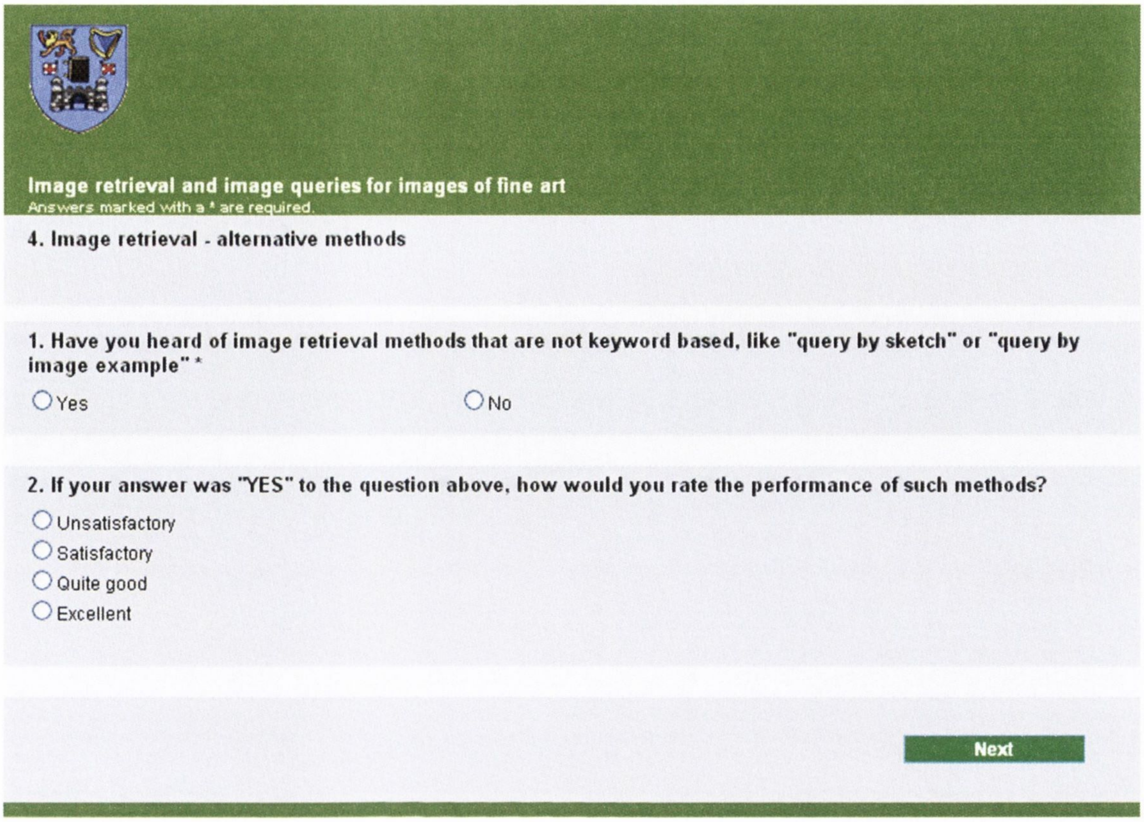

Image retrieval and image queries for images of fine art

Answers marked with a * are required.

4. Image retrieval - alternative methods

1. Have you heard of image retrieval methods that are not keyword based, like "query by sketch" or "query by image example" *

○ Yes ○ No

2. If your answer was "YES" to the question above, how would you rate the performance of such methods?

○ Unsatisfactory
○ Satisfactory
○ Quite good
○ Excellent

Next

Figure C.4 – Page 4: Image retrieval - alternative methods

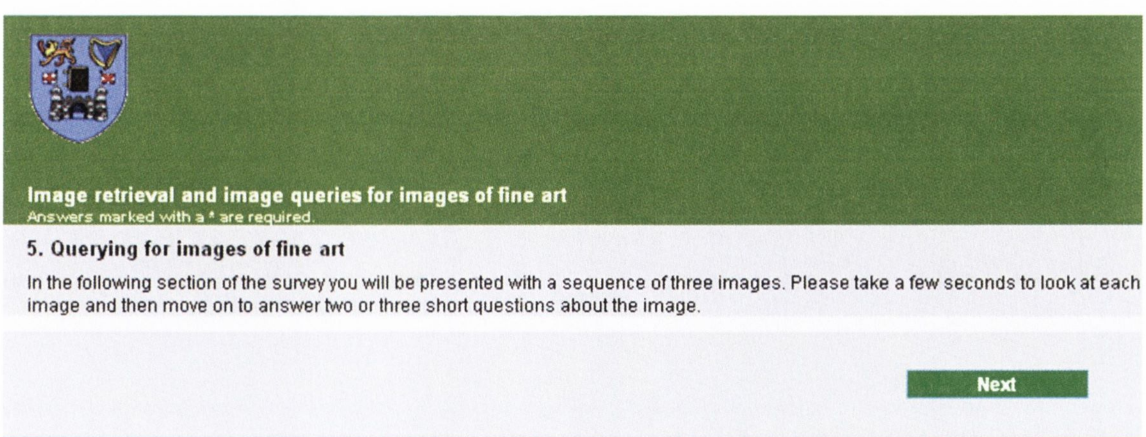

Image retrieval and image queries for images of fine art

Answers marked with a * are required.

5. Querying for images of fine art

In the following section of the survey you will be presented with a sequence of three images. Please take a few seconds to look at each image and then move on to answer two or three short questions about the image.

Next

Figure C.5 – Page 5: Querying for images of fine art

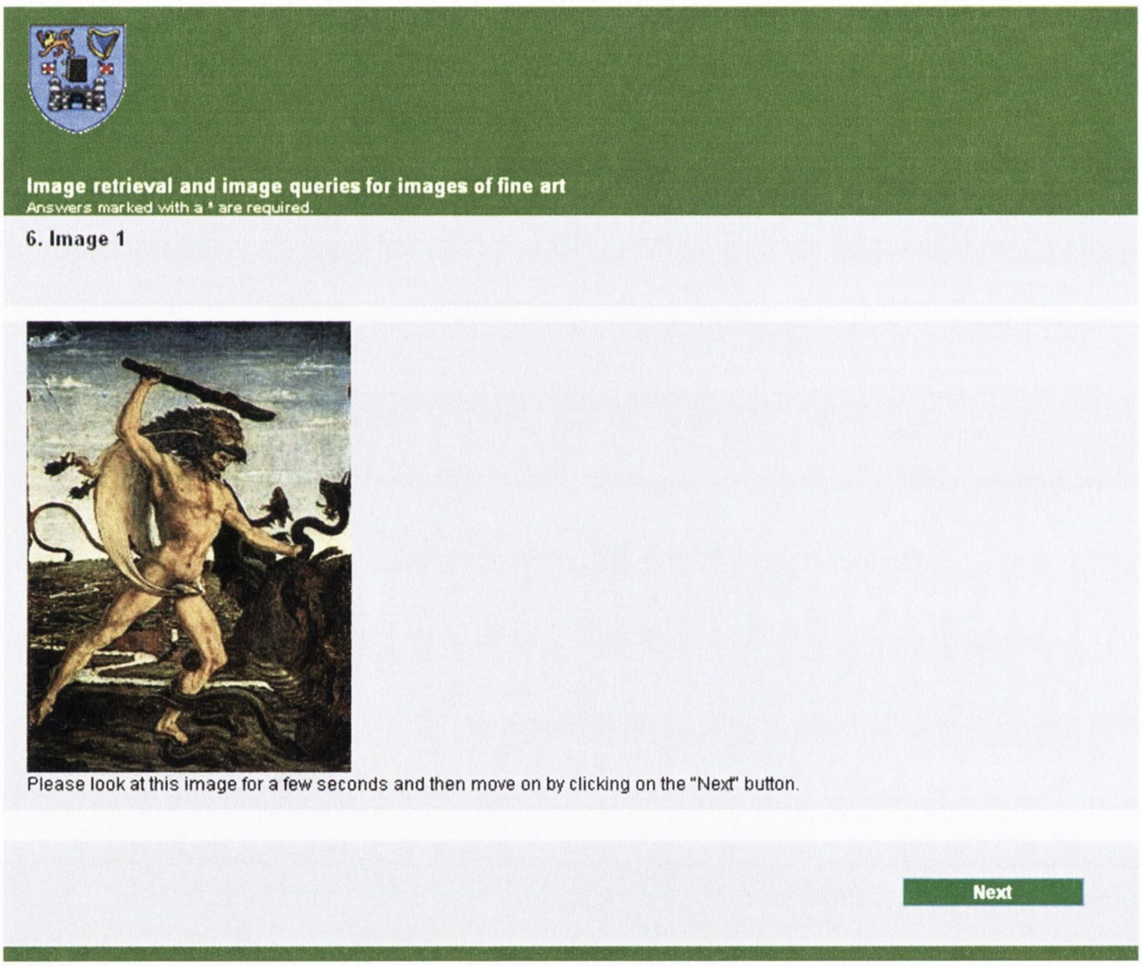

Figure C.6 – Page 6: Image 1

Image retrieval and image queries for images of fine art

Answers marked with a * are required.

7. Image 1 - Questions

1. What keywords would you use if you were looking for this image online? *

2. Have you seen a similar image before? *

○ Yes ○ No

3. If your answer was "YES" to the question above, indicate in what way the image was similar? (multiple answers possible)

- ☐ It was similar in composition
- ☐ It showed the same characters
- ☐ It had similar colours
- ☐ It was art historically related to this image
- ☐ It depicted a similar scene
- ☐ It illustrated the same story
- ☐ Other (Please Specify)

Next

Figure C.7 – Page 7: Image 1 - Questions

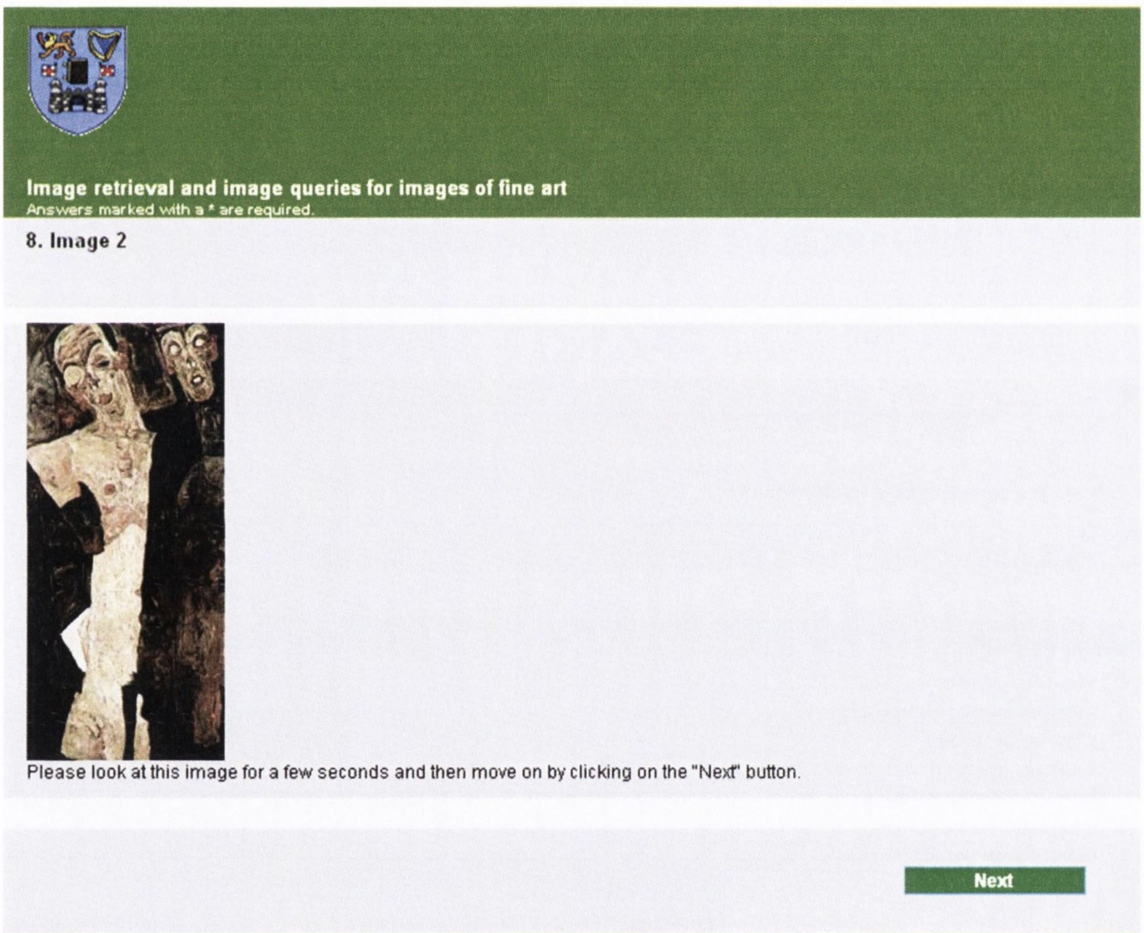

Figure C.8 – Page 8: Image 2

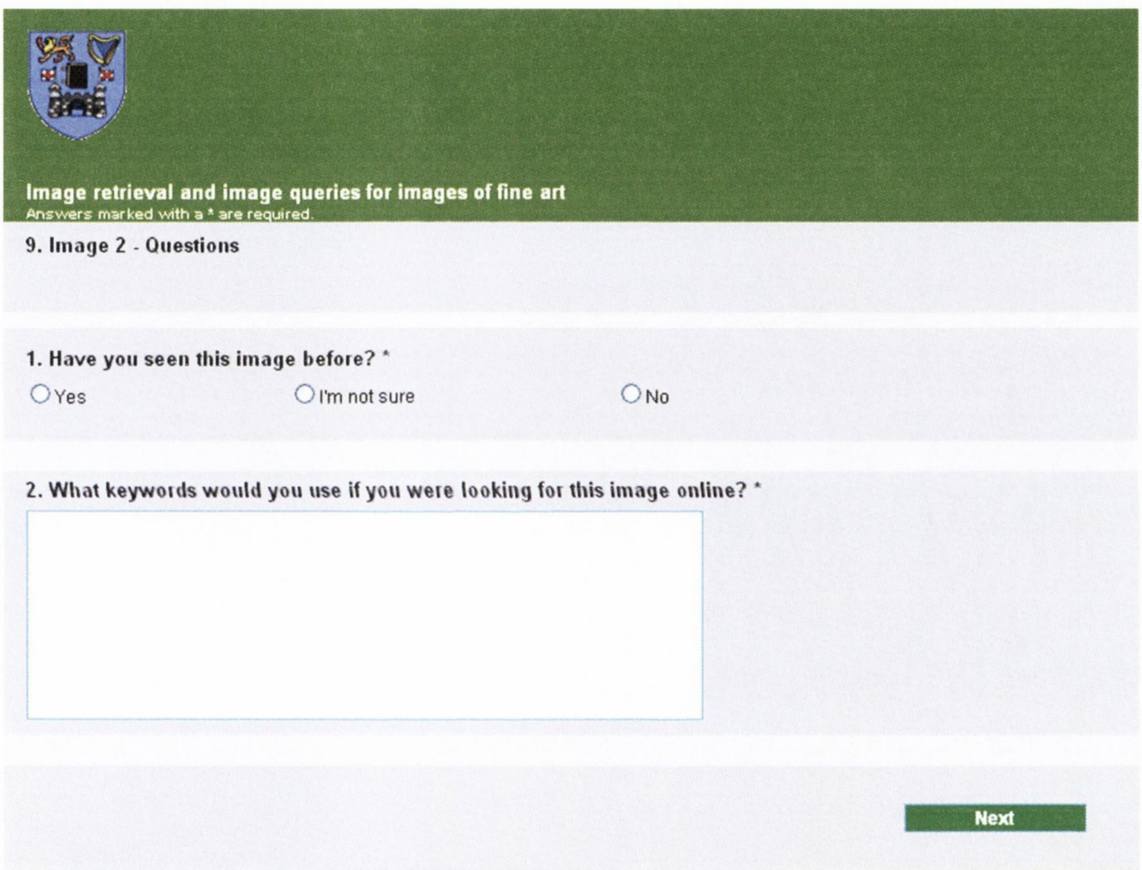

Image retrieval and image queries for images of fine art
Answers marked with a * are required.

9. Image 2 - Questions

1. Have you seen this image before? *

○ Yes ○ I'm not sure ○ No

2. What keywords would you use if you were looking for this image online? *

Next

Figure C.9 – Page 9: Image 2 - Questions

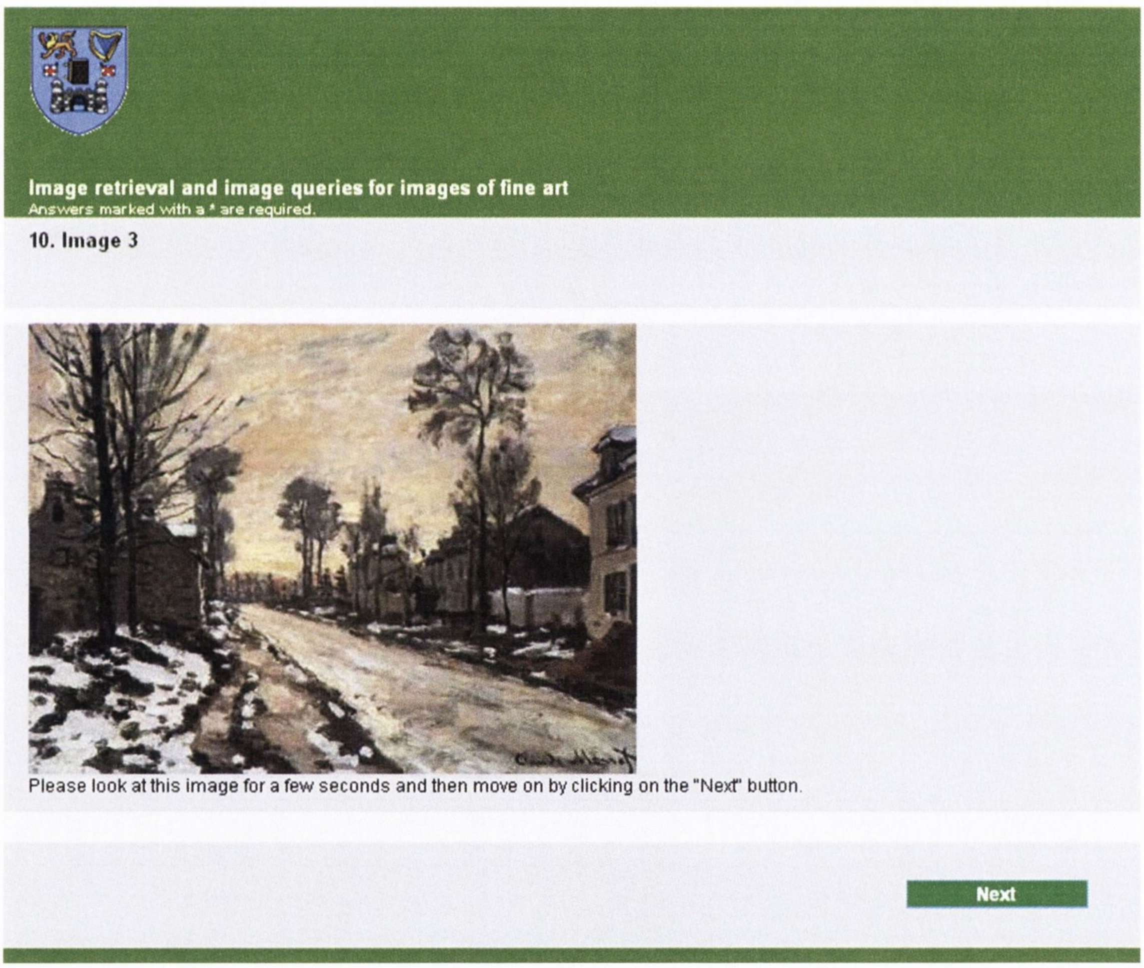

Figure C.10 – Page 10: Image 3

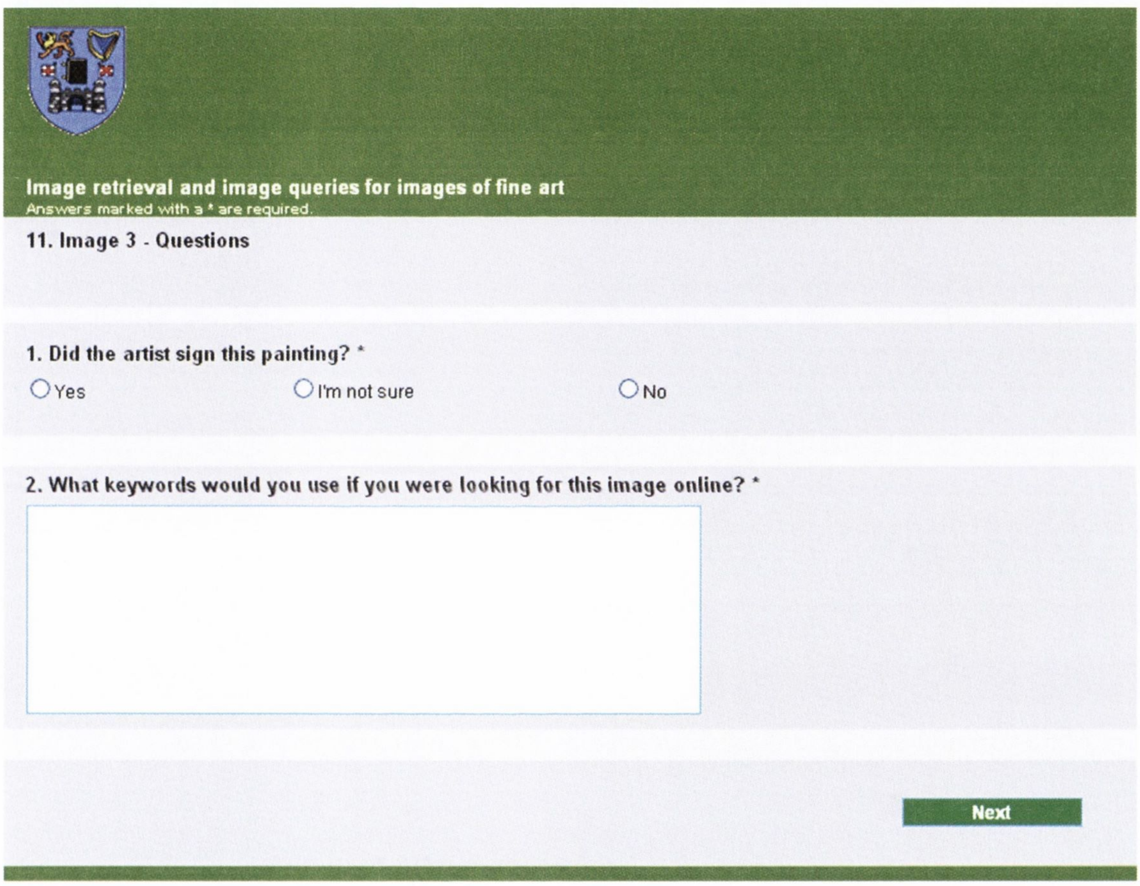

Figure C.11 – Page 11: Image 3 - Questions

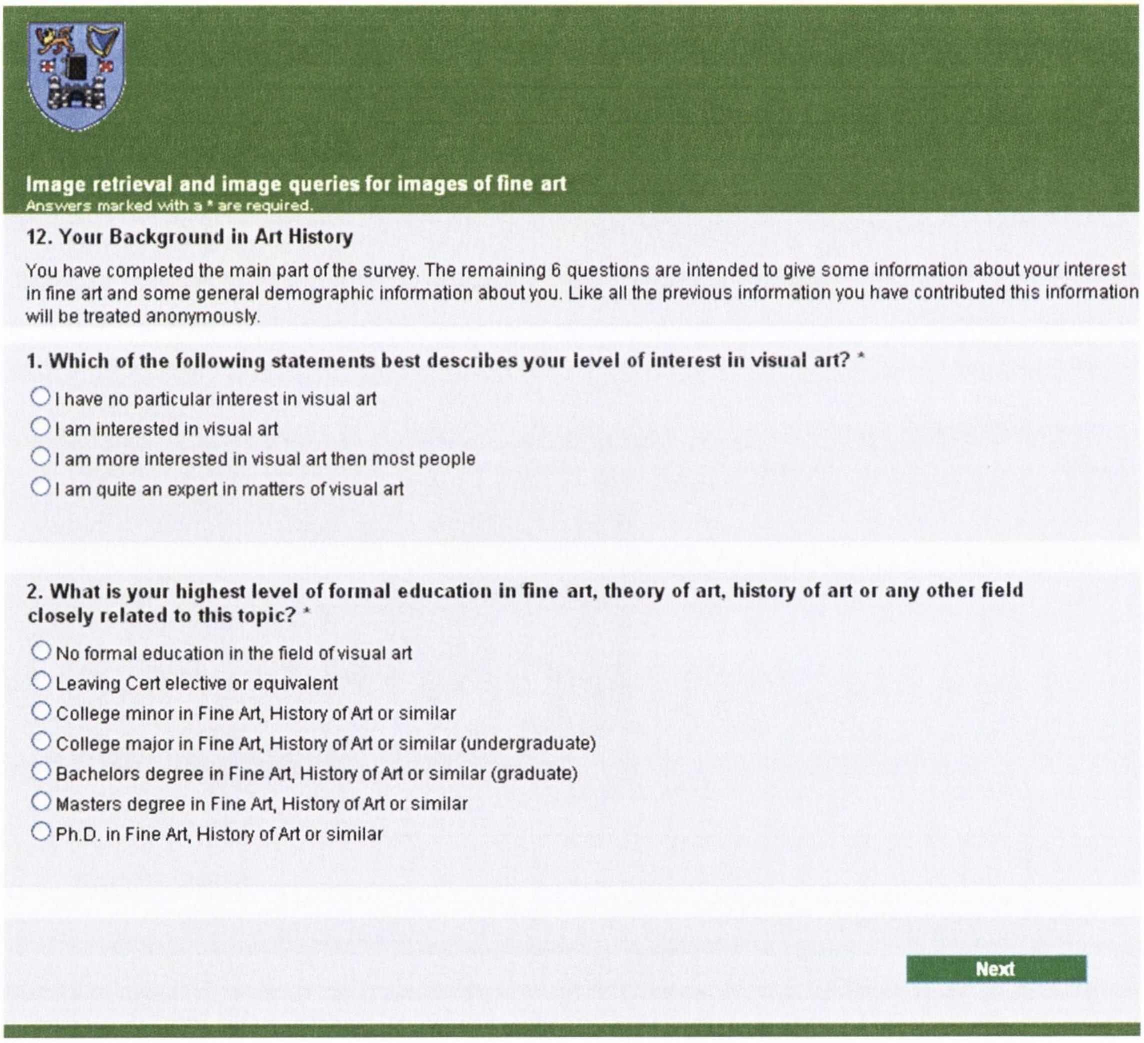

Image retrieval and image queries for images of fine art
Answers marked with a * are required.

12. Your Background in Art History

You have completed the main part of the survey. The remaining 6 questions are intended to give some information about your interest in fine art and some general demographic information about you. Like all the previous information you have contributed this information will be treated anonymously.

1. Which of the following statements best describes your level of interest in visual art? *

- I have no particular interest in visual art
- I am interested in visual art
- I am more interested in visual art then most people
- I am quite an expert in matters of visual art

2. What is your highest level of formal education in fine art, theory of art, history of art or any other field closely related to this topic? *

- No formal education in the field of visual art
- Leaving Cert elective or equivalent
- College minor in Fine Art, History of Art or similar
- College major in Fine Art, History of Art or similar (undergraduate)
- Bachelors degree in Fine Art, History of Art or similar (graduate)
- Masters degree in Fine Art, History of Art or similar
- Ph.D. in Fine Art, History of Art or similar

Next

Figure C.12 – Page 12: Your Background in Art History

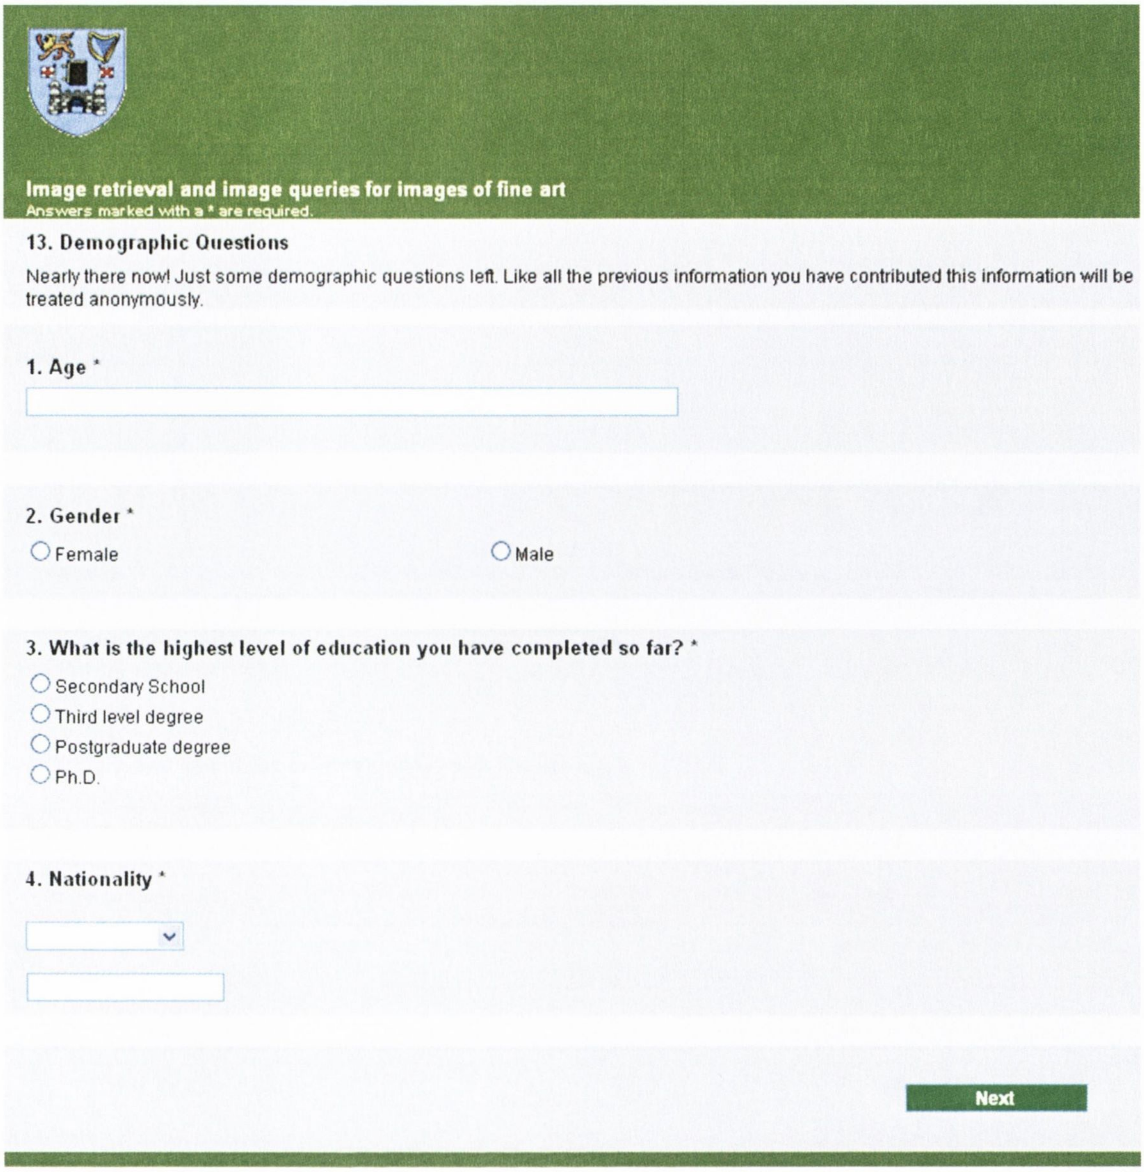

Image retrieval and image queries for images of fine art

Answers marked with a * are required.

13. Demographic Questions

Nearly there now! Just some demographic questions left. Like all the previous information you have contributed this information will be treated anonymously.

1. Age *

2. Gender *

○ Female ○ Male

3. What is the highest level of education you have completed so far? *

○ Secondary School
○ Third level degree
○ Postgraduate degree
○ Ph.D.

4. Nationality *

Next

Figure C.13 – Page 13: Demographic Questions

Image retrieval and image queries for images of fine art

Answers marked with a * are required.

14. Thank you very much for completing this survey!

Please click "Finished" to submit your information.

Finished

Figure C.14 – Page 14: Thank you very much for completing this survey!

Appendix D

Personal Communication

This appendix contains a partial record of an email exchange between the author of this thesis and Bernhard Schiemann, who was, at the time, responsible for maintaining the OWL implementation of the CIDOC Conceptual Reference Model (CRM) by the University of Erlangen-Nürnberg.[1] The exchange mostly covered modelling issues in the Erlangen CRM or the CRM in general. The emails are listed in chronological order. In 2011, a GOOGLE group was set up to facilitate discussions such as the ones covered in this appendix.[2]

```
From: "Daniel Isemann" <isemandi@cs.tcd.ie>
Subject: OWL version of CRM
Date: Fri, October 24, 2008 5:16 pm
To: goerz@informatik.uni-erlangen.de,schiemann@informatik.uni-erlangen.de
```

```
Dear Prof Görz, dear Dipl-Ing Schiemann,

my name is Daniel Isemann. I am a PhD student at Trinity College Dublin
and I am currently trying to build up a small OWL DL ontology about
paintings and drawings. I would like to use the CIDOC CRM as top-level
ontology to embed my domain ontology in. In that context I came across
your OWL DL versions of the CRM at

http://www8.informatik.uni-erlangen.de/IMMD8/Services/cidoc_crm/
cidoc_crm_4.2.4_owl_dl.owl
http://www8.informatik.uni-erlangen.de/IMMD8/Services/cidoc-crm
/2008/09/01/cidoc-crm-4.2.4.owl#
http://www8.informatik.uni-erlangen.de/IMMD8/Services/cidoc-crm
/2008/09/01/cidoc-crm-4.2.4.LessExpressive.owl

and also a version which was called cidoc-crm-4.2.4.weak.owl which I
downloaded earlier this week.

So far I have been using Aldo Gangemi's OWL rendering of the CRM which
can be found at:
```

[1] For the current version of the Erlangen-Nürnberg CRM implementation see http://erlangen-crm.org/current-version (last accessed 19/03/2012).

[2] The GOOGLE group can be found at http://groups.google.com/group/erlangen-crm/topics?hl=en-GB&gvc=2 (last accessed 19/03/2012).

http://cidoc.ics.forth.gr/OWL/cidoc_v4.2.owl

I have a few questions concerning your OWL versions. Even though all are labeled as OWL DL only the last two without cardinality restrictions are classified as DL by the Protege "Determine OWL Sublanguage" tool. The others come out as OWL Full. Are they actually OWL Full or is there a problem with the Protege tool?

Secondly I was wondering whether you could point me to any "best practice" in terms of modelling information about works of art under CRM. If I understand the CRM correctly then in order to express when a painting was created one is supposed to introduce the following individuals:

- the painting
- the production event of the painting
- the time span of the production event
- the time primitive of the time span
- and then finally following the "container" representation in OWL you suggested a data type instance that corresponds/expresses the time primitive

This seems to lead to a great proliferation of individuals if you deal with many works of art. Was that the approach you took in building the domain ontology about sculpture in the GNM?

Any feedback you could give me on these topics would be very appreciated.

With best regards
Daniel

From: "Bernhard Schiemann" <Bernhard.Schiemann@informatik.uni-erlangen.de
>
Subject: Re: OWL version of CRM
Date: Fri, October 31, 2008 11:22 am
To: "Daniel Isemann" <isemandi@cs.tcd.ie>
Cc: "Guenther Goerz" <Guenther.Goerz@informatik.uni-erlangen.de>," Guenther Goerz" <guenther.goerz@gmail.com>

Dear Daniel
> Dear Prof Görz, dear Dipl-Ing Schiemann,
My name is Bernhard.

> http://www8.informatik.uni-erlangen.de/IMMD8/Services/cidoc-crm
/2008/09/01/cidoc-crm-4.2.4.LessExpressive.owl
> and also a version which was called cidoc-crm-4.2.4.weak.owl which I
> downloaded earlier this week.
The weak version is old. These two versions have identical content, but
it seems that the term "weak" could be misunderstood by non-logicans.
>
> So far I have been using Aldo Gangemi's OWL rendering of the CRM which
> can be found at:
> http://cidoc.ics.forth.gr/OWL/cidoc_v4.2.owl
This OWL is based on a older text specification (June 2007).
>
> I have a few questions concerning your OWL versions. Even though all
are
> labeled as OWL DL only the last two without cardinality restrictions
are
> classified as DL by the Protege "Determine OWL Sublanguage" tool. The
> others come out as OWL Full. Are they actually OWL Full or is there a
> problem with the Protege tool?
Just a Protege error. Which version do you use? We prefer the 3.4Beta
Build >500. To validate the OWL files you can use:
http://www.mygrid.org.uk/OWL/Validator
>
> Secondly I was wondering whether you could point me to any "best
> practice" in terms of modelling information about works of art under
> CRM. If I understand the CRM correctly then in order to express when a
> painting was created one is supposed to introduce the following
> individuals:
>
> - the painting
> - the production event of the painting
> - the time span of the production event
> - the time primitive of the time span
> - and then finally following the "container" representation in OWL you
> suggested a data type instance that corresponds/expresses the time
> primitive

This is correct based on the definitions given by the text specification
of the CIDOC CRM. Could be easier, but isn't yet (one of the open
discussions in the CRM SIG). From a computer scientist (my) perspective
this model is not a real problem, because you just have the datetime
(timeprimitive datatype, W3C XSD) data once. Other instances of e.g.
other time_spans then are allowed to point to this data. So, they can be
seen as pointers as e.g. LISP, JAVA,...
>
> This seems to lead to a great proliferation of individuals if you deal
> with many works of art.
A database stores eg. ten thousand times the date 2008-10-31. We store
it only one time and then create ten thousand pointers to it. Note that
the XSD type "datetime" is just one proposal to encode dates. The CRM
itself is compatible to ALL formats encoded.

>Was that the approach you took in building the
> domain ontology about sculpture in the GNM?
Yes.

kind regards,
Bernhard

From: "Daniel Isemann" <isemandi@cs.tcd.ie>
Subject: Re: OWL version of CRM
Date: Mon, January 26, 2009 8:01 pm
To: Bernhard.Schiemann@informatik.uni-erlangen.de

Dear Bernhard,

a happy new year to you and those close to you. Following on from last years email, I have a few more questions about the Erlangen implementation of the CRM.

1) First of all I was delighted to see that you added the begins-at and ends-at relations to the CRM_5_TQ version. I was already using similar self-defined relations to model time intervals and will now change to your suggested ones. Why is the range of both properties the union of Time-Span and Time_Primitive and not just Time_Primitive? (Not that it makes a big difference for my purposes, but I was just wondering what the rationale behind it was?)

2) Based on the CRM document a number of properties seem to have qualities which could be modeled in OWL-DL. E.g. P10_falls_within and P89_falls_within really look transitive to me. And of course the transitivity would allow nice inferences/save a lot of "redundancy" in specifying relations. Is there any doubt that these properties are intended as transitive or is there another reason why you did not model them as transitive?

3) Similarily I tend to interpret P9I_forms_part_of and P88I_forms_part_of as subproperties of P10 and P89 respectively. This seems less justified, since the CRM specification does make subproperty relations explicit in other cases. Was it discussed in the CRM SIG to make these properties subproperties at one point?

4) You mention in [1] that properties of the form PX.1 in the CIDOC CRM specification are not subproperties of PX (because effectively PX.1 is a tertiary relation between elements of the sets domain(PX), E55_Type and range(PX)). But surely for a fixed E55_Type argument we get a subproperty of PX? So partially, i.e. for special cases that might be important for a certain application, PX.1 could be modelled by a number of subproperties of PX: PX.1(as E55_Type type0), PX.1(as E55_Type type1), ... etc.?

[1]
http://www8.informatik.uni-erlangen.de/IMMD8/Services/cidoc-crm/docu/crm_owl_cidoc2008.pdf

Best regards
Daniel

From: "Daniel Isemann" <isemandi@cs.tcd.ie>
Subject: Object properties in the Erlangen CRM
Date: Tue, July 14, 2009 11:37 am
To: Bernhard.Schiemann@informatik.uni-erlangen.de

Dear Bernhard,

I am interested in the OWL implementation of the CIDOC CRM that you helped to develop. I am trying to use it for "ontological mark up" of a database of paintings and artists. As you might or might not remember we have already been in touch by email earlier this year.

I recently revisited the current version of the Erlangen CRM with respect to the characteristics of object properties (symmetry etc.). I found some of the assigned property characterisations counterintuitive (see the third group below). Maybe you could help in clarifying this. I am using

http://www8.informatik.uni-erlangen.de/IMMD8/Services/cidoc-crm/erlangen-crm_090330_5_0_1.owl

in which the following characteristics for CRM object properties are defined:

Functional: P22, P99
Inverse functional: P22I, P99I, P3, P57, P81, P82
Symmetric: P114, P121, P122, P132, P133, P69
Transitive: P25, P26, P27, P89, P89I

In some cases the defined characteristics are very intuitive or follow directly from stipulations in the CRM document (e.g. P69). These cases are in my opinion:

P114, P121, P122, P132, P133, P69, P89, P89I

In other cases the property characteristics restrict the interpretation in ways that are (to me) less obvious, albeit understandable:

P22 transfered_title_to (functional): ok, if joint ownership is always modelled as ownership by a group
P99 dissolved (functional): every Dissolution can be the dissolution of only one group. So there can never be a single dissolution event for several groups at the same time (e.g. two groups of vulcanologists meeting in a crater at eruption time? :-) i'm sure one could envisage more plausible scenarios as well)
P22I (inv func): see P22 above.
P99I: see above.
P3 has_note (inv func): Why can items not share a note, e.g. "Part of a

group of five marble busts by..."

In a third group of cases I find it hard to see the rationale behind the specified property characterisations:

P25 moved (transitive): domain and range of this property are necessarily disjoint. So the transitivity never fires (granted, it doesn't do any harm either, but it might be confusing to some).
P57 has number of parts (inv func): Two chess sets that are made up of 33 pieces each are not necessarily identical. Was the intention maybe to make P57 functional? Even that might be a bit strong but inverse functional doesn't seem to make sense.
P26 moved_to (transitive): an issue similar to P25 except that this time domain and range are not technically required to be disjoint. But is there really an E53.Place that is at the same time an E9.Move event? It doesn't seem to make a lot of sense.
P27 moved_from (transitive): same same as P26
P81 ongoing_throughout (inv func): If two time spans are ongoing throughout the same time primitive (an interval say) they shouldn't have to be identical. All that might be known about the life spans of two ancient dignitaries is that they were both alive throughout a given year (say 55 BC). But that doesn't mean they were born and died at the same time. Was the intention maybe to make P81 functional to ensure that there is only one inner boundary for time spans?
P82 at_some_time_within (inv func): argument for P81 applies accordingly.

Sorry for the lengthy email. I really appreciate the work you and your colleagues are doing on the OWL-CRM. I'm just trying to understand the modelling decisions you made. It would be nice if you could say sth especially about the third group of property characteristics above.

If you are not in charge for these questions please let me know who the right person to talk to is or if there is a better place to post questions and feedback like this (mailing list etc.).

Thanks for your help in advance
Daniel Isemann

Trinity College Dublin

From: "Bernhard Schiemann" <Bernhard.Schiemann@informatik.uni-erlangen.de
>
Subject: Re: Object properties in the Erlangen CRM
Date: Tue, July 14, 2009 4:57 pm
To: "Daniel Isemann" <isemandi@cs.tcd.ie>

Dear Daniel,
> I am interested in the OWL implementation of the CIDOC CRM that you
> helped to develop. I am trying to use it for "ontological mark up" of a
> database of paintings and artists. As you might or might not remember
we
> have already been in touch by email earlier this year.
Yep, I remember. Fine project!
> I recently revisited the current version of the Erlangen CRM with
> respect to the characteristics of object properties (symmetry etc.). I
> found some of the assigned property characterisations counterintuitive
> (see the third group below). Maybe you could help in clarifying this. I
> am using
> http://www8.informatik.uni-erlangen.de/IMMD8/Services/cidoc-crm/
erlangen-crm_090330_5_0_1.owl
Yes, thank you for your feedback! I will try to answer your questions
carefully.
>
> In some cases the defined characteristics are very intuitive or follow
> directly from stipulations in the CRM document (e.g. P69). These cases
> are in my opinion:
> P114, P121, P122, P132, P133, P69, P89, P89I
ok.
> In other cases the property characteristics restrict the interpretation
> in ways that are (to me) less obvious, albeit understandable:
>
> P22 transfered_title_to (functional): ok, if joint ownership is always
> modelled as ownership by a group
Yes, that's the reason why E39 is used instead of E21. You can interpret
the corresponding E8 as an "atomic transaction".
> P99 dissolved (functional): every Dissolution can be the dissolution of
> only one group. So there can never be a single dissolution event for
> several groups at the same time (e.g. two groups of vulcanologists
> meeting in a crater at eruption time? :-) i'm sure one could envisage
> more plausible scenarios as well)
;-) super example, but the death of all members doesn't imply that the
groups are dissolved. Hm, the examples in the CIDOC CRM aren't very
helpful, but I can think of German "history" example: The SED dissolved
any other (more than one) political party by one single event
(proclamation). So, I think you are right that it is not a functional
property. I will correct this in the 501 version now.
> P22I (inv func): see P22 above.
> P99I: see above.

Yes.
> P3 has_note (inv func): Why can items not share a note, e.g. "Part of a
> group of five marble busts by..."
This is set because of the scope note of p3: "but [b]a[/b] note is attached to [b]a[/b] specific item."

> In a third group of cases I find it hard to see the rationale behind the
> specified property characterisations:
>
> P25 moved (transitive): domain and range of this property are
> necessarily disjoint. So the transitivity never fires (granted, it
> doesn't do any harm either, but it might be confusing to some).
The "confusing" argument is convincing me! I delete this one.
> P57 has number of parts (inv func): Two chess sets that are made up of
> 33 pieces each are not necessarily identical. Was the intention maybe to
> make P57 functional? Even that might be a bit strong but inverse
> functional doesn't seem to make sense.
No it doesn't make sense, I think this is a failure from a previous version. I delete this one too!
> P26 moved_to (transitive): an issue similar to P25 except that this time
> domain and range are not technically required to be disjoint. But is
> there really an E53.Place that is at the same time an E9.Move event? It
> doesn't seem to make a lot of sense.
> P27 moved_from (transitive): same same as P26
See P25. I have not found any internal documentation about these transitivities, but I remember a discussion, that a combined relation is transitive. I delete these two ones too.
> P81 ongoing_throughout (inv func): If two time spans are ongoing
> throughout the same time primitive (an interval say) they shouldn't have
> to be identical. All that might be known about the life spans of two
> ancient dignitaries is that they were both alive throughout a given year
> (say 55 BC). But that doesn't mean they were born and died at the same
> time. Was the intention maybe to make P81 functional to ensure that
> there is only one inner boundary for time spans?
Yes, but It is possible to restrict that on the time span level. We have indeed an ongoing discussion about the time spans, because they are defined to be fuzzy (unclear beginning & end) but make use of Allen's time relations, which for my point of view doesn't work well. Functional and inverse functional properties should be used, if one have to guarantee the "one-to-one" relationship. So I delete this one also (doesn't hurt), and write me a note that i should think of a restriction on Time-Span/E2's.
> P82 at_some_time_within (inv func): argument for P81 applies accordingly.
Thank you again!

```
>
> Sorry for the lengthy email. I really appreciate the work you and your
> colleagues are doing on the OWL-CRM. I'm just trying to understand the
> modelling decisions you made. It would be nice if you could say sth
> especially about the third group of property characteristics above.
As you can see, we have to thank you for your detailed feedback. I
updated the website already.
>
> If you are not in charge for these questions please let me know who the
> right person to talk to is or if there is a better place to post
> questions and feedback like this (mailing list etc.).
;-)


Thank you!
Ben

>
> Thanks for your help in advance
> Daniel Isemann
>
> Trinity College Dublin
>
```

From: "Daniel Isemann" <isemandi@cs.tcd.ie>
Subject: Re: Object properties in the Erlangen CRM
Date: Thu, July 16, 2009 6:31 pm
To: Bernhard.Schiemann@informatik.uni-erlangen.de

Hi Bernhard,

thanks for your very fast response. I didn't expect a new version the same day. That was great!

I like your SED example. It's spot on and (just slightly :-) less contrived then my vulcanologists...

>> P3 has_note (inv func): Why can items not share a note, e.g. "Part of
a
>> group of five marble busts by..."
>>
> This is set because of the scope note of p3: "but [b]a[/b] note is
> attached to [b]a[/b] specific item."
>
I actually thought that this scope note gave rise to the setting. I suppose I don't like the scope note then :-) I don't have strong feelings about this one because I'm probably not going to use the property for my work. But I think "semantic overcommitment" could restrict usability. Maybe some curators would like to use a single note for several objects (a nice feature of digital notes). Maybe you can bring it up in the CRM SIG. But as I said, no big deal as far as I'm concerned.

There is another thing I meant to ask you. At first I was delighted when I saw that you introduced P79a.begins_at and P80a.ends_at in the TQ versions. But then I started to wonder about some details. Domain and range of the two properties are set as Time-Span and (Time-Span u Time_Primitive) respectively. First of all I wondered why a Time-Span would signify the beginning or end of another Time-Span. Is the intention here to model initial and final segments of Time-Spans? Secondly (and more importantly) I feel that P79a and P80a somewhat bypass the inner and outer boundary properties P81.ongoing_throughout and P82.at_some_time_within. If I want to specify a beginning and an end of a Time-Span as a Time Primitive (say a time point) what do I mean then by saying some Time-Span begins at point x and ends at point y? The exact beginning and end? But then I have to make the kind of precise commitment which P.81 and P.82 afford me to avoid. Of course I could (re-)interpret P79a as "has definitly been going on at point..." and P80a as "was definitly still going on at point..." thereby capturing the inner boundary idea (P.81) or similarly for the outer temporary boundary (P.82). But if I want to use P79a and P80a consistently I cannot (according to my current understanding) capture both ideas. Therefore I am currently using a different solution (if you are interested I can send you details, it's very simple). Can you vaguely sympathise with what I said or have I misunderstood the intended meaning of P79a and P80a ?

One last thing: You did not change the ontology URI and the default namespace in the newly uploaded versions. Was that on purpose?

Best whishes,
Daniel

From: "Bernhard Schiemann" <Bernhard.Schiemann@informatik.uni-erlangen.de
>
Subject: Re: Object properties in the Erlangen CRM
Date: Fri, July 17, 2009 12:01 pm
To: "Daniel Isemann" <isemandi@cs.tcd.ie>

Dear Daniel,
> thanks for your very fast response. I didn't expect a new version the
> same day. That was great!
Thx.
> I like your SED example. It's spot on and (just slightly :-) less
> contrived then my vulcanologists...
;-) ;-) ;-) Live long and prosper.
>
> I actually thought that this scope note gave rise to the setting. I
> suppose I don't like the scope note then :-) I don't have strong
> feelings about this one because I'm probably not going to use the
> property for my work. But I think "semantic overcommitment" could
> restrict usability. Maybe some curators would like to use a single note
> for several objects (a nice feature of digital notes). Maybe you can
> bring it up in the CRM SIG. But as I said, no big deal as far as I'm
> concerned.
Hm, Please take into account, that the corresponding datatype property
is not functional, which allows to reduce strictness of the model of
single notes (what you call a "semantic overcommitment").
e.g.: a1 instanceOf E1, a2 instanceOf E1, b1 instanceOf E62, b2
instanceOf E62, b1 has_primitiveString "Test.", b2 has_primitiveString
"Test.", b1 != b2 (not equivalent), which should show that b1 and b2 are
"functional connected" to a1 resp. a2, but not the datatype:Strings. A
note "identifies" the corresponding instance of E1 and not the not
parsed datatype:String.
>
> There is another thing I meant to ask you. At first I was delighted
when
> I saw that you introduced P79a.begins_at and P80a.ends_at in the TQ
> versions. But then I started to wonder about some details. Domain and
> range of the two properties are set as Time-Span and (Time-Span u
> Time_Primitive) respectively. First of all I wondered why a Time-Span
> would signify the beginning or end of another Time-Span. Is the
> intention here to model initial and final segments of Time-Spans?
This is the CIDOC CRM way of defining that. We have an ongoing
discussion in the CIDOC SIG about this issue. Think of the E52
"Plaistozaen" which could be a time span about 5 mill. years duration,
about 5 mill. years in the past. Now, you can model a E52 time-span for
the beginning of the "Plaistozaen", which has no beginning with a date,
but an approximated start time-span (start-start) and approximated end
of the start time-span, and so on. I do prefer E61.Time_primitives as
well.

> Secondly (and more importantly) I feel that P79a and P80a somewhat
> bypass the inner and outer boundary properties P81.ongoing_throughout
> and P82.at_some_time_within. If I want to specify a beginning and an
end
> of a Time-Span as a Time Primitive (say a time point) what do I mean
> then by saying some Time-Span begins at point x and ends at point y?
The
> exact beginning and end? But then I have to make the kind of precise
> commitment which P.81 and P.82 afford me to avoid. Of course I could
> (re-)interpret P79a as "has definitly been going on at point..." and
> P80a as "was definitly still going on at point..." thereby capturing
the
> inner boundary idea (P.81) or similarly for the outer temporary
boundary
> (P.82). But if I want to use P79a and P80a consistently I cannot
> (according to my current understanding) capture both ideas. Therefore I
> am currently using a different solution (if you are interested I can
> send you details, it's very simple). Can you vaguely sympathise with
> what I said or have I misunderstood the intended meaning of P79a and
P80a?

P81 and P82 allow vague intervals. P79a and P80a allow sharp borders for
vague intervals. A simple example: My Lifespan (L1) is instanceOf E52.
L1 -P81-> minInterval1, L1 -P82-> maxInterval1, minInterval1 -P79a->
1976, minInterval1 -P80a-> today, maxInterval1 -P79a-> 1976,
maxInterval1 -P80a-> 31.12.2009. If no P79a and no P80a connections are
allowed, how do you would express these dates?

> One last thing: You did not change the ontology URI and the default
> namespace in the newly uploaded versions. Was that on purpose?

?

```
"<rdf:RDF
    xmlns:xsp="http://www.owl-ontologies.com/2005/08/07/xsp.owl#"

xmlns="http://www8.informatik.uni-erlangen.de/IMMD8/Services/cidoc-crm/
erlangen-crm_090714_5_0_1_TQ.owl#"
    xmlns:swrlb="http://www.w3.org/2003/11/swrlb#"
    xmlns:swrl="http://www.w3.org/2003/11/swrl#"
    xmlns:protege="http://protege.stanford.edu/plugins/owl/protege#"
    xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
    xmlns:xsd="http://www.w3.org/2001/XMLSchema#"
    xmlns:rdfs="http://www.w3.org/2000/01/rdf-schema#"
    xmlns:owl="http://www.w3.org/2002/07/owl#"

xml:base="http://www8.informatik.uni-erlangen.de/IMMD8/Services/cidoc-crm
/erlangen-crm_090714_5_0_1_TQ.owl">
  <owl:Ontology rdf:about="">
    <rdfs:comment xml:lang="en">The Erlangen CIDOC CRM "
from
```

http://www8.informatik.uni-erlangen.de/IMMD8/Services/cidoc-crm/erlangen-crm_090714_5_0_1_TQ.owl

kind regards
Ben

From: "Daniel Isemann" <isemandi@cs.tcd.ie>
Subject: Re: Object properties in the Erlangen CRM
Date: Mon, July 27, 2009 12:25 pm
To: Bernhard.Schiemann@informatik.uni-erlangen.de

Dear Bernhard,

thanks for your email. It took me a while to digest all the information

> Hm, Please take into account, that the corresponding datatype property
> is not functional, which allows to reduce strictness of the model of
> single notes (what you call a "semantic overcommitment").
> e.g.: a1 instanceOf E1, a2 instanceOf E1, b1 instanceOf E62, b2
> instanceOf E62, b1 has_primitiveString "Test.", b2 has_primitiveString
> "Test.", b1 != b2 (not equivalent), which should show that b1 and b2
are
> "functional connected" to a1 resp. a2, but not the datatype:Strings. A
> note "identifies" the corresponding instance of E1 and not the not
> parsed datatype:String.
Ok

> This is the CIDOC CRM way of defining that. We have an ongoing
> discussion in the CIDOC SIG about this issue. Think of the E52
> "Plaistozaen" which could be a time span about 5 mill. years duration,
> about 5 mill. years in the past. Now, you can model a E52 time-span for
> the beginning of the "Plaistozaen", which has no beginning with a date,
> but an approximated start time-span (start-start) and approximated end
> of the start time-span, and so on. I do prefer E61.Time_primitives as
> well.
I remember that one of my professors in logic used to say that
three-valued logic, while in principle just as arbitrary as classical
two-valued logic was well suited to natural language analysis, because
in practical terms most people can agree on a core extension (true) and
an anti-extension (false) which leaves a grey area ("maybe") in between.
I believe time spans and especially time spans in a cultural context
could be a good example for this. My understanding was that P.81 and
P.82 are instruments for modelling a core time (ongoing throughout) and
a core time plus grey area (at some time within).
I suppose the use of time-spans to mark the beginning and end of other
time spans you describe could be seen as modelling such a grey area as
well?
If you iterate this idea however don't you end up very quickly with
highly hypothetical constructs that even expert practitioners wouldn't
use anymore? I mean is it not hard to have a clear idea of the start of
the start of the start of the Pleistocene?

> P79a and P80a allow sharp borders for vague intervals.
You have in mind the case when P79a and P80a are used to specify a

time-primitive not a (vague) time-span?

> A simple example: My Lifespan (L1) is instanceOf E52.
> L1 -P81-> minInterval1, L1 -P82-> maxInterval1, minInterval1 -P79a->
> 1976, minInterval1 -P80a-> today, maxInterval1 -P79a-> 1976,
> maxInterval1 -P80a-> 31.12.2009.

minInterval1 and maxInterval1 are instances of E61 Time Primitive
because of the range of P81 and P82 respectively, at the same time they
are instances of E52 Time-Span because of the domain associated with
P79a (and P80a). I know it's not inconsistent to be an E52 and an E61 at
the same time, but is the modelling as you describe it intended?

> If no P79a and no P80a connections are
> allowed, how do you would express these dates?

Currently I would model your life span as follows:
L1 --P81--> minInterval
L1 --P82--> maxInterval

minInterval --beginsAt--> minBeginning --hasPrimitiveTime-->
1.1.1977;00:00:00
minInterval --endsAt--> minEnd --hasPrimitiveTime--> 17.7.2009;12:01:09
(assuming you were still alive when you sent your last email :-)

maxInterval --beginsAt--> maxBeginning --hasPrimitiveTime-->
1.1.1976;00:00:00
maxInterval --endsAt--> maxEnd --hasPrimitiveTime--> 1.1.2176;00:00:00
(people don't live to be 200 :-)

Here min- and maxInterval are instances of my home-made "TimeInterval"
which is a subclass of E61.TimePrimitive. min- and maxBeginning and min-
and maxEnd are instances of "TimePoint" (meaning essentially time
interval of length 1sec) which is also a subclass of
E61.TimePrimitive. beginsAt and endsAt are "home-made" relations with
domain TimeInterval and range TimePoint. Then I use your relation
hasPrimitiveTime to link TimePoints to their actual xsd:dateTime
representation.

Does that make sense from your point of view or would you have
objections against it? (Is it a misunderstanding or "misuse" of sorts of
CRM modelling instruments? :-)

> ?
> "<rdf:RDF
>   xmlns:xsp="http://www.owl-ontologies.com/2005/08/07/xsp.owl#"
>
> xmlns="http://www8.informatik.uni-erlangen.de/IMMD8/Services/cidoc-crm/
erlangen-crm_090714_5_0_1_TQ.owl#"
>   xmlns:swrlb="http://www.w3.org/2003/11/swrlb#"
>   xmlns:swrl="http://www.w3.org/2003/11/swrl#"
>   xmlns:protege="http://protege.stanford.edu/plugins/owl/protege#"

\>   xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
\>   xmlns:xsd="http://www.w3.org/2001/XMLSchema#"
\>   xmlns:rdfs="http://www.w3.org/2000/01/rdf-schema#"
\>   xmlns:owl="http://www.w3.org/2002/07/owl#"
\>
\> xml:base="http://www8.informatik.uni-erlangen.de/IMMD8/Services/cidoc-crm/erlangen-crm_090714_5_0_1_TQ.owl">
\>   <owl:Ontology rdf:about="">
\>   <rdfs:comment xml:lang="en">The Erlangen CIDOC CRM "
\> from
\> http://www8.informatik.uni-erlangen.de/IMMD8/Services/cidoc-crm/erlangen-crm_090714_5_0_1_TQ.owl

Funny. I could have sworn that the first release still had the old URI and namespace. I even found the file with new name and changes but old URI/NS on my harddrive. I must have messed it up somehow :-)

Finally, is it save to assume that P10.falls_within is transitive (as opposed to P9 which is probably not transitive)?
I always think of P9I as a subproperty of P10. Is there a reason why they are unrelated in CRM?

Thanks!

Daniel

From: "Bernhard Schiemann" <Bernhard.Schiemann@informatik.uni-erlangen.de
>
Subject: Re: Object properties in the Erlangen CRM
Date: Mon, July 27, 2009 4:26 pm
To: "Daniel Isemann" <isemandi@cs.tcd.ie>

Dear Daniel,
>> This is the CIDOC CRM way of defining that. We have an ongoing
>> discussion in the CIDOC SIG about this issue. Think of the E52
>> "Plaistozaen" which could be a time span about 5 mill. years duration
,
>> about 5 mill. years in the past. Now, you can model a E52 time-span
for
>> the beginning of the "Plaistozaen", which has no beginning with a date
,
>> but an approximated start time-span (start-start) and approximated end
>> of the start time-span, and so on. I do prefer E61.Time_primitives as
>> well.
> I remember that one of my professors in logic used to say that
> three-valued logic, while in principle just as arbitrary as classical
> two-valued logic was well suited to natural language analysis, because
> in practical terms most people can agree on a core extension (true) and
> an anti-extension (false) which leaves a grey area ("maybe") in between
.
> I believe time spans and especially time spans in a cultural context
> could be a good example for this. My understanding was that P.81 and
> P.82 are instruments for modelling a core time (ongoing throughout) and
> a core time plus grey area (at some time within).
> I suppose the use of time-spans to mark the beginning and end of other
> time spans you describe could be seen as modelling such a grey area as
> well?
Yes could be seen as model alternative for the "grey" areas, but I think
the importants issue is that the give grey areas "sharp" boundaries.
> If you iterate this idea however don't you end up very quickly with
> highly hypothetical constructs that even expert practitioners wouldn't
> use anymore? I mean is it not hard to have a clear idea of the start of
> the start of the start of the Pleistocene?
Yep. As I said, this is an ongoing discussion on the CRM SIG mailing
list. I think you summed some of my arguments in this discussion up.
Therefore: agreed! We need P79a and P80a, ;-)
>
>> P79a and P80a allow sharp borders for vague intervals.
> You have in mind the case when P79a and P80a are used to specify a
> time-primitive not a (vague) time-span?
Yep.
>
>> A simple example: My Lifespan (L1) is instanceOf E52.
>> L1 -P81-> minInterval1, L1 -P82-> maxInterval1, minInterval1 -P79a->

```
>> 1976, minInterval1 -P80a-> today, maxInterval1 -P79a-> 1976,
>> maxInterval1 -P80a-> 31.12.2009.
> minInterval1 and maxInterval1 are instances of E61 Time Primitive
> because of the range of P81 and P82 respectively, at the same time they
> are instances of E52 Time-Span because of the domain associated with
> P79a (and P80a). I know it's not inconsistent to be an E52 and an E61
at
> the same time, but is the modelling as you describe it intended?
This is a two step model: minInterval1, maxInterval1 should be seen as
instance Of a Timespan. Yes they cannot be instanceOf E52 and instanceOf
E61 at the same. The CRM document says that both E52 and E61 are time
intervals.
>
>> If no P79a and no P80a connections are
>> allowed, how do you would express these dates?
> Currently I would model your life span as follows:
> L1 --P81--> minInterval
> L1 --P82--> maxInterval
> minInterval --beginsAt--> minBeginning --hasPrimitiveTime-->
> 1.1.1977;00:00:00
> minInterval --endsAt--> minEnd --hasPrimitiveTime--> 17.7.2009;12:01:09
> (assuming you were still alive when you sent your last email :-)
> maxInterval --beginsAt--> maxBeginning --hasPrimitiveTime-->
> 1.1.1976;00:00:00
> maxInterval --endsAt--> maxEnd --hasPrimitiveTime--> 1.1.2176;00:00:00
> (people don't live to be 200 :-)
there is no endsAt and beginsAt if no p79a and p80a is allowed!
>
> Here min- and maxInterval are instances of my home-made "TimeInterval"
> which is a subclass of E61.TimePrimitive. min- and maxBeginning and min
-
> and maxEnd are instances of "TimePoint" (meaning essentially time
> interval of length 1sec) which is also a subclass of E61.TimePrimitive.
> beginsAt and endsAt are "home-made" relations with domain TimeInterval
> and range TimePoint. Then I use your relation hasPrimitiveTime to link
> TimePoints to their actual xsd:dateTime representation.
>
> Does that make sense from your point of view or would you have
> objections against it? (Is it a misunderstanding or "misuse" of sorts
of
> CRM modelling instruments? :-)
THe home made is my problem, these two consists mainly of the same
semantics of p79a and p80a.
>
>> ?
>> "<rdf:RDF
>>     xmlns:xsp="http://www.owl-ontologies.com/2005/08/07/xsp.owl#"
>>
>> xmlns="http://www8.informatik.uni-erlangen.de/IMMD8/Services/cidoc-crm
/erlangen-crm_090714_5_0_1_TQ.owl#"
```

>>
>> xmlns:swrlb="http://www.w3.org/2003/11/swrlb#"
>> xmlns:swrl="http://www.w3.org/2003/11/swrl#"
>> xmlns:protege="http://protege.stanford.edu/plugins/owl/protege#"
>> xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
>> xmlns:xsd="http://www.w3.org/2001/XMLSchema#"
>> xmlns:rdfs="http://www.w3.org/2000/01/rdf-schema#"
>> xmlns:owl="http://www.w3.org/2002/07/owl#"
>>
>> xml:base="http://www8.informatik.uni-erlangen.de/IMMD8/Services/cidoc-crm/erlangen-crm_090714_5_0_1_TQ.owl">
>>
>> <owl:Ontology rdf:about="">
>> <rdfs:comment xml:lang="en">The Erlangen CIDOC CRM "
>> from
>> http://www8.informatik.uni-erlangen.de/IMMD8/Services/cidoc-crm/erlangen-crm_090714_5_0_1_TQ.owl
>>
> Funny. I could have sworn that the first release still had the old URI
> and namespace. I even found the file with new name and changes but old
> URI/NS on my harddrive. I must have messed it up somehow :-)

Maybe a early bird beta release cached? Browsercache? If its fine now.... ;-)

>
> Finally, is it save to assume that P10.falls_within is transitive (as
> opposed to P9 which is probably not transitive)?
> I always think of P9I as a subproperty of P10. Is there a reason why
> they are unrelated in CRM?

The relation between P9i and P10 is probably stated in the scope note:"This property describes an instance of E4 Period, which falls within the E53 Place and E52 Time-Span of another. The difference with P9 consists of (forms part of) is subtle. Unlike P9 consists of (forms part of), P10 falls within (contains) does not imply any logical connection between the two periods and it may refer to a period of a completely different type. Example: the Great Plague (E4) falls within The Gothic period (E4)." So the problem was, how to formalise this statement. Any suggestions?

>
> Thanks!

Thank You!!! ;-)

Ben

>
> Daniel
>

From: "Daniel Isemann" <isemandi@cs.tcd.ie>
Subject: Re: Object properties in the Erlangen CRM
Date: Sun, August 2, 2009 1:34 am
To: Bernhard.Schiemann@informatik.uni-erlangen.de

Dear Ben,

> Yes could be seen as model alternative for the "grey" areas, but I
think
> the importants issue is that the give grey areas "sharp" boundaries.
>
Which ones give the grey area sharp boundaries? P81 and P82 (I would
think). But you were talking about P79a and P80a here, weren't you?

>> If you iterate this idea however don't you end up very quickly with
>> highly hypothetical constructs that even expert practitioners wouldn't
>> use anymore? I mean is it not hard to have a clear idea of the start
of
>> the start of the start of the Pleistocene?
>>
> Yep. As I said, this is an ongoing discussion on the CRM SIG mailing
> list. I think you summed some of my arguments in this discussion up.
> Therefore: agreed! We need P79a and P80a, ;-)
>
Yes, if you specify time primitives straight away. But by allowing
time-spans in the range as well you allow for those iterations I
mentioned above. As in
(time-span of) pleistocene (E52) --P79a--> (time-span of) beginning of
pleistocene (E52 again!) --P79a--> (time-span of) beginning of beginning
of pleistocene (E52 again!!) and so on...

>>> A simple example: My Lifespan (L1) is instanceOf E52.
>>> L1 -P81-> minInterval1, L1 -P82-> maxInterval1, minInterval1 -P79a->
>>> 1976, minInterval1 -P80a-> today, maxInterval1 -P79a-> 1976,
>>> maxInterval1 -P80a-> 31.12.2009.
>>>
>> minInterval1 and maxInterval1 are instances of E61 Time Primitive
>> because of the range of P81 and P82 respectively, at the same time
they
>> are instances of E52 Time-Span because of the domain associated with
>> P79a (and P80a). I know it's not inconsistent to be an E52 and an E61
at
>> the same time, but is the modelling as you describe it intended?
>>
> This is a two step model: minInterval1, maxInterval1 should be seen as
> instance Of a Timespan. Yes they cannot be instanceOf E52 and
instanceOf
> E61 at the same. The CRM document says that both E52 and E61 are time

> intervals.
>
I'm not sure if I follow. My original point was, that in your model above, minInterval1 (same for maxInt of course) is a time-primitive (E61) in virtue of appearing on the right-hand side of P81. That is part of the CRM spec and also of the OWL file. But at the same time it appears on the left hand side of P79a which has domain E52 time-span. So this, if I am not mistaken, enforces that minInterval1 is E52 and E61 at the same time. Not strictly speaking an inconsistency, since disjointness of E52 and E61 is not stated in the OWL, but somewhat counterintuitive. If you agree that the intervals shouldn't be E52 and E61 how would you model your life-span then?

> Here min- and maxInterval are instances of my home-made "TimeInterval"
> which is a subclass of E61.TimePrimitive. min- and maxBeginning and min
-
> and maxEnd are instances of "TimePoint" (meaning essentially time
> interval of length 1sec) which is also a subclass of E61.TimePrimitive.
> beginsAt and endsAt are "home-made" relations with domain TimeInterval
> and range TimePoint. Then I use your relation hasPrimitiveTime to link
> TimePoints to their actual xsd:dateTime representation.
>
> Does that make sense from your point of view or would you have
> objections against it? (Is it a misunderstanding or "misuse" of sorts
of
> CRM modelling instruments? :-)
>

THe home made is my problem, these two consists mainly of the same semantics of p79a and p80a.

I think what might be at the bottom of this, is the following: The examples given in the CRM documentation for E61 Time-Primitive suggest that the class contains primitive values that denote time intervals (in an informal sense, not E52 or sth) and possibly time points. Your, in the implementation context very useful datatype property "has_PrimitiveTime" links E61 (domain) to xsd:dateTime (range). But xsd:dateTime only describes very particular time intervals (if you like) of length one second (date;hh:mm:ss). So for the Time-Primitive "1994-1997" (an example from the CRM doc) I wouldn't know which xsd:dateTime to assign. So how can I use P81 and P82 to relate my E52 Time Spans to E61 Time Primitives (which will usually denote long intervals) and at the same time link the latter via some datatype property to an xsd:dateTime for the beginning and another one for the end. Currently I achieve this with my beginsAt and endsAt relations which link "longer" Time Primitives to descriptors for their beginning (second) and their end (last second) (my properties should probably be called "beginningIsDenotedBy" and "endIsDenotedBy" or sth like that). Maybe not the most elegant solution, I admit. A possible alternative would be to have two datatype properties "hasPrimitiveBeginningTime" and "hasPrimitiveEndTime" with domain E61 and range xsd:dateTime.

My "problem" with P79a and P80a is that they establish a link between E52's and E61's which bypasses the connection between these two classes defined by the CRM doc (namely P81 and P82). How do I formally say in CRM-OWL: "Renaissance was an E4.Period, that started sometime between 1300 and 1400 and ended sometime between 1600 and 1700" (where 1300 etc. are ints or xsd:dateTimes or some "real" datatype as opposed to E61's)? That's basically all I want :-)

> Finally, is it save to assume that P10.falls_within is transitive (as
> opposed to P9 which is probably not transitive)?
> I always think of P9I as a subproperty of P10. Is there a reason why
> they are unrelated in CRM?
>

The relation between P9i and P10 is probably stated in the scope note:"This property describes an instance of E4 Period, which falls within the E53 Place and E52 Time-Span of another. The difference with
P9 consists of (forms part of) is subtle. Unlike P9 consists of (forms
part of), P10 falls within (contains) does not imply any logical connection between the two periods and it may refer to a period of a completely different type. Example: the Great Plague (E4) falls within
The Gothic period (E4)." So the problem was, how to formalise this statement. Any suggestions?

I've looked at that scope note a number of times and I've always thought of it as follows (to give an example):

Early Renaissance --P9i--> Renaissance (Early Renaissance is a "subperiod" of Renaissance, as opposed to something which just coincidentally happened at the same time and place)

Leonard da Vinci takes his last [breath] before he starts painting ML --P10--> Renaissance (not quite a subperiod, but happens within the geospatial confines of the Renaissance, apparently all that is required for P10)

But of couse also...
Early Renaissance --P10--> Renaissance (forming part of the Renaissance the Early Renaissance falls (necessarily) into the geotemporal scope of the Renaissance)

So in short, my reading would be:
P9i is a subproperty of P10
P9 is a subproperty of P10i and
P10 and P10i are transitive (P9 probably not).

Maybe the CRM doc doesn't specify the subproperty relationship, because the two properties are stated "in opposite direction" which makes it more awkward to express it in natural language (that's just speculation on my part).

That's it for today :-)

Best whishes,
Daniel

From: "Bernhard Schiemann" <Bernhard.Schiemann@informatik.uni-erlangen.de
>
Subject: Re: Object properties in the Erlangen CRM
Date: Mon, August 3, 2009 10:10 am
To: "Daniel Isemann" <isemandi@cs.tcd.ie>

Dear Daniel,
>> Yes could be seen as model alternative for the "grey" areas, but I
think
>> the importants issue is that the give grey areas "sharp" boundaries.
> Which ones give the grey area sharp boundaries? P81 and P82 (I would
> think). But you were talking about P79a and P80a here, weren't you?
Yes. Please take into account, that from my point of view a vague
time-span needs sharp boundaries. If not give, you easily design an
infinite recuring/iterate definition. (see attached graphic)
Just try to model a large, vague Time-span just with these two P81 and
P82 ones.

> Yes, if you specify time primitives straight away. But by allowing
> time-spans in the range as well you allow for those iterations I
> mentioned above. As in
> (time-span of) pleistocene (E52) --P79a--> (time-span of) beginning of
> pleistocene (E52 again!) --P79a--> (time-span of) beginning of
beginning
> of pleistocene (E52 again!!) and so on...
(see attached graphic), the reason to allow E52's in the range of P79a
and P80a is to make such models possible. We see E61 ones as
Time-Primitives (maybe exact time points) AND the minimal time-intervals
(minimum defined by the granularity of the project modeled).

>
>>>> A simple example: My Lifespan (L1) is instanceOf E52.
>>>> L1 -P81-> minInterval1, L1 -P82-> maxInterval1, minInterval1 -P79a->
>>>> 1976, minInterval1 -P80a-> today, maxInterval1 -P79a-> 1976,
>>>> maxInterval1 -P80a-> 31.12.2009.
> I'm not sure if I follow. My original point was, that in your model
> above, minInterval1 (same for maxInt of course) is a time-primitive
> (E61) in virtue of appearing on the right-hand side of P81. That is
part
> of the CRM spec and also of the OWL file. But at the same time it
> appears on the left hand side of P79a which has domain E52 time-span.
So
> this, if I am not mistaken, enforces that minInterval1 is E52 and E61
at
> the same time. Not strictly speaking an inconsistency, since
> disjointness of E52 and E61 is not stated in the OWL, but somewhat
> counterintuitive. If you agree that the intervals shouldn't be E52 and
> E61 how would you model your life-span then?

My failure, the given example is not illustrative. startInterval1 -P79a-> 1.1.1976, startInterval1 -P80a-> 31.12.1976, endInterval1 -P79a-> today, endInterval1 -P80a-> 31.12.2009, L1 -P79a-> startInterval1, L1 -P80a-> endInterval1, I would leave the P81 and P82 out, but would add P83 and P84.

> THe home made is my problem, these two consists mainly of the same
> semantics of p79a and p80a.
> I think what might be at the bottom of this, is the following: The
> examples given in the CRM documentation for E61 Time-Primitive suggest
> that the class contains primitive values that denote time intervals (in
> an informal sense, not E52 or sth) and possibly time points. Your, in
> the implementation context very useful datatype property
> "has_PrimitiveTime" links E61 (domain) to xsd:dateTime (range). But
> xsd:dateTime only describes very particular time intervals (if you like
)
> of length one second (date;hh:mm:ss). So for the Time-Primitive
or one day (date) or ms (dateThh:mm:ss.msmsms) ,... As the granularity
of the modeled project indicates.
> "1994-1997" (an example from the CRM doc) I wouldn't know which
In our point of view this example should not be an E61 but an E52
(because a nearby same example exists in the scope note of E52).
> xsd:dateTime to assign. So how can I use P81 and P82 to relate my E52
> Time Spans to E61 Time Primitives (which will usually denote long
> intervals) and at the same time link the latter via some datatype
> property to an xsd:dateTime for the beginning and another one for the
> end. Currently I achieve this with my beginsAt and endsAt relations
> which link "longer" Time Primitives to descriptors for their beginning
> (second) and their end (last second) (my properties should probably be
> called "beginningIsDenotedBy" and "endIsDenotedBy" or sth like that).
> Maybe not the most elegant solution, I admit. A possible alternative
> would be to have two datatype properties "hasPrimitiveBeginningTime"
and
> "hasPrimitiveEndTime" with domain E61 and range xsd:dateTime.
>
> My "problem" with P79a and P80a is that they establish a link between
> E52's and E61's which bypasses the connection between these two classes
> defined by the CRM doc (namely P81 and P82). How do I formally say in
> CRM-OWL: "Renaissance was an E4.Period, that started sometime between
> 1300 and 1400 and ended sometime between 1600 and 1700" (where 1300 etc
.
> are ints or xsd:dateTimes or some "real" datatype as opposed to E61's)?
> That's basically all I want :-)

E4("Renaissance"), E4("vagueStartofRenaissance"), E4("vagueStartofRenaissance"), "Renaissance" -P116I-> "vagueStartofRenaissance", "Renaissance" -P115I-> "vagueEndofRenaissance", "vagueStartofRenaissance" -P4-> "1300-1400", "vagueEndofRenaissance" -P4-> "1600-1700"
and:
a) "1300" -P79a-> 1.1.1300, "1300" -P80a-> 31.12.1300, "1400" -P79a->

1.1.1400, "1400" -P80a-> 31.12.1400, "1600" -P79a-> 1.1.1600, "1600"
-P80a-> 31.12.1600, "1700" -P79a-> 1.1.1700, "1700" -P80a-> 31.12.1700,
"1300-1400" -P79a-> "1300", "1300-1400" -P80a-> "1400",
"1600-1700" -P79a-> "1600", "1600-1700" -P80a-> "1700"
(where "Names" are instancesOf E52)
or (more easier)
b) "1300-1400" -P79a-> 1.1.1300, "1300-1400" -P80a-> 31.12.1400,
"1600-1700" -P79a-> 1.1.1600, "1600-1700" -P80a-> 31.12.1700,
"1300-1400" -P79-> ""ca."", "1300-1400" -P80-> ""ca./vague/by rule of
thumb/..."",
"1600-1700" -P79-> ""ca."", "1600-1700" -P80-> ""ca.""

a) & b) should be flexible enough. I prefer a) because it addresses
proper years as time-spans. The P79 and P80 can be added to a) as well,
but I think they are needed in b). Is this ok for you?

> The relation between P9i and P10 is probably stated in the scope
> note:"This property describes an instance of E4 Period, which falls
> within the E53 Place and E52 Time-Span of another. The difference
with
> P9 consists of (forms part of) is subtle. Unlike P9 consists of (
forms
> part of), P10 falls within (contains) does not imply any logical
> connection between the two periods and it may refer to a period of a
> completely different type. Example: the Great Plague (E4) falls
within
> The Gothic period (E4)." So the problem was, how to formalise this
> statement. Any suggestions?
> I've looked at that scope note a number of times and I've always
thought
> of it as follows (to give an example):
> Early Renaissance --P9i--> Renaissance (Early Renaissance is a
> "subperiod" of Renaissance, as opposed to something which just
> coincidentally happened at the same time and place)
> Leonard da Vinci takes his last [breath] before he starts painting ML
> --P10--> Renaissance (not quite a subperiod, but happens within the
> geospatial confines of the Renaissance, apparently all that is required
> for P10)
> But of couse also...
> Early Renaissance --P10--> Renaissance (forming part of the Renaissance
> the Early Renaissance falls (necessarily) into the geotemporal scope of
> the Renaissance)
> So in short, my reading would be:
> P9i is a subproperty of P10
> P9 is a subproperty of P10i and
> P10 and P10i are transitive (P9 probably not).
> Maybe the CRM doc doesn't specify the subproperty relationship, because
> the two properties are stated "in opposite direction" which makes it
> more awkward to express it in natural language (that's just speculation

> on my part).
Well,...,yes. I think I see the point. Let's say that if we do not
formalize the subproperty-relationship you are allowed to model:
Early Renaissance -P9i-> Renaissance and Early Renaissance -P10->
Renaissance, but it is not automatically given. I will discuss the
P9i/P10 relationship with my colleagues as well to find out if we have
an idea what is this "direction" about.
>

Have a nice day,
Ben

From: "Daniel Isemann" <isemandi@cs.tcd.ie>
Subject: Questions about CRM
Date: Wed, March 31, 2010 7:29 pm
To: Bernhard.Schiemann@informatik.uni-erlangen.de

Dear Bernhard,

while I am developing the ontology-based application some of the
modelling questions are still open. I know my last email was
terrifyingly long and you are certainly very busy, but just in case you
merely overlooked my email I thought I'll run it against you again ;-)
It would be great to hear your opinion on some of these points,
especially section (I). But if you don't have the time don't worry about
it.

Thanks,
Daniel

Daniel Isemann wrote:
> Dear Bernhard,
>
> I hope you had a good Christmas holiday and I wish you a happy new
> year! Sorry for interrupting our discussion about the CRM for so long.
> It was partly due to reasons beyond my control, but for the most part
> simply my fault. Apologies.
>
> I still have open questions about the CRM. Some of these concern
> topics we have already covered, others we haven't talked about so far.
> Roughly my questions fall into three categories.
>
> I) Modelling time in CRM (with datatypes)
> II) Properties of the mereological relations on E4.Period.
> III) Modelling location in CRM
>
> I am going to use the following two references to distinguish between
> the CRM definition and your implementation:
> [1] Crofts et al., Definition of the Cidoc CRM, retrieved on 15/1/2010
> from http://cidoc.ics.forth.gr/docs/cidoc_crm_version_5.0.1_Nov09.pdf
> [2] Erlangen CRM / OWL 2009-12-17, retrieved on 15/1/2010 from
> http://erlangen-crm.org/091217
>
>
> ad I)
> In our last communication I asked you to flesh out a modelling for
> "Renaissance was an E4.Period, that started sometime between 1300 and
> 1400 and ended sometime between 1600 and 1700". Thank you very much
> for the two detailed alternatives you presented. They differ slightly

> from how I am using the model currently. There are two points in
> particular I would like to hear your opinion on: Firstly my use of
> E52.Time-Spans is less "numerical" than yours, i.e. I am more
> interested in time-spans like "duration of the Ming Dynasty" than
> "From 12-17-1993 to 12-8-1996" (both examples are taken from [1] as
> you might have noticed, but I admit that the lonely Ming Dynasty is
> outnumbered 4 to 1 by examples of the "numerical" kind.). Secondly
> your modelling examples ended at the level of E52.Time-Spans or
> E61.Time_Primitives (depending on whether you see 1.1.1300 e.g. as an
> E52 or E61). But I would like to model all the way through to
> datatypes I can reason with. Currently I use 4 datatype properties
> that bypass E61's (to avoid complications :-) The statement about the
> Renaissance above looks like this:
>
> E4("Renaissance")
> E52("Time-Span_of_Renaissance")
> Renaissance --P4.has_time-span--> Time-Span_of_Renaissance
> Time-Span_of_Renaissance --beganAfter--> xsd:int(1300)
> Time-Span_of_Renaissance --beganBefore--> xsd:int(1400)
> Time-Span_of_Renaissance --endedAfter--> xsd:int(1600)
> Time-Span_of_Renaissance --endedBefore--> xsd:int(1700)
>
> My time points are ints because Pellet and Fact++ support int-ranges
> and comparisons and the like. That limits the granularity to one year
> (which is ok for my purposes), but in the long run they should be
> xsd:date or xsd:dateTime of course. The 4 datatype properties
> basically capture what I would like to express. A more roundabout way
> to say the same thing using E61s would be:
>
> E4("Renaissance")
> E52("Time-Span_of_Renaissance")
> Renaissance --P4.has_time-span--> Time-Span_of_Renaissance
> E61("1300-1700")
> E61("1400-1600")
> E61("1300")
> E61("1400")
> E61("1600")
> E61("1700")
> Time-Span_of_Renaissance --P81.ongoing_throughout--> "1400-1600"
> Time-Span_of_Renaissance --P82.at_some_time_within--> "1300-1700"
> "1400-1600" --begins_at--> "1400"
> "1400-1600" --ends_at--> "1600"
> "1300-1700" --begins_at--> "1300"
> "1300-1700" --ends_at--> "1700"
> "1300" --has_PrimitiveTime_int--> xsd:int(1300)
> "1400" --has_PrimitiveTime_int--> xsd:int(1400)
> "1600" --has_PrimitiveTime_int--> xsd:int(1600)
> "1700" --has_PrimitiveTime_int--> xsd:int(1700)
>
> Again the ints just stand in until reasoners can handle xsd:date

> and/or xsd:dateTime. Then "has_PrimitiveTime_int" could be replaced by
> "has_PrimitiveTime" which is already implemented in [2]. The reason
> I'm not using the second model (apart from the fact that queries get
> more complicated) is that I am not sure yet what E61s really are. Is
> E61 just a placeholder class for something like xsd:dateTime (possibly
> including intervalls)? Then I could probably go ahead with the second
> modelling. Or is E61 a bag of all sorts of strings and expressions
> that can refer to time primitive data? In short, are "2000-01-02" and
> "2000-01-01 + 1 day" two different names for the same E61 individual
> (I think they should be) or are they two separate E61s? Can you shed
> some light on this? Also let me know what you think of my time
> modelling if you can.
>
>
> ad II)
> Just before Christmas I had a systematic look again at the four
> relations in CRM with domain and range E4.Period. These are P9, P10,
> P132 and P133. The last three are purely mereological in nature (in
> the sense of a formal part-whole relationship elaborated upon e.g.
> here http://plato.stanford.edu/entries/mereology/). Would you agree?
> P9 on the other hand seems to combine a mereological notion ("part
> of") with something else.
>
> In the following I will list what formal relationships I think should
> hold between these relations and their inverses. I will give arguments
> for my views in brackets.
>
> P9I subproperty of P10 (precisely because P9I expresses "part of" and
> sth else it is stronger than P10 which only seems to express the "part
> of" bit. Therefore P9I is a subproperty of the weaker P10)
>
> P9 subproperty of P10I (this is equivalent to the above statement,
> hence the same argument applies)
>
> P10 subproperty of P132 (being part of something is an extreme case of
> overlapping with that something)
> P10I subproperty of P132 (having something as a part is also an
> extreme case of overlapping with that something)
>
> P10 is transitive (transitivity is a basic property of mereological
> "part of" relations, even in the weakest of mereologies; A part of B,
> B part of C therefore A part of C. Also compare the passage on page xv
> of [1] from which the two quotes below are taken which strongly
> suggests that P10 is actually regarded as a transitive property in [1])
.
> P10I is transitive (equivalent to the above)
>
> P132 is disjoint from P133 (In mereological systems one of the two
> relations is usually defined as the negation of the other; the wording
> in [1] also strongly supports this. Now I don't think this can be

> expressed in OWL 2.0, but at the very least one can say that they are
> disjoint)
>
> P132 is symmetric (I am just mentioning this for completeness. It is
> already implemented in [2].)
> P133 is symmetric (This is also implemented in [2], however, unlike
> P132, P133 is not listed as it's own inverse. This is no real problem,
> more of a "cosmetic" inconsistency between the two properties. Since
> symmetric implies self-inverse maybe delete the self-inverse statement
> from P132?)
>
> Which of the above statements would you agree/disagree with? If you
> disagree, why?
>
> Finally there are some statements in [1] which say something about the
> properties of properties that I find a little confusing. Especially
> this one:
>
> "Properties that have identical domain and range are either symmetric
> or transitive." [1, p. xv].
>
> In general this is certainly not the case as mathematical
> counterexamples abound. So the statement can only be interpreted as
> meaning that all the properties in the CRM with identical domain and
> range are (more or less coincidentally) either symmetric or
> transitive. Now even if that were true the statement is quite
> dangerous because if a new property was added to the CRM the statement
> might be invalidated and have to be removed (a source of error).
> Instead of talking about symmetry and transitivity of certain
> properties in a preamble I think it would be best to list these
> characteristics for each property separately in the property
> description. To make matters worse however I think CRM already
> contains at least one property which conflicts with the above
> statement, namely P9. As a corollary to the statement above P9 is
> transitive (because unlike the symmetric properties it has a
> parenthetical form that relates the meaning of the inverse direction):
>
> "Transitive asymmetric properties, such as E4 Period. P9 consist of
> (forms part of): E4 Period, have a parenthetical form that relates to
> the meaning of the inverse direction." [1, p. xv].
>
> That P9 is asymmetric seems to be mentioned as an added bonus without
> justification. But since it is in line with what I think the semantics
> of P9 are I have no problem here :-) Transitivity of P9 however
> strikes me as questionable. If
>
> Middle Minoan --P9I.forms_part_of--> Cretan Bronze Age
> and
> Cretan Bronze Age --P9I.forms_part_of--> European Bronze Age (say)
>

```
> I don't see why
>
> Middle Minoan --P9I.forms_part_of--> European Bronze Age
>
> should hold. We're not talking about mere spatio-temporal containment
> after all but about being a subsidiary period in some (historical or
> art historical) sense. Now perhaps I have a flawed understanding of
> art history and its categorisations or of the intended semantics of
> P9, but [1] contains additional information about P9 that indicates
> that something is at odds here I think:
>
> "P9 [...] Quantification: one to many, (0,n:0,1)." [1, p. 39]
>
> In the language of OWL this means that the property P9 is inverse
> functional. Accordingly P9I is functional. If P9 is transitive so is
> P9I. So P9I is functional and transitive at the same time. But this is
> quite a potent mix. In the example above one of our premisses was
>
> Middle Minoan --P9I--> Cretan Bronze Age
>
> and transitivity together with the other premiss gives
>
> Middle Minoan --P9I--> European Bronze Age.
>
> Hence by functionality of P9I:
>
> European Bronze Age = Cretan Bronze Age!!!
>
> And of course the same thing occurs in every structurally identical
> case. So either there are no P9I chains of length greater than 3 (in
> which case there is no point in demanding transitivity anyway because
> it will never fire) or else I think one of the two characteristics
> functionality and transitivity has to go (else undesirable collapses
> like European Bronze Age = Cretan Bronze Age occur). Functionality is
> a strong claim but not in conflict with what I think the intended
> semantics of P9I are. Transitivity on the other hand doesn't feel
> right for P9I, so I would advocate that it should go. I do realise of
> course that neither of the two characteristics are currently
> implemented in [2]. The whole last paragraph only refers to the CRM
> definition [1]. If you have the time it would be great to learn what
> of the above you agree with and what you disagree with and what you
> think is in need of further clarification (maybe from the CRM SIG).
>
>
> ad III)
> Seeing that my email has by now reached encyclopaedic dimensions I
> will keep the last point very short. I would like to model location
> and place of origin of works of art in the CRM. The "granularity" I
> have in mind are countries and major cities. I suppose the central
> category for this is E53.Place. Do you have any words of advice for
```

```
> me? How do I link present-day Italy to Leonardo da Vinci's Italy? Can
> you recommend an existing geographical ontology/vocabulary? Have you
> ever combined [2] or a predecessor with such a vocabulary (e.g. Getty
> Thesaurus of Geographic Names:
> http://www.getty.edu/research/conducting_research/vocabularies/tgn/)?
>
> Sorry again for the enormous proportions of this email. Partial
> answers are of course ok if you think all this is too much.
>
> Thanks for your help
> Daniel
>
```

Appendix E

Evaluation Slides

This appendix contains a reproduction of the presentation slides that were shown to participants of the system evaluation sessions (see Section 5.5) as an introduction and as a guide throughout the session.

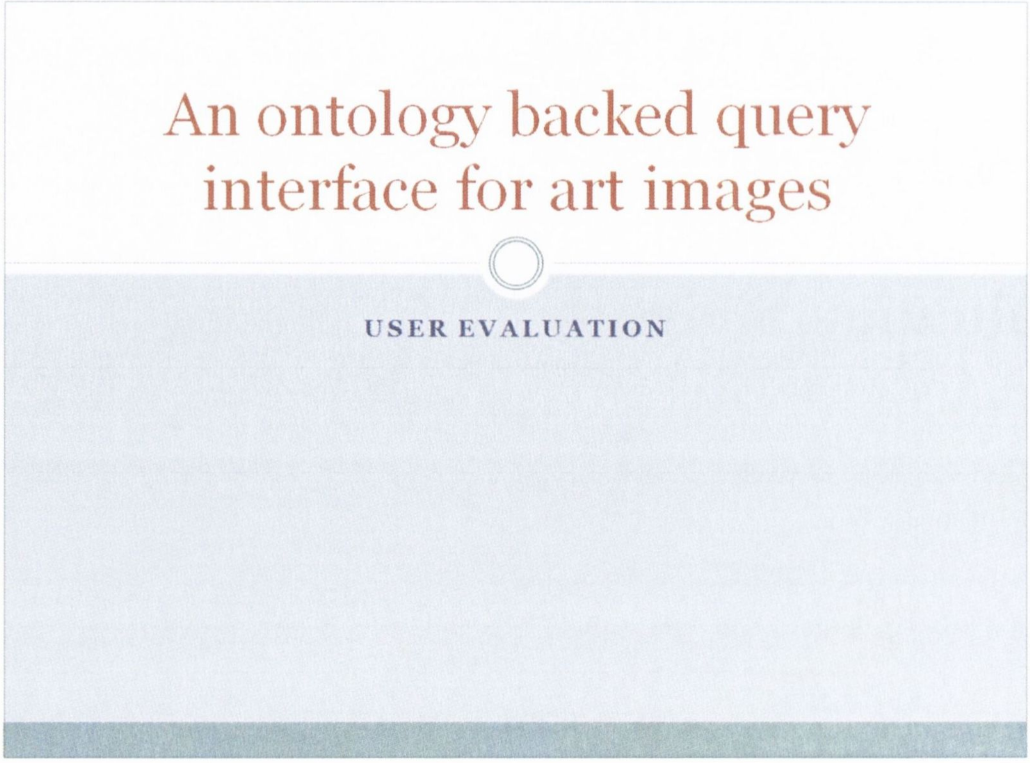

Figure E.1 – Slide 1: Title

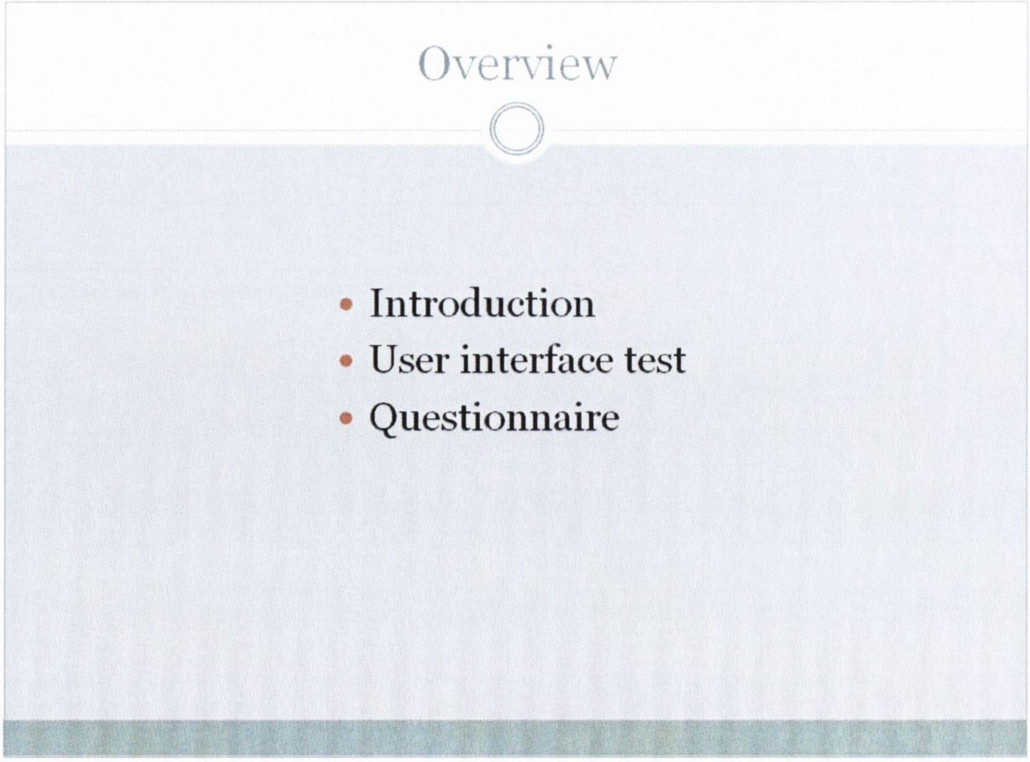

Figure E.2 – Slide 2: Overview

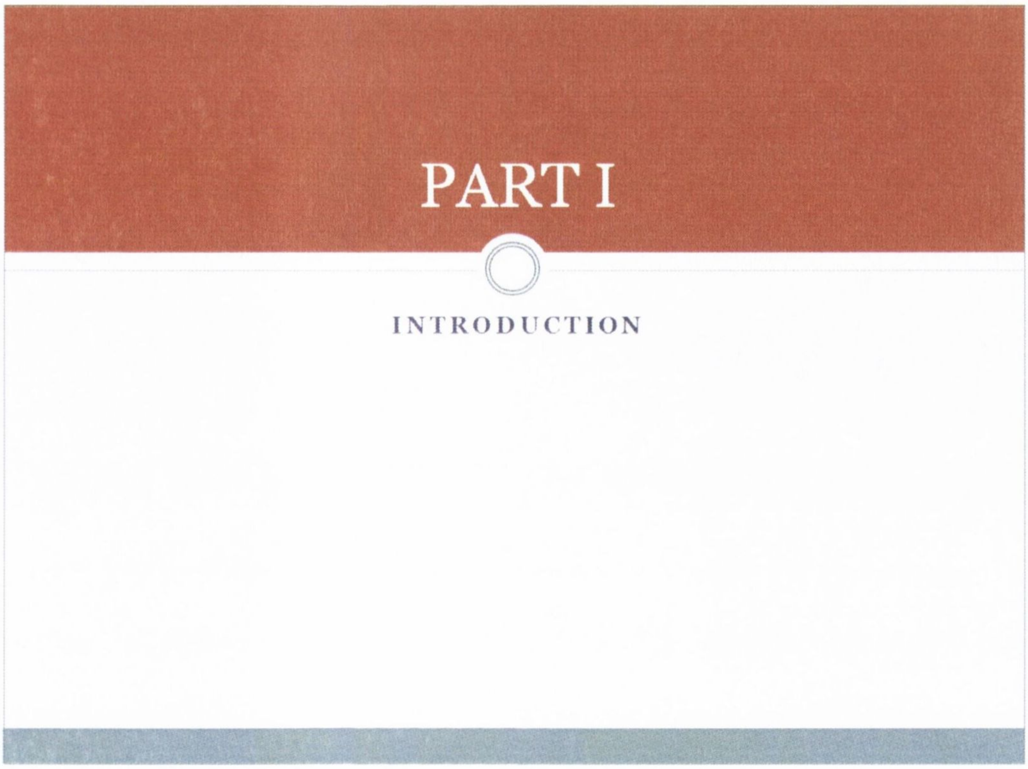

Figure E.3 – Slide 3: Part 1 – Title

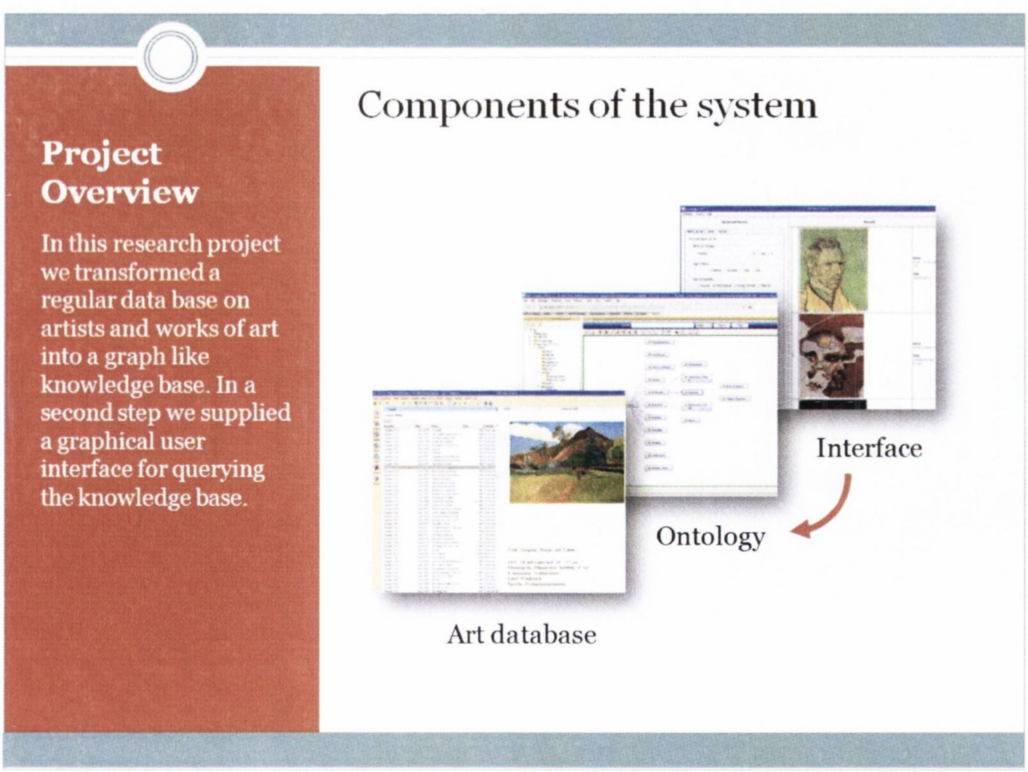

Figure E.4 – Slide 4: Components of the system

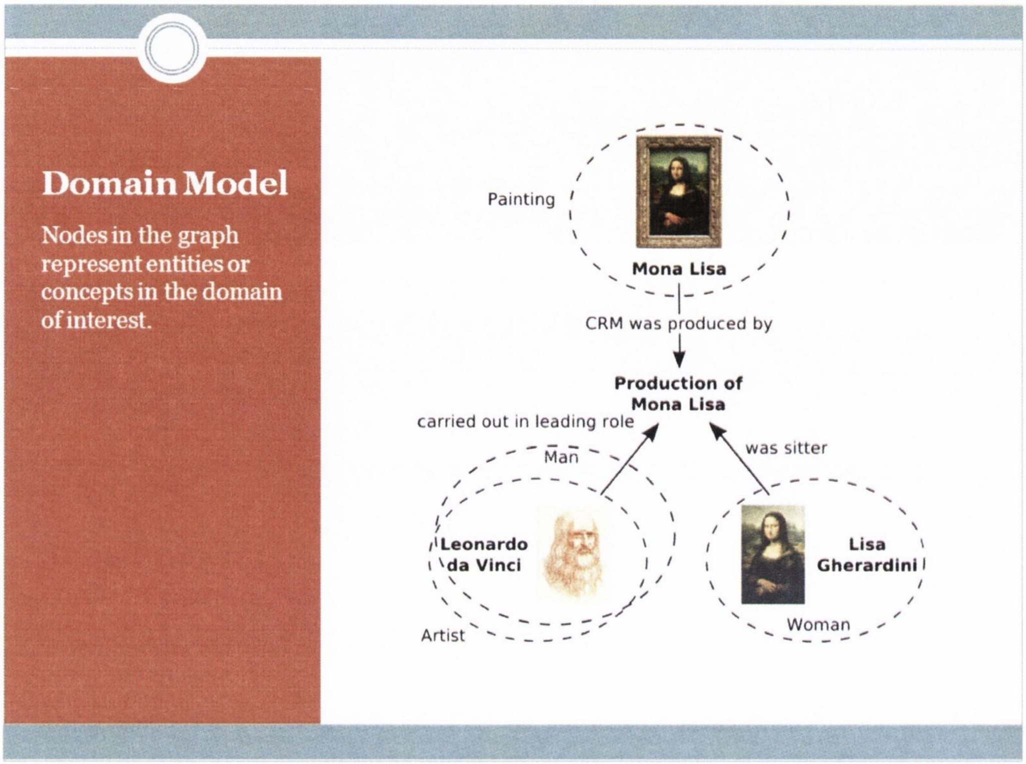

Figure E.5 – Slide 5: A modelling example

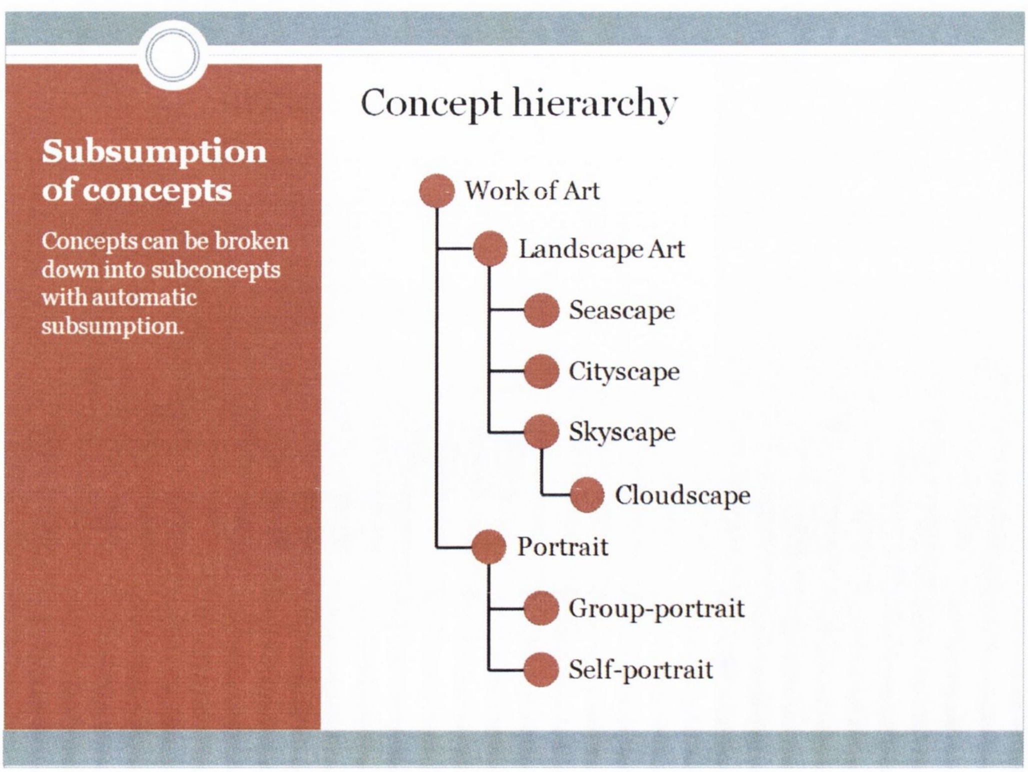

Figure E.6 – Slide 6: Concept hierarchy

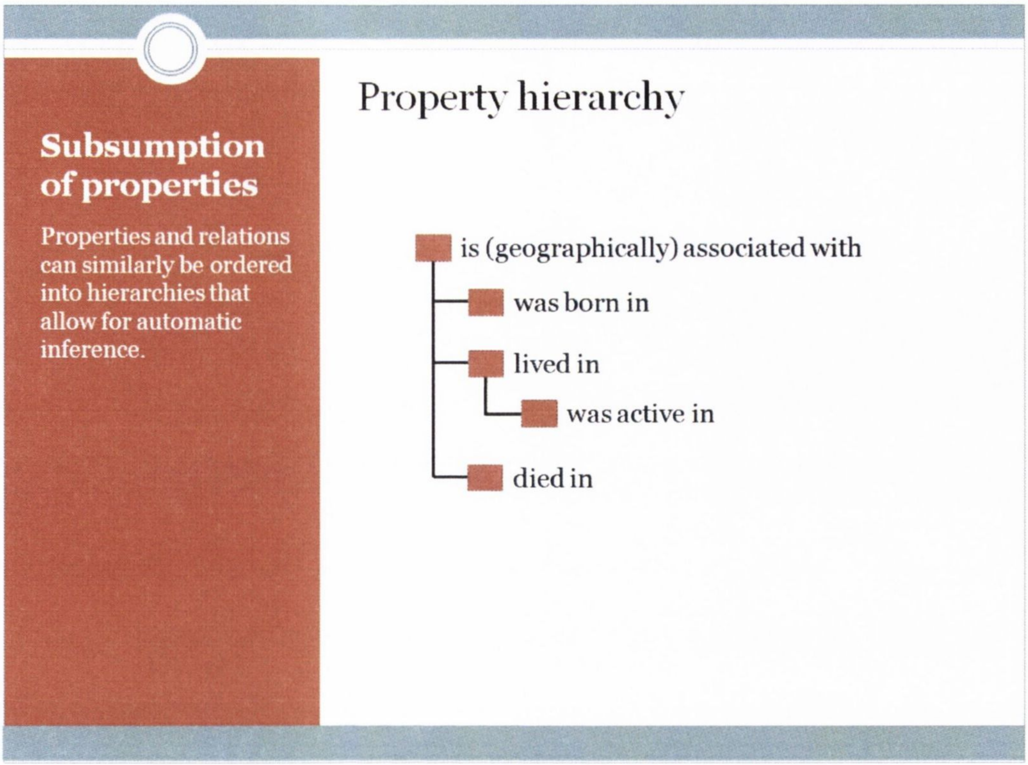

Figure E.7 – Slide 7: Property hierarchy

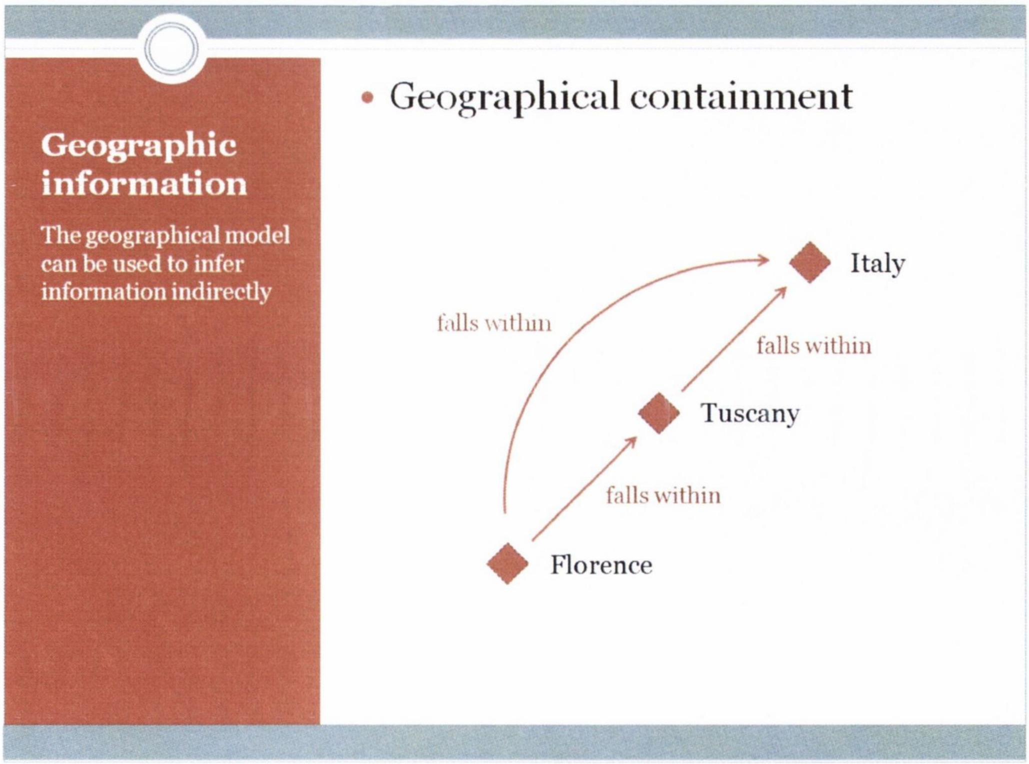

Figure E.8 – Slide 8: Transitivity: Geographical containment

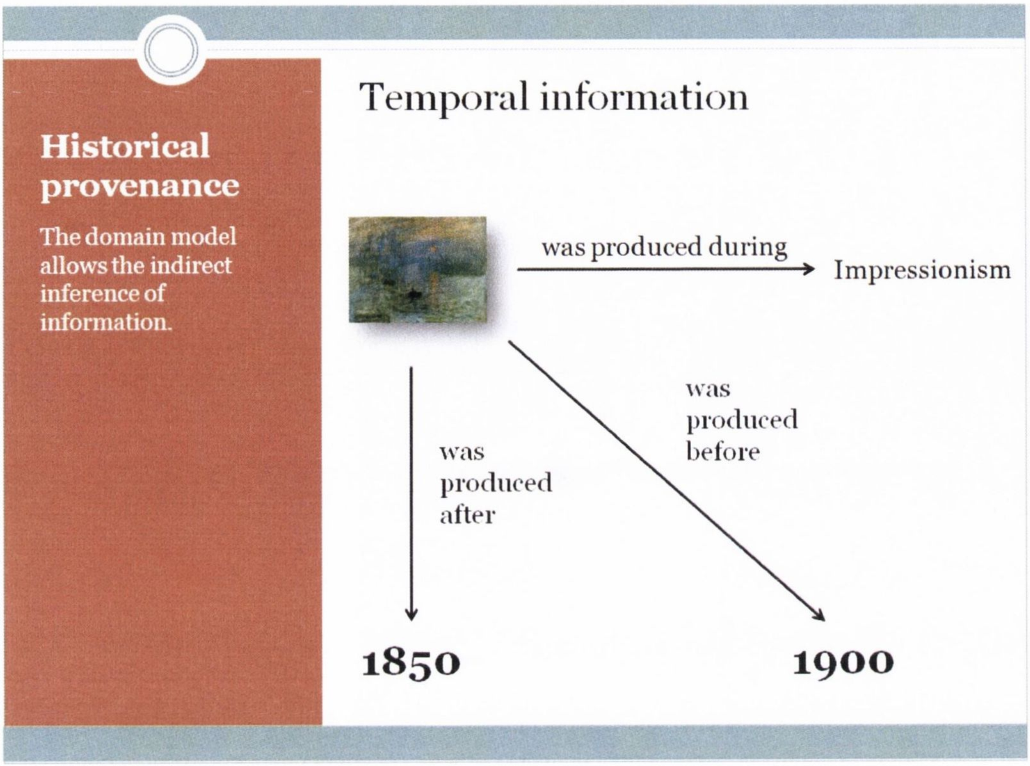

Figure E.9 – Slide 9: Temporal inference

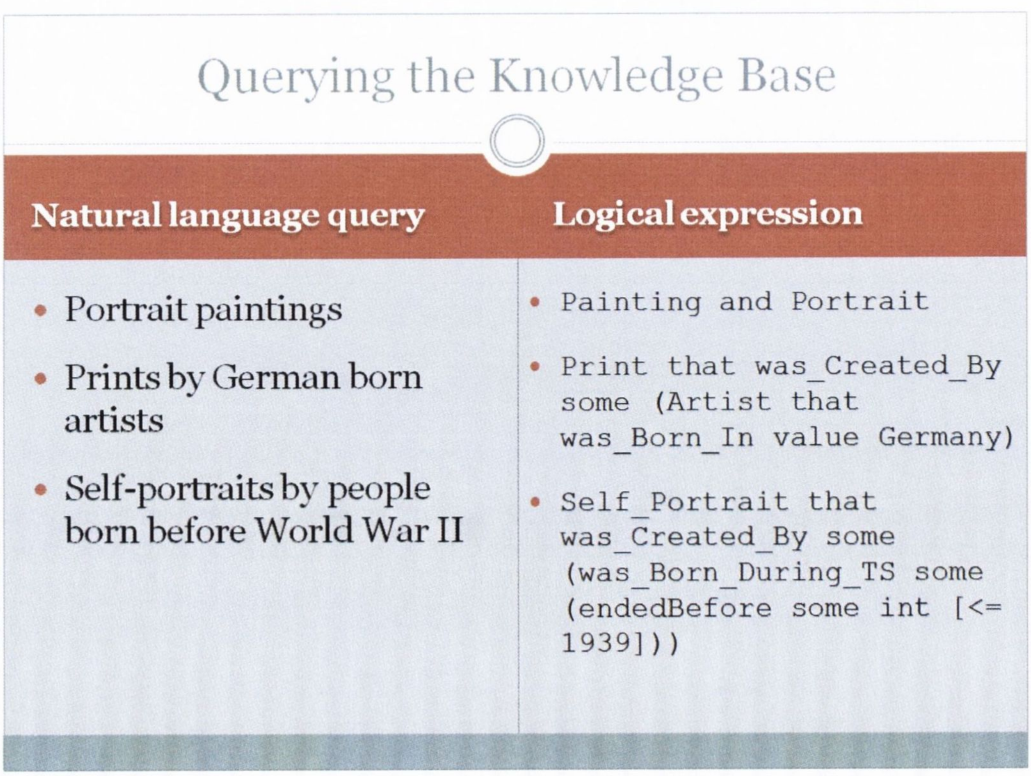

Figure E.10 – Slide 10: Example queries in natural language and in Manchester OWL syntax

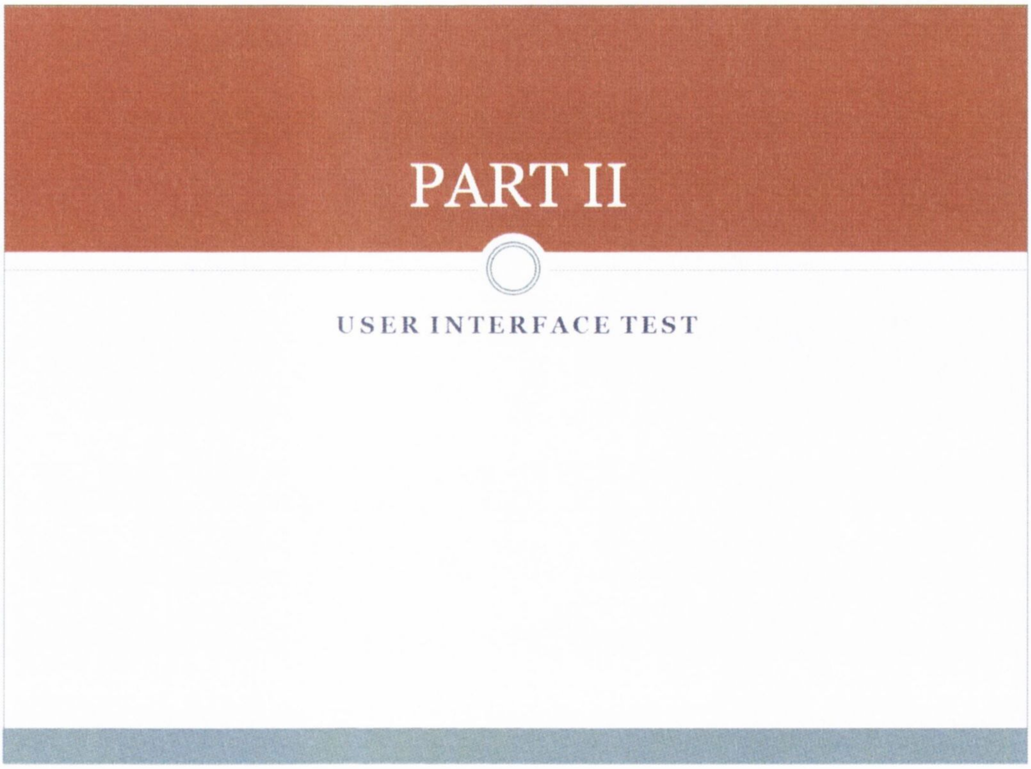

Figure E.11 – Slide 11: Part 2 – Title

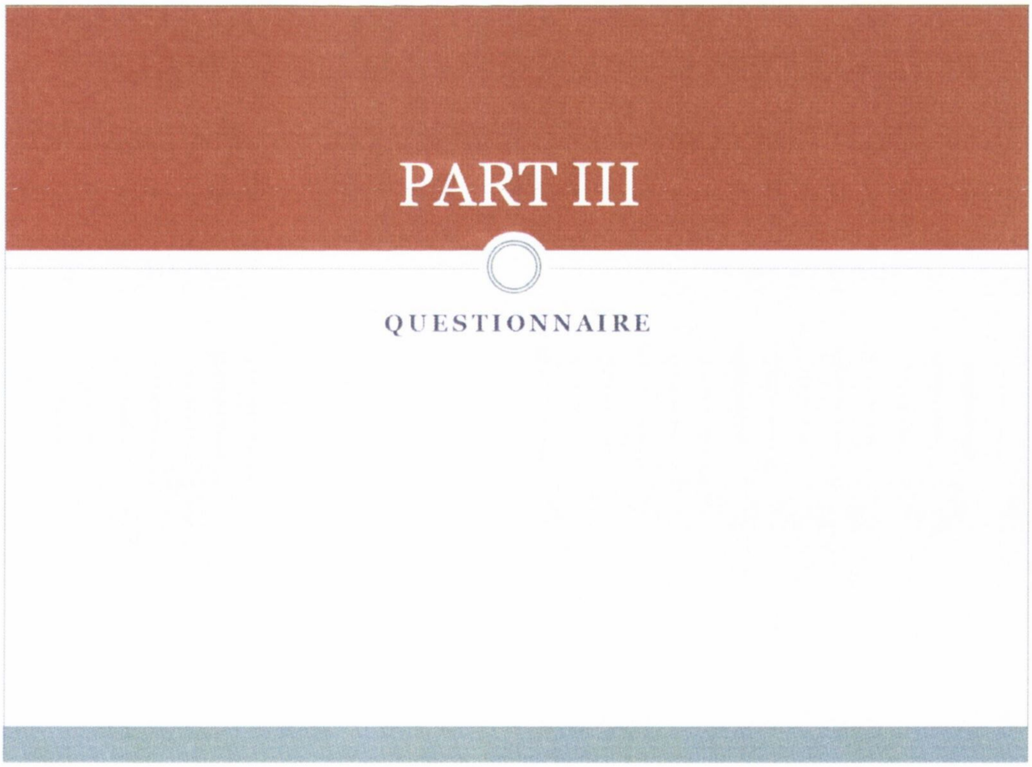

Figure E.12 – Slide 12: Part 3 – Title

Figure E.13 – Slide 13: Thank you

Appendix F

Evaluation Questionnaire

This appendix contains screenshots of the 12 pages of the Web-based questionnaire that was part of the system evaluation discussed in Section 5.5. The screenshots are cropped to only show the main content of each page. The first page additionally shows the advertisements placed by the *eSurveyspro*[1] platform which was used to conduct the survey.

[1] `http://www.esurveyspro.com` (last accessed 08/08/2011).

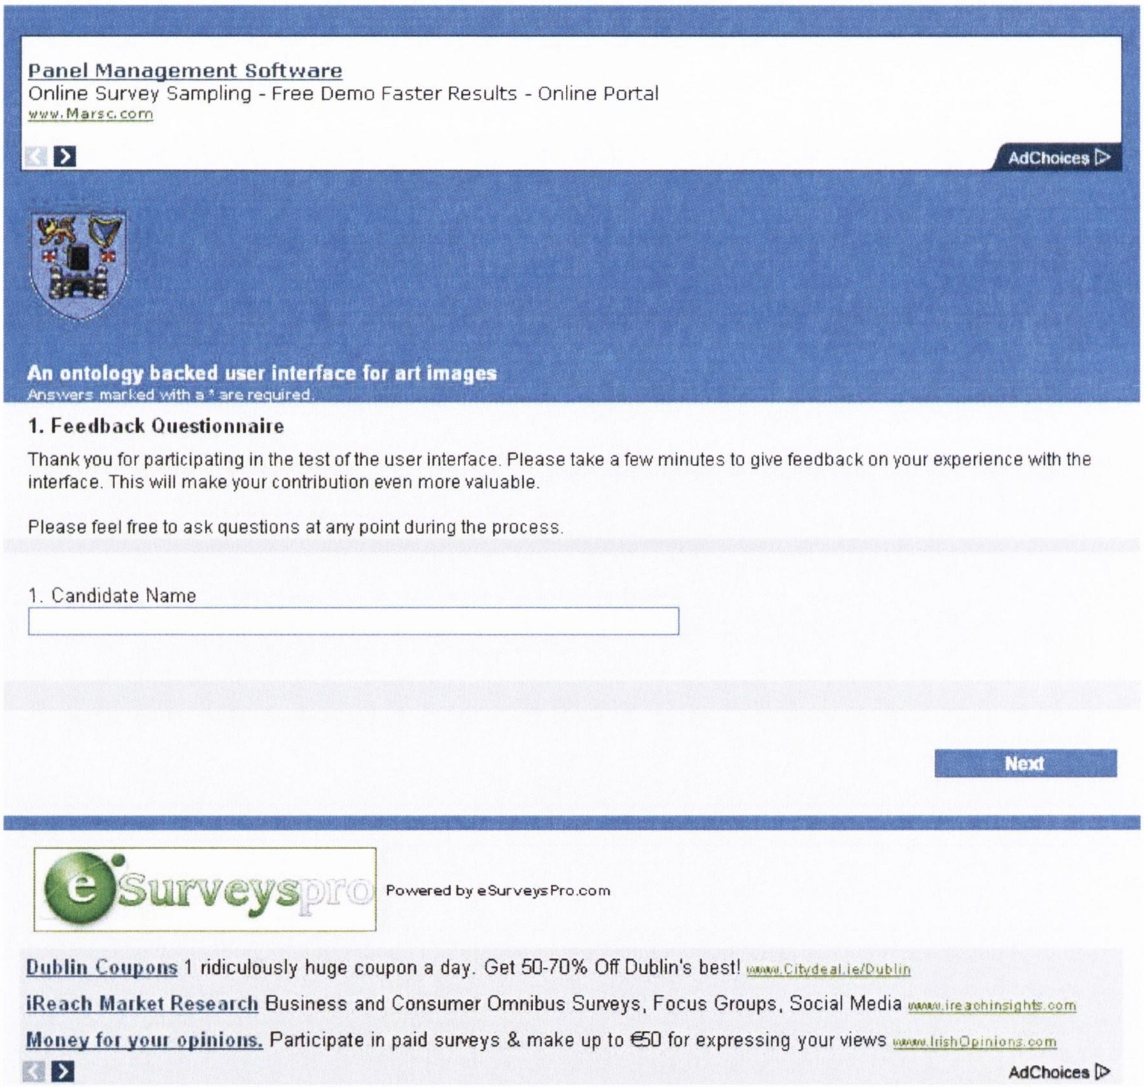

Figure F.1 – Page 1

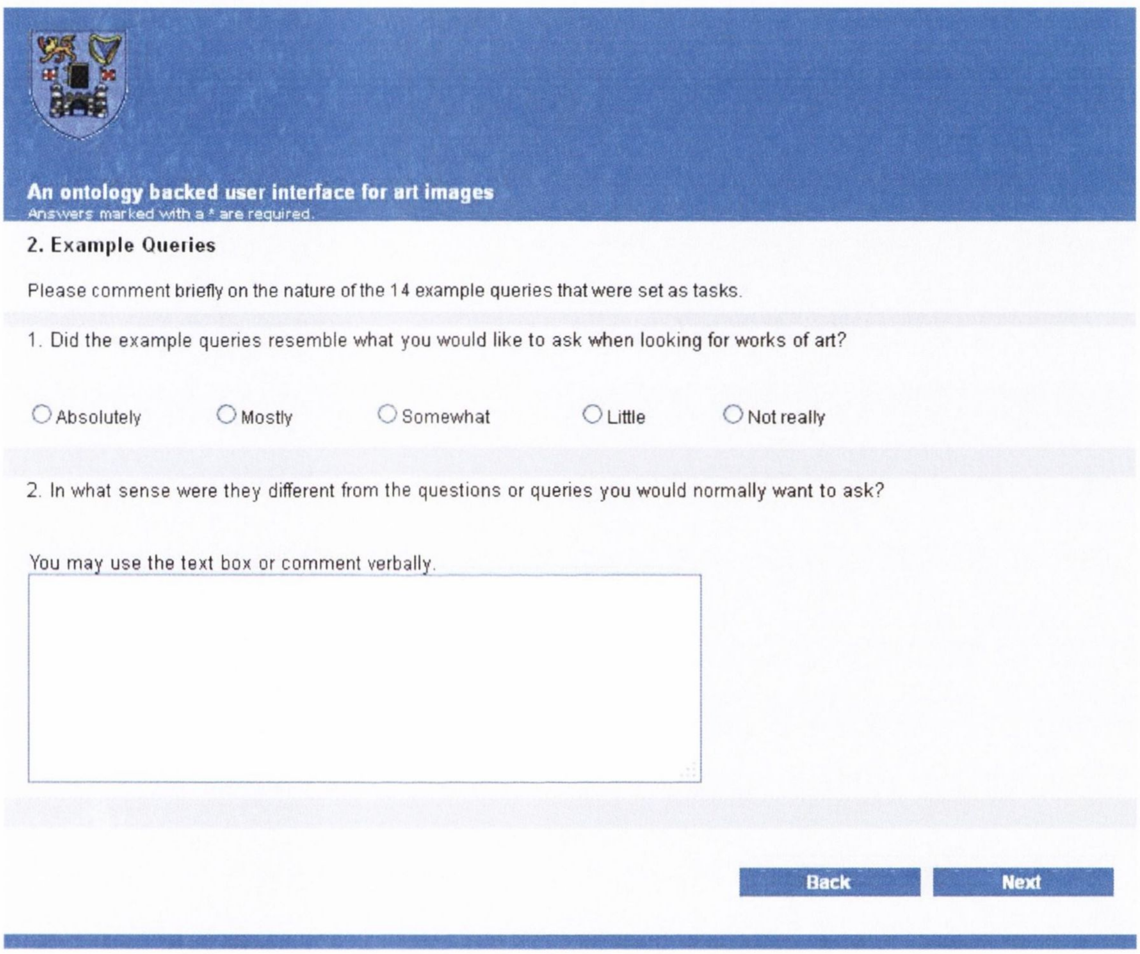

An ontology backed user interface for art images

Answers marked with a * are required.

2. Example Queries

Please comment briefly on the nature of the 14 example queries that were set as tasks.

1. Did the example queries resemble what you would like to ask when looking for works of art?

○ Absolutely ○ Mostly ○ Somewhat ○ Little ○ Not really

2. In what sense were they different from the questions or queries you would normally want to ask?

You may use the text box or comment verbally.

Back Next

Figure F.2 – Page 2

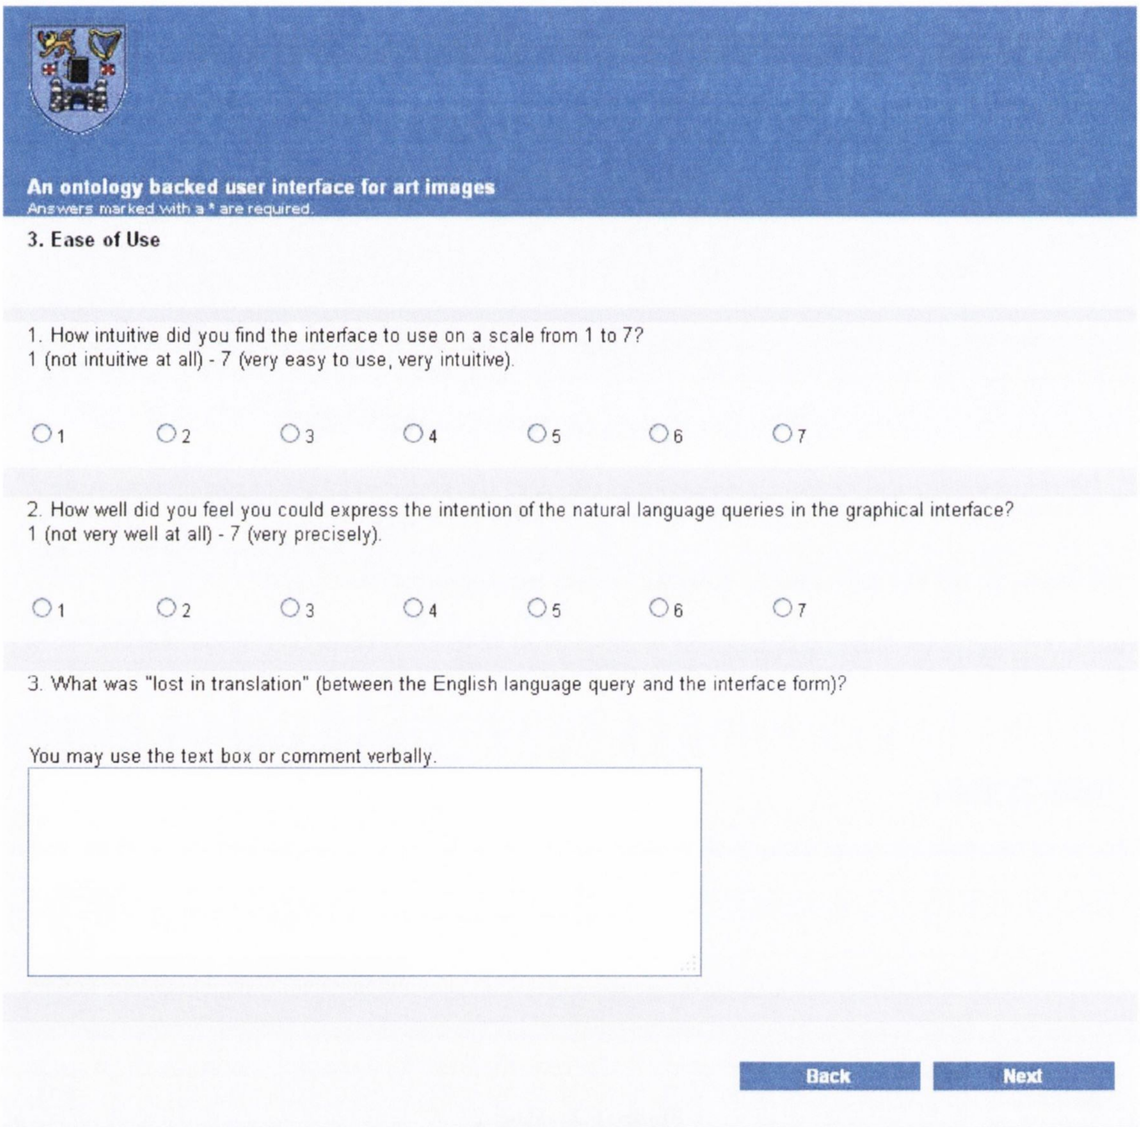

An ontology backed user interface for art images

Answers marked with a * are required.

3. Ease of Use

1. How intuitive did you find the interface to use on a scale from 1 to 7?
1 (not intuitive at all) - 7 (very easy to use, very intuitive).

○1 ○2 ○3 ○4 ○5 ○6 ○7

2. How well did you feel you could express the intention of the natural language queries in the graphical interface?
1 (not very well at all) - 7 (very precisely).

○1 ○2 ○3 ○4 ○5 ○6 ○7

3. What was "lost in translation" (between the English language query and the interface form)?

You may use the text box or comment verbally.

Back Next

Figure F.3 – Page 3

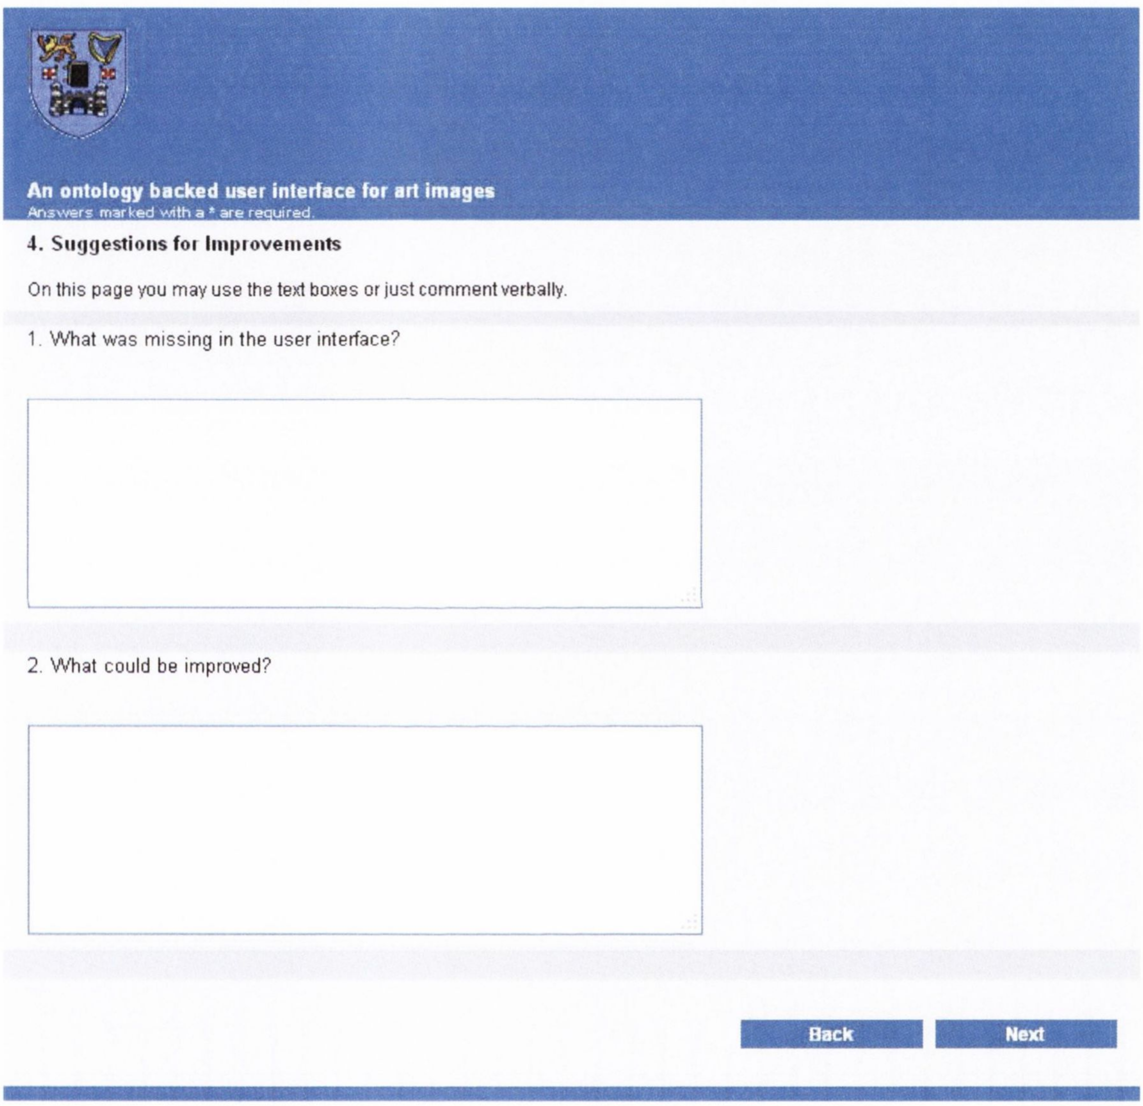

An ontology backed user interface for art images
Answers marked with a * are required.

4. Suggestions for Improvements

On this page you may use the text boxes or just comment verbally.

1. What was missing in the user interface?

2. What could be improved?

Back Next

Figure F.4 – Page 4

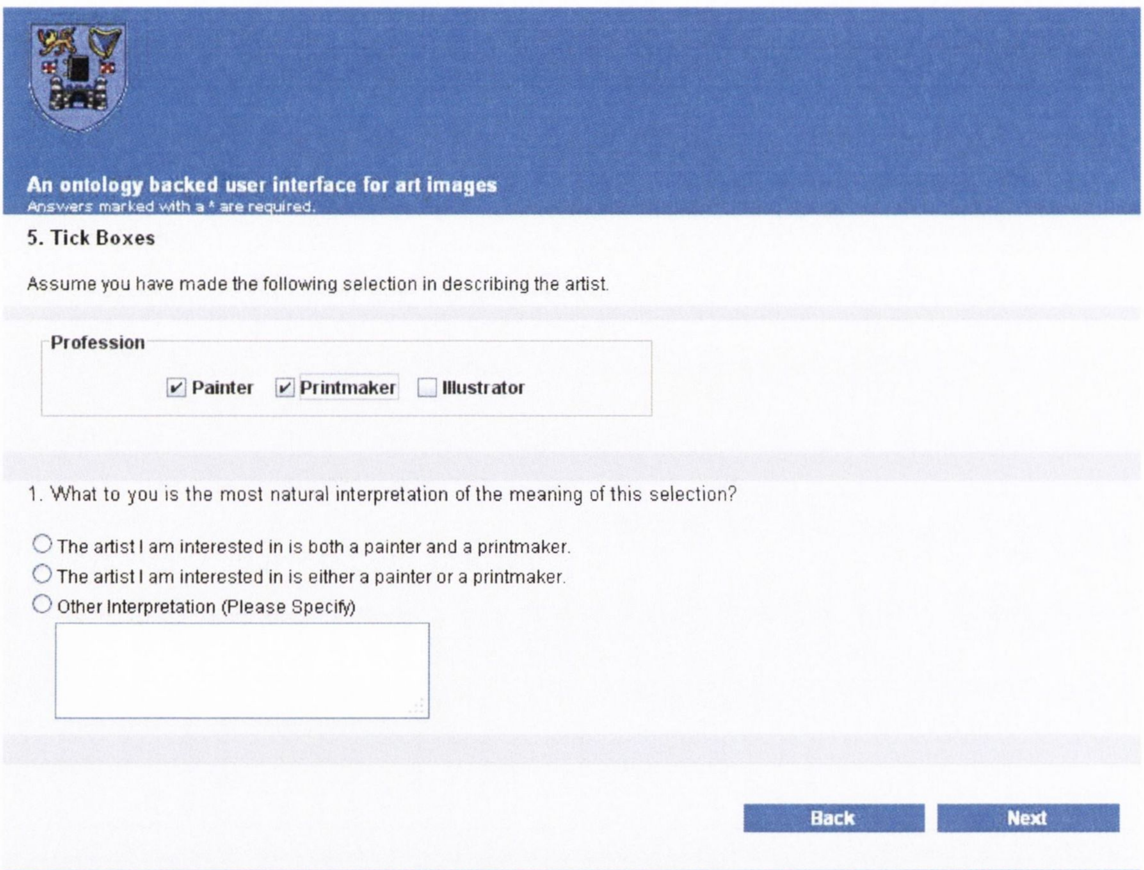

Figure F.5 – Page 5

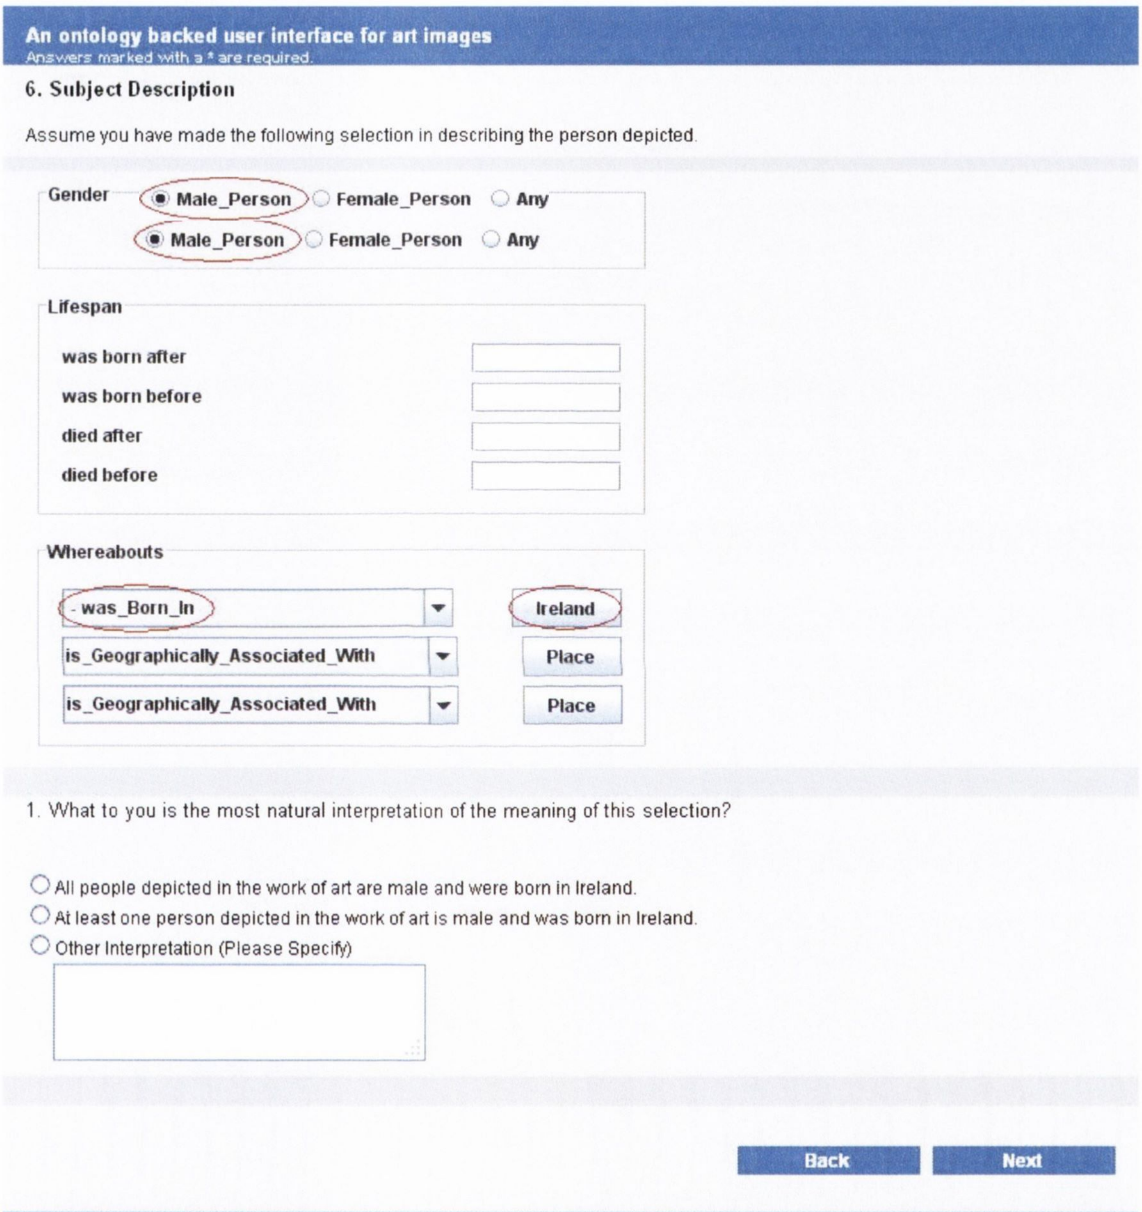

Figure F.6 – Page 6

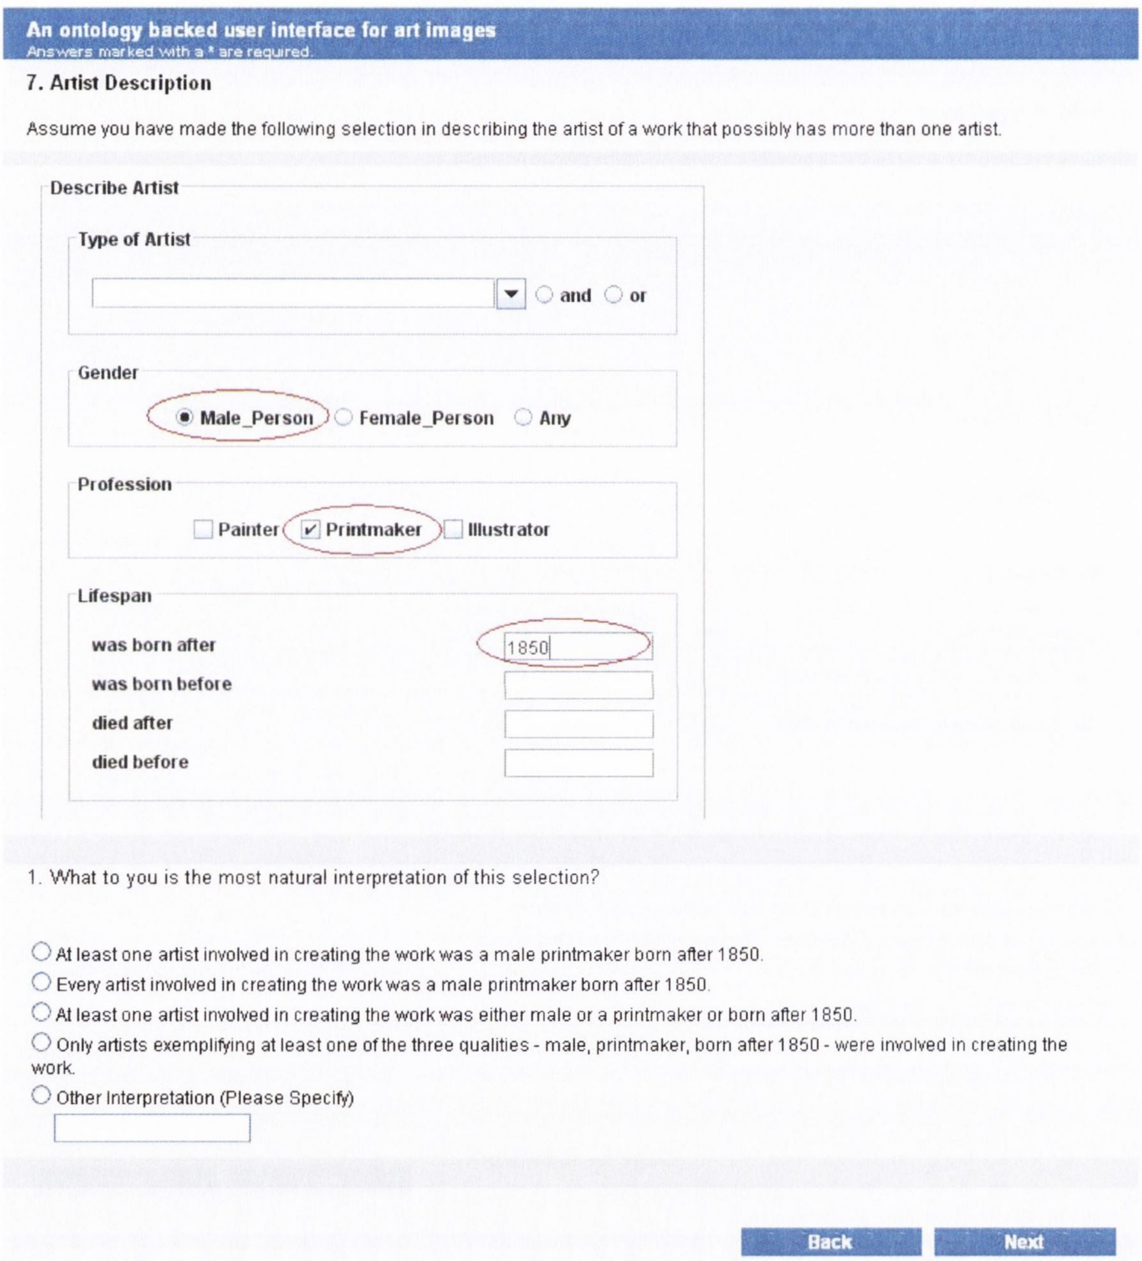

An ontology backed user interface for art images

Answers marked with a * are required.

7. Artist Description

Assume you have made the following selection in describing the artist of a work that possibly has more than one artist.

Describe Artist

Type of Artist

○ and ○ or

Gender

● Male_Person ○ Female_Person ○ Any

Profession

☐ Painter ☑ Printmaker ☐ Illustrator

Lifespan

was born after 1850

was born before

died after

died before

1. What to you is the most natural interpretation of this selection?

○ At least one artist involved in creating the work was a male printmaker born after 1850.

○ Every artist involved in creating the work was a male printmaker born after 1850.

○ At least one artist involved in creating the work was either male or a printmaker or born after 1850.

○ Only artists exemplifying at least one of the three qualities - male, printmaker, born after 1850 - were involved in creating the work.

○ Other Interpretation (Please Specify)

Back Next

Figure F.7 – Page 7

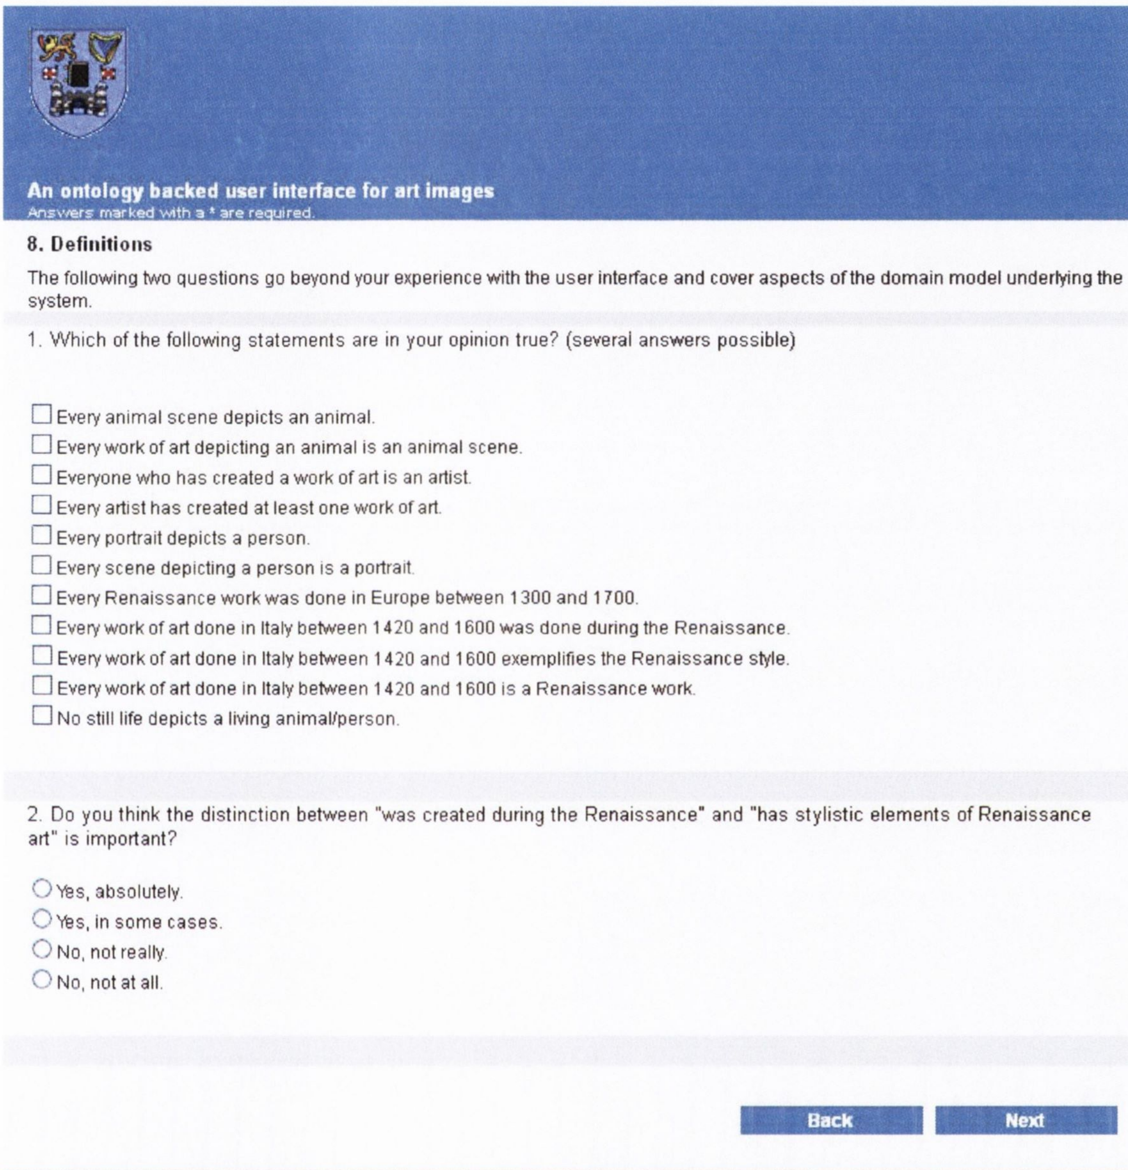

An ontology backed user interface for art images
Answers marked with a * are required.

8. Definitions

The following two questions go beyond your experience with the user interface and cover aspects of the domain model underlying the system.

1. Which of the following statements are in your opinion true? (several answers possible)

- ☐ Every animal scene depicts an animal.
- ☐ Every work of art depicting an animal is an animal scene.
- ☐ Everyone who has created a work of art is an artist.
- ☐ Every artist has created at least one work of art.
- ☐ Every portrait depicts a person.
- ☐ Every scene depicting a person is a portrait.
- ☐ Every Renaissance work was done in Europe between 1300 and 1700.
- ☐ Every work of art done in Italy between 1420 and 1600 was done during the Renaissance.
- ☐ Every work of art done in Italy between 1420 and 1600 exemplifies the Renaissance style.
- ☐ Every work of art done in Italy between 1420 and 1600 is a Renaissance work.
- ☐ No still life depicts a living animal/person.

2. Do you think the distinction between "was created during the Renaissance" and "has stylistic elements of Renaissance art" is important?

- ○ Yes, absolutely.
- ○ Yes, in some cases.
- ○ No, not really.
- ○ No, not at all.

Back | Next

Figure F.8 – Page 8

An ontology backed user interface for art images
Answers marked with a * are required.

9. Your experience with search engines

You have completed the main part of the questionnaire. The remaining questions are intended to give information about your background and experience in art history and some general demographic information about you.

1. How often do you use regular search engines (Google, Yahoo etc.) for image search?

- ○ Almost never.
- ○ Less than once a week.
- ○ More than once a week but not every day.
- ○ Almost every day.

2.
How often do you use advanced image search interfaces on museum web pages and art portals for image search?

- ○ Never.
- ○ Less than once a month.
- ○ More than once a month but less than once a week.
- ○ More than once a week but not every day.
- ○ Almost every day.

3. Have you ever used dedicated art image search systems or collection management software?

- ○ Yes
- ○ No
- ○ Maybe

4. If the answer to the previous question was YES:

Would you describe yourself as experienced in using these systems?

- ○ Yes, very experienced
- ○ Yes, somewhat experienced
- ○ No, not very much experienced

Back Next

Figure F.9 – Page 9

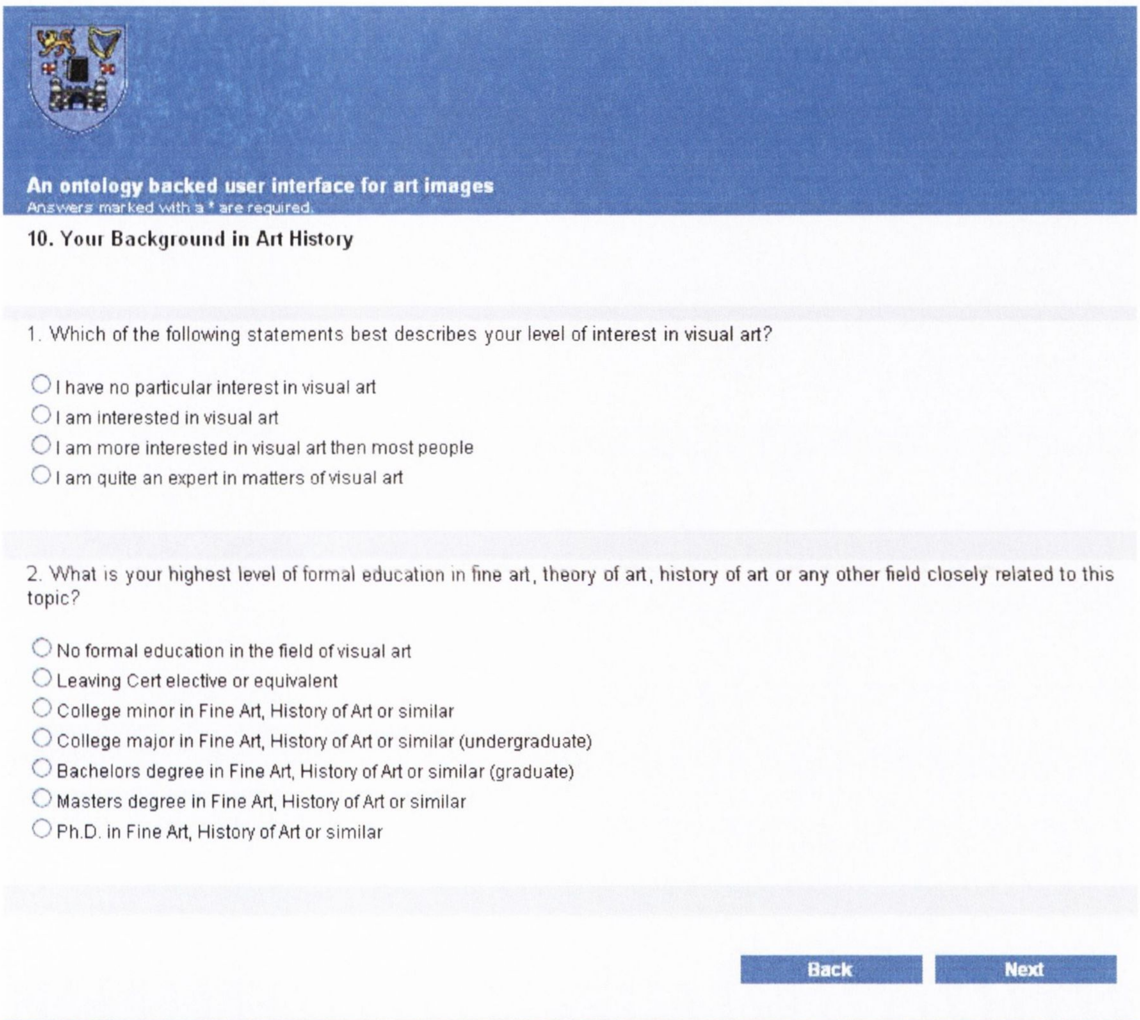
An ontology backed user interface for art images
Answers marked with a * are required.

10. Your Background in Art History

1. Which of the following statements best describes your level of interest in visual art?

- ◯ I have no particular interest in visual art
- ◯ I am interested in visual art
- ◯ I am more interested in visual art then most people
- ◯ I am quite an expert in matters of visual art

2. What is your highest level of formal education in fine art, theory of art, history of art or any other field closely related to this topic?

- ◯ No formal education in the field of visual art
- ◯ Leaving Cert elective or equivalent
- ◯ College minor in Fine Art, History of Art or similar
- ◯ College major in Fine Art, History of Art or similar (undergraduate)
- ◯ Bachelors degree in Fine Art, History of Art or similar (graduate)
- ◯ Masters degree in Fine Art, History of Art or similar
- ◯ Ph.D. in Fine Art, History of Art or similar

Back Next

Figure F.10 – Page 10

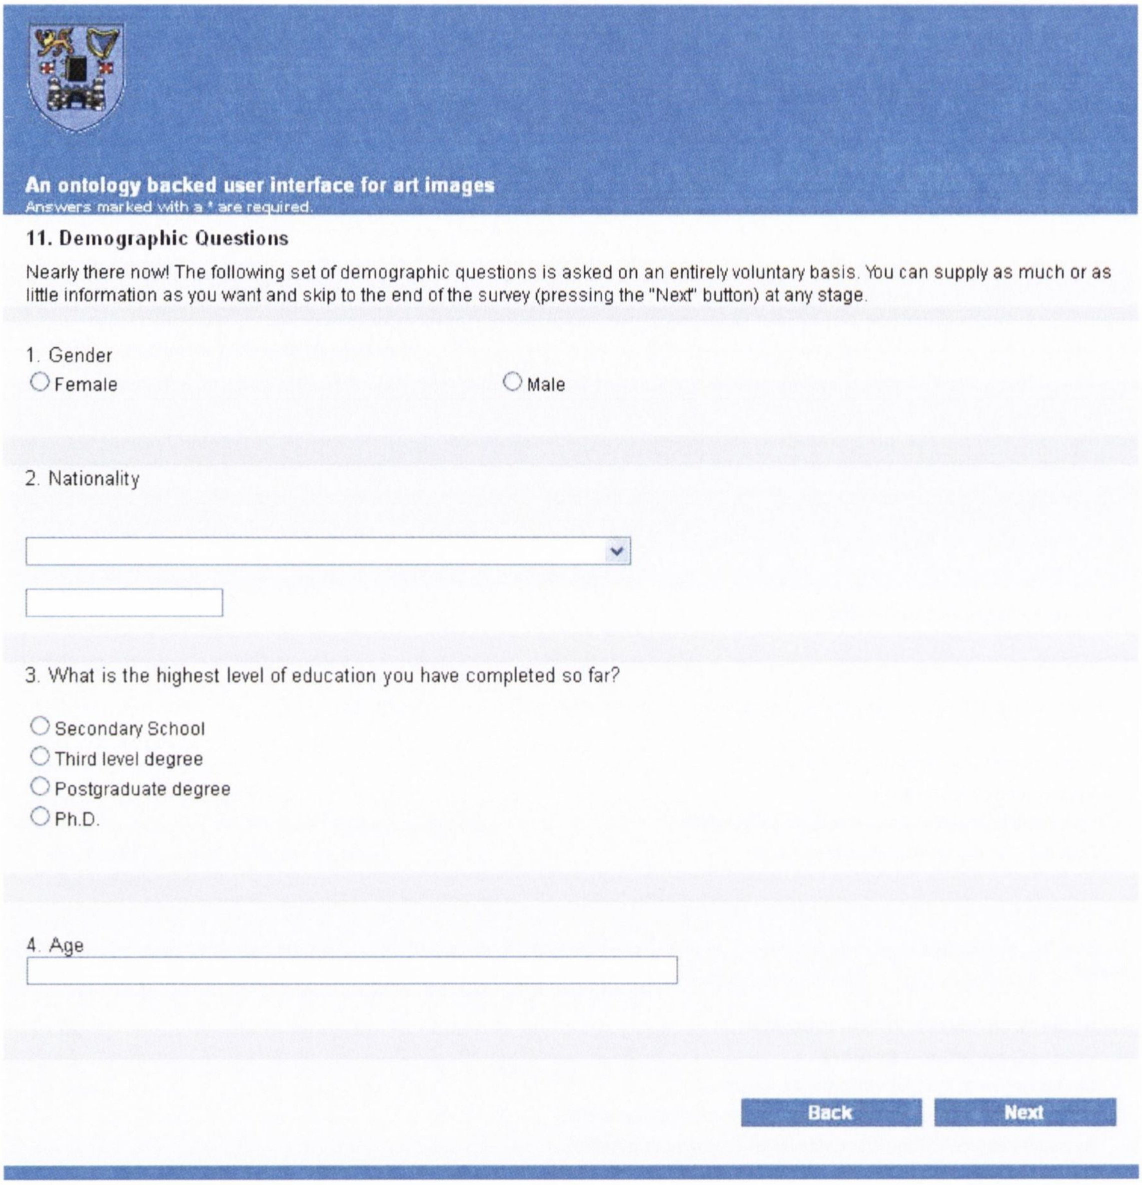

An ontology backed user interface for art images
Answers marked with a * are required.

11. Demographic Questions

Nearly there now! The following set of demographic questions is asked on an entirely voluntary basis. You can supply as much or as little information as you want and skip to the end of the survey (pressing the "Next" button) at any stage.

1. Gender
   - ○ Female
   - ○ Male

2. Nationality

3. What is the highest level of education you have completed so far?
   - ○ Secondary School
   - ○ Third level degree
   - ○ Postgraduate degree
   - ○ Ph.D.

4. Age

Back | Next

Figure F.11 – Page 11

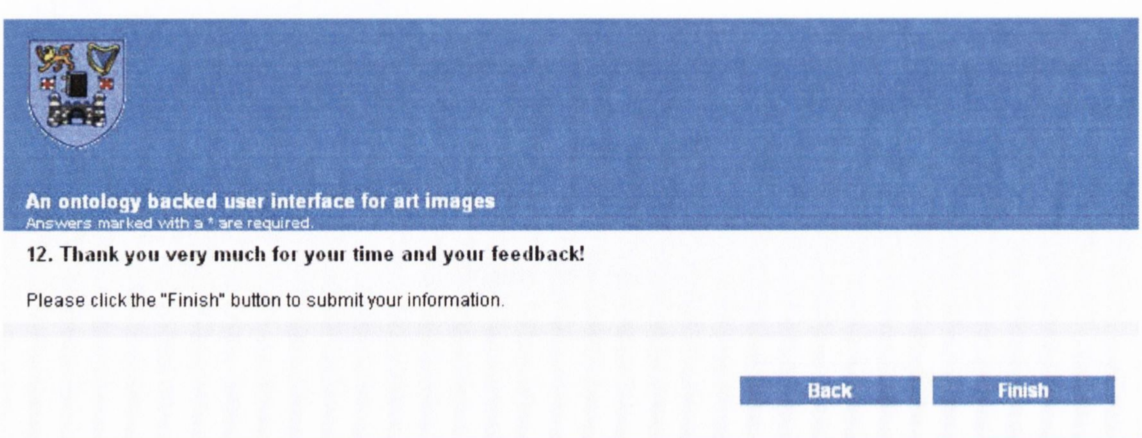

An ontology backed user interface for art images
Answers marked with a * are required.

12. Thank you very much for your time and your feedback!

Please click the "Finish" button to submit your information.

Back | Finish

Figure F.12 – Page 12

Bibliography

[1] ADDIS, M., MARTINEZ, K., LEWIS, P., STEVENSON, J., AND GIORGINI, F. New ways to search, navigate and use multimedia museum collections over the web. In *Museums and the Web 2005: Proceedings* (Toronto, March 31 2005), J. Trant and D. Bearman, Eds., Archives & Museum Informatics.

[2] ALLEMANG, D., AND HENDLER, J. *Semantic Web for the Working Ontologist.* Morgan Kaufmann, London, New York, Oxford, 2008.

[3] BAADER, F., AND HANSCHKE, P. A schema for integrating concrete domains into concept languages. In *Proc. of the 12th Int. Joint Conf. on Artificial Intelligence (IJCAI'91)* (1991), pp. 452–457.

[4] BAADER, F., AND HANSCHKE, P. Extensions of concept languages for a mechanical engineering application. In *Proc. of the 16th German Workshop on Artificial Intelligence (GWAI'92)* (Berlin, 1992), Lecture Notes in Computer Science LNCS 671, Springer, pp. 132–143.

[5] BAADER, F., HORROCKS, I., AND SATTLER, U. Description logics. In *Handbook on Ontologies*, S. Staab and R. Studer, Eds., second ed. Springer, Berlin, Heidelberg, 2009, pp. 21–43.

[6] BAKER, K., Ed. *Philosophical Dictionary.* Washington, 1972. An Englisch translation of [15].

[7] BENJAMINS, V. R., CASANOVAS, P., BREUKER, J., AND GANGEMI, A., Eds. *Law and the Semantic Web*, vol. LNAI 3369 of *Lecture Notes in Artificial Intelligence.* Springer, Berlin, Heidelberg, 2005.

[8] BEREZHNOY, I., POSTMA, E., AND VAN DEN HERIK, J. Computer analysis of Van Gogh's complementary colours. *Pattern Recogn. Lett. 28* (April 2007), 703–709.

[9] BERGER, J. *Ways of Seeing.* British Broadcasting Corporation and Penguin Books, London, 1972.

[10] BERGER, J. *About Looking.* Vintage Books, New York, September 1991. Originally published by Pantheon Books, 1980.

[11] BERNARAS, A., LARESGOITI, I., AND CORERA, J. Building and reusing ontologies for electrical network applications. In *Proceedings of the European Conference on Artificial Intelligence (ECAI'96)* (Chichester, United Kingdom, 1996), W. Wahlster, Ed., John Wiley and Sons, pp. 298–302.

[12] Berners-Lee, T., Hendler, J., and Lassila, O. The Semantic Web. *Scientific American 284*, 5 (May 2001), 34–43.

[13] Bizer, C., Lehmann, J., Kobilarov, G., Auer, S., Becker, C., Cyganiak, R., and Hellmann, S. DBpedia — a crystallization point for the web of data. *Journal of Web Semantics: Science, Services and Agents on the World Wide Web 7* (2009), 154–165.

[14] Bower, J. M., and Baca, M. *Union List of Artist Names*. G.K. Hall, New York, 1994.

[15] Brugger, W., Ed. *Philosophisches Wörterbuch*. 1967.

[16] Bunnin, N., and Yu, J., Eds. *The Blackwell Dictionary of Western Philosophy*. Blackwell Publishing, Malden, Oxford, Carlton, 2004.

[17] Burford, B., Briggs, P., and Eakins, J. P. A taxonomy of the image: on the classification of content for image retrieval. *Visual Communication 2*, 2 (2003), 123–161.

[18] Chandrasekaran, B., Josephson, J. R., and Benjamins, V. R. What are ontologies, and why do we need them? *IEEE Intelligent Systems 14*, 1 (January/February 1999), 20–26.

[19] Chen, H.-L. An analysis of image queries in the field of art history. *Journal of the Am. Soc. for Information Science and Technology 52*, 3 (2001), 260–273.

[20] Chen, H.-L. An analysis of image retrieval tasks in the field of art history. *Journal of the Am. Soc. for Information Science and Technology 37* (2001), 701–720.

[21] Chen, H.-L. A socio-technical perspective of museum practitioners' image-using behaviors. *Electronic Library 25*, 1 (2007), 18–35.

[22] Cimiano, P. *Ontology Learning and Population from Text*. Springer, New York, 2006.

[23] Couprie, L. D. Iconclass, a device for the iconographical analysis of art objects. *Museum 30*, 3/4 (1978), 194–198. Published by Unesco.

[24] Craig, E. Ontology. In *Routledge Encyclopedia of Philosophy*, E. Craig, Ed., vol. 7. London, New York, 1998, pp. 117–118.

[25] Crofts, N., Doerr, M., Gill, T., Stead, S., and Stiff, M. Definition of the CIDOC Conceptual Reference Model. Available at `http://www.cidoc-crm.org/docs/cidoc_crm_version_5.0.2.pdf`, January 2010. Version 5.0.2 of the CRM reference document (last accessed 01/08/2011).

[26] Crofts, N., Doerr, M., Gill, T., Stead, S., and Stiff, M. Definition of the CIDOC Conceptual Reference Model. Available at `http://www.cidoc-crm.org/docs/cidoc_crm_version_5.0.3.pdf`, May 2011. Version 5.0.3 of the CRM reference document (last accessed 20/03/2012).

[27] Crofts, N., Doerr, M., Gill, T., Stead, S., and Stiff, M. Definition of the CIDOC Conceptual Reference Model. Available at `http://www.cidoc-crm.org/docs/cidoc_crm_version_5.0.4.pdf`, November 2011. The current version 5.0.4 of the CRM reference document (last accessed 20/03/2012).

[28] CUTZU, F., HAMMOUD, R., AND LEYKIN, A. Distinguishing paintings from photographs. *Computer Vision and Image Understanding 100* (December 2005), 249–273.

[29] DAVEY, B. A., AND PRIESTLEY, H. A. *Lattices and Order*, second ed. Cambridge University Press, Cambridge, 2002.

[30] DE RYNCK, P. *How to Read a Painting: Decoding, Understanding and Enjoying the Old Masters*. Thames and Hudson, 2004.

[31] DEWEY, J. Ontology. In *Dictionary of Philosophy and Psychology*, J. M. Baldwin, Ed., vol. 2. London, 1902, pp. 203–204.

[32] DILLMAN, D. A. *Mail and Internet Surveys*, second ed. John Wiley & Sons, New York, 2000.

[33] DIMITROV, D. A., PUNDALEEKA, R., QASEM, A., AND HEFLIN, J. Isens: A system for information integration, exploration, and querying of multi-ontology data sources. In *2009 IEEE International Conference on Semantic Computing* (2009), pp. 330–335.

[34] DOERR, M. The CIDOC CRM — An ontological approach to semantic interoperability of metadata. *AI Magazine 24*, 3 (2003).

[35] DOERR, M. Ontologies for cultural heritage. In Staab and Studer [142], pp. 463–486.

[36] DONDIS, D. A. *A primer of visual literacy*. MIT Press, August 1974.

[37] DUYGULU, P., BARNARD, K., DE FREITAS, J. F. G., AND FORSYTH, D. A. Object recognition as machine translation: Learning a lexicon for a fixed image vocabulary. In *ECCV '02: Proceedings of the 7th European Conference on Computer Vision – Part IV* (2002), pp. 97–112.

[38] EAKINS, J. P. Towards intelligent image retrieval. *Pattern Recognition 35*, 1 (January 2002), 3–14.

[39] ENSER, P. The evolution of visual information retrieval. *Journal of Information Science 34*, 4 (2008), 531–546.

[40] ENSER, P. G. B., AND MCGREGOR, C. G. *Analysis of visual information retrieval queries*. British Library Research and Development Report, 6104. British Library, 1992.

[41] FELLBAUM, C., Ed. *WordNet — an electronic lexical database*. MIT Press, Cambridge, Massachusetts and London, England, 1998.

[42] FENSEL, D., Ed. *The semantic Web and its languages*, vol. 15 of *Intelligent Systems, IEEE*. Nov/Dec 2000, pp. 67–73.

[43] FENSEL, D., VAN HARMELEN, F., HORROCKS, I., MCGUINNESS, D. L., AND PATEL-SCHNEIDER, P. F. OIL: an ontology infrastructure for the Semantic Web. *Intelligent Systems, IEEE 16*, 2 (2001), 38–45.

[44] FIDEL, R. The image retrieval task: implications for the design and evaluation of image databases. *New Review of Hypermedia and Multimedia 3*, 1 (1997), 181–199.

[45] FLEW, A., Ed. *A Dictionary of Philosophy*. London, 1979.

[46] FÜRST, F., AND TRICHET, F. Heavyweight ontology engineering. In *On the Move to Meaningful Internet Systems 2006: OTM 2006 Workshops* (2006), R. Meersman, Z. Tari, and P. Herrero, Eds., Lecture Notes in Computer Science LNCS 4277, Springer, pp. 38–39.

[47] GANGEMI, A., GUARINO, N., MASOLO, C., OLTRAMARI, A., OLTRAMARI, R., AND SCHNEIDER, L. Sweetening ontologies with DOLCE. In *Proceedings of the 13th International Conference on Knowledge Engineering and Knowledge Management, EKAW '02. Ontologies and the Semantic Web.* (2002), Springer, pp. 166–181.

[48] GANTER, B., AND WILLE, R. *Formal Concept Analysis: Mathematical Foundations.* Springer, Berlin, 1999.

[49] GLASER, B., AND STRAUSS, A. *The discovery of grounded theory.* Aldine, Chicago, 1967.

[50] GÓMEZ-PÉREZ, A., FERNÁNDEZ-LÓPEZ, M., AND CORCHO, O. *Ontological Engineering.* Springer, London, 2004.

[51] GOMBRICH, E. H. *The Story of Art*, sixteenth ed. Phaidon, London, 1995.

[52] GRANT, A., AND COWTON, J. *SPECTRUM: UK Museum Documentation Standard.* Museum Documentation Association, 1997.

[53] GRAU, B. C., HORROCKS, I., MOTIK, B., PARSIA, B., PATEL-SCHNEIDER, P., AND SATTLER, U. OWL 2: The next step for OWL. *Web Semantics: Science, Services and Agents on the World Wide Web 6*, 4 (November 2008), 309–322.

[54] GRÜNINGER, M., AND FOX, M. S. The role of competency questions in enterprise engineering. In *Proceedings of the IFIP WG5.7 Workshop on Benchmarking – Theory and Practice* (Trondheim, Norway, 1994).

[55] GRÜNINGER, M., AND FOX, M. S. Methodology for the design and evaluation of ontologies. In *Proceedings of the IJCAI'95 Workshop on Basic Ontology Issues in Knowledge Sharing* (Montreal, 1995).

[56] GRUBER, T. R. A translation approach to portable ontology specifications. *Knowledge Acquisition 5*, 2 (1993), 199–220.

[57] GRUBER, T. R. Toward principles for the design of ontologies used for knowledge sharing. *Human-Computer Studies 43*, 5-6 (Nov./Dec. 1995), 907–928.

[58] GRUBER, T. R. Ontology. In *Encyclopedia of Database Systems*, L. Liu and M. T. Özsu, Eds. Springer, 2009.

[59] GÖRZ, G., SCHIEMANN, B., AND OISCHINGER, M. An implementation of the CIDOC Conceptual Reference Model (4.2.4) in OWL-DL. In *CIDOC 2008* (Athens, September 2008). Available at `http://www8.informatik.uni-erlangen.de/IMMD8/Services/cidoc-crm/docu/crm_owl_cidoc2008.pdf` (last accessed 20/01/2009).

[60] GUARINO, N., OBERLE, D., AND STAAB, S. What is an ontology? In Staab and Studer [142], pp. 1–17.

[61] GUARINO, N., AND WELTY, C. Evaluating ontological decisions with OntoClean. *Communications of the ACM 45*, 2 (2002), 61–65.

[62] GUARINO, N., AND WELTY, C. An overview of OntoClean. In Staab and Studer [142], pp. 201–220.

[63] HAARSLEV, V., HIDDE, K., MÖLLER, R., AND WESSEL, M. The RacerPro knowledge representation and reasoning system. *Semantic Web 1* (2011), 11 pages. Available at `http://users.encs.concordia.ca/~haarslev/publications/SW2011.pdf` (last accessed 4/05/2011).

[64] HARE, J. S., LEWIS, P. H., ENSER, P. G. B., AND SANDOM, C. J. Mind the gap: Another look at the problem of the semantic gap in image retrieval. In *Multimedia Content Analysis, Management and Retrieval 2006* (San Jose, California, USA, January 2006).

[65] HASTINGS, S. Query categories in a study of intellectual access to digitized art images. In *Proceedings of the 58th ASIS Annual Meeting* (Chicago, 9-12 October 1995), T. Kinney, Ed., pp. 3–8.

[66] HASTINGS, S. Evaluation of image retrieval systems: Role of user feedback. *Library Trends 48*, 2 (1999), 438–452.

[67] HEBELER, J., FISHER, M., BLACE, R., AND PEREZ-LOPEZ, A. *Semantic Web Programming*. Wiley Publishing Inc., Indianapolis, 2009.

[68] HENDLER, J. Agents and the Semantic Web. *IEEE Inteligent Systems 16*, 2 (March/April 2001), 30–37.

[69] HENSS, J., KLEB, J., GRIMM, S., AND BOCK, J. A database backend for OWL. In *Proceedings of the 6th International Workshop on OWL: Experiences and Directions (OWLED 2009)* (2009).

[70] HEPP, M. Ontologies: State of the art, business potential, and grand challenges. In *Ontology Management: Semantic Web, Semantic Web Services, and Business Applications*, M. Hepp, P. D. Leenheer, A. de Moor, and Y. Sure, Eds. Springer, New York, 2007, pp. 3–22.

[71] HOHMANN, G., AND SCHOLZ, M. Recommendation for the representation of the primitive value classes of the CRM as data types in RDF/OWL implementations. Available at `http://erlangen-crm.org/docs/crm-values-as-owl-datatypes.pdf` (last accessed 13/08/2011), February, 24 2011.

[72] HOLLINK, L., SCHREIBER, A. T., WIELEMAKER, B., AND WIELINGA, B. Semantic annotation of image collections. In *Proceedings of the KCAP'03 Workshop on Knowledge Markup and Semantic Annotation* (Florida, USA, October 2003).

[73] HOLLINK, L., SCHREIBER, A. T., WIELINGA, B. J., AND WORRING, M. Classification of user image descriptions. *International Journal of Human-Computer Studies 61*, 5 (March 2004), 601–621.

[74] HOLLINK, L., SCHREIBER, G., AND WIELINGA, B. Patterns of semantic relations to improve image content search. *Journal of Web Semantics 5* (2007), 195–203.

[75] HORRIDGE, M., AND BECHHOFER, S. The OWL API: A Java API for working with OWL 2 ontologies. In *Proceedings of OWL: Experiences and Directions 2009 (OWLED 2009)* (October 23–24 2009). Available at `http://www.webont.org/owled/2009/papers/owled2009_submission_29.pdf` (last accessed 8/09/2010).

[76] HORRIDGE, M., AND PATEL-SCHNEIDER, P. F. OWL 2 Web Ontology Language Manchester Syntax. Available at `http://www.w3.org/TR/2009/NOTE-owl2-manchester-syntax-20091027`, October 2009. (last accessed 3/04/2011).

[77] HORROCKS, I. DAML+OIL: a description logic for the Semantic Web. *IEEE Data Engineering Bulletin 25* (2002), 4–9.

[78] HORROCKS, I., FENSEL, D., BROEKSTRA, J., DECKER, S., ERDMANN, M., GOBLE, C., VAN HARLEMEN, F., KLEIN, M., STAAB, S., STUDER, R., AND MOTTA, E. The Ontology Inference Layer OIL. Tech. rep., Vrije Universiteit Amsterdam, 2000. Available at `http://xml.coverpages.org/OIL-inference.pdf` (last accessed 16/05/2011).

[79] HORROCKS, I., KUTZ, O., AND SATTLER, U. The even more irresistible $\mathcal{SROIQ}$. In *Proc. of the 10th Int. Conf. on Principles of Knowledge Representation and Reasoning (KR2006)* (June 2006), pp. 57–67.

[80] HORROCKS, I., AND SATTLER, U. A tableaux decision procedure for $\mathcal{SHOIQ}$. In *Proc. of 19th Int. Joint Conference on Artificial Intelligence (IJCAI 2005)* (Los Altos, 2005), Morgan Kaufmann.

[81] HORROCKS, I., SATTLER, U., AND TOBIES, S. Practical reasoning for expressive description logics. In *Proc. of the 6th Int. Conf. on Logic for Programming and Automated Reasoning (LPAR'99)* (1999), H. Ganzinger, D. McAllester, and A. Voronkov, Eds., Lecture Notes in Artificial Intelligence LNAI 1705, Springer, pp. 161–180.

[82] HYVÖNEN, E. Semantic portals for cultural heritage. In Staab and Studer [142], pp. 757–778.

[83] HYVÖNEN, E., MÄKELÄ, E., SALMINEN, M., VALO, A., VILJANEN, K., SAARELA, S., JUNNILA, M., AND KETTULA, S. MuseumFinland — Finnish museums on the semantic web. *Journal of Web Semantics 3*, 2 (2005), 224–241.

[84] HYVÖNEN, E., STYRMAN, A., AND SAARELA, S. Ontology-based image retrieval. In *Towards the semantic web and web services, Proceedings of XML Finland 2002 Conference* (October 21–22 2002), pp. 15–27.

[85] ISEMANN, D., AND AHMAD, K. Navigating cultural heritage in style. Sketching an ontological representation of metadata: the example of artistic period and style. *Museum Ireland 19* (2009), 149–155.

[86] ISEMANN, D., AND AHMAD, K. Art ontology: Experiences with CRM-based organization of heritage images. In *Proceedings of CIDOC 2011 – Knowledge management and museums* (Sibiu, Romania, 4th-9th September 2011).

[87] ISEMANN, D., AND AHMAD, K. Query terms for art images: A comparison of specialist and layperson terminology. In *Prodeedings of the 25th BCS Conference on Human-Computer Interaction* (Newcastle, UK, 4th-8th July 2011).

[88] ISEMANN, D., AN AIRCHINNIGH, M. M., AND AHMAD, K. Navigating heritage on well-grounded terminology and robust ontology. In *Specialized Language in Global Communication – Proceedings of the XVIth European Symposium on Language for Special Purposes (LSP), Hamburg (Germany) 2007* (Frankfurt am Main, 2010), W. von Hahn and C. Vertan, Eds., Sprache in der Gesellschaft. Beiträge zur Sprach- und Medienwissenschaft. Bd. 30, Peter Lang, pp. 78–89.

[89] JAIMES, A., AND CHANG, S.-F. A conceptual framework for indexing visual information at multiple levels. In *IS&T/SPIE Internet Imaging 2000* (San Jose, CA, January 2000).

[90] JEON, J., AND MANMATHA, R. Using maximum entropy for automatic image annotation. In *Image and Video Retrieval: Third International Conference, CIVR 2004, Dublin, Ireland, July 21-23, 2004, Proceedings* (2004), P. G. B. Enser, Y. Kompatsiaris, N. E. O'Connor, A. F. Smeaton, and A. W. M. Smeulders, Eds., Lecture Notes in Computer Science, 3115, Springer, pp. 24–32.

[91] JÖRGENSEN, C. Attributes of images in describing tasks. *Information Processing & Management 34*, 2/3 (1998), 161–174.

[92] KAMMERER, P., LETTNER, M., ZOLDA, E., AND SABLATNIG, R. Identification of drawing tools by classification of textural and boundary features of strokes. *Pattern Recogn. Lett. 28* (April 2007), 710–718.

[93] LAKOFF, G. *Women, fire, and dangerous things*. University of Chicago Press, Chicago, 1987.

[94] LAM, T., AND SINGH, R. Semantically relevant image retrieval by combining image and linguistic analysis. In *ISVC 2006: International Symposium on Visual Computing*, R. Boyle, B. Parvin, D. Koracin, A. Nefian, G. Meenakshisundaram, V. Pascucci, J. Zara, J. Molineros, H. Theisel, and T. Malzbender, Eds., vol. 4292 of *Lecture Notes in Computer Science*. Springer Berlin / Heidelberg, 2006, pp. 770–779.

[95] LASSILA, O., AND MCGUINNESS, D. L. The role of frame-based representation on the Semantic Web. Tech. Rep. KSL-01-02, Knowledge Systems Laboratory Stanford University, January 2001. Available at `http://www.ksl.stanford.edu/people/dlm/etai/lassila-mcguinness-fbr-sw.html` (last accessed 30/09/2010).

[96] LEE, H.-J., AND NEAL, D. A new model for semantic photograph description combining basic levels and user-assigned descriptors. *Journal of Information Science 36*, 5 (2010), 547–565.

[97] LENAT, D. B., AND GUHA, R. V. *Building Large Knowledge-Based Systems: Representation and Inference in the Cyc Project*. Addison-Wesley, January 1990.

[98] LEW, M. S., SEBE, N., DJERABA, C., AND JAIN, R. Content-based multimedia information retrieval: State of the art and challenges. *ACM Transactions on Multimedia Computing, Communications, and Applications 2*, 1 (Feb. 2006), 1–19.

[99] LIU, Y., ZHANG, D., LU, G., AND MA, W.-Y. A survey of content-based image retrieval with high-level semantics. *Pattern Recognition 40* (2007), 262–282.

[100] LOPES, D. M. Painting. In *The Routledge Companion to Aesthetics*, B. Gaut and D. M. Lopes, Eds. Routledge, London, New York, 2005, pp. 625–637.

[101] LOTZ, J. B. Ontologie. In Brugger [15]. An English translation can be found in [6].

[102] LOWE, E. J. Ontology. In *The Oxford Companion to Philosophy*, T. Honderich, Ed., second ed. Oxford University Press, Oxford, 2005, pp. 670–671.

[103] LUTZ, C. NEXPTIME-complete description logics with concrete domains. In *Proc. of the Int. Joint Conf. on Automated Reasoning (IJCAR 2001)* (Berlin, 2001), Lecture Notes in ArtiÞcial Intelligence LNAI 2083, Springer, pp. 45–60.

[104] LYU, S., ROCKMORE, D., AND FARID, H. A digital technique for art authentication. *Proceedings of the National Academy of Sciences of the United States of America 101*, 49 (December 7 2004), 17006–17010.

[105] MACINTYRE, A. Ontology. In *The Encyclopedia of Philosophy*, P. Edwards, Ed., vol. 5. Macmillan Publishers Ltd, New York, London, 1967, pp. 542–543.

[106] MAEDCHE, A., STAAB, S., STOJANOVIC, N., STUDER, R., AND SURE, Y. SEmantic portAL: The SEAL approach. In *Spinning the Semantic Web: Bringing the World Wide Web to Its Full Potential*, D. Fensel, J. Hendler, H. Lieberman, and W. Wahlster, Eds. MIT Press, 2005, pp. 317–359.

[107] MÄKELÄ, E., HYVÖNEN, E., AND SAARELA, S. Ontogator — a semantic view-based search engine service for web applications. In *Proceedings of the 5th International Semantic Web Conference (ISWC 2006)* (November 2006).

[108] MAUTNER, T. *The Penguin dictionary of philosophy*. Penguin, London, 1997.

[109] MCLELLAN, E., MACQUEEN, K. M., AND NEIDIG, J. L. Beyond the qualitative interview: Data preparation and transcription. *Field Methods 15*, 63 (2003), 63–84.

[110] MERRILL, G. H. Engineering a development platform for ontology-enhanced knowledge applications. In *Ontologies: A Handbook of Principles, Concepts and Applications in Information Systems*, R. Sharman, R. Kishore, and R. Ramesh, Eds. Springer, New York, 2007, pp. 777–822.

[111] MERRILL, G. H. Ontological realism: Methodology or misdirection? *Applied Ontology 5* (2010), 79–108.

[112] MOTIK, B., GRAU, B. C., AND SATTLER, U. Structured objects in OWL: Representation and reasoning. WWW 2008, April 21-25, 2008, Beijing, China.

[113] MOTIK, B., PATEL-SCHNEIDER, P. F., AND PARSIA, B. OWL 2 Web Ontology Language structural specification and functional-style syntax. Available at `http://www.w3.org/TR/2009/REC-owl2-syntax-20091027`, October 2009. (last accessed 24/03/2011).

[114] MOTIK, B., SHEARER, R., AND HORROCKS, I. Hypertableau reasoning for description logics. *Journal of Artificial Intelligence Research 36* (2009), 165–228.

[115] NECHES, R., FIKES, R., FININ, T., GRUBER, T., PATIL, R., SENATOR, T., AND SWARTOUT, W. R. Enabling technology for knowledge sharing. *AI Magazine 12*, 3 (1991), 36–56.

[116] NIELSEN, J. *Usability Engineering*. Morgan Kaufmann, San Francisco, 1993.

[117] NOWAK, S., AND DUNKER, P. A consumer photo tagging ontology: Concepts and annotations. In *THESEUS/ImageCLEF Pre-Workshop 2009, Co-located with the Cross-Language Evaluation Forum (CLEF) Workshop and 13th European Conference on Digital Libraries ECDL* (Corfu, Greece, 2009).

[118] NOY, N. F., AND MCGUINNESS, D. L. Ontology development 101: A guide to creating your first ontology. Tech. rep., Stanford Knowledge Systems Laboratory Technical Report KSL-01-05 and Stanford Medical Informatics Technical Report SMI-2001-0880, March 2001. Available at `http://www-ksl.stanford.edu/people/dlm/papers/ontology-tutorial-noy-mcguinness.pdf` (last accessed 8/09/2010).

[119] O'CONNOR, M. J., AND DAS, A. K. SQWRL: a query language for OWL. In *Proceedings of OWL: Experiences and Directions (OWLED), Fifth International Workshop* (Chantilly, VA, 2009). Available at `http://bmir.stanford.edu/file_asset/index.php/1538/BMIR-2009-1395.pdf` (last accessed 20/08/2011).

[120] ÖHGREN, A., AND SANDKUHL, K. Towards a methodology for ontology development in small and medium-sized enterprises. In *IADIS International Conference on Applied Computing 2005* (Algarve, Portugal, 2005).

[121] PANOFSKY, E. *Studies in Iconology*. Harper & Row, New York, 1972.

[122] PATEL-SCHNEIDER, P. F., HAYES, P., AND HORROCKS, I. OWL Web Ontology Language semantics and abstract syntax. Available at `http://www.w3.org/TR/2004/REC-owl-semantics-20040210`, February 2004. (last accessed 24/03/2011).

[123] PETERSEN, T., Ed. *Art & Architecture Thesaurus*. Oxford University Press, New York, Oxford, 1990.

[124] RECTOR, A., WELTY, C., NOY, N., AND WALLACE, E. Simple part-whole relations in OWL ontologies. Available at `http://www.w3.org/2001/sw/BestPractices/OEP/SimplePartWhole/simple-part-whole-relations-v1.5.html`, August 2005. (last accessed 24/03/2011).

[125] REED, P. A. *CIDOC Relational Data Model: A Guide*. ICOM/CIDOC, 1995.

[126] REESE, W. L. *Dictionary of Philosophy and Religion — Eastern and Western Thought*. Humanities Press, Atlantic Highlands, N. J., 1980.

[127] REYNOLDS, D., SHABAJEE, P., AND CAYZER, S. Semantic information portals. In *Proceedings of the 13th International Conference on World Wide Web - Alternate Track Papers & Posters* (New York, USA, May 17-20 2004), S. I. Feldman, M. Uretsky, M. Najork, and C. E. Wills, Eds., pp. 290–291.

[128] ROHRBAUGH, G. Ontology of art. In *The Routledge Companion to Aesthetics*, B. Gaut and D. M. Lopes, Eds. Routledge, London, New York, 2005, pp. 241–253.

[129] ROSCH, E. Principles of categorization. In *Cognition and Categorization*, E. Rosch and B. B. Lloyd, Eds. Lawrence Erlbaum, Hillsdale, New Jersey, 1978, pp. 27–48.

[130] SCHMIDT-SCHAUSS, M., AND SMOLKA, G. Attributive concept descriptions with complements. *Artificial Intelligence 48*, 1 (1991), 1–26.

[131] SCHREIBER, G. Web science: The digital-heritage case. In *SOFSEM 2010: Theory and Practice of Computer Science*, J. van Leeuwen, A. Muscholl, D. Peleg, J. Pokorný, and B. Rumpe, Eds., vol. 5901 of *Lecture Notes in Computer Science*. Springer Berlin / Heidelberg, 2010, pp. 108–111.

[132] SEGARAN, T., EVANS, C., AND TAYLOR, J. *Programming the Semantic Web*. O'Reilly, Sebastopol, CA, 2009.

[133] SEIDENBERG, J., AND RECTOR, A. Representing transitive propagation in OWL. In *Conceptual Modeling - ER 2006*, D. Embley, A. Olivé, and S. Ram, Eds., vol. 4215 of *Lecture Notes in Computer Science*. Springer Berlin / Heidelberg, 2006, pp. 255–266.

[134] SHATFORD, S. Analysing the subject of a picture: a theoretical approach. *Cataloging and Classification Quarterly 6*, 3 (1986), 39–62.

[135] SHEARER, R., MOTIK, B., AND HORROCKS, I. HermiT: A highly-efficient OWL reasoner. In *Proc. of the 5th Int. Workshop on OWL: Experiences and Directions (OWLED 2008 EU)* (Karlsruhe, Germany, October 26-27 2008).

[136] SHNEIDERMAN, B. Web science: a provocative invitation to computer science. *Commun. ACM 50* (June 2007), 25–27.

[137] SIDHU, A. S., DILLON, T. S., AND CHANG, E. Current status of biomedical ontologies: Developments in 2007. In *2nd IEEE International Conference on Digital Ecosystems and Technologies* (Phitsanulok, 26-29 February 2008), pp. 538–542.

[138] SIRIN, E., PARSIA, B., GRAU, B. C., KALYANPUR, A., AND KATZ, Y. Pellet: A practical OWL-DL reasoner. *Journal of Web Semantics 5*, 2 (2007), 51–53.

[139] SMEULDERS, A. W. M., WORRING, M., SANTINI, S., GUPTA, A., AND JAIN, R. Content-based image retrieval at the end of the early years. *IEEE Transactions on Pattern Analysis and Machine Intelligence 22*, 12 (Dec 2000), 1349–1380.

[140] SMITH, B. Ontology. In *The Blackwell Guide to the Philosophy of Computing and Information*, L. Floridi, Ed. Blackwell Publishing, Malden, MA and Oxford, 2004.

[141] SOWA, J. F. *Knowledge Representation: Logical, Philosophical, and Computational Foundations*. Brooks Cole Publishing Co., Pacific Grove, CA, 2000.

[142] STAAB, S., AND STUDER, R., Eds. *Handbook on Ontologies*, second ed. International Handbooks on Information Systems. Springer, Berlin, Heidelberg, 2009.

[143] STAAB, S., STUDER, R., SCHNURR, H.-P., AND SURE, Y. Knowledge processes and ontologies. *IEEE Intelligent Systems 16*, 1 (January/February 2001), 26–34.

[144] SUGUMARAN, V., AND STOREY, V. C. Ontologies for conceptual modeling: their creation, use, and management. *Data & Knowledge Engineering 42*, 3 (September 2002), 251–271.

[145] SUMMERS, D. *Real Spaces: World Art History and the Rise of Western Modernism*. Phaidon, New York, 2003.

[146] SWARTOUT, B., PATIL, R., KNIGHT, K., AND RUSS, T. Toward distributed use of large-scale ontologies. In *AAAI'97 Spring Symposium on Ontological Engineering* (Stanford University, California, 1997), A. Farquhar, M. Grüninger, A. Gómez-Pérez, M. Uschold, and P. van der Vet, Eds., pp. 138–148.

[147] SWARTOUT, W., AND TATE, A. Ontologies. *IEEE Intelligent Systems 14*, 1 (January/February 1999), 18–19.

[148] TARSKI, A. Der Wahrheitsbegriff in den formalisierten Sprachen. *Studia Philosophica 1* (1935), 261–405. Translation from Polish by L. Blaustein.

[149] THOMASSON, A. L. The ontology of art. In *The Blackwell Guide to Aesthetics*, P. Kirvy, Ed. Blackwell, Malden, MA, 2004, pp. 78–92.

[150] TOBIES, S. *Complexity results and practical algorithms for logics in knowledge representation*. PhD thesis, LuFG Theoretical Computer Science, RWTH-Aachen, Germany, 2001.

[151] TSARKOV, D., AND HORROCKS, I. Fact++ description logic reasoner: System description. In *Proc. of the Int. Joint Conf. on Automated Reasoning (IJCAR 2006)* (2006), Springer, pp. 292–297.

[152] USCHOLD, M. Building ontologies: Towards a unified methodology. In *Proceedings of Expert Systems '96, Cambridge* (Cambridge, December 16-18 1996).

[153] USCHOLD, M., AND GRÜNINGER, M. Ontologies: Principles, methods and applications. *Knowledge Engineering Review 11*, 2 (1996), 93–136.

[154] USCHOLD, M., AND GRÜNINGER, M. Ontologies and semantics for seamless connectivity. *SIGMOD Record 33*, 4 (December 2004), 58–64.

[155] USCHOLD, M., AND KING, M. Towards a methodology for building ontologies. In *Workshop on Basic Ontological Issues in Knowledge Sharing, held in conjunction with IJCAI-95* (1995).

[156] VAILAYA, A., JAIN, A., AND ZHANG, H. J. On image classification: City images vs. landscapes. *Pattern Recognition 31*, 12 (1998), 1921–1935.

[157] VAN HAGE, W. R., MALAISÉ, V., SEGERS, R. H., HOLLINK, L., AND SCHREIBER, G. Design and use of the Simple Event Model (SEM). *Journal of Web Semantics 9*, 2 (2011).

[158] VAN RIJSBERGEN, C. J. *Information Retrieval*, second ed. Butterworth, 1979.

[159] VELTMAN, K. Towards a Semantic Web for Culture. *Journal of Digital Information 4*, 4 (March 2004). Article No. 255, 2004-03-15.

[160] VROCHIDIS, S., DOULAVERAKIS, C., GOUNARIS, A., NIDELKOU, E., MAKRIS, L., AND KOMPATSIARIS, I. A hybrid ontology and visual-based retrieval model for cultural heritage multimedia collections. *International Journal of Metadata, Semantics and Ontologies 3*, 3 (2008), 167–182.

[161] WANG, H., LIU, S., AND CHIA, L.-T. Does ontology help in image retrieval?: a comparison between keyword, text ontology and multi-modality ontology approaches. In *Proceedings of the 14th annual ACM international conference on Multimedia* (New York, NY, USA, 2006), MULTIMEDIA '06, ACM, pp. 109–112.

[162] ZOGLAUER, T. *Einführung in die formale Logik für Philosophen*, fourth ed. Vandenhoeck & Ruprecht, Göttingen, 2008.