



DACURA: A NEW SOLUTION TO DATA HARVESTING AND KNOWLEDGE EXTRACTION FOR THE HISTORICAL SCIENCES

Journal:	<i>Historical Methods: A Journal of Quantitative and Interdisciplinary History</i>
Manuscript ID	VHIM-2017-0034.R2
Manuscript Type:	Original Article
Keywords:	data harvesting, RDF triplestore, database ontology, database metamodels, data curation

SCHOLARONE™
Manuscripts

1
2
3 **DACURA: A NEW SOLUTION TO DATA HARVESTING AND**
4
5 **KNOWLEDGE EXTRACTION FOR THE HISTORICAL SCIENCES**
6
7
8
9

10 Peter N. Peregrine, Lawrence University, 711 E. Boldt Way, Appleton WI 54911
11
12 and Santa Fe Institute, 1399 Hyde Park Road, Santa Fe NM, 87501
13
14 (peter.n.peregrine@lawrence.edu).[†]
15

16
17 Rob Brennan, ADAPT & Knowledge and Data Engineering Group, School of
18
19 Computer Science and Statistics, Trinity College Dublin, Ireland
20
21 (rob.brennan@scss.tcd.ie)
22

23
24 Thomas Currie, Department of Biosciences, University of Exeter—Penryn
25
26 Campus, Cornwall, TR10 9FE, UK (t.currie@exeter.ac.uk)
27

28
29 Kevin Feeney, Knowledge and Data Engineering Group, School of Computer
30
31 Science and Statistics, Trinity College Dublin, Ireland
32
33 (kevin.feeney@scss.tcd.ie)
34

35
36 Pieter François, School of Humanities, De Havilland Campus, University of
37
38 Hertfordshire, Hatfield, AL10 9EU, UK and Institute of Cognitive and
39
40 Evolutionary Anthropology, Oxford University, Oxford OX4 1QH, UK
41
42 (pieter.francois@stb.ox.ac.uk)
43

44
45 Peter Turchin, Department of Ecology and Evolutionary Biology, University of
46
47 Connecticut, 75 N. Eagleville Road, Storrs, CT 06269-3042
48
49 (peter.turchin@uconn.edu)
50

1
2
3 Harvey Whitehouse, Institute of Cognitive and Evolutionary Anthropology,
4
5 Oxford University, Oxford OX4 1QH, UK.

6
7 (harvey.whitehouse@anthro.ac.uk)
8
9

10
11
12 †Corresponding author
13
14

15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

For Peer Review Only

ABSTRACT

New advances in computer science address problems historical scientists face in gathering and evaluating the now vast data sources available through the Internet. As an example we introduce Dacura, a dataset curation platform designed to assist historical researchers in harvesting, evaluating, and curating high-quality information sets from the Internet and other sources. Dacura uses semantic knowledge graph technology to represent data as complex, inter-related knowledge allowing rapid search and retrieval of highly specific data without the need of a lookup table. Dacura automates the generation of tools to help non-experts curate high quality knowledge bases over time and to integrate data from multiple sources into its curated knowledge model. Together these features allow rapid harvesting and automated evaluation of Internet resources. We provide an example of Dacura in practice as the software employed to populate and manage the Seshat databank.

KEYWORDS

data harvesting; RDF triplestore; database ontology; database metamodels; data curation

1
2
3 Current developments in computer science provide new ways of
4
5 harvesting, storing, and retrieving data from the Internet that have the potential to
6
7 transform how literature reviews and data harvesting are undertaken in the
8
9 historical sciences. Dacura is a data curation platform that uniquely reflects three
10
11 of these developments: (1) a data-model based on a knowledge graph (as opposed
12
13 to the standard column and row data structure), (2) the use of Web Ontology
14
15 Language (OWL) for data definition, and (3) an automated process based on
16
17 semantic reasoning for weeding out the thousands of on-line and database hits not
18
19 directly related to a problem of interest and/or of dubious accuracy. Dacura was
20
21 built in tandem with the Seshat databank, which is being constructed to coordinate
22
23 quantitative historical and archaeological data to statistically test models of
24
25 historical dynamics (Turchin et al. 2015). We introduce both Dacura and Seshat
26
27 here as concrete examples of how the advances in computer science noted above
28
29 might be employed by historical researchers.
30
31
32
33
34

35 We begin with the basic problem the Dacura data curation platform is
36
37 intended to address: the overabundance of unevaluated information available to
38
39 researchers. As an example, consider a researcher who wants to gather
40
41 quantitative data to answer a specific question, let's say whether population
42
43 increased on island of Hawaii prior to the emergence of states ca. 1500 CE, as
44
45 would be predicted by a simple population-pressure model (e.g. Diamond 1997,
46
47 284-292). If she were to simply type "ancient Hawaii population" into Google,
48
49 she would obtain nearly 250,000 results (some discussing modern demographics)
50
51 with no easy way of knowing which of the many thousands of results on ancient
52
53
54
55
56
57
58
59
60

1
2
3 Hawaii would provide the information she needs, nor which of them would
4
5 provide reliable information (the Wikipedia page on “Ancient Hawaiian
6
7 Population”, for example, provides only high estimates and apparently from only
8
9 one source; the inability to clearly identify the source of the data is itself a serious
10
11 problem). If this researcher were to use Google Scholar instead, the results would
12
13 be fewer (around 165,000), and although she could expect somewhat better
14
15 quality, there would remain the daunting task of identifying papers and books
16
17 directly relevant to her interests. Even JSTOR, with quality-ensured content,
18
19 would proffer around 60,000 articles to churn through.
20
21
22

23
24 The example above illustrates a central problem in contemporary
25
26 historical research: the Internet and open-access publishing provide researchers
27
28 abundant information on virtually any topic of interest, but there is no quality
29
30 assurance, and even where quality can be assumed (as in peer-reviewed open-
31
32 access publications), the amount of information is often overwhelming. What is
33
34 needed is a tool that allows a researcher to build a dataset containing high-quality
35
36 structured information which provides specific pieces of information needed to
37
38 answer such questions. Such a search tool requires a carefully designed
39
40 hierarchical structure (ontology) to allow a scholar to easily dig down through
41
42 results to those that are directly relevant to his or her research. This search tool
43
44 also requires detailed indexing across result domains so that a search for Hawaii
45
46 population estimates, as our hypothetical researcher might perform, not only
47
48 recovers all information on Hawaii population estimates but does not also retrieve
49
50 unrelated information on other demographics or other locations. In short, such a
51
52
53
54
55
56
57
58
59
60

1
2
3 search tool must be able to apply an integrated thesaurus or set of thesauri as part
4
5 of the basic search routine.
6

7
8 There are quite a few extant search tools that provide this functionality:
9
10 rapid retrieval of specific, quality information across domains. Considering only
11
12 archaeological search tools (the first author is an archaeologist), a good example
13
14 is eHRAF (Human Relations Area Files; hraf.yale.edu) which provides two
15
16 archives of documents (ethnographic and archaeological, respectively) organized
17
18 using detailed ontologies (the *Outline of World Cultures* and *Outline of*
19
20 *Archaeological Traditions*) and employing a rich thesaurus (the *Outline of*
21
22 *Cultural Materials*). Individual paragraphs from nearly three-quarters of a
23
24 million pages of archaeological and ethnographic primary and secondary source
25
26 documents are indexed in eHRAF and can be easily searched and retrieved at
27
28 varying levels of detail using hierarchical and Boolean search strategies. The
29
30 results are specific, of excellent quality and specificity, and manageable in
31
32 number. However, the range of results is limited to the documents that have been
33
34 included in the eHRAF archives. The reason eHRAF provides such excellent
35
36 information retrieval is that the information has been extensively pre-processed to
37
38 the extent that every document has been individually placed into the ontology and
39
40 every paragraph in every document individually indexed by Ph.D.-holding
41
42 anthropologists. In short, a huge amount of work is required to make search and
43
44 retrieval of high-quality data easy, and that means the data provided by eHRAF
45
46 grows slowly and eHRAF must charge a user fee.
47
48
49
50
51
52
53
54
55
56
57
58
59
60

1
2
3 An alternative model of a search tool providing rapid retrieval of specific,
4 high-quality information across domains is tDAR (the Digital Archaeological
5 Record; www.tdar.org). Like eHRAF, entire documents (including raw datasets,
6 shapefiles, and the like) are available through tDAR, and are organized within a
7 basic ontology. Unlike eHRAF, these documents are not processed by tDAR staff
8 (although there is review of the processing to ensure it has been done correctly),
9 but rather the individuals who submit documents complete a metadata form which
10 is attached to the document (Watts 2011). This allows the number of documents
11 in tDAR to increase relatively rapidly, and also allows free retrieval of
12 information from tDAR (there are modest fees for contributing documents).
13
14 However, because contributors provide the ontological and indexing information
15 themselves, the level of detail and accuracy vary, meaning that searches may not
16 retrieve all relevant documents. And, like eHRAF, the available information is
17 limited to the documents within the database.

18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35 Open Context (www.opencontext.org) is another excellent data repository
36 that is similar to tDAR, but which provides several additional features that expand
37 its range beyond archaeological data. Like tDAR, archaeological data are
38 contributed for a modest fee. Unlike tDAR, Open Source editors work with
39 contributors to create the metadata and clean the data sources for publication on
40 the web, and the data sources themselves are evaluated for their importance; that
41 is, not all data sources are published, but only those that peer reviewers think will
42 be of use to the broader field. Once incorporated into Open Context, data sources
43 are linked to related data sources on the web by Linked Data standards (Kansa
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

1
2
3 2010). This allows Open Context to expand beyond the archived data,
4
5 overcoming a limitation of both eHRAF and tDAR while providing an assurance
6
7 of data quality for those data contributed directly to Open Context, though not to
8
9 linked data, which we see as a serious limitation.
10

11
12 We present here what we argue is a more comprehensive approach to the
13
14 problem of retrieving specific, high-quality information across domains than the
15
16 three outlined above (and there are many other excellent programs and data
17
18 repositories we could have cited)—a set of data harvesting, evaluation, assembly,
19
20 and output processes that has been implemented in Dacura (dacura.cs.tcd.ie) and
21
22 which itself is being employed as the managing software for the Seshat databank
23
24 (seshatdatabank.info) which we introduce below. By being developed and
25
26 implemented in tandem with a data-heavy research initiative, Dacura has
27
28 benefitted from the ongoing identification of problems and shortcomings that
29
30 gathering and managing large and complex historical data entail, and thus serves
31
32 as a good example of a resource that would be of use to academic researchers.
33
34
35
36

37
38 Knowledge graph technology is increasingly being used by companies
39
40 such as Google (N.d.) and Facebook (N.d.) to manage and structure the vast and
41
42 diverse sources of information they curate. Traditional SQL storage solutions
43
44 based on tables and rows are insufficiently expressive to capture the structure or
45
46 semantics of the information that they manage and the complexity of the
47
48 relationships between things. However, to date, knowledge graph technology
49
50 remains very technical and difficult to use except by very large and technically
51
52 savvy corporations. Dacura's design goal is to help make knowledge graph
53
54
55
56
57

1
2
3 technology available to researchers in such a way that it does not require
4
5 tremendous expertise to use. Dacura achieves this by automating much of the
6
7 software required to harvest and curate data from the semantic model, providing
8
9 users with easy to use tools and interfaces which do not require them to
10
11 understand the underlying technology.
12
13

14 We do not intend this article to be simply an advertisement for Dacura, but
15
16 rather we use Dacura to illustrate an approach to harvesting, evaluating, and
17
18 retrieving data from the Internet or any “big data” source that has been made
19
20 possible by new advances in computer science and that we believe will have
21
22 profound impact on the historical sciences.
23
24
25
26
27

28 **DACURA**

29
30
31
32

33 Dacura is a data curation platform designed to assist historical researchers
34
35 in creating and curating high-quality datasets in the form of rich semantic
36
37 knowledge graphs [Endnote 1]. The basic idea is simple—the researcher starts by
38
39 defining the structure of the data that they would like to collect. The system uses
40
41 this information to support the user in discovering, harvesting, filtering, correcting,
42
43 refining and analyzing information from the Internet in order to compile the
44
45 highest quality information possible. The details provided by the researcher
46
47 include basic information such as the definition of the fundamental entities of
48
49 interest (e.g. Hawaii), the properties of those entities in which she or he is
50
51 interested (e.g. population estimates), the datatypes and desired units of each
52
53
54
55
56
57
58
59
60

1
2
3 property (e.g. numeric), and relationships with other entities both within and
4
5 beyond the dataset itself (e.g. Hawaii is in Polynesia).
6

7
8 The process of defining the structure of the desired data is one of the
9
10 strengths of this approach. All historical scientists know that data must be
11
12 property questioned before they can speak. Because Dacura requires the precise
13
14 shape of the desired data to be specified before performing a search, Dacura
15
16 encourages researchers to think carefully about the nature of the desired data and
17
18 the means by which they will be questioned before starting to collect those data.
19
20 Such preparation before data collection saves time and effort in what is often the
21
22 most difficult task in historical research—identifying useful information sources.
23
24 Dacura’s simple and user friendly interface (discussed below) makes the process
25
26 of data definition easy, and, because Dacura provides a flexible search structure,
27
28 this process can be iterative, changing as data are interrogated and as questions
29
30 become more focused.
31
32
33
34

35 Dacura encodes the structure of the dataset defined by the researcher as a
36
37 semantic web ontology according to the Web Ontology Language (OWL)
38
39 standard of the World Wide Web Consortium (W3C), the main international
40
41 standards body for the web. OWL is a rich and flexible language which allows a
42
43 wide variety of constraints and inference rules to be specified on the data to be
44
45 collected (e.g. the population of a town should not be greater than the population
46
47 of the region that it is in). In contrast with the unstructured natural language
48
49 strings that drive most search engine results, the highly structured and precisely
50
51 specified nature of ontological dataset specifications can be exploited by the
52
53
54
55
56
57
58
59
60

1
2
3 computer to provide much greater specificity in results. The richer the structural
4
5 specification, the easier it is for the system to automate the harvesting of data and
6
7 the generation of useful tools with which to analyze, improve and curate it over
8
9 time.
10

11
12 Dacura is based on semantic web technology. At its core is a Resource
13
14 Description Framework (RDF) triplestore, a specific form of graph database (as
15
16 opposed to a two-dimensional column and row database used in most
17
18 spreadsheets) in which data are identified by a subject-predicate-object
19
20 combination like “Hawaii is Polynesia”, “Hawaii has Island”, or “Polynesia has
21
22 Island” (www.w3.org/TR/rdf11-concepts/). The subject-predicate-object
23
24 structure can be understood as nodes-edges-properties within a three-dimensional
25
26 graph which represents and stores data. The graphic structure of an RDF
27
28 triplestore allows for index-free adjacency, meaning every subject-predicate-
29
30 object triple directly links to related subject-predicate-object triples so that no
31
32 index lookups are necessary. In the example above, Polynesia, Hawaii, and Island
33
34 are all linked so that no indexed search is required to identify Hawaii as a
35
36 Polynesian Island.
37
38
39
40

41
42 OWL ontologies are used in Dacura to enable semantic reasoning in
43
44 quality control and data harvesting; that is, if there are conflicts between triples
45
46 Dacura identifies and marks them as conflicts for further evaluation (see
47
48 dacura.scss.tcd.ie/ontologies/dacura-130317.ttl). Dacura is designed to produce
49
50 and consume data in line with the linked open data principles. This makes it easy
51
52 to import information from existing structured information sources and to enrich
53
54
55
56
57

1
2
3 curated datasets by interlinking them with publicly available Linked Data sources
4 (e.g. DBpedia or wikidata, the linked data versions of Wikipedia) and datasets
5
6 curated by Dacura can similarly be easily linked. A video example of Dacura's
7
8 operation is provided at <https://youtu.be/AEb1wF3jAgk>. A key aspect of Dacura
9
10 is that harvested data, including those harvested through Linked Data, are
11
12 evaluated for quality through both automated and human evaluators as part of the
13
14 system's workflow. Thus Dacura not only harvests data quickly and easily, it also
15
16 evaluates it for quality.
17
18
19
20

21 The Dacura workflow breaks the process of data creation and curation
22
23 down into 4 stages, as illustrated in Figure 1. The first stage is data harvesting:
24
25 identifying sources of high-quality information with which to populate the dataset.
26
27 Dacura supports a number of approaches to data harvesting: from identifying
28
29 relevant data in known public data sources, to deploying automated agents to
30
31 search the Internet, to manual specification of information sources by researchers.
32
33 The goal of the system is to automate, as much as possible, the identification of
34
35 the sources of information that will be needed in order to populate the dataset. In
36
37 this stage, the goal is not to find documents about the entities in which one is
38
39 interested, but to find specific sources of information which can populate the
40
41 properties and relationships that a researcher has defined in their data
42
43 specification.
44
45
46
47
48

49 [Figure 1 about here]

50
51 The second stage in the Dacura data creation and curation process is
52
53 knowledge extraction. This involves extracting the precise information from
54
55
56
57

1
2
3 harvested sources into the structure required by the researcher's data specification.
4
5 Although Natural Language Processing and other artificial intelligence
6
7 technologies continue to improve all the time, they remain error prone and thus, in
8
9 order to ensure high quality data, some human input is normally required to filter
10
11 out false positives. Dacura employs tools to support both human users and
12
13 automated agents in screening, filtering, improving, annotating and interlinking
14
15 candidate records to produce knowledge reports; that is, authoritative accounts of
16
17 the relevant knowledge contained in a source, enriched through links into the web
18
19 of data.
20
21
22
23

24 The third stage in the Dacura process is perhaps the most important for
25
26 ensuring data quality: expert analysis. One of Dacura's strengths is its focus on
27
28 data quality, providing both automated and manual tools to ensure that harvesting
29
30 captures accurate and complete information that conforms to the researcher's data
31
32 specifications. Initial data evaluation is performed through automated tools,
33
34 which use semantic consistency checking and validity testing to reconcile various
35
36 data points into a composite account that represent the tools' best estimate of
37
38 authoritative data which accurately represents reality. These composite accounts
39
40 are reviewed by experts in the data domain (like our hypothetical researcher
41
42 interested in Hawaii population estimates), allowing the expert to correct
43
44 misinterpretations and identify disagreements between the expert and the
45
46 automated tools. Experts can create their own personal interpretations (for
47
48 example, by specifying that only particular sources should be trusted) and overlay
49
50
51
52
53
54
55
56
57
58
59
60

1
2
3 this on the dataset to produce a custom dataset, representing their view on what
4
5 the data should be.
6

7
8 In the case of Seshat experts are solicited by the Seshat management team
9
10 based on their publication record in the area(s) they are asked to review. All have
11
12 Ph.D.s in their field and the vast majority have academic appointments. Experts
13
14 are periodically evaluated to identify individuals whose input appears inconsistent
15
16 with harvested data on a regular basis. One might be concerned that developing a
17
18 pool of experts by solicitation might not be effective, but the quality of those
19
20 currently volunteering can be seen in the list at [http://seshatdatabank.info/seshat-](http://seshatdatabank.info/seshat-about-us/contributor-database/)
21
22 [about-us/contributor-database/](http://seshatdatabank.info/seshat-about-us/contributor-database/). The number (77 at the time of this writing) and
23
24 range of expertise of these volunteers illustrates that it is quite feasible to
25
26 incorporate expert evaluation into a data harvesting system like Dacura. Most
27
28 projects employing systems like Dacura will not likely be as large in scope as
29
30 Seshat, and only one expert may be required to evaluate the data—perhaps only
31
32 the researcher herself.
33
34
35
36

37
38 Finally, Dacura supports a variety of output tools to make data available to
39
40 third parties in a range of formats. Dacura publishes its curated data as Linked
41
42 Data so other users and platforms (such as Open Context, discussed above) can
43
44 readily access it. Dacura also provides a SPARQL endpoint—a query language
45
46 for RDF graphs—which supports sophisticated filtering and retrieval of data.
47
48 This allows intelligent applications to interact with the data in unforeseen ways.
49
50 For human users, Dacura can produce graphs, charts, maps and other
51
52 visualizations to provide users with easy-to-understand insights into the data.
53
54
55
56
57

1
2
3 Data for graphs or other outputs can be browsed, searched, and selected providing
4
5 users with the ability to access the sections of datasets they find most useful.
6

7
8 Dacura also allows datasets or subsets of them to be exported in a wide range of
9
10 formats for external analysis, including geographic information systems and
11
12 statistical packages such as SPSS and R.
13
14

15 16 17 **IMPLEMENTING DACURA: THE SESHAT META-MODEL** 18

19
20
21 As an example of how Dacura works in practice, Figure 2 shows the meta-model
22
23 being used to implement Seshat: Global Historical Databank (Turchin et al. 2015,
24
25 Francois et al., 2016; see also dacura.scss.tcd.ie/ontologies/seshat-130317.ttl).
26
27 Seshat (seshatdatabank.info) was designed to bring together into one place a
28
29 comprehensive body of knowledge about human history and prehistory for the
30
31 purpose of empirically testing hypotheses about cultural evolution, including the
32
33 possible role of religion, ritual, warfare, agriculture and other clusters of variables
34
35 in the rise of social complexity. To date Seshat has been used to identify a single
36
37 dimension of complexity that explains about three-quarters of the variation in
38
39 human social organization (Turchin et al. 2018) and has also been used to
40
41 demonstrate that the hypothesized socio-political transformations of the Axial
42
43 Age took place at varied times across Eurasia, encompassing more than a
44
45 millennium (Mullins et al. 2018). Exploring such large-scale questions with
46
47 appropriate statistical techniques requires data that are both valid and reliable; that
48
49
50
51
52
53
54
55
56
57
58
59
60

1
2
3 is, data that define what the researcher thinks they define and that are measured in
4
5 the same manner across cases.
6

7
8 Dacura, also under development when Seshat was being planned, seemed
9
10 an excellent platform for Seshat. The two groups of researchers decided to
11
12 develop Dacura and Seshat in tandem, so that both would inform the other.
13
14 Because computer scientists and historical scientists do not regularly work
15
16 together, the cooperation between the Dacura and Seshat research teams proved
17
18 extremely fruitful, allowing the Dacura team to understand the needs of historical
19
20 scientists, and the Seshat team to understand both the possibilities and limitations
21
22 of triplestore data harvesting and curation.
23
24

25
26 [Figure 2 about here]
27

28
29 There are two fundamental pieces of information upon which Seshat cases
30
31 are based: a **Location** and a **Duration**. A location is a point or polygon anywhere
32
33 on the earth's surface, and defines an entity called a **Territory**. Three entity
34
35 classes of Territory have been defined in Seshat (more may be defined later as
36
37 Seshat expands):
38

- 39
40 (1) *Natural-Geographic Areas* (NGA), which are a contiguous area roughly
41
42 100 by 100 kilometers encompassing a reasonably homogenous ecological
43
44 region.
45
46 (2) *Biomes*, which encompass a contiguous biotic region or region of similar
47
48 climatic conditions.
49
50 (3) *World Regions*, which may be pre-defined entities such as nations or states,
51
52 or can be defined by other specific criteria.
53
54
55
56
57
58
59
60

1
2
3 A Duration can be a single date or a date range. Adding a Duration to a Territory
4 entity class defines one of two temporally bounded entities: (1) a **Human**
5 **Population**, which is group of humans in a defined territory during a specified
6 period of time; and (2) an **Event**, defined as an occurrence taking place in a
7 specific territory in a specific period of time.
8
9

10
11
12
13
14 Seshat provides the ability to create entity classes within Human
15 Populations and Events for specific research questions. Within the Human
16 Population entity, current entity classes are:
17
18
19

20
21 (1) *Tradition*, which is defined as a human population “sharing similar
22 subsistence practices, technology, and forms of socio-political
23 organization that are spatially contiguous over a relatively large area and
24 which endure temporally for a relatively long period of time” (Peregrine
25 and Ember, 2001, ix). For this entity class there is a formal sampling
26 universe for selecting cases, the *Outline of Archaeological Traditions*
27 (hereafter OAT) (hraf.yale.edu/online-databases/ehraf-archaeology/outline-of-archaeological-traditions-oat/) and a formal
28 thesaurus for coding data, the *Outline of Cultural Materials* (hereafter
29 OCM) (hraf.yale.edu/online-databases/ehraf-world-cultures/outline-of-cultural-materials/).

30
31
32
33
34
35
36
37
38
39
40 (2) *Cultural Group*, which is a human population sharing norms, beliefs,
41 behaviors, values, attitudes, etc. The primary sampling universe for this
42 entity class is the *Outline of World Cultures* (hereafter OWC)(Murdock,
43 1983) and the thesaurus is the OCM.
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

1
2
3 (3) *Polity*, which is a human population that is a politically independent unit
4
5 with a shared system of governance. This is an example of an entity class
6
7 created for a specific research project. The sample consists of 30 cases
8
9 selected for characteristics of sociopolitical organization and geographic
10
11 location (Turchin et al., 2015). The primary thesaurus for this entity class
12
13 is the OCM.
14
15

16
17 (4) *Settlement*, is a human population in a physical location and material
18
19 facilities ranging in size and complexity from a temporary camp to a great
20
21 metropolis. Because of the great range of settlements that could be coded,
22
23 there is no defined sampling universe for the entity class. The primary
24
25 thesaurus is again the OCM.
26
27

28
29 (5) *Identity Group*, which is a human population with a shared sense of being
30
31 part of the same group. Like *Polity*, this entity class was created for a
32
33 specific set of research projects and the sample is opportunistic (see
34
35 Whitehouse, Francois, and Turchin, 2015). There is no formal thesaurus,
36
37 though the OCM is used for some domains.
38
39

40
41 (6) *Linguistic Group*, which is a human population with a common language.
42
43 The sampling universe for this entity class is *Ethnologue*
44
45 (www.ethnologue.com), but there is no formal thesaurus (again, the OCM
46
47 is being used for some domains).
48

49 In addition, subclasses can be added to entity classes to provide for more specific
50
51 sets of data. Figure 3 shows entity subclasses that have been created for the
52
53 current entity classes listed above.
54
55

[Figure 3 about here]

The Event entity obviously encompasses an almost infinite range of possible entity classes and subclasses. To maintain some order the event class in DBpedia is used (mappings.dbpedia.org/server/ontology/classes/) as a basic ontology. As shown in Figure 2, the current entity classes for the Event entity include:

- (1) *Inter-group Conflict*, such as a war, a battle, a feud, or the like.
- (2) *Socio-Natural Disaster*, such as a famine, or epidemic.
- (3) *Natural Disaster*, such as a drought, a flood, an infestation, a volcanic eruption, etc.
- (4) *Societal Collapse*
- (5) *Transition Ritual*, such as a marriage, a coronation, or an initiation.
- (6) *Social Movement*, including physical movements like migration, but also social movements such as revitalization, millenarianism, strikes, etc.
- (7) *Technological*, such as inventions, discoveries, innovations, and the like.

Populating Seshat: The Dacura Workflow Model

As an example of how an historical scientist might employ Dacura to populate a dataset, Figure 4 illustrates how data for the Tradition entity class is incorporated into the Seshat databank through Dacura. The area in the blue rectangle can be entirely automated, while the area outside the blue rectangle requires automated analysis and experts to ensure the Seshat data are valid and reliable. Starting at the top of the blue rectangle, a Human Population entity is defined by a Duration

1
2
3 within a Territory. The characteristics of the Human Population entity are then
4
5 classified through the OAT thesaurus to define a Tradition entity class. Data
6
7 mining begins by automated searching of the Internet for Cultural Domain
8
9 information as classified through the OCM thesaurus. At this point a researcher
10
11 can also have Dacura search for Cultural Domain information through both
12
13 Internet and print sources. Dacura compares information on a specific Cultural
14
15 Domain, identified in Figure 4 as Archaeological Data, with values in DBpedia to
16
17 determine if linked values should be included from other sources, and then
18
19 evaluated by an automated analyst for consistency. Inconsistent data (data with
20
21 semantic or value conflicts) are output to the researcher or an expert on the Cultural
22
23 Group or Cultural Domain for evaluation. The researcher or expert either decides
24
25 upon a canonical value for the Cultural Domain or, if there are conflicts that
26
27 cannot be resolved, a non-canonical value is given. Regardless, all the harvested
28
29 values are included in Seshat and marked as either canonical or non-canonical so
30
31 that other researchers can return to those values and revise or augment their
32
33 interpretation. Canonical values are also exported to DBpedia to assist other
34
35 researchers and future searches.

36
37
38
39
40
41
42 [Figure 4 about here]

43
44
45 Researchers may also input their own data, including images, media, and
46
47 shapefiles. Incorporating new data allows other researchers to access it through
48
49 Linked Data, making it widely available. Perhaps more importantly, data
50
51 incorporated into Seshat or other Dacura-generated dataset can be continuously
52
53 improved as Dacura allows researchers to comment on and re-evaluate previously
54
55

1
2
3 incorporated data. We suggest this allows Dacura to not only create and curate
4 large datasets, but also to ensure the continuous improvement of data quality.
5
6
7
8
9

10 *Using Seshat: Outputs from Dacura*

11
12
13

14 Our researcher interested in Hawaii population estimates would be able to quickly
15 identify accurate and fully referenced estimates through Seshat (or her own
16 Dacura-generated dataset). She would open Seshat through Dacura (a mock-up is
17 provided at <http://dacura.scss.tcd.ie/seshat/>) [Endnote 2], select the Natural-
18 Geographic Area for Hawaii, select the Polity sub-class of the Human Population
19 inhabiting Hawaii for the time period of interest, and then select the population
20 variable. A video illustrating this procedure is available at
21 <https://youtu.be/tZNCqNq2Mp0>. The data on population she would obtain in this
22 case would be taken from the Seshat data repository created with Dacura through
23 the data harvesting and verification process described above. But our researcher
24 could also create a new ontology using Dacura to conduct her own unique search,
25 as discussed above and illustrated in the video mentioned earlier at
26 <https://youtu.be/AEb1wF3jAgk>.
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43

44 Our researcher would have a wide range of possible outputs from her
45 search, whether in Seshat or with data she employed Dacura to harvest from other
46 sources. As noted earlier, Dacura publishes datasets as Linked Data and employs
47 SPARQL for output. SPARQL is a query language for RDF graphs which can
48 produce documents and raw data sets but also graphs, charts, maps and other
49
50
51
52
53
54
55
56
57

1
2
3 visualizations. Important for historical researchers, SPARQL works with
4
5 GeoSPARQL to allow data integration into geographic information system using
6
7 well-understood OGC query standards (GML, WKT, etc.). Raw textual, media,
8
9 or numeric data produced through Dacura can be browsed, searched, and selected,
10
11 allowing our researcher the ability to access the sections of texts, images, media,
12
13 or datasets she finds most useful. Dacura also allows harvested or input material
14
15 (or subsets of them) to be exported in a wide range of formats for external
16
17 analysis. For example our researcher might want numerical data on population
18
19 estimates as output for statistical analysis. Dacura would produce a comma-
20
21 delimited file that could be ported directly into a spreadsheet or statistical package
22
23 and our researcher could then run any analysis she required to answer her
24
25 question. Figure 5 shows a simple line graph of Hawaii population estimates
26
27 derived through Dacura and Seshat with data output to a csv file and graphed
28
29 using an Excel spreadsheet.
30
31
32
33
34

35 The answer to the researcher's initial question would be that population
36
37 indeed increased prior to the first state system on the island of Hawaii, as
38
39 expected under the population-pressure model. This is not a particularly
40
41 impressive result in itself (and would in fact be quite simplistic—see Field,
42
43 Ladefoged, and Kirch, 2011), but consider that our researcher would have been
44
45 able to compile these data in a matter of minutes, be confident of their quality,
46
47 and have access to all the metadata attached to them.
48
49
50

51 [Figure 5 about here]
52
53
54
55
56
57

CONCLUSIONS

The Internet provides historical scientists abundant information, but often the information is too abundant, and usually lacks quality control. Dacura was designed to address these problems. It provides a way to harvest information from the Internet easily, with an assurance of quality, and with a manageable body of results. Dacura's carefully designed ontology (dacura.scss.tcd.ie/ontologies/dacura-130317.ttl) allows researchers to readily identify and retrieve information directly relevant to their research. Dacura's integrated thesauri and RDF triplestore structure removes the need for detailed indexing across result domains so that all information on a given subject, even information that might not be obviously related or indexed as related, is retrieved. And Dacura offers a wide range of possible outputs, from texts to visualizations to spreadsheets. Dacura is not the only data harvesting and curation package available, but because it has been developed hand-in-hand with the Seshat databank, it provides a unique model for new computer-based methods of historical and archaeological data handling.

In this way, Dacura represents a set of important new tools for the historical sciences. As Kintigh et al. (2015, 3) have recently pointed out “[historians and] archaeologists are increasingly challenged as they acquire, manage, and analyze large volumes of disparate data.” Dacura illustrates several responses to this problem. Specifically, Dacura incorporates (1) a semantic knowledge graph technology based on an RDF triplestore, (2) the use of Web

1
2
3 Ontology Language (OWL) for data definition, (3) semantic reasoning as the
4 foundation of an automated process for data evaluation, and (4) output following
5
6 Linked Data standards.
7
8

9
10 OWL data definition and knowledge graph technology allow the
11 enormous volume of data available to historical researchers to be quickly and
12 easily reduced to the most important information on a given question, and then
13 output in a variety of useful formats. Semantic reasoning provides mechanisms
14 for rapid data evaluation and ongoing curation. And Linked Data standards allow
15 other researchers to readily access data that has already been harvested and
16 evaluated.
17
18

19
20 In addition, the recent advances in computer science available in platforms
21 like Dacura offer a way to provide useful and accurate historical data to scholars
22 who are not historical scientists. It has long been a frustration among historical
23 researchers that data which can provide both a diachronic record of cultural
24 stability and change and empirical examples of practices that have been
25 successful or unsuccessful in human societies, has not been widely used outside
26 of archaeology and history. But it is also not surprising, as historical data can be
27 hard to access and hard to understand by non-professionals (Kintigh et al. 2015,
28 2). By providing an automated means of harvesting, evaluating, and exporting
29 historical data that has been evaluated for accuracy, platforms like Dacura offer
30 both a means and a model for economists, political scientists, ecologists,
31 geographers, and others to access and explore the rich and valuable record of
32 human history.
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

ACKNOWLEDGEMENTS

This work was supported by a John Templeton Foundation grant to the Evolution Institute, entitled "Axial-Age Religions and the Z-Curve of Human Egalitarianism," a Tricoastal Foundation grant to the Evolution Institute, entitled "The Deep Roots of the Modern World: The Cultural Evolution of Economic Growth and Political Stability," an ESRC Large Grant to the University of Oxford, entitled "Ritual, Community, and Conflict" (REF RES-060-25-0085), a grant from the European Union Horizon 2020 research and innovation program (grant agreement No 644055 [ALIGNED, www.aligned-project.eu]), and an European Research Council Advanced Grant to the University of Oxford, entitled "Ritual Modes: Divergent modes of ritual, social cohesion, prosociality, and conflict."

The authors wish to thank the participants of a workshop held at the Santa Fe Institute May 4-6, 2015 during which the needs for harvesting and integrating quality information was discussed and the Seshat meta-model developed. We gratefully acknowledge the contributions of our team of research assistants, post-doctoral researchers, consultants, and experts. Additionally, we have received invaluable assistance from our collaborators. Please see the Seshat website (www.seshatdatabank.info) for a comprehensive list of private donors, partners, experts, and consultants and their respective areas of expertise. Finally, we want to thank the anonymous reviewers whose insightful comments allowed us to substantially improve this paper.

ENDNOTES

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

1. The Dacura software was developed as part of the European ALIGNED Horizon 2020 project. All of the software, as well as other useful tools for managing semantic datasets and knowledge graphs, is available with an open source license through the project's web site at <http://aligned-project.eu/open-source-tools/>. However, this still requires users to configure and install their own knowledge-graph server, which remains a complex undertaking. It is our goal to also make the system available to researchers through the web as service in such a way that no technical knowledge is required to use it. We anticipate releasing a pilot version of this service in the middle of 2018. In the meantime, for any researchers who are particularly interested in seeing and using the web service, we are running an ongoing series of trials with Seshat and other collaborators and are open to new research collaborations. For updates check the Dacura website at <http://dacura.cs.tcd.ie/> or by email at dacura@scss.tcd.ie.

2. The version of Dacura shown on the site is a mock-up that lacks full functionality and is intended solely to give readers a sense of what the Dacura interface and output might look like when installed on their institutional computer. This version searches only within a portion of the Seshat databank and only returns simple html output. As discussed in the article, the fully-functional version of Dacura does much more than the mock-up provided here.

REFERENCES CITED

Diamond, J. 1997. *Guns, germs, and steel*. New York: W.W. Norton.

Facebook, N.d. Introducing graph. <https://www.facebook.com/graphsearcher/>.
Accessed 1/3/2018.

Field, J. S., T. Ladefoged, and P. Kirch. 2011. Household expansion linked to agricultural intensification during emergence of Hawaiian archaic states. *PNAS*, 108(18), 7327-7332.

François, P., J. Manning, H. Whitehouse, R. Brennan, et al. 2016. A macroscope for global history: Seshat Global History Databank. *Digital Humanities Quarterly* 10(4).

Google. N.d.. The knowledge graph.
<https://www.google.com/intl/bn/insidesearch/features/search/knowledge.html>.
Accessed 1/3/2018.

Mullins, D.A., D. Hoyer, C. Collins, T. Currie, et al. 2018. A systematic assessment of “Axial Age” proposals using global comparative historical evidence. *American Sociological Review* (in press).

1
2
3
4
5 Kansa, E.C. 2010. Open Context in context: Cyberinfrastructure and distributed
6 approaches to publish and preserve archaeological data. *SAA Archaeological*
7
8 *Record* 10(5):12-16.
9
10

11
12
13
14 Kintigh, K.W., J. Altschul, A. Kinzig, W.F. Limp, et al.. 2015. Cultural
15 dynamics, deep time, and data. *Advances in Archaeological Practice* 3:1-15
16
17

18
19
20
21 Murdock, G.P. 1983. *Outline of world cultures, 6th edition*. New Haven, CT.:
22 Human Relations Area Files.
23
24

25
26
27
28 Peregrine, P.N. and M. Ember (editors). 2001. *Encyclopedia of prehistory, 9 vols.*
29
30 New York: Kluver Academic / Plenum Publishers.
31
32

33
34
35 Turchin, P., R. Brennan, T. Currie, K. Feeney, et al. 2015. Seshat: The Global
36 History Databank. *Cliodynamics* 6: 77-107
37
38

39
40
41
42 Turchin, P, T. Currie, H. Whitehouse, P. Francois, et al. 2018. Quantitative
43 historical analysis uncovers a single dimension of complexity that structures
44 global variation in human social organization. *PNAS* 115(2): E144-E151.
45
46
47
48
49 <https://doi.org/10.1073/pnas.1708800115>.
50
51

1
2
3 Watts, J. 2011. Building tDAR: Review, reduction, and ingest of two reports
4 series. *Reports in Digital Archaeology* 1:1-15.
5
6
7

8
9
10 Whitehouse, H., P. François, and P. Turchin. 2015. The role of ritual in the
11 evolution of social complexity: Five predictions and a drum roll. *Cliodynamics*
12 6(2):199-216.
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

For Peer Review Only

Figure 1: The four stages of the Dacura data curation process.

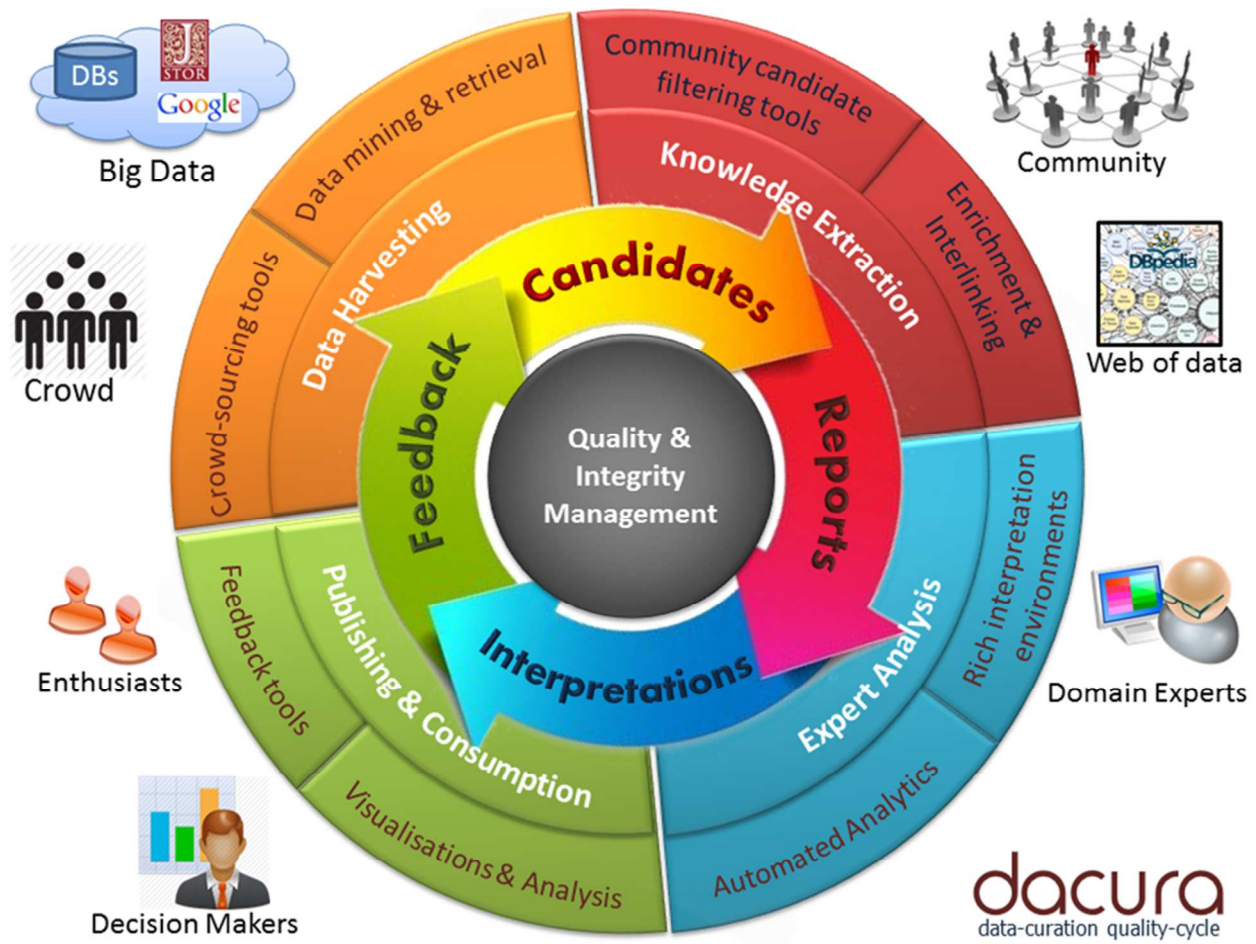
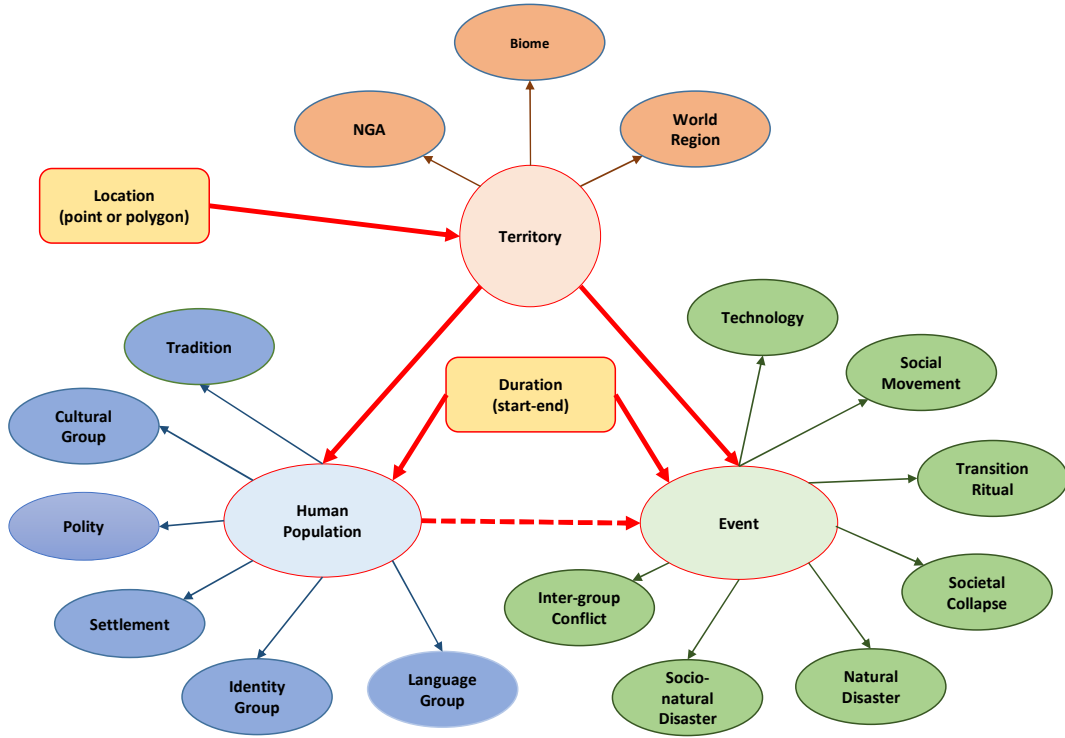
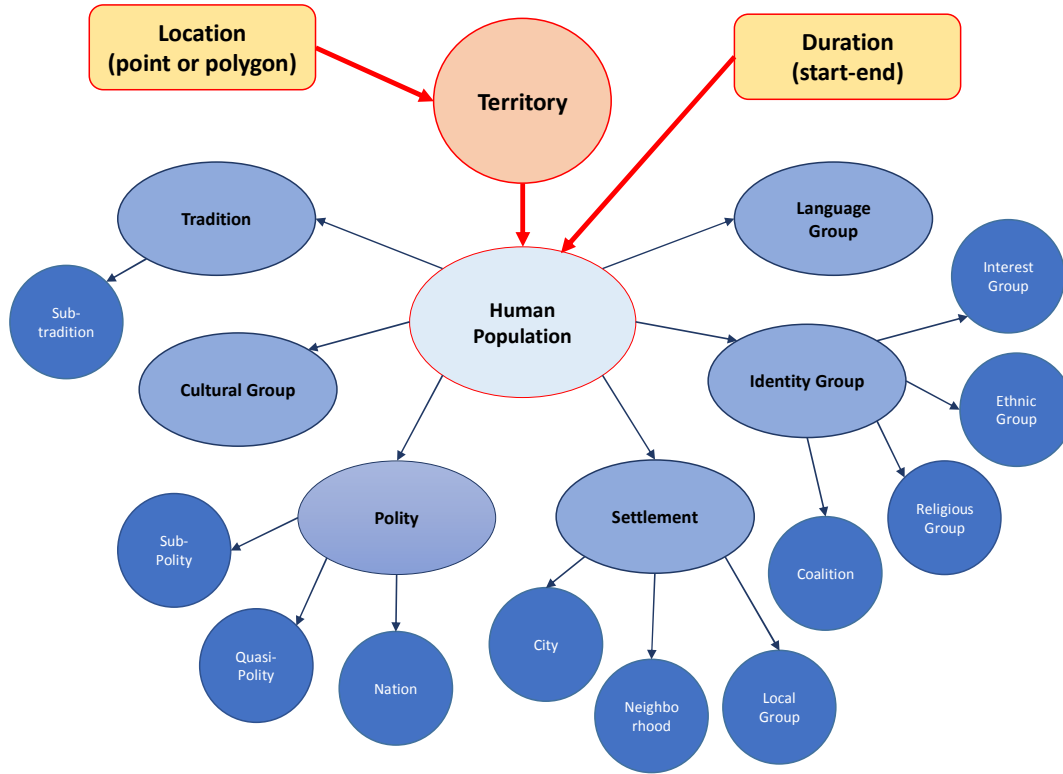


Figure 2: Metamodel for Seshat: The Global History Databank.



Only

Figure 3: Detail of the Human Population entity, showing current entity classes and subclasses.



Only

Figure 4: Workflow for the incorporation of numerical data for the Archaeological Tradition entity class into Seshat through Dacura. Human characters marked “A” represent automated data evaluators; the character marked “E” represents a human scholar or “expert” evaluator.

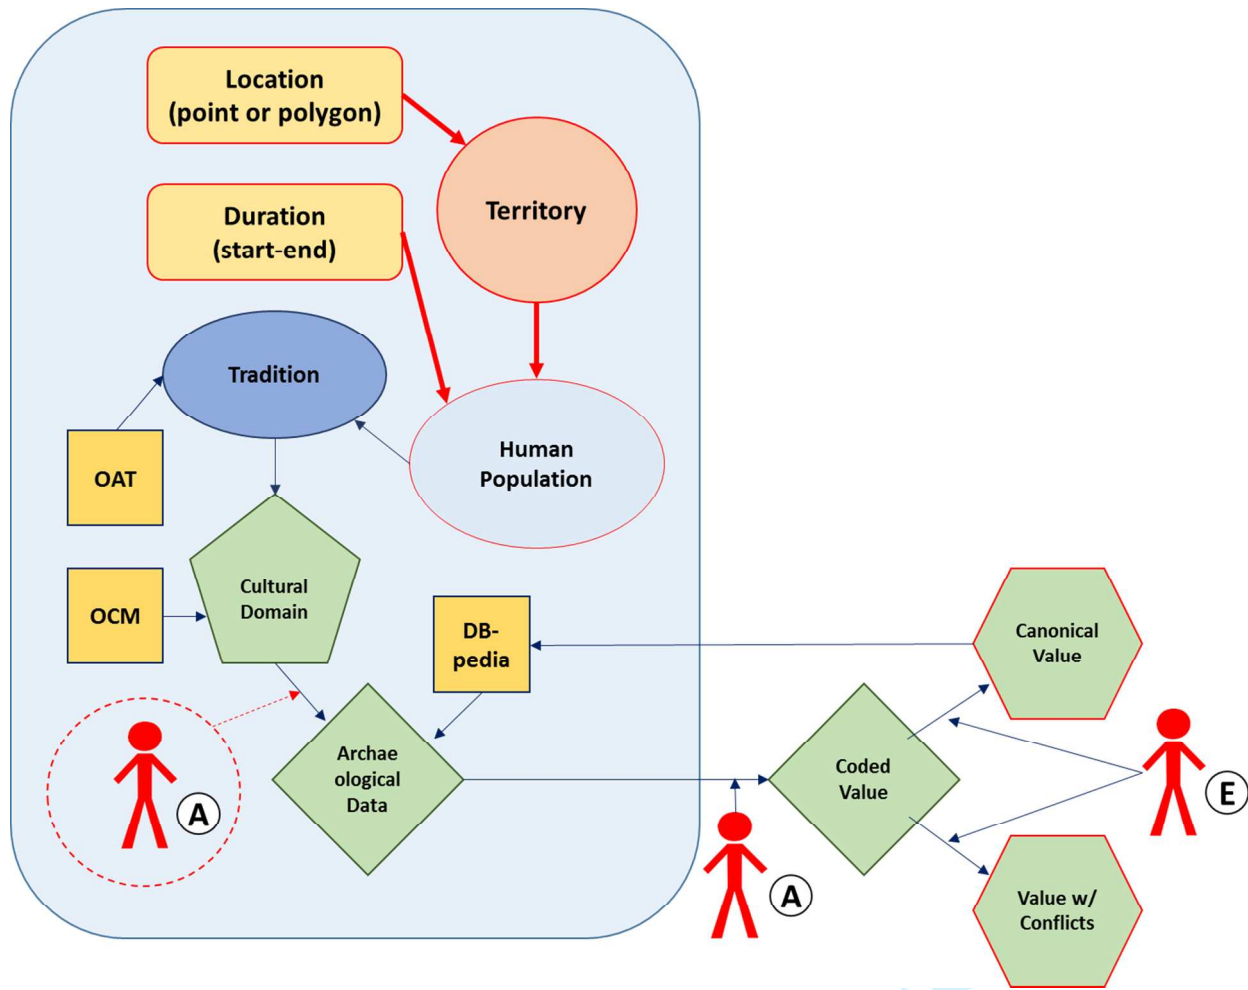


Figure 5: Population dynamics on Big Island Hawaii from 1200 to 1700 CE.

