Service-Oriented E-Learning Platforms

From Monolithic Systems to Flexible Services

The learning management system (LMS) has dominated Internet-based education for the past two decades. However, the traditional LMS is failing to keep pace with advances in Internet technologies and social interactions online. To support technological diversity, current frameworks such as the E-Learning Framework (ELF), the IMS Abstract Framework, and the Open Knowledge Initiative (OKI) have defined the initial steps toward service-oriented e-learning platforms. Next-generation platforms will be based on these service-oriented visions. Here, the authors discuss LMS evolution and present core challenges that must be addressed to achieve information interoperability in next-generation e-learning platforms.

Declan Dagger,
Alexander O'Connor,
Séamus Lawless,
Eddie Walsh,
and Vincent P. Wade
Trinity College Dublin

As Internet technologies proliferate into our daily lives, we come closer to realizing new and exciting online opportunities. One such opportunity is in e-learning, in which more dynamic platforms are emerging and replacing traditional, passive ones. Active e-learning employs a broad range of Internet technologies, such as personalization, simulation, and mobility, to achieve pedagogic scenarios otherwise inaccessible to traditional forms of learning.[1] Thus, today's e-learning platforms must deal with an increasing set of requirements.

The demand for modularized and personalizable e-learning platforms is growing.[2] Traditional platforms can't support architectural flexibility due to their monolithic designs. E-learning vendors are addressing this demand by providing toolkits that support customization or by making their source code available for modification under various open source licenses. This indicates an emerging shift from generic solutions to specific applications. Future e-learning platforms will support a wider range of needs by providing interoperability architectures for various existing and emergent services. These needs include federated exchange among services (information and control), various levels of interoperability (intradomain and interdomain), and service composition (orchestration and choreography). Howev-

er, these next-generation platforms also introduce wide-ranging issues from numerous research areas, including the Semantic Web, adaptive hypermedia, dynamic services, and federated modeling.

Here, we explore e-learning platforms' evolution and illustrate some key challenges to information interoperability in next-generation platforms.

E-Learning Platforms

Traditional e-learning platforms, or *learning management systems* (LMSs), provide holistic environments for delivering and managing educational experiences. They present suites of tools that support online course creation, maintenance, and delivery, student enrollment and management, education administration, and student performance reporting. We can group LMSs into two main categories:

- Open source initiatives include Moodle (www. moodle.org), Sakai (www.sakaiproject.org), ATutor (www.atutor.ca), and Whiteboard (http:// whiteboard.sourceforge.net).
- Proprietary solutions include WebCT/Black board (www.blackboard.com), Gradepoint (www. gradepoint.net), Desire2Learn (www.desire2 learn.com), and Learn.com (www.learn.com).

Open source LMSs are typically built on extensible frameworks that let implementers adjust and modify the systems to suit their specific needs. Although the proprietary sector hasn't widely adopted it, this approach is emerging through such initiatives as WebCT's PowerLinks kit and Blackboard's Building Blocks. These provide software developers with "hooks" to tie third-party software into the LMS.

First Generation

The first generation of e-learning platforms (from roughly 1993 on) provided, in essence, black-box solutions. For the most part, these systems used proprietary formats to manage courses directly. You could commonly find a one-to-one mapping between systems and courses, with very limited user tracking, if any. These platforms focused on the delivery and interoperability of content designed for a specific purpose, such as a particular course. During this time, a range of standards emerged – Dublin Core (www.dublincore.org), IMS Learning Resource Metadata (www.imsglobal.org/ specifications.html), IEEE Learning Object Metadata (http://ltsc.ieee.org/wg12/), and others that are still in use today. These standards described content that compliant LMSs could share at the asset level; however, this practice wasn't widespread. Instead, these standards made provisions for interoperability between different systems at the content level via communication channels such as AICC Computer Managed Instruction (CMI; www. aicc.org/pages/aicc_ts.htm). AICC described a common communication API to be included in all content that supported such functionalities as notifications to the LMS indicating the commencement and completion of learning content. First-generation platforms provided for future evolution via their support for sharing content in an interoperable way. Examples of these early e-learning platforms include the first versions of WebCT and Blackboard.

Second Generation

The second, or current, generation of e-learning platforms (from roughly 1999 on) expanded on their predecessors' successes and began addressing their failures. Examples of these second-generation platforms include WebCT/Blackboard, Moodle, and Sakai. In terms of e-learning evolution, they provided a shift toward modular architectural designs and recognized a need for semantic exchange. They began to focus not only on sharing content but also on sharing learning objects, sequences of learning objects, and learner information. Standards and specifications emerging during this time, such as the Shareable Content Object Reference Model (SCORM; www.adlnet.gov), IMS Content Packaging, and IMS Learning Design (www.ims global.org/specifications.html), supported the ability of standards-compliant platforms to share courses or parts of courses via an import–export paradigm, while IMS Tool Interoperability provided guidelines for interoperating tools between different LMSs. Another significant development was that second-generation platforms began to embrace the "services" principle, exposing certain aspects of their functionality externally. As designs became more modularized, it was easier for platforms to integrate new functionality as it arose. In some cases, such as with Sakai, Web services were designed to expose a limited set of functionality supported within. The LMS community made an increased move toward separating content from tools, and the learner information became more distinguished. However, these systems aren't entirely learner-centric: they still focus strongly

on learning administration (course management) rather than on the learner.

Evolution

As these changes continue, they forge a path toward next-generation e-learning platforms. This, in essence, involves applying service frameworks to the platforms' modular design. Separating LMS and *learning content management system* (LCMS) functionality will provide support for greater interoperability, in which systems can not only share content and learning scenarios but also exchange tools, functionalities, semantics, and control seamlessly and dynamically. This also involves a much wider range of information — such as user information, context, sequencing, workflow, and control — which services can use and reason across in e-learning platforms. This will let users build custom e-learning platforms from a vast range of e-learning services, for their specific needs, at specific times, and in dynamic ways. The LMS vendor of old will no longer sell monolithic, one-size-fits-all solutions, but rather interoperable platforms and a range of e-learning services, letting consumers choose the right combination of services for their requirements.

Challenges for Flexibility and Interoperability

Next-generation e-learning platforms introduce several key challenges for information interoperability — in part due to the shift toward semantic exchange — that we must address going forward. Information will be not only exportable and importable across different environments but also exchangeable among heterogeneous environments. Current e-learning platforms, for example, can send user information (knowledge state, skill level, preferences, and so on) or content information (metadata about subject area, technical aspects of the content, and so on) between environments. However, there is no support for these environments to understand such information's semantics, how to reason across it, or how different environments can use it. For this level of interoperability, environments must exchange both the information's syntax and its semantics. This can't occur through what some call "shared" semantics — that is, single global semantic models. In reality, people (the core ingredient of information systems) are adversely different in, for example, how they describe objects.[3] Thus, flexible solutions must

identify and support dynamic semantic mappings to support true semantic interoperability.[4]

Another key aspect to such semantic exchange relates to control. We can't assume that e-learning services are simply dumb content. They have and will continue to have their own internal representations, control flows, and, in some cases, tracking mechanisms. Thus, future e-learning platforms must support the exchange of control between interoperating services.

A further challenge is to support higher levels of interoperability. One aspect of *intradomain* interoperability involves creating frameworks and standards to support plugability, which numerous initiatives and standards bodies have already begun. *Interdomain* interoperability, however, is a much newer field — it involves the interoperability of context, a topic we describe further in a later section.

Service composition will let these e-learning platforms dynamically discover and assemble e-learning services to achieve a given user's specific purpose. This will include the realization of service orchestration and service choreography,[5] topics that merit lengthy discussion but are outside this article's scope. Research initiatives in this area are beginning to address all these challenges, but most work is in its early stages.

Standards and Technologies

Several standards and technologies support the interoperability requirements for next-generation e-learning platforms, as Figure 1 illustrates. Existing and emerging methodologies evolve around modularization and separation of concerns. This, in essence, means that functionality is divided into modules, which can then be combined to provide an integrated e-learning platform. Service-oriented architectures (SOAs) describe an architectural concept that defines the expression of processes and logic as individual services, which in turn publish or expose facets of their functionality in a standardized way, letting other services access and use this functionality in a flexible manner.

Component-oriented architectures (COAs), on the other hand, are more finely grained and tightly coupled than SOAs. Changes to individual components typically impact the software those components access, making COAs less flexible and extensible than SOAs. A COA is primarily concerned with a local component, its properties, and its methods.

In middleware architectures, a software layer

between applications, or services facilitates interoperability, helping to provide services such as identification, authentication, authorization, information exchange, and security. Middleware is especially applicable to modern knowledge-based systems developed on XML, SOAP, and SOAs. Indeed, middleware is often considered the "plumbing" between Web services.

The standards community has been pioneering a variety of frameworks, specifications, and guidelines for service-oriented e-learning platforms for some time. The IMS Abstract Framework (www.imsglobal.org/specifications.html) loosely identifies and represents the core components and interfaces of an e-learning system. The E-Learning Framework (ELF; www.elframework.org) illustrates e-learning systems' common functionalities. Similarly, the Open Knowledge Initiative (OKI; www.okiproject.org) defines service layers for developing e-learning platforms. The common approach among these emergent standards is to modularize functionality, usually defining the following groupings:

- sets of applications (such as LMSs);
- application services (finer-grained services, such as quizzes and simulations, with which the user directly interacts);
- educational services (usually revolving around education administration such as course management and scheduling);
- common services (functionality that the user isn't directly exposed to but that is essential, such as authentication, file sharing, logging, and database management); and
- infrastructure (the backbone of the services, including HTTP, SOAP, and XML).

The frameworks, specifications, and guidelines then define layered approaches to constructing e-learning systems from these collections of previously defined services. These specifications provide representations for person and group information (IMS Enterprise), a person's profile and learning history (IMS Learner Information Package and ePortfolio), assessment (IMS Question and Test Interface), groupings of learning content (IMS Content Package and SCORM), dynamic content sequencing (IMS Simple Sequencing), learner competencies (IMS Reusable Definition for Competence and Educational Objectives), learning activities (IMS Learning Design), searching across federated databases
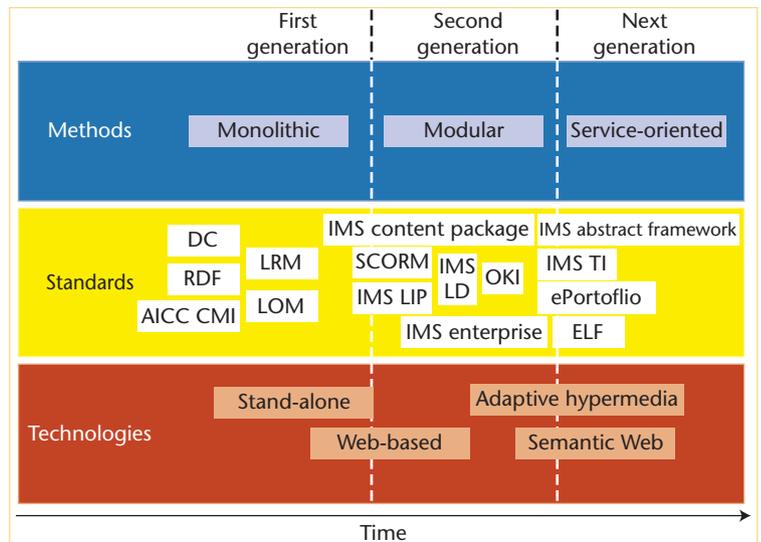

Figure 1. Generations of learning management systems. First-generation systems supported content-only interoperation. Second-generation systems began to take account of users and their associated profiles, whereas next-generation systems will support targeted personalization.

(IMS Digital Repositories Interoperability), and connecting diverse eLearning tools (IMS Tool Interoperability; see www.imsglobal.org/specifications.html for more about these specifications). At a fine-grained level, these standards and specifications describe the syntax that the various services should implement to present information externally. However, a crucial component of interoperability is missing — the ability to dynamically use the information meaningfully. This is where the Semantic Web community makes a vital contribution.

A key assumption with machine-readable information and services is that they can interoperate and negotiate on the fly. The Semantic Web began with RDF and DAML+OIL and has since progressed toward the Web Ontology Language (OWL; www.w3.org/TR/owl-features/), which recently became a W3C standard. OWL can be used to express *ontologies* — knowledge bases of concepts that services can query to retrieve information. In particular, ontologies can help produce new knowledge, in that inference can occur within the knowledge base to achieve a goal, independently of information input specifically by the ontology's creator. This ability extends to the realm of Web services with OWL for Services (OWL-S).

At a finer-grained level, developers can use W3C's Web Service Description Language (WSDL; www.w3.org/TR/wsdl) to describe a Web service's functionality in terms of its *inputs*, *outputs*, *pre-*

conditions, and *effects* (IOPE). WSDL provides a syntax through which a calling system or service can access exposed functionality without concerning itself with its inner workings.

At a broader level of granularity, the Semantic Web community is developing specifications for services organization and workflow. Web Services Business Process Execution Language (WSBPEL; http://docs.oasis-open.org/wsbpel/2.0/wsbpel-specification-draft.html) supports service organization and flow within SOAs. Enterprise Service Bus (ESB)[6] also illustrates service organization and orchestration.

Adaptive hypermedia[7] systems form the second tranche of technology for interoperation. Adaptive systems benefit from a different approach because they tend to be open-ended information systems with a highly developed sense of their domain. This means that external sources can add new information relatively easily and that a small amount of additional information will considerably improve performance. Adaptive services, taking advantage of OWL-S in addition to core OWL, provide a rich, highly expressive framework for interoperating services.

The main link between the Semantic Web and adaptive hypermedia is Web service technology and composition. SOAP and REST allow remote systems to use Web-based communication to create complex systems built from atomic services.

E-Learning Services

E-learning services can represent and manage any technical aspect of an e-learning scenario. Initiatives such as ELF, the IMS Abstract Framework, and OKI have defined a wide variety of coarse-grained functionalities that e-learning services can provide. E-learning services include traditional functionalities such as authentication, tracking, course management, scheduling, activities, tools, and assessment, as well as emerging functionalities such as personalization, resource harvesting, context management, federated exchange, simulation, games, wiki, blogging, podcasting, and so on. Thus, an e-learning platform must support this expanding range of coarse-grained and emerging finer-grained services to provide real flexibility. The complexities and variations in these second-generation standards will lead to more diverse implementations than with first-generation (content-related) standards. Additionally, e-learning platforms must be able to support loose coupling, thus enabling interoperability of services' semantics.

Service Interoperability

Service interoperability concerns the seamless creation, deployment, consumption, and orchestration of Web-based services. Most research in this field focuses on service syntax interoperation — that is, the ability to create a service on any platform in any environment and interact with it from a different platform and environment. Through specifications such as OKI, the e-learning community is specifying common framework-based APIs for certain e-learning application areas. However, these APIs provide support for tightly knit collections of services, grouped according to functionality. Going forward, looser connections should provide greater potential for interoperability and flexibility, IMS Tool Interoperability being a positive step in this direction. Instead of specifying APIs for communications in a common syntax, we will focus on semantic-exchange and control-flow management between services. The semantic exchange supports a more detailed view of what a service can do, what it produces, how to manage it, and its uses. It also paves the way for generating and interpreting other information, such as context, trust, and security.

Context-informed services. Pure ubiquitous computing research is based on two key concepts: pervasive computing and calm computing.[8] Pervasive, calm systems help users by managing their complex information profiles in an easy, blended way. Context-informed e-learning aims to lower the information burden on users and let them concentrate on their learning tasks. Numerous projects (see www.cs.cmu.edu/~aura/ for one) demonstrate how various cooperating systems can enhance users' experience by placing a much wider variety of information at their immediate disposal.

Educational applications represent a natural application domain for this type of work, in part because of their immediate familiarity for academic researchers, but more importantly because of students' rich information needs. Projects such as ActiveClass[9] address domain-specific concerns while maintaining the generalized ubiquitous computing approach.

Typically, current systems[10] represent context via physical properties such as location. Next-generation systems will be able to leverage rich

information sources made available by a service-oriented approach. Services will provide a great deal of information about learners and their activities, including metadata to describe that information. This richer information space lets systems incorporate a much wider amount of information, much of it located and integrated at runtime.

One approach we're currently investigating involves integrating adaptive systems with semantically described services.[11] This approach uses a semantic overlay network to create links between the knowledge residing in external services and the e-learning application. The informed environment can then use these links to transfer information known by external services to the e-learning application based on users' situations and needs (that is, their context).

By generating a *shared semantic view* of the user's context, the e-learning platform can include a wide variety of extra information relevant to his or her task without having to model the types of useful information a priori. The semantic approach provides a way to tailor adaptive systems' knowledge and behavior to exploit extra available knowledge in a dynamic fashion. The platform designer needn't specify the type of information in advance; rather, the informed environment discovers services and finds links based on the descriptions that accompany those services. Initiatives designed to add rich descriptions to services and their information (such as the Semantic Web) present a clear opportunity to greatly enhance the e-learning process.

Flexible service interoperability. To illustrate flexible service interoperability in an e-learning environment, consider the following example. David is a site engineer working for the Tech Corporation, a large software services and hardware vendor. As part of his work, David travels to client premises to perform upgrades and maintenance on a wide variety of software and servers that Tech Corp provides.

In this particular scenario, David is on site with Widgets Inc. and needs to upgrade the client's servers. The task involves some new technologies that David isn't familiar with, so he needs to use Tech Corp's knowledge base for assistance.

Tech Corp maintains a *technical knowledge base* (TKB), illustrated in Figure 2, to assist employees of different levels and areas in learning new skills and performing specific tasks. In addition to the TKB, Tech Corp maintains a human resources (HR)
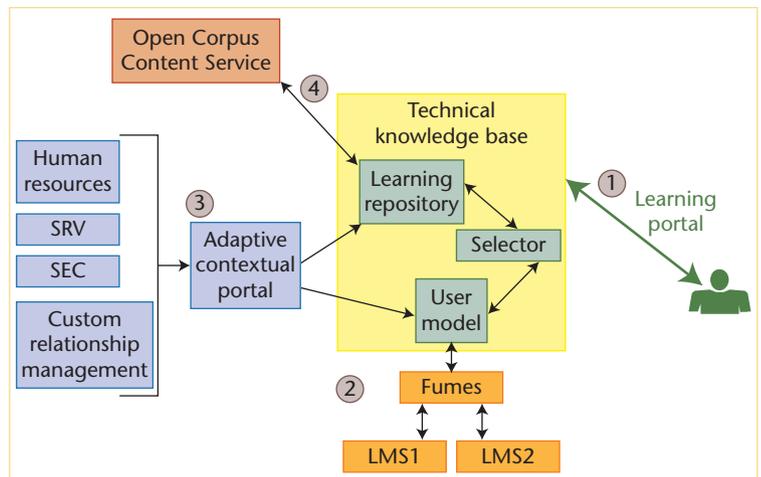

Figure 2. Dynamic personalized e-learning scenario. Loosely coupled services dynamically interoperate to personalize and contextualize the content and activities of an e-learning experience.

database with details of individuals and their positions. This is complemented by a customer relationship management (CRM) database containing each customer's deployment details. Finally, the site that Widgets Inc. maintains includes security- and server-management services.

First, David connects to the Tech Corp TKB, through a learning portal (such as an LMS or a corporate Web site), which authenticates and manages his session and exchanges information as required with the ancillary services. Next, the TKB retrieves David's learning preferences through a *federated user modeling exchange service* (FUMES), which extracts them from LMS services (LMS1 and LMS2) that he's previously used. This provides a semantically rich description of David's past e-learning experiences.

The TKB is a context-informed adaptive service suite. The *adaptive contextual portal* (ACP)[11] that manages it is aware of the services situated both on site and in David's office. The ACP maintains a shared semantic view of David's context and can tailor his learning experience, based on information held in the CRM and HR services. The system can guide David in his choice of learning resources and specify additional parameters for personalizing his experience. The system chooses to select technical, descriptive learning resources because David is an on-site engineer who needs practical advice on task completion.

The TKB maintains an archive of content that Tech Corp believes is both necessary and beneficial to its staff's training and professional development. However, a task assigned to an employee

might occasionally involve technologies that the TKB doesn't currently support. In that case, the TKB could use a complementary content-harvesting service called the *Open Corpus Content Service* (OCCS). Research is ongoing into the development of this service, which can leverage open corpus content for use in e-learning environments.[12]

The OCCS employs a focused Web crawler to traverse the Web and many digital repositories, creating an archive of all the relevant content it encounters. The OCCS indexes this archive, and the TKB can request content from it. The TKB extracts content requirements from the semantic interpretation of David's task, as derived by the ACP from his current learning preferences and past learning experiences (FUMES).

The OCCS must address some interoperability issues in order to leverage the available Web content, including problems with content that has insufficient or, in some cases, no associated metadata descriptions. Problems also arise when incompatible metadata standards are implemented on various pieces of content. To overcome these issues, the OCCS applies semantic mappings to a canonical metadata model using a fixed taxonomy of terms. Semiautomatic or automatic content annotation also generates descriptions of content with no associated metadata.

For this context-informed approach, the TKB participates in the following process. Each service located for the ACP registers an ontology of the information that it can share. A reasoner suite then examines the concepts known to each service and builds a shared semantic view as a topic map within the system. As the TKB performs the operations needed to build the course for David, hooks within the TKB share the current knowledge of the service with the ACP. The external services enhance this knowledge with their knowledge via the ACP's shared semantic view and "push" it back to the TKB. David then engages the offering from the TKB (step 5 in Figure 2) and successfully completes the current task. The context-informed approach means that the TKB's designers needed only to define the points at which the TKB could use context and allow the aware environment to push context-enhanced models to the system.

As innovative Internet technologies and applications continue to grow, people are placing more aspects of their lives online. Passive consumer models are rapidly being outpaced by more active producer models. Defining rigid frameworks and boundaries for e-learning platforms doesn't fit with the free model of the Internet. Wrapping e-learning as a deliverable is a failed approach. In reality, every contribution to the Web has a learning-related value associated with it. Tomorrow's e-learning platforms will deliver knowledge when, where, and how you want it.

EU projects such as iClass, PROLearn, and Unfold are pioneering European efforts to create sustainable e-learning services in a range of educational sectors. The iClass project, in which we're participating, is responsible for researching, developing, and disseminating e-learning services for a self-regulated personalized e-learning environment for the primary and secondary (K12) school market.

Much of the current work in this field is concerned with interoperation and exchange. However, only a few services are generally available for consumption. Without a critical mass of such services, we risk hindering the evolution of next-generation LMSs. ⬚

References

1. T. Mayes and S. de Freitas, "Review of E-Learning Theories, Frameworks, and Models," *Joint Information Systems Committee (JISC) e-Learning Models Desk Study*, 2004; www.jisc.ac.uk/uploaded_documents/Stage%202%20Learning%20Models%20(Version%201).pdf.
2. P. Trafford, "PLEs as Environments for Personal and Personalised Learning," position paper as part of the RAMBLE project; www.weblearn.ox.ac.uk/site/asuc/oucs/staff/pault/research/ramble/present/uploads/PaulT_Mini%20PLEPaper1.0.pdf.
3. A. Mathes, "Folksonomies — Cooperative Classification and Communication through Shared Metadata," *Computer-Mediated Communication Magazine*, 2004; www.adammathes.com/academic/computer-mediated-communication/folksonomies.html.
4. D. O'Sullivan and D. Lewis, "Semantically Driven Service Interoperability for Pervasive Computing," *Proc. 3rd ACM Int'l Workshop on Data Engineering for Wireless and Mobile Access*, ACM Press, 2003, pp. 80–92.
5. C. Feier et al., "Towards Intelligent Web Services: The Web Service Modeling Ontology (WSMO)," *Proc. Int'l Conf.*

Intelligent Computing (ICIC), 2005; http://platon.escet.urjc.es/~axel/publications/feie-etal-2005.pdf.

6. D.A. Chappell, *Enterprise Service Bus*, O'Reilly Media, 2004.

7. P. Brusilovsky, "Methods and Techniques of Adaptive Hypermedia," *User Modeling and User-Adapted Interaction*, vol. 6, nos. 2–3,1996, pp. 87–129.

8. M. Weiser, "The Computer for the 21st Century," *Scientific Am.*, vol. 9, Sept. 1991, pp. 933–940.

9. M. Ratto et al., "The ActiveClass Project: Experiments in Encouraging Classroom Participation," *Proc. Int'l Conf. Computer Support for Collaborative Learning 2003*, Computer-Supported Collaborative Learning Series, vol. 2, B. Wasson, S. Ludvigsen, and U. Hoppe, eds., Kluwer Academic, 2003, pp. 477–486.

10. P. Lonsdale et al., "Context Awareness for MOBIlearn: Creating an Engaging Learning Experience in an Art Museum," *Proc. MLearn 2004: Learning Anytime, Everywhere, Learning and Skills Development Agency (LSDA)*, 2004, pp. 115–118.

11. A. O'Connor and V. Wade, "Informing Context to Support Adaptive Services," *Proc. Conf. Adaptive Hypermedia and Adaptive Web-Based Systems*, LNCS 4018, Springer-Verlag, 2006, pp. 366–369; doi 10.1007/11768012_52.

12. S. Lawless and V. Wade, "Dynamic Content Discovery, Harvesting, and Delivery from Open Corpus Sources, for Adaptive Systems" *Proc. 4th Int'l Conf. Adaptive Hypermedia and Adaptive Web-based Systems*, V. Wade, H. Ashman, and B. Smyth, eds., LNCS 4018, Springer-Verlag, 2006, pp. 445–451; www.cs.tcd.ie/~slawless/papers/ah2006.pdf.

Declan Dagger is a research fellow in the Knowledge and Data Engineering Group in Trinity College Dublin's Department of Computer Science. His research interests include user-centric adaptivity and personalization, composition methodologies for knowledge-driven adaptive systems, and user-driven interoperability. Dagger has a PhD in computer science from Trinity College Dublin. Contact him at declan.dagger@cs.tcd.ie.

Alexander O'Connor is a PhD candidate in the Knowledge and Data Engineering Group in Trinity College Dublin's Department of Computer Science. His research interests include semantic interoperation, e-learning, and ubiquitous computing. O'Connor has a masters degree in networks and distributed systems from Trinity College Dublin. He is a member of the ACM. Contact him at oconnoat@cs.tcd.ie.

Séamus Lawless is a PhD candidate in the Knowledge and Data Engineering Group in Trinity College Dublin's Department of Computer Science. His research interests include open corpus content, information retrieval, and personalized e-learning systems. Lawless has an honors degree, BSc, in computer science from Trinity College Dublin. Contact him at slawless@cs.tcd.ie.

Eddie Walsh is a PhD Candidate in the Knowledge and Data Engineering Group in Trinity College Dublin's Department of Computer Science. His research interests include e-learning, user modeling, and semantic interoperation. Walsh has an MSc (research) in e-learning from the University of Limerick. Contact him at walshe8@cs.tcd.ie.

Vincent P. Wade is research director at the Intelligent Systems Laboratory in Trinity College Dublin's Department of Computer Science. His research interests include knowledge-driven adaptive systems, service composition and content personalization, knowledge-driven management and control systems, knowledge-driven integration and interoperability, knowledge elicitation and machine learning, and user-centric knowledge and data engineering. Wade has a PhD in computer science from Trinity College Dublin. He is a member of the IEEE, the ACM, and the AACE. Contact him at vincent.wade@cs.tcd.ie.