

Opportunities for AI in Intelligent Web-based Technology-Supported Learning

Carsten Ullrich, Jian Wang, Ruimin Shen
The Department of Computer Science and Engineering,
Shanghai Jiaotong University, Shanghai 200030, China
{ullrich_c, jwang, rmshen}@sjtu.edu.cn

Frank Quosdorf
Totuba Inc.
Shanghai, China
frank.quosdorf@totuba.com

Abstract

The Web is still changing at a rapid rate. Principles that define the Web 2.0 are now better understood than a few years ago, but continue to evolve. The Semantic Web is coming of age, resulting, for instance, in a significant amount of data now being available as Linked Open Data. Advanced Web services make sophisticated functionality such as entity extraction and location-data available for free. Against this background, we feel that the potential of today's Web for intelligent technology-supported learning is under-exploited. Especially for AI exciting research opportunities arise from building on top of available data and services. In this paper, we sketch the current state of the art of Web data and services, highlight potential future development and point out opportunities and challenges for Web-based technology-supported learning.

1. Introduction

The amount of new development in today's Web, with new Web sites, services and standards being published with high frequency, raises the question which of these developments are meant to last and worth investigating from a research perspective. This question is especially relevant since some of today's Web developments have the potential to advance significantly the current state of the art in AI for technology-supported learning. We feel, however, that the potential is still under-exploited, especially with regard to data and services that are offered by today's Web 2.0 and Semantic Web.

In Section 2 we analyze the current trends of both Web 2.0 and Semantic Web and highlight those characteristics that we think are of relevance today, especially for technology-supported learning: an architecture of assembly allows new ways of integrating and combining services; the Web is becoming more centered on the individual user; and semantics improve existing and allow powerful new services. These sections are not a review of state-of-the-art *research*, but of *results* of prior research—that is, of results available for usage today and that provide excellent opportunities that research in AI for Education can exploit today. We illustrate

these trends with an example in Section 3 and conclude in Section 4.

2. Trends

The Web is evolving at a fast rate and it is unclear which of the current trends will significantly determine the shape of the Web and which will simply fade away. In this section, we will analyze those trends in the current development of the Web we consider being highly relevant for the future research on technology-supported learning today. We start with clarifying what precisely we mean with Web 2.0 and Semantic Web, since often quite different ideas are associated with these terms.

Especially “Web 2.0” is not particularly well defined. In this paper, we follow the conceptualization made popular by [1], who use it to characterize applications that have in common the stimulation of individual creativity, the massive scale of contributions and of available data, and technical aspects such as an architecture of assembly that enables easy access to data and functionality of, thereby allowing other services to build on existing functionality.

The vision of the Semantic Web (also called “Web of Data” and sometimes even “Web 3.0”), a Web of data understandable by machines, was sketched in [2]. About half a decade later, the basic technologies of the SW have been developed and standardized [3], and one impressive sign of the coming of age of the Semantic Web is the rapidly growing amount of data available in RDF.

2.1. Architecture of Assembly

The “architecture of assembly” is a term we use to characterize the fact that software is becoming increasingly interoperable, i. e., functionality of a service is made available to other services. The concept itself is not a new one, as middleware gained popularity in the 1980s, and today Web services are an active area of work of the World Wide Web Consortium (W3C).

However, assembly and sharing of functionality has reached new heights in the Web 2.0. First of all, machine to machine access is provided by APIs (Application Programming Interfaces) that allows to add, change, and retrieve the

data available in a service. In difference to traditional Web-services, the APIs are characterized by pragmatic solutions and lightweight formats, thus are straightforward to use. Secondly, Web 2.0 offers previously unavailable vast amounts of diverse data. Photos, geo-data, bookmarks, videos and other kinds of resources are available in the Web and often can be reused for remixing, i. e., adding, extending and modifying of resources. Since most of this content is produced by users, Web 2.0 services have an interest to stimulate the production of content by making their services as easy to use as possible as well as by making them open. Users want to be able to access their and other users' content in a multitude of ways, through the services Web interface but also through other means.

A good example is Twitter, a micro-blogging site (<http://twitter.com>). Twitter allows its users to post short messages that are distributed to their friends. The traffic from Twitter's API is 10 times the regular Web traffic. This traffic is created by tools written by third parties: desktop applications, visualization tools and other services that Twitter by itself would not have been able to create due to limited manpower [4].

In addition to APIs, exchange of data is facilitated by RSS (Really Simple Syndication). RSS is a family of standards (RSS/Atom) originally developed by Netscape for publishing a list (feed) of new and updated items such as news articles. Users can subscribe to RSS feeds in RSS readers (aggregators) and thereby be informed about updates without visiting the Web site itself. Today, news articles are just one type of data which is syndicated by RSS; other include photos, blog entries, bookmarks on social bookmarking sites. Having such a common format for information about data greatly simplifies the architecture of assembly. For instance, an RSS feed of the latest Twitter feeds can be combined with a feed of photos from the photo community Flickr in order to automatically add pictures to the updates. Since both sites syndicate their content as RSS feeds, combining the data is very easy.

Additionally, existing Web 2.0 services often disseminate their functionality by plug-in modular components, so called widgets [5]. This allows integrating the service on a given Web-site by drag&drop or adding only a few lines of code.

Combining existing services to achieve new functionality has become mainstream, too. An example of a generic assembly platforms is the mash-up engine Yahoo! Pipes (<http://pipes.yahoo.com>). This service allows a drag&drop construction of applications that integrate functionality from existing services without any programming knowledge. Semantic versions are currently researched [6] but not yet ready for daily application.

A previous project [7] taught us that the architecture of assembly is not just a marketing term of Web 2.0 advocates, but that its techniques facilitate the construction of more complex services significantly. In this project, we were able to increase the interactivity of a language learning course

by integrating a micro-blogging service. Due to open APIs we were able to completely reuse the services of an existing service.

Several opportunities for AI, especially for application in education, arise from the architecture of participation: access to huge amounts of data that can be retrieved in order to support the learning process; analyzing data in a multitude of formats produced by learners in order to make estimations about their knowledge/competencies; and finding services and combining of services to help the student to train specific skills.

2.2. User-Centered Web

The first trend, the architecture of assembly adds to the Web of documents a layer of specific, reusable services. The next big trend, the user-centered Web, focuses on the individual who is using these documents and services.

Today's Web is mostly a Web of connected *documents*. The emergence of social networks such as Facebook shows a new trend, to a Web that connects *people*, resulting in a social graph. Facebook is a fine example that shows how this social graph might change how we interact with the Web: a feature of Facebook called "life streaming" presents a summary of the activities of a user's friends, e. g., their postings, travels, usage of applications, books they read, etc. With one glance, a user is updated about all actions within his or her network. This adds a social layer on top of the existing Web.

The user-centered Web will require the integration of data that is currently dispersed over several networks. Several standards and organizations were created to meet the demand of improved data management. OpenID [8] provides a decentralized single-sign-on mechanism; information exchange between services can be handled by OAUTH (<http://oauth.net/>). Additional efforts to improve data management and portability exist, notably Google's OpenSocial (<http://code.google.com/apis/opensocial/>) that provides a set of APIs for common functionality for social applications.

The data about the user that will be relevant for the user-centered, personalized Web of the future is not restricted to his social graph. Personalization requires information about the activities of the individual, a fact long known in the user modeling community (for a recent overview, see [9]). The Web 2.0 "rediscovery" of the user model is called *Attention Data*. Attention data includes browsed Web sites, read blog entries, music listened to, in short all activity that is collected by some service in order to capture the users' interests (and their dislikes). This "click-stream" can be used to compile a set of interests. Again, a standard is arising for representing this compiled attention data, namely the Attention Profiling Markup Language (APML, <http://www.apml.org>). This data can already today be used to present a personal view of the Web.

Personal information is not restricted to interests, but also includes the current position of a user. Today, this information is readily available. Google's browser extension Google Gears and mobile location services such as Brightkite (gears.google.com, <http://brightkite.com>) provide information about the position of the user to Web applications at no cost.

Usage of such data for personalization and adaptivity was pioneered in the fields of Adaptive Hypermedia [10] and e-learning. Advanced learning environments such as ACTIVE MATH [11] have a representation of the learners' competencies, goals and preferences and use it to select those educational resources best suited for the individual learner. However, currently such systems are restricted to the activities that happen within the environment, since they cannot track the learners' activities outside the system. Now, sufficient "outside" data is available to make it attractive to research how to use social information as captured by attention data and attention profile, such as the books bought via Amazon, the membership and the frequency of visits of groups in social networks etc., in order to provide an improved personalized learning experience. This is an ideal playing ground for research in AI, for instance recommender systems.

This trend will even be supported by the browser itself. The Mozilla foundation in charge of the development of the Firefox browser envisions Firefox as an information broker that analyzes the data within documents and then suggests suitable actions how to process the data [12]. For instance, address information in a Web page can be exported to the address book or plotted on a map. This process is supported by metadata, i. e., annotations of parts of the content using formats such as RDFa [13]. This information will be used by the browser to further enhance the browsing, e. g., by looking up additional information in Wikipedia or Amazon. This is a fundamental shift, since then a browser will propose sensible actions about the object that is describe in a page. The shift in focus of the Web, from documents (text) to what the texts describe (semantics), will be discussed in detail in the following section.

2.3. Semantics

Thanks to the Web it is not longer a problem to access information, rather it is problematic to find the specific piece of information that is relevant for the current situation, or, even more difficult, to combine existing pieces to provide an answer. Semantics, i. e., the possibility of the computer to understand what a document is about, promises a solution to this problem.

Over the last years, the Linked Open Data Initiative [14] has led to a significant number of datasets available in RDF format. The datasets include Wikipedia, DBLP RKB Explorer, the CIA World Fact Book, OpenCyc and other,

and is growing exponentially. The data is interlinked and provides a never seen before database that can be exploited by intelligent applications.

Similarly, advanced semantic services become publicly available. The most prominent example is Calais (<http://opencalais.com>), a company and service owned by Thomson Reuters. OpenCalais performs entity extraction, based on a rich set of entities build from Reuters news archive. Calais offers a Web service that takes a document as input and extracts a set of metadata that describe the document as a result. The usage is for free, even for commercial application and is fast enough to be used in Web applications on the fly.

Semantic is a means to an end. The hope is that by making machines "understand" what a resource is about, the quality of services that process the resource can be improved, resulting in a better categorization, recommendation, automatic linking with other resources, etc. [15].

3. Example

In this section, we present the example of a prototype application that is based on some of the services discussed above. This application, called the Totuba Toolkit (henceforth TT), provides services which were extremely expensive and time consuming to achieve only a few years ago. Today, thanks to the trends shown in the previous section, we were able to build a first prototype in a few person-weeks effort. The point we want to make for this paper is that the "basic" functionality we achieved by reusing existing services and data would have been prohibitively expensive only a few years ago. Today, we can reuse the functionality and thus face the new research questions that arise in such an advanced setting.

The goal of the TT is to support scientific research of novices as well as experts. Based on user input (keywords, but also whole sentences and paragraphs of text), the TT actively suggests resources that provide the users with relevant information about their current research topic. The suggestions include resources from different information sources and related topics.

The TT functionality is achieved by employing named entity recognition [16] and querying linked open data [14]. Named entity recognition is used to identify the main topics from the input text in case the users input more than a single keyword, and the type of these topics. Web services that offer such a functionality are Yahoo! Term Extraction and OpenCalais. While the former returns only a list of terms, OpenCalais returns very detailed information, including the types of the identified entities, relations between entities, and also a link to the equivalent entity in DBpedia.

Once the TT has identified the relevant topics, it looks up related topics. This is done by querying a DBpedia SPARQL endpoint (SPARQL is the query language for the Semantic Web endorsed by the W3C; a SPARQL endpoint is a Web

service that processes SPARQL queries [17]). For each topic t , the query retrieves all outgoing and incoming relations of t , as well as the corresponding topics. The TT user interface renders these relations as links. Thus, the user can decide which of the relations to explore.

It is important to realize that the existence of DBPedia and similar data sources is very much a game changing fact. For a long time, AI in general suffered from the lack of freely available data about the common world. Today, the situation is quite contrary—there exist several data sources that fill that gap, e. g., Freebase (<http://www.freebase.com/>), umbel (<http://www.umbel.org/>), and DBPedia (dbpedia.org/). At the time the project was done, we selected DBPedia due to its open SPARQL endpoint that allowed us to query its content without having to set-up our own data mirror. On the negative side, the DBPedia data set is not particularly clean, for instance differently named relations exist that actually are semantically equivalent.

Furthermore, the TT retrieves scientific articles that provided additional information about the identified topics from scholarly databases. It then shows the links to the articles in the user interface. For that purpose, the TT queries one of the RKBExplorer repositories [18] for articles whose titles contain the topics. Another suited data source is the DBLP Bibliography Database (<http://dblp.l3s.de/d2r/>), but this server truncates results over a certain limit and thus could not be used for our purpose. Again, we could have downloaded the datasets and set up our own database, but the existence of the SPARQL endpoints saved us this work while developing the prototype.

The above described system achieves a functionality for free that was hard to achieve a few years ago: it semantically analyzes a given text, and presents related topics and articles. Once we set up that basic system, the “real” questions emerged: which of the many relations represented in DBPedia are of interest to the current goal of the user? Which of the many resources from the cloud of linked open data that are linked to the current topics should be shown to the user? How to present them in the user interface? We do not argue that no work has been done on these topics before, quite the contrary.

Significant research has been invested into determining how to support users while learning. One example of such a service is PAIGOS [19]. It is embedded in the learning environment for mathematics ACTIVEMATH. PAIGOS automatically generates courses adapted to the competencies, preferences and learning goal of the student. The courses include learning materials as well as services such as a concept mapping tool. The generation is based on explicitly represented pedagogical knowledge that is applied in an AI planning framework. PAIGOS illustrates the potential that lies in semantics: due to its knowledge about resources and tools, it can generate in-depth courses, courses suited for rehearsal, exam preparations etc. However, the content that

PAIGOS uses is stored within the ACTIVEMATH database, filled with learning objects that were very time consuming to author. Now, thanks to advances in the Semantic Web and Linked Open Data in particular, a significant amount of semantic data is suddenly available freely in the Web. Thus, today is the time where such research is no longer only within the reach of very specialized research groups, but are now within the reach of those who previously had not the technical competencies or means to build their own semantic services.

Obviously, some minor problems still remain. We did not mirror the data in our own databases but accessed the open SPARQL endpoints. Thus, we depend on the services being online and often the response times were quite high.

4. Conclusion

In this paper, we gave an overview of current trends in the Web and on data and services that are ready to be used today.

We believe that in the future the Web 1.0, Semantic Web, Web 2.0 and the mobile Web (which we did not cover in this article) will coexist: the static pages of the Web 1.0 will not vanish, since they are easy to set up and maintain. Web 2.0 will continue to grow, with a focus on a socially centered Web. The Semantic Web will be invisible, a hidden layer in the Web with which users will not come in direct contact, but they will experience a significant improvement of functionality, wherever it is used as an underlying technique.

The example of the TT showed that not long ago ground breaking services and data have now become basic services. And we argue that this is a level that deserves more attention from today’s AI research in technology-enhanced learning.

During our work we also faced a severe obstacle: most of the existing services are designed for the English language. OpenCalais for instance, will process other languages only a later milestone. This gives research groups that have to work in languages other than English a huge disadvantage, which in the long run will hurt research in this area. This also clearly shows the importance of making research results available to other parties, either as a Web service or as open source.

Acknowledgment

This article is sponsored by the National Natural Science Foundation of China (grant no. 60672066), National High-tech Research and Development Program of China (grant no. 2007AA01Z157), by the European Commission within the Role IP (Grant no.: 231396), and by Totuba Inc..

References

- [1] T. O'Reilly, "What is Web 2.0: Design Patterns and Business Models for the next generation of software," O'Reilly Media Inc., 2005, last accessed 18.10.2007. [Online]. Available: <http://www.oreillynet.com/pub/a/oreilly/tim/news/2005/09/30/what-is-web-20.html>
- [2] T. Berners-Lee, J. Hendler, and O. Lassila, "The semantic web," *Scientific American*, vol. 284, no. 5, pp. 34–43, 2001.
- [3] N. Shadbolt, T. Berners-Lee, and W. Hall, "The semantic web revisited," *IEEE Intelligent Systems*, vol. 21, no. 3, pp. 96–101, 2006.
- [4] S. Ammirati, "Biz Stone—co-founder Twitter," 2007, this is an electronic document. Date of publication: September 5, 2007. Date retrieved: February 13, 2008. [Online]. Available: <http://readwritetalk.com/2007/09/05/biz-stone-co-founder-twitter/>
- [5] M. Caceres and A. van Kesteren, "Widgets 1.0," W3C, W3C Working Draft, Oct. 2007, <http://www.w3.org/TR/2007/WD-widgets-20071013/>.
- [6] D. L. Phuoc, A. Polleres, G. Tummarello, C. Morbidoni, and M. Hauswirth, "Rapid semantic web mashup development through semantic web pipes," in *Proceedings of the 18th World Wide Web Conference (WWW2009)*. ACM Press, 2009, pp. 581–590.
- [7] C. Ullrich, K. Borau, H. Luo, X. Tan, L. Shen, and R. Shen, "Why Web 2.0 is Good for Learning and for Research: Principles and Prototypes," in *Proceedings of the 17th International World Wide Web Conference*. ACM, 2008, pp. 705–714. [Online]. Available: <http://www.carstenullrich.net/pubs/Ullrich08Why.pdf>
- [8] D. Recordon and D. Reed, "OpenID 2.0: a platform for user-centric identity management," in *DIM '06: Proceedings of the second ACM workshop on Digital identity management*. New York, NY, USA: ACM, 2006, pp. 11–16.
- [9] P. Brusilovsky and E. Millán, "User models for adaptive hypermedia and adaptive educational systems," in *The Adaptive Web*, ser. Lecture Notes in Computer Science, P. Brusilovsky, A. Kobsa, and W. Nejdl, Eds., vol. 4321. Springer, 2007, pp. 3–53.
- [10] P. Brusilovsky, "Adaptive hypermedia," *User Modeling and User Adapted Interaction*, vol. 11, no. 1/2, pp. 87–110, 2001.
- [11] E. Melis, G. Gogvadze, M. Homik, P. Libbrecht, C. Ullrich, and S. Winterstein, "Semantic-aware components and services of ActiveMath," *British Journal of Educational Technology*, vol. 37, no. 3, pp. 405–423, 2006. [Online]. Available: <http://www.carstenullrich.net/pubs/Melisetal-SemanticAware-BJET-2005.pdf>
- [12] A. Faaborg, "Microformats - part 0: Introduction," 2006, this is an electronic document. Date of publication: November 11, 2006. Date retrieved: February 15, 2008. [Online]. Available: <http://blog.mozilla.com/faaborg/2006/12/11/microformats-part-0-introduction/>
- [13] M. Birbeck and B. Adida, "RDFa primer," W3C, W3C Working Draft, Oct. 2007, <http://www.w3.org/TR/2007/WD-xhtml-rdfa-primer-20071026/>.
- [14] C. Bizer, T. Heath, K. Idehen, and T. Berners-Lee, "Linked data on the web (ldow2008)," in *WWW '08: Proceedings of the 17th international conference on World Wide Web*. New York, NY, USA: ACM, 2008, pp. 1265–1266.
- [15] L. Sauermaun, A. Bernardi, and A. Dengel, "Overview and outlook on the semantic desktop," in *Proceedings of the 1st Workshop on The Semantic Desktop at the ISWC 2005 Conference*, S. Decker, J. Park, D. Quan, and L. Sauermaun, Eds., 2005. [Online]. Available: <http://www.dfki.uni-kl.de/~sauermaun/papers/Sauermaun+2005d.pdf>
- [16] H. Cunningham, Y. Wilks, and R. J. Gaizauskas, "Gate: a general architecture for text engineering," in *Proceedings of the 16th conference on Computational linguistics*. Morristown, NJ, USA: Association for Computational Linguistics, 1996, pp. 1057–1060.
- [17] A. Seaborne and E. Prud'hommeaux, "SPARQL query language for RDF," W3C, W3C Recommendation, Jan. 2008, <http://www.w3.org/TR/2008/REC-rdf-sparql-query-20080115/>.
- [18] H. Glaser, I. Millard, and A. Jaffri, "Rkbexplorer.com: a knowledge driven infrastructure for linked data providers," in *European Semantic Web Conference*, vol. 5021/2008. Springer, June 2008, pp. 797–801. [Online]. Available: <http://eprints.ecs.soton.ac.uk/15152/>
- [19] C. Ullrich, *Pedagogically Founded Courseware Generation for Web-Based Learning – An HTN-Planning-Based Approach Implemented in PAIGOS*, ser. Lecture Notes in Artificial Intelligence. Springer, September 2008, no. 5260, iSBN 978-3-540-88213-8. [Online]. Available: <http://www.springerlink.com/content/k604618p5351>