Adaptive Navigation through Semantic Annotations and Service Descriptions

Ruben Verborgh

supervised by
Rik Van de Walle

Ghent University – iMinds – Multimedia Lab
Gaston Crommenlaan 8 bus 201, 9050 Ghent, Belgium
ruben.verborgh@ugent.be

Abstract. While hyperlinks are absolutely crucial to the Web's success, they are currently uni-directional, as information is augmented with controls from the perspective of the information publisher. However, it is the user who needs those links to navigate—and the publisher cannot know how any user might want to interact with the information. Therefore, the most relevant links for a user might be omitted, severely limiting the applicability of hypertext. In this paper, I outline a plan to tackle this problem as part of my doctoral research, by explaining the research questions, the underlying hypotheses, and the approach, in which semantic technologies play a crucial role.

Keywords: affordance, hypermedia, Semantic Web, Web services

1 Problem Statement

The hyperlink-driven information model of the World Wide Web [3] has introduced humanity to a novel way of information consumption. Information has become *actionable* [11], in contrast to the passive medium it used to be. Although hypertext was envisioned long before [15], the Web was the first hypertext system that worked on a global scale. Still, the initial visions of hypermedia featured a much richer arsenal of link mechanisms, such as non-breaking n-way links and transclusion [16].

The main problem of hyperlinks on the Web is that the responsibility of link creation lies entirely with the publisher. Indeed, when creating a Web site or application, it is the publisher of the information who decides what actions the consumers of that information can perform. This poses a threat to the said actionability of the information, because it can only be called "actionable" to the extent the consumer can actually perform the actions he or she desires. If the publisher omits the hyperlinks that afford these actions, then the hypertext document becomes as passive as any pre-Web medium, defeating its purpose.

The problem statement of my doctoral research is therefore how we can enhance the controls in hypermedia documents on the Web in a personalized way,

such that they afford the relevant actions for each consumer. I want to look at this problem from the standpoint of both human visitors of websites and machine clients of Web APIs, as they each have unique challenges. The aim is to provide these controls in an automated way with the least possible amount of human intervention.

In the next section, I will explain the relevancy of this problem, followed by a summary of related work in Section 3. Section 4 poses several research questions, for which I formulate hypotheses in Section 5. My approach is presented in Section 6, followed by a reflection in Section 7. Section 8 discusses the evaluation and preliminary results are shown in Section 9.

2 Relevancy

It is crucial to realize that hyperlinks—and hypermedia controls in general—are not *enablers* but *affordances* [11,17], *i.e.,* they offer the information consumer an action possibility, but the action itself is also achievable through other means. For instance, if a document refers to another page without a hyperlink, that page might still be accessible (*e.g.,* through a search engine). However, this catapults us back to the age of paper documents, as the document does not directly contain the needed affordances, defeating the purpose of hypertext. Furthermore, such indirect ways are more time- and resource-consuming, while the ability to browse quickly in a hypermedia system is critical to its usability [2]. Thus, having the right hypermedia controls in place is necessary for efficient Web use. Three reasons in particular make the need for more relevant controls very actual.

Continued expansion The Web is growing at an ever increasing rate, which means that if the average number of links per page does not increase, the connectedness of the Web is decreasing. In 1999, the size of the Web was estimated at 800 million documents and the average document distance already at 19 clicks [19]—and it is not hard to imagine that this distance will only go up. This means that the trade-off between *completeness* (a publisher offering all relevant links) and *efficiency* (a consumer quickly finding the link she needs) becomes increasingly difficult to manage.

Aging documents Even if publishers could somehow strike an ideal balance between completeness and efficiency at the time of publication, it is highly unlikely that their choice of links will remain optimal as the document ages. Since the Web does not implement the concept of backlinks, the fact that new documents can link to older ones does not improve the affordance of the older documents. Furthermore, links to certain pages break if documents move or cease to exist [14]. Therefore, the hypermedia controls found on older documents are not the most relevant ones for a user. For instance, given the current trend of online social networks, many recent sites provide interaction controls with those networks. When a user browses older pages, these controls can be missed, especially by users who lack the necessary skills to perform these actions without the help of such designated controls.

Mobile growth In the past few years, mobile use of the Web has increased tremendously and will continue to do so in the coming years [10]. The nature of mobile devices makes the Web browsing experience different: average screen sizes are smaller and physical keyboards are either miniaturized or not present. The small screen size leaves less space for links, so the few that appear must be relevant. Also, the absence of a full-scale keyboard makes it more difficult to reach a goal in indirect ways (such as entering keywords in a search engine) if a direct link that leads to the user's desired goal is missing.

3 Related Work

3.1 Adaptive Hypermedia

The personalization of hypermedia documents is part of the research field of *adaptive hypermedia* [6]. Within adaptive hypermedia research, *adaptive navigation support* [8] is concerned with personalizing hypermedia controls to match the intentions or goals of the user. Most adaptive navigation support systems are *a)* operating on a closed corpus, *b)* focused on linking to related information, and *c)* used in a specific context such as learning. In contrast, we want to approach the problem statement *a)* on the open corpus of the entire Web, *b)* with a focus on performing actions, and *c)* for day-to-day usage. In fact, *open-corpus adaptive hypermedia* has been identified as an important challenge in the field [7], but it has not been tackled intensively. Semantic Web technologies were listed as a candidate to help overcome the open-corpus problem on the Web [9].

3.2 Social Interaction Widgets

In response to the sudden rise of many social networks, publishers started adding so-called *widgets* to their sites, small snippets of code that provide user interactions. In contrast with simple hyperlinks, which connect one document to another, these widgets typically allow the user to perform an *action* on the current document, such as sharing it within a social circle or adding a comment. Examples include the Facebook *Like* button and the Twitter *Tweet* button [5]. However, as more social networks emerge, it becomes impractical for publishers to manually add widgets for each of them.

In order to cope with this increasing diversity in social networks, aggregated widgets such as AddThis [1] were created. AddThis is a single widget that gives access to sharing options on various social networks through a list that every user can personalize. The benefit on the publisher side is that he does not need to know about the user's preferred network, nor must provide a sharing link. Additionally, the user is not bothered by non-relevant sharing links, because social networks that she does not use are not needlessly displayed by the widget. However, the usage of AddThis is limited to sharing actions, and it is thereby not a generalized solution for personalized action links on the Web.

3.3 Web Intents

A solution that does support a wider range of actions is Web Intents [4]. The idea derives from a concept on phones with the Android operating system, where applications can indicate their *intent* to support a certain action, such as calling or sending a message. The specification defines several standard actions, such as sharing, editing, and viewing, which can be supported by Web applications. Content publishers should indicate which actions their users can perform. However, this still implies the publisher must "predict" what type of action the user might want to do [21]. In contrast, we want to determine the action through the user's preferences, which are highly personal and can change over time.

4 Research Questions

The main question in my doctoral research is:

> *How can we enhance hypermedia controls in a personalized way, such that they complement a piece of information with the affordance a user requires to perform the next steps he or she needs?*

This question gives rise to two others. On the one hand, there is the human aspect:

> *How does such enhanced affordance help users browsing the Web in achieving their goals, and how can we achieve maximum effectiveness in this regard?*

And on the other hand, there is the machine aspect:

> *How does enhanced affordance help machine clients consume Web APIs, and can it lead to true serendipitous reuse [23] of services?*

This last question is inspired by the concept *hypermedia as the engine of application state* [12], which aims to achieve loose conversational coupling [18] by augmenting representations with controls, also for machine clients. However, as is the case with humans surfing the Web, publishers of information are unaware of the goals of the information consumer, and it is therefore hard for them to provide the necessary affordance [21].

5 Hypotheses

The main hypothesis, related to the main research question, is:

> *Current semantic technologies are sufficiently flexible and powerful to connect a piece of information and matching actions at runtime, while providing loose coupling at design time.*

In addition to the feasibility, the following hypothesis relates to the necessity:

> *Semantic technologies offer an added value to the creation process of enhanced affordance.*

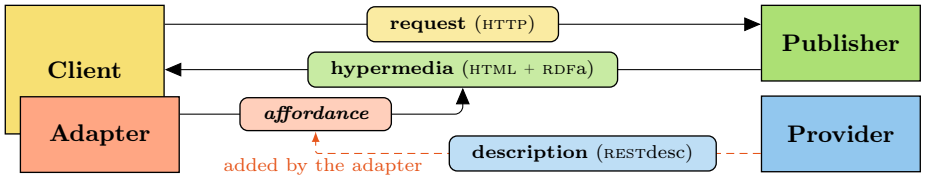

Fig. 1. An adapter at the *client* side adds affordance to the hypermedia representation, based on the semantic annotations the latter contains (*e.g.*, RDFa or HTML5 microdata).

My hypotheses regarding the effectiveness of hypermedia documents enhanced with personalized affordance are:

> *Users can browse the Web faster and more efficient when the relevant affordance is in place.*
>
> *Machine clients will be more functional and more resilient to change if they receive messages with enhanced affordance.*

6 Approach

My approach to address the research questions is to develop a technology and architecture for what I call *distributed affordance* [21]. The core idea is that publishers add *semantic annotations* to the hypermedia documents they serve to a client, which are matched at runtime by the client to *semantic service descriptions* that describe the *functionality* offered by providers of the user's choice. Figure 1 shows a client making a request to an information publisher, who replies with an HTML document that has been enriched with RDFa markup. Earlier, a client-side adapter has accessed functional descriptions in RESTdesc [22] format, which are now instantiated with the RDFa annotations to generate affordances that can be added to the HTML document. These affordances will lead to actions that operate directly on the resources inside the page.

This addresses the main research question, and will also put the main hypothesis to the test. The proposed platform is loosely coupled, as the publisher, provider, and adapter do not need to know about each other. Instead, the publisher and provider offer sufficient semantics for the adapter to interpret what combinations are possible. This contrasts with current hyperlinks on the one hand, where the publisher has to know about the action provider, and with closed-corpus adaptation on the other hand, where the adapter has to know about the publisher and/or the action provider.

To address the research questions and hypothesis about user and machine usage of the distributed affordance platform, I will create a corpus consisting of websites with annotations and Web services with descriptions. The goal is to incorporate as many real-world examples as possible, in order to have a realistic testing environment, as well as several use cases wherein the technology proves its added value.

7 Reflections

The main difference in my approach with existing work on adaptive navigation support [8], is that I specifically want to perform adaptation on an *open* corpus, *i.e.,* the entire Web, instead of a controlled subset. Open corpus adaptive hypermedia has in fact been identified as an important challenge [7], and semantic technologies have been hinted at as a potential solution [9], although no concrete systems have emerged yet. However, I believe that the Web only recently is becoming ready for this, as it is only in the past few years that we see sufficient semantic annotations appear, despite the technologies (such as RDFa) being around for a longer time. Furthermore, my previous experimentation with functional description of REST APIs [22] gives me the confidence that this technology is sufficiently mature to apply it on automated action creation using even limited semantic annotations.

A second difference, as indicated in Figure 1, is that the adaptation happens at the *client* side and therefore is fully scalable, instead of classical adaptation systems that rely on a central adaptation component. Furthermore, whereas the majority of work on adaptive navigation focuses on linking static documents together, my goal is to connect *dynamic* documents, *i.e.,* generate links towards actions on the current document. This allows for the creation of much richer interactions.

8 Evaluation Plan

For the evaluation, there are three main lines of interest.

user studies As the main focus is on creating affordance that will help users browse more efficiently, it is of utmost importance to conduct user studies that follow people's browsing behavior as they make use of the developed platform. Several tests should be conducted in a (double-)blind setting, where the participant (and possibly the experiment conductor) are unaware whether the platform is active. I will analyze qualitative parameters on the one hand, such as the user's impression of effectiveness, and quantitative parameters on the other hand, such as the time to complete a task.

performance evaluation In a platform that should manipulate web pages in real-time, speed will be crucial. Therefore, various aspects of the platform should be tested for performance, especially the semantic matching of content and services, which can become complex quickly. In addition to that, the whole pipeline must be tested, and optimized so that it stays under the threshold that is deemed acceptable in the user studies.

client code complexity Finally, as I also want to focus on automated clients of Web APIs, the code of such clients should be less complex [13] as a result of the enhanced affordance. This makes it necessary to compare the implementations of clients with and without the use of (additional) hypermedia controls in the server's response.

9 Preliminary Results

Previously, I have evaluated the performance of Web API matching and composition with RESTdesc [20], which led me to conclude two things. First, it is possible to create relevant composition chains in a few milliseconds. Second, this level of performance can be maintained even with thousands of different descriptions of Web site actions that are potential matches. Together, this indicates that finding the few API descriptions that match a given resource out of a large description set is possible in a reasonable amount of time.

Additionally, I have started a user study in collaboration with researchers from the human-computer interaction field, in which we observe users as they perform tasks on the Web with and without the distributed affordance platform. The first results seem to suggest that users navigate indeed more efficiently when the affordance has been optimized for their needs. The distributed affordance platform itself is currently under development, the progress of which can be followed at `http://distributedaffordance.org/`.

10 Conclusion

This paper has presented the outline of my doctoral research, in which I want to focus on personalized affordance created from distributed sources. My approach is to build a platform that works on the client side, enhancing hypermedia representations returned by the server with hyperlinks and controls of the user's preference. Semantic technologies enable a loose coupling between publishers of information and action providers, which allows the platform to have a truly distributed nature. The evaluation of this work will consist of user studies, performance evaluations and code complexity comparisons. My goals are to make browsing the Web more efficient for users, and to enable a more serendipitous reuse of services for machines.

Acknowledgements The described research activities were funded by Ghent University, the Institute for the Promotion of Innovation by Science and Technology in Flanders (IWT), the Fund for Scientific Research Flanders (FWO Flanders), and the European Union.

References

1. AddThis, `http://www.addthis.com/`
2. Akscyn, R.M., McCracken, D.L., Yoder, E.A.: KMS: a distributed hypermedia system for managing knowledge in organizations. Communications of the ACM 31(7), 820–835 (Jul 1988)
3. Berners-Lee, T., Cailliau, R., Groff, J.F.: The world-wide web. Computer Networks and ISDN Systems 25(4–5), 454–459 (1992), `http://www.sciencedirect.com/science/article/pii/016975529290039S`
4. Billock, G., Hawkins, J., Kinlan, P.: Web Intents. W3C working draft (Jun 2012), `http://www.w3.org/TR/web-intents/`
5. Bodle, R.: Regimes of sharing. Information, Communication and Society 14(3), 320–337 (Apr 2011)

6. Brusilovsky, P.: Methods and techniques of adaptive hypermedia. User Modeling and User-Adapted Interaction 6(2–3), 87–129 (1996), http://dx.doi.org/10.1007/BF00143964
7. Brusilovsky, P.: Adaptive hypermedia. User Modeling and User-Adapted Interaction 11(1–2), 87–110 (2001), http://dx.doi.org/10.1023/A%3A1011143116306
8. Brusilovsky, P.: Adaptive navigation support. In: Brusilovsky, P., Kobsa, A., Nejdl, W. (eds.) The Adaptive Web, pp. 263–290. Springer-Verlag (2007), http://dl.acm.org/citation.cfm?id=1768197.1768207
9. Brusilovsky, P., Henze, N.: Open corpus adaptive educational hypermedia. In: Brusilovsky, P., Kobsa, A., Nejdl, W. (eds.) The Adaptive Web, pp. 671–696. Springer-Verlag, Berlin, Heidelberg (2007), http://dl.acm.org/citation.cfm?id=1768197.1768223
10. Cisco: Visual networking index: Global mobile data traffic forecast update, 2012–2017 (Feb 2013), http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-520862.pdf
11. Fielding, R.T.: REST APIs must be hypertext-driven. Untangled – Musings of Roy T. Fielding (Oct 2008), http://roy.gbiv.com/untangled/2008/rest-apis-must-be-hypertext-driven
12. Fielding, R.T., Taylor, R.N.: Principled design of the modern Web architecture. Transactions on Internet Technology 2(2), 115–150 (May 2002)
13. Lanning, D., Khoshgoftaar, T.: Modeling the relationship between source code complexity and maintenance difficulty. Computer 27(9), 35–40 (1994)
14. Markwell, J., Brooks, D.W.: Broken links: the ephemeral nature of educational WWW hyperlinks. Journal of Science Education and Technology 11(2), 105–108 (2002)
15. Nelson, T.: Complex information processing: a file structure for the complex, the changing and the indeterminate. In: Proceedings of the ACM 20th National Conference. pp. 84–100. ACM, New York, NY, USA (1965)
16. Nelson, T.: Dream machines. self-published (1974)
17. Norman, D.A.: The Design of Everyday Things. Doubleday, New York (1988)
18. Pautasso, C., Wilde, E.: Why is the Web loosely coupled? – A multi-faceted metric for service design. In: Proceedings of the 18th international conference on World Wide Web. pp. 911–920. ACM, New York (2009), http://doi.acm.org/10.1145/1526709.1526832
19. Reka, A., Jeong, H., Barabasi, A.L.: Diameter of the World-Wide Web. Nature 401(6749), 130–131
20. Verborgh, R., Haerinck, V., Steiner, T., Van Deursen, D., Van Hoecke, S., De Roo, J., Van de Walle, R., Gabarró Vallés, J.: Functional composition of sensor Web APIs. In: Proceedings of the 5th International Workshop on Semantic Sensor Networks (Nov 2012), http://ceur-ws.org/Vol-904/paper6.pdf
21. Verborgh, R., Hausenblas, M., Steiner, T., Mannens, E., Van de Walle, R.: Distributed affordance: An open-world assumption for hypermedia. In: Proceedings of the Fourth International Workshop on RESTful Design (May 2013), http://distributedaffordance.org/publications/ws-rest2013.pdf
22. Verborgh, R., Steiner, T., Van Deursen, D., Coppens, S., Gabarró Vallés, J., Van de Walle, R.: Functional descriptions as the bridge between hypermedia APIs and the Semantic Web. In: Proceedings of the Third International Workshop on RESTful Design. pp. 33–40. ACM (Apr 2012), http://www.ws-rest.org/2012/proc/a5-9-verborgh.pdf
23. Vinoski, S.: Serendipitous reuse. Internet Computing 12(1), 84–87 (2008), http://steve.vinoski.net/pdf/IEEE-Serendipitous_Reuse.pdf