The Rise of the Verb

Paul Groth

p.t.groth@vu.nl
VU University Amsterdam
De Boelelaan 1081a,
1081 HV, Amsterdam, The Netherlands

Abstract. In the next 10 years, we will see a Semantic Web that is infused with a richer set of verbs: the ability not just to represent knowledge about static datasets but the ability to use knowledge to perform actions or operations. We argue that there are three trends that make this outcome likely, namely, demand from current web applications (e.g. Facebook's Like), the ubiquity of Javascript and the increasing instrumentation of the real world.

To Siri (now): "Move my meeting from 3 to 4"
To Siri (10 years): "Mow the lawn"

One of the most exciting things when you use Apple's Siri for the first time is not just when it finds the correct information but when it actually does something, for example, rescheduling a meeting or setting a reminder. In this short paper, I will argue that the next 10 years of the Semantic Web will focus on "doing stuff" or, put another way, the *verb* will dominate. In the past 10 years, we have seen a massive explosion in semantically structured data, from the world ontology wikipedia to the structured mark-up of content through languages like RDFa and Cascading Style Sheets. This introduction of semantics to the Web has primarily improved search - whether that is improving the ability to perform complex queries across datasets or providing a more coherent view of common concepts. While there has been a large amount of effort in Semantic Web Services [6] even going as far as developing a standard for describing those services [2], we have not seen a corresponding take-up in using these language to enable the execution of actions either on the Web or in the real world. This is about to change.

1 The Need for Lightweight Actions

Increasingly, we see the need to able to exchange knowledge on the Web that correspond to verbs (i.e. operations or actions). The canonical example is Facebook's Like button, which communicates the action of "liking" from a given website to Facebook. Other examples of such actions include sending a tweet or subscribing to an RSS feed. Tantek provides a fairly comprehensive cataloging of current web verbs [7]. The need to invoke or communicate actions between

websites is also seen in proposals such as Web Intents[1] that describe how to both advertise available verbs from a website and allow other sites to discover and invoke them. For those familiar with the Semantic Web services literature, this may sound strikingly familiar. However, the key difference is that these proposals are driven by the demand of existing applications and we are already see examples in use. Furthermore, the kinds of operations being assigned semantics are extremely simple or at least their semantics are often summarized in a single verb (e.g. share, edit). This simplicity parallels the developments that sparked the growth of the use of semantic mark-up on the Web (e.g. the creation of micro formats/data). This organic emergence is one reason that I believe the next 10 years will focus on how to both communicate and instantiate such verbs.

2 The Web is Ready for Verbs

Another reason for the potential focus on verbs and actions is that we now have a ubiquitous substrate to build upon, namely, Javascript. Just as the pervasiveness of the hyperlink made Linked Data possible, the now pervasiveness of Javascript will make it possible to communicate in an interoperable fashion about verbs. The Web is now infused, from both the client side (Javascript-enabled web browsers) and the server side (Node.js), with a common programming language for communicating actions. This poses challenges about how best to use or interact with Javascript to describe operations. Clearly, it is intractable to understand an arbitrary Javascript program but as we have seen intractability may not matter in realistic environments [8]. Additionally, it poses issues about what the environment of the browser means in terms of actions that can be performed and the semantics of those actions. Thus, the investigation of verbs in particular with respect to Javascript poses exciting research problems. We already see some investigation into better understanding of behavioral web pages and Javascript through research into AJAX crawling [5, 4].

If we can describe the first decade of Semantic Web research as being driven by 1) the availability of a real world lab (the web) 2) exciting research questions about distributed knowledge representation and 3) the meeting of the web and database research, then we can see a strong parallel to today's environment but with the focus on verbs. We have the availability of a real world lab (the web with Javascript), exciting research questions about defining the semantics of actions, and a meeting of the web and programming language research.[2] This is further evidence that the Semantic Web's next step is towards operations.

3 Semantics Grounded in the Real World

The final sign that the verb will be central in the further development of the Semantic Web is the ongoing instrumentation of the real world. Increasingly,

[1] http://webintents.org/

[2] See the First Workshop on Programming the Semantic Web http://iswc2012.semanticweb.org/workshops/psw12.inf.puc-rio.br

we see a world in which we not only have a plethora of rich sensors but also of actuators all connected to the Web. This is the vision of the Internet of Things that is being actively implemented [1]. Concretly, we can recreate a 3D representations of the Colosseum just based from images on the Web [3]. With today's Web APIs, we can already make a phone vibrate[3] or send a physical letter.[4] Going forward, we need not just have a Semantic Web grounded in mathematical definitions but we can have a Semantic Web that is referent to reality. This is particularly important for verbs as we may want to perform our operations not just upon information objects but also on real world objects. Thus, I see the massive movement towards an instrumented world as supporting research into operations on the Semantic Web.

4 Conclusion

Here, I have tried to argue that we are perfectly situated to begin to enrich the Semantic Web with the ability *to do*. In 10 years, the Semantic Web will not just be a connected set of datasets but will be intertwined with verbs that are grounded in operations that can operate in bot the real and virtual worlds. In 10 years, I expect Siri will have the knowledge, provided using Semantic Web technology, to know how to get my lawn mowed.

References

1. Luigi Atzori, Antonio Iera, and Giacomo Morabito. The internet of things: A survey. *Computer Networks*, 54(15):2787 – 2805, 2010.
2. Mark Burstein, Jerry Hobbs, Ora Lassila, Drew Mcdermott, Sheila Mcilraith, Srini Narayanan, Massimo Paolucci, Bijan Parsia, Terry Payne, Evren Sirin, Naveen Srinivasan, and Katia Sycara. OWL-S: Semantic Markup for Web Services. Website, November 2004.
3. David Crandall and Noah Snavely. Modeling people and places with internet photo collections. *Commun. ACM*, 55(6):52–60, June 2012.
4. Cristian Duda, Gianni Frey, Donald Kossmann, and Chong Zhou. Ajaxsearch: crawling, indexing and searching web 2.0 applications. *Proc. VLDB Endow.*, 1(2):1440–1443, August 2008.
5. Ali Mesbah and Arie van Deursen. Invariant-based automatic testing of ajax user interfaces. In *Proceedings of the 31st International Conference on Software Engineering*, ICSE '09, pages 210–220, Washington, DC, USA, 2009. IEEE Computer Society.
6. Jinghai Rao and Xiaomeng Su. A survey of automated web service composition methods. In Jorge Cardoso and Amit Sheth, editors, *Semantic Web Services and Web Process Composition*, volume 3387 of *Lecture Notes in Computer Science*, pages 43–54. Springer Berlin / Heidelberg, 2005.
7. Celik Tantek. Web actions: Identifying a new building block for the web, 2012. http://tantek.com/2011/220/b1/web-actions-a-new-building-block.

[3] http://www.w3.org/TR/vibration/

[4] http://www.postalmethods.com/postal-api

8. Jacopo Urbani, Spyros Kotoulas, Jason Maassen, Frank van Harmelen, and Henri E. Bal. Owl reasoning with webpie: Calculating the closure of 100 billion triples. In *The Semantic Web: Research and Applications, 7th Extended Semantic Web Conference, ESWC 2010*, pages 213–227, 2010.