

Creating a Common Ground for URI Meaning Using Socially Constructed Web sites

John Black^{*}
2100 Wisteria Drive
Charlottesville, Virginia 22911
JohnBlack@kashori.com

ABSTRACT

The semantic web proposes to inject machine meaningful data into the existing human language oriented web. As part of this effort, on the semantic web, URIs are used to identify entities. But there is currently no standard way to specify what it is that any given URI is to identify, or to whom, or when. Recent work in linguistics offers ideas for a solution to this lack. It focuses on the pragmatics of actual language use among ensembles of people. Also, the World Wide Web provides a set of technologies, in the form of socially constructed web sites, that could be employed to provide a solution. In this paper, I suggest how such socially constructed web sites could be used to address the problem of establishing common ground among a community of machines of the referent of a URI used on the semantic web. The result is a proposal to automate social meaning by creating societies of machines that share knowledge representations identified by URIs.

1. INTRODUCTION

In this paper, I presuppose that the original vision of the semantic web is a goal which should be pursued. That vision is of a web where *machines* can themselves act on the *meaning* of the content of messages on the web[2]. Thus I will focus on machine agents, but most of what is said would apply to humans, given the development of suitable human user interfaces.

In the study of language there are problems of structure and problems of use. The problems of structure have been dealt with extensively. We have detailed accounts of syntactic structures, their formal semantics, and reasoning over those structures. But the problems of use, how animated agents employ language to accomplish social purposes, are less well-defined. Herbert H. Clark, in his book Using Language[4], following work that gained traction with Austin[1], has created an extensive theory of the pragmatic use of language. I will attempt to apply some of those ideas to the issue of social meaning on the semantic web.

One of the ideas in this work is that of a performative, an instance of speech that, in and of itself, forms an action that accomplishes some goal. One such performative is a

naming. In an act of naming, it is the utterance of certain phrases in certain ways in certain situations that actually creates a name.

There are at least three elements of such a naming when applied to the semantic web. First, there is what I will call the fingerprint of an entity to be named. This is a digital signature of some canonical form of a knowledge representation of that entity. Second, a URI that will function as the name. And third, a public posting of the intention of those agents that will use that name to denote that entity identified by that fingerprint. Such a speech act[1] creates both a name and a language community around that name. This community can range in size from two individuals to as many agents as post acceptance of the creation of the name. What they all have in common, their common ground, is at a minimum, a shared knowledge of what that name identifies.

An outline of the rest of the paper follows. In section 2, I propose that web sites meeting certain criteria can function as the basis for the common ground of a naming. In section 3, I argue that URI reference happens as part of a joint activity[4] between users of a URI and requires involvement by both (or all) parties. In section 4, I suggest directions for future work. In section 5, I conclude with a summary.

2. A BASIS FOR COMMON GROUND

A URI used as a web name must be supported by common knowledge[8] or common ground. In Using Language[4], on page 94, Clark asserts that ‘*Common ground (shared basis)*’ can be defined like this:

proposition p is common ground for members of community C if and only if:

1. every member of C has information that basis b holds;
2. b indicates to every member of C that every member of C has information that b holds;
3. b indicates to members of C that p

2.1 The Grounding of Common Ground

The web site del.icio.us forms a basis, b , as defined in this definition. It is like a public square, where naming authorities baptize[7] tags in front of the crowd. The set of propositions, p , are of the form, ‘this URL is usefully associated with this tag’. The collective users of del.icio.us form C , the members of the community. These socially constructed web sites could be used by machines as well. The same conditions necessary to form such a basis for common ground exist for

^{*}John Black works at Deltek Systems, Inc. However, all of the ideas and opinions expressed here are his own and do not necessarily express the views of Deltek Systems, Inc.

automated systems. By viewing their own posts, each agent verifies that they are part of and known to the community. Thus there is self-awareness. Every agent can verify that every other agent is able to access the site as well, because they receive their posts. Finally, naming propositions, true because they were performed properly, are downloadable from the site. Thus this technology can be used for establishing the referent of a URI in a community of machine agents.

These sites form a shared basis for the common ground of names. On such a web site, every URI forms a language community. More complex communities can be built by adding more names and more agents. A large community of machines, each of which has posted its perception of every one of a large number of name sense pairs, becomes an environment for complex joint activities based on the semantics of the URIs so created. It becomes the common ground for semantic activities by all those members who have adopted the same terms.

2.2 The Three Parts of a Naming

First, for establishing an entity in the world that will be named, a unique fingerprint is needed. Second, a URI is created for naming the entity. And third, the digital identifier of the namer is included for identifying the creator of the name for the entity. These three parts, the fingerprint, the name, and the identifier of the namer are all posted to the social web site together as a triple:

(namer-agent, URI-names, KR-fingerprint)

At once, this triple is available to view by any user of the site. Additional users complete the naming. They do this by signing on and adding their name to the list of adoptees of that naming.

(adopter-agent, URI-names, KR-fingerprint)

Now it is known by the namer and all adoptees that all have knowledge of that name for that entity.

2.2.1 Grounding in the World

Using a fingerprint, that is, something that a machine can input from the world and match to an internal version is necessary to connect statements on the web to the world. Without a connection to the world, names end up naming other names in an endless hall of mirrors. There are various proposals for the material to use for grounding names, depending on the type of the entity. For individual concrete entities, something unique about their physical nature should be used. For humans, literally fingerprints could be used, but for other entities, such as an open source archive download, the MD5 signature can be used. The bar coded package serial number of the kind that is pasted on the outside of a package for a shipment of some kind is another example of an entity fingerprint.

For more abstract entities, the signature of other forms of knowledge representations, put in a canonical form, could be stored. See, for example, the Web Proper Names[6] proposal. Even more complex entities can be named. For example, using named graphs[3], arbitrary collections of facts could be named.

2.2.2 URIs as Names

On the semantic web, URIs are used to identify things. From the W3C's RDF Primer, 'RDF is based on the idea of identifying things using Web identifiers (called Uniform Resource Identifiers, or URIs)'. Everything, in fact, should

be identified with URIs. Not just individual humans, but their properties, web resources, categories of thought, and possessions of all kinds. On the semantic web, URIs are used for the digital identities of everything.

2.2.3 The Namer

The first namer is essential because it is his, hers, or it's agency that makes the posting of the triple a performative, a naming. Otherwise it is just another web posting, which could be experimental, example, error, farce, fraud or anything else. A performative is an action with a purpose, a goal, and agents are, by definition, animated entities with a purpose.

3. JOINT ACTIVITY

But naming is not solely the responsibility of the namer. Suppose no one ever visits the site? Did a name get created? Not until another user adopts that name for that entity has there been a naming. Clark[4] thoroughly develops the idea that language use is a joint activity. It is like dancing. It cannot be done alone. Even if you could perform the very same movements as you did while dancing, it would not be the same activity without the coordinated actions of your partner.

Without the public uptake by adoptees, just publishing such a triple would amount to no more than a 'private language'. As Wittgenstein says of it,

'Why can't my right hand give my left hand money? –My right hand can put it into my left hand. My right hand can write a deed of gift and my left hand a receipt. –But the further practical consequences would not be those of a gift. When the left hand has taken the money from the right, etc., we shall ask: Well, and what of it? And the same could be asked if a person had given himself a private definition of a word; I mean, if he has said the word to himself and at the same time has directed his attention to a sensation'

- Wittgenstein[9], *Philosophical Investigations*, 268.

The second namer, and all further namers, consummate the performative and form a community around the common knowledge of the referent of that name.

This is a critical factor. *The mutually verifiable adoption by members of the community is essential.* With a new name, there is no common ground until it is established in just such a public naming. This explains some of the effectiveness of del.icio.us. It is due not so much to the technology of tagging, but to the technology of social software.

3.1 Social Meaning Requires Societies

The social meaning created in this manner is not stored anywhere. Specifically, it is not the referent of a URI. Like the V-shape of a flock of geese migrating south for the winter, social meaning is what will emerge spontaneously and ephemerally among communities of machines that coordinate activity around their common ground. The creation of meaning in this way is a *joint activity* serving a *joint project* of two or more agents[4].

3.2 A Calculus of Communities

The common ground created by the users of such a site would not be a fixed region. In particular, for many URIs,

it will not be shared globally. It would be more like a set of concentric and overlapping circles that range in extent from large portions of the population down to two communicators. Any two communicators could establish their own name for an entity and this local name could preempt that name used by the broader community. This creates the potential for micro-senses. But it also records the common ground to use to disambiguate these micro-senses. Any group of communicators could engage in grounding acts that established a new, local common ground that is shared between them alone. Thus factions can emerge, each group posting adherence to a different representation for a URI.

There is no need for there to be only one of these socially constructed ontology sites. In fact, there is an obvious path for distribution of this functionality through the emphasis on different content. In human language use, Clark points out that common ground centers around various attributes of the language users. Among the attributes around which constellations of common ground accumulate are

Nationality, Residence, Education, Occupation, Employment, Hobby, Language, Religion, Politics, Ethnicity, Subculture, Cohort, Gender.

- Clark[4], page 103.

3.3 Google vs. del.icio.us

It might be objected that there is no need for such web sites. That the architecture of the web already includes all these elements. That contributors to existing web sites, publishing ontology documents using existing web technology, combined with aggregating engines such as Swoogle[5] will produce the same results more effectively. But joint activity is different than the aggregation or analysis of independent activity. You can use Google to search for able sellers, and you can use it to search for willing buyers, and you may be able to find a precise match through analysis of the two sets. But this would not create a contract for the sale of goods. Such a contract must be both offered and accepted by the respective parties. Both agents must sign it. So it is with social meaning. No amount of aggregating, reasoning, or merging can, by itself, turn independently created propositions into common knowledge.

4. FUTURE WORK

Here are a few directions I would like to pursue in the future.

4.1 Consensus Building by Machine Agents

Such socially based web sites could perform a function similar to what the W3C or OASIS perform for human vocabulary builders. It could act like a machine oriented Wikipedia, forming a place to record the moves that are public record, such as the Wikipedia history of edits. Such a ground would allow for maintenance and repair as well, which fosters further acceptance.

As enabled by del.icio.us, when I view a tag, I have the option of copying it. On a socially constructed naming web site, this would become a vote for the naming. To view a name and move on would be to vote against it.

4.2 Socially Oriented Queries

These web sites will allow for new queries that are not currently possible:

- what is the *URI-name* adopted by agent *Agent-ID* for this *KR-ID*?
- which *URI-name* is used by the most agents for this *KR-ID*?
- what are all the *URI-names* that have been jointly adopted by the set *y* of agents?

Where *KR-ID* is the fingerprint or identifying knowledge representation as described in section 2.2.

Many of these queries, or something similar to them, are already in use on socially constructed web sites and are much of what makes them so satisfying to use for human participants.

4.3 Current State of the Activity

A common ground for ontology terms mimics language conventions[8], but in actual use, each utterance must also become a part of the common ground. Thus a means of tracking the current state of the activity is needed. This discourse both depends upon the common ground and becomes a part of it.

5. CONCLUSION

In this paper I have argued for a proposal to continue progress on the semantic web by attempting to automate the kind of social meaning that is so important for the use of language in human communities. I have shown how it may be possible to achieve this by using web technologies currently in use by socially constructed web sites such as del.icio.us.

6. REFERENCES

- [1] J. L. Austin. *How to Do Things with Words*. Harvard University Press, Cambridge, Massachusetts, 1962.
- [2] T. Berners-Lee, J. Hendler, and O. Lassila. The semantic web. *Scientific American*, May 2001.
- [3] J. J. Carroll, P. Hayes, C. Bizer, and P. Stickler. Named graphs, provenance and trust. <http://www2005.org/cdrom/docs/p613.pdf>, 2005.
- [4] H. H. Clark. *Using language*. Cambridge University Press, Cambridge, 1996.
- [5] L. Ding, T. Finin, A. Joshi, R. Pan, R. S. Cost, Y. Peng, P. Reddivari, V. C. Doshi, and J. Sachs. Swoogle: A Search and Metadata Engine for the Semantic Web. In *Proceedings of the Thirteenth ACM Conference on Information and Knowledge Management*. ACM Press, November 2004.
- [6] H. Halpin and H. S. Thompson. Web proper names: Naming referents on the web. <http://www.webpropernames.org/paper/>, 2005.
- [7] S. A. Kripke. *Naming and Necessity*. Harvard University Press, Cambridge, Massachusetts, 1972.
- [8] D. Lewis. *Convention - A Philosophical Study*. Blackwell Publishers, Malden, MA, 2002.
- [9] L. Wittgenstein. *Philosophical Investigations*. (Translated by Anscombe, G.E.M.). Basil Blackwell, Oxford, 1953.