

Metcalfe's Law, Web 2.0, and the Semantic Web

James Hendler¹
Jennifer Golbeck²

¹ Rensselaer Polytechnic Institute
² University of Maryland, College Park
hendler@cs.rpi.edu, jgolbeck@umd.edu

***Abstract:** The power of the Web is enhanced through the network effect produced as resources link to each other with the value determined by Metcalfe's law. In Web 2.0 applications, much of that effect is delivered through social linkages realized via social networks online. Unfortunately, the associated semantics for Web 2.0 applications, delivered through tagging, is generally minimally hierarchical and sparsely linked. The Semantic Web suffers from the opposite problem. Semantic information, delivered through ontologies of varying amounts of expressivity, is linked to other terms (within or between resources) creating a link space in the semantic realm. However, the use of the Semantic Web has yet to fully realize the social schemes that provide the network of users. In this article, we discuss putting these together, with linked semantics coupled to linked social networks, to deliver a much greater effect.*

Introduction

In talking about the Web, whether the original model, the so-called "Web 2.0", or the emerging Semantic Web (aka Web 3.0), one of the most important things to keep in mind is the *network effect*. The power of the Web emerges through the link space realized between Web pages. This is evidenced in a number of pieces of work, most famously the PageRank algorithm (Brin and Page, 1998) that was behind the early success of Google. Unlike traditional information retrieval algorithms, which were solely based on the information content of the individual pages, PageRank takes into effect how Web pages are linked to each other. By coupling this information with traditional indexing schemes, the system was able to outperform its competitors.

The network effect describes the value of a service to a user that arises from the number of people using the service. At its core, it captures that value increases as the number of users increases, because the potential links increase for every user as a new person joins. This is best quantified by what has come to be known as Metcalfe's Law. This proposition developed by Bob Metcalfe in the early 1980s, was originally defined to better explain to his customers why they needed more Ethernet boards than they were buying¹. Metcalfe hypothesized that while the cost of the network grew linearly with the number of connections, the value was proportional to the square of the number of users. For example, given n users of ethernet cards, the number of possible connections that can be made is $n(n-1) = O(n^2)$.

¹ Bob Metcalfe, personal communication, June 2007.

Metcalfe's law has been used to explain the growth of many technologies ranging from phones, cell phones, and faxes to web applications and social networks, especially online social networks. The intuition clearly holds that as the number of people in the network grows, the connectivity increases, and if people can link to each other's content, the value grows at an enormous rate.

Recently, there has been some interesting debate with respect to the validity of Metcalfe's law. On the low end, in a 2006 column in *IEEE Spectrum*, Brisco et. al (2006) opined that value in a network grows more like $O(n \log n)$ arguing that not all connections are of equal value. At the other extreme, in a 2001 article in *Harvard Business Review*, Reed (2001) claimed that the value of the network grew exponentially in the number of connections. His argument is essentially that in a largely connected network, such as a social networking Web site, the value is in the creation of subgroups and the number of these subgroups (i.e. the subnetworks of size 2, size 3, ... size n) grows exponentially with n . While none of these effects have been validated in practice, it is clear that the network effect is quite real, and even the most pessimistic view still provides for significant value as the number of connections in the network grows.

There is a corollary of Metcalfe's law that is sometimes missed: for the network effect to happen, linking must be present. The Web, if it were simply a collection of pages of content, would not have the value it has today. It is precisely because every Web page can, in principle, link to any other page that the Web has grown as it has. Without this linking, information would get cut and pasted onto larger and larger individual pages; instead of the Web, we would have a large number of disconnected pages and little or no index.

In this paper, we look at Web 2.0 and Semantic Web applications from the point of view of the linked spaces being created – where does the network effect come from? The social nature of Web 2.0 sites primarily allows linking between people, not content, thus creating large, and valuable, social networks, but with impoverished semantic value among the tagged content. Conversely, the Semantic Web is able to take advantage of significant linking in semantic space, and while it can represent social networks, it does not have social constructs that lead to linking between users. Furthermore, many production level Semantic Web applications are not exploring how to create links between different ontologies. We will look at how a combination of these could be designed to take advantage of the joint network effects of links in social space with links in the semantic space. By combining the social networks of Web 2.0 with the (small "s") semantic networks of the Semantic Web, a tremendous value is promised.

Web 2.0 as a social phenomenon

Much is made of the incredible success of so-called "Web 2.0" applications, even though there is no widely agreed upon definition of what makes something one. In a widely cited web article, Tim O'Reilly, who is generally considered to have coined the term, discusses the many aspects of Web 2.0 (O'Reilly, 2005). The discussion includes exploring the technologies of AJAX, Web Services and other means for making Web

content more dynamic. In this view, Google Maps is considered the prototypical Web 2.0 application, even though it does not include interaction between users. He also discussed that tagging sites, like flickr² and Del.icio.us, are archetypes of Web 2.0 as they allow users to create content easily. (This is often joined with an argument, sometimes attributed to Clay Shirky (2003), that "folksonomy" will magically answer many of the traditional problems of knowledge representation and create what others have called the "small s semantic web."³)

The idea that users can create content is considered a critical aspect of Web 2.0. Blogs, Wikipedia⁴, and other sites that are considered successes of the new approach focus on this aspect. However, in discussing the difference between blogs and home pages, O'Reilly makes it clear that content creation is not enough. Rather, RSS, permalinks and other trackback technologies are considered critical. These, he states, are what contribute to the link space that enables the network effect to work in the dynamic content space of blogs and the like.

In the discussion of Web 2.0, O'Reilly tends to focus on the technologies and not as much on the social phenomena underlying Web 2.0 applications. In the past few years, however, it has become increasingly clear with the growth of sites such as MySpace and Facebook that the social networking construct is critical to the success of Web 2.0 applications. The fact that sharing of content can be enhanced by personal connections, rather than primarily via search or other query techniques, has emerged as a major, and perhaps defining, aspect of successful Web 2.0 applications.

As an example of this, consider YouTube⁵, another successful modern Web application. YouTube allows users to upload video content to the Web, and provides a number of mechanisms for letting users share this content. Interestingly, email and blogging has proven to be one of the crucial aspects of the YouTube phenomenon. Pointers to the videos on the site are often shared and that has become the primary way in which videos become successful. Once a video has "made it," getting many thousands of views, it can become a popular node in the network of videos, which are linked by a number of metadata features (who they are by, what the main subject is, where the content originated, etc.) Search in YouTube is primarily enhanced by the social context, not by the "semantic content" of what is in the videos (Marcus, Perez, 2007). While automated technologies to create indexes of these videos are being sought, the primary indexing comes through the social overlay of the site.

This, we argue, is actually true of almost all of the successful Web 2.0 applications. For example, while the English version of Wikipedia is a clear success of the new generation of Web technologies, it is less clear why so many other Wiki sites have fallen flat. What one sees when examining Wikipedia, and other successful sites, is the social construct being critical. As Jimmy Wales, developer of Wikipedia, stated in his (2005) talk at the

² <http://flickr.com>

³ A term that is usually attributed to either Rohit Khaare or Tantek Celik.

⁴ <http://www.wikipedia.org/>

⁵ <http://youtube.com>

Doors of Perception Conference, "Wikipedia is not primarily a technological innovation, but a social and design innovation."

In fact, if one looks at some of the early Web 2.0 successes in this light it becomes clear that the success of tagging has far more to do with the social interactions it allows than with the semantic vocabularies it creates (Marlow et al., 2006). There is significant evidence pointing to this. For example, on flickr photo sharing appears to be most successful for two different sets of users. One group is those who attend a uniquely named event or who, at some event, determine (out of band) what keyword will be used by those who want to upload and share their photos. Where a clear and unique keyword exists, the search capabilities of flickr work fine, but where there isn't one, the flatness of the tag space (which is not hierarchical) and the lack of links make it more difficult to find the content one desires.

The second and more socially successful use of flickr is within known communities where specific tags can have some meaning. For example, if you search flickr for the string "pi" you will find over seventy thousand photos which include the substring in the tags. On the other hand, if you are a member of the community of users that can access the photos posted by Jennifer Golbeck and her social network, you would see that "Pi" brings up pictures of a specific dog named, not surprisingly, Pi. This is not unusual, many common tags on flickr include terms like "dad" (80,000+ photos), "Fred" (90,000+ photos) and "My (something)" (over 8,000,000 photos). Clearly these terms are not very useful outside of specific contexts, but are very meaningful within them. Similar effects are seen in Amazon, where tags like "dad's favorite" are common, and in del.icio.us⁶ and many other tagging sites that allow users to create tags within contexts other than the globally shared one.

A problem for many Web 2.0 sites, in fact, is that tags do not create much of a link space. Even if one postulates that the multiple tags put on a single item create a graph (i.e. all items sharing a tag are considered linked to each other), this graph is very sparse. Most items typically have a very small number of tags associated, and many of the terms used are ambiguous or context dependant. Thus, attempts to use statistics to cluster in tag space have not been very successful (and many sites, such as flickr, have removed the clustering features from their primary page views), and page-rank-like algorithms have not been successful. Search in such sites doesn't work well, as it is basically traditional IR used on large numbers of documents with small numbers of keywords, and browsing in the impoverished graph is not very rewarding.

Returning to our earlier discussion of Metcalfe's law, it becomes clear that in many, if not most, of the Web 2.0 sites that use tagging, the network effect is not primarily coming from links between content and tags. Rather, we argue that given the prevalence of the social constructs within these sites, that value of the network effect is coming from the links between people arising from the interactions using these sites. For social networking sites like mySpace and Facebook, it is obvious that the social network graph is denser and more connected than that of the content space. For sites like flickr and

⁶ <http://del.icio.us>

YouTube, this effect is less obvious, but it is clear, as we have argued above, that it is still the primary value source. The success comes from the rapidly growing social network and the value growth driven by Metcalfe's law operating over the social links.

The Semantic Web Graph

Some of the original motivations for the Semantic Web came from the very same failures in early Web applications that cause the problems for search and browsing in Web 2.0 applications. Latent Semantics, the attempt to "mine" meaning from the words in Web content, is always problematic due to ambiguity and polysemy (the many meanings of a single word such as "run" or "left"). Also problematic are the class and subclass relations that are crucial to language use. For example, a search for information about "dog"s won't find a picture of Pi unless you know that Pi is a dog. Similarly, raw statistics are not terribly successful for determining that dogs are meat eaters, snails are vegetarians (but meat when consumed), etc. This problem is made even worse as sometimes whether something is a member of some class is dependent on a specific context. For example, the term "chattel" is used in law to refer to certain kinds of personal property. Whether Pi is chattel or not depends on the specific context of her ownership by Golbeck. Similarly, whether a particular gene is a "cancer gene," whether a particular airplane flight is an "on time flight" and many other class memberships are dependent on complex relationships that are not easily mined from textual content.

The situation is even worse for non-textual data. It is an old cliché that "a picture is worth a thousand words." Unfortunately, if this is true, then understanding the content of a particular picture would require long paragraphs to be written describing it, not something that happens often. Worse, a video is essentially a collection of photos, consider how many words it takes to describe, as completely as possible, what is going on in even a short video. While automated understanding of photos and videos is an active area of research, its realization is still far off, and thus using text-based approaches to search and browsing of video, without some sort of semantic annotation, remains a distant promise. Data is also a non-textual form, and again, searching and browsing data without some kind of organizing schema is beyond current capabilities.

Semantic Web technologies were developed in part to address these faults. For applications that wanted to share information that was not yet in textual form, or was in a form where the textual information was hard to extract, it was clear that some form of knowledge representation was needed. This was not a new observation, it had been realized in fields like Natural Language Processing and machine translation years earlier. What was new in the Semantic Web technologies was an attempt to do knowledge representation in a form that was web embedded, that is, where terms and relationships were assigned persistent URIs and linking between these terms, and between these terms and other Web resources, was easy to do. The key was to create another web graph, this time a graph between semantic terms and between these terms and what they described.

The Semantic Web languages RDF, RDFS and OWL are all based on a model in which terms are assigned specific URIs. While much is made about the representational

capabilities of these languages, and their ability to express certain relationships, a much more critical aspect is that they can be used to provide common referents. Some of the most used Semantic Web vocabularies, like the Friend of a Friend (FOAF) ontology, get their primary value not from the terms they express but, as Metcalfe's law predicts, from the many instances linked to each other through the common (and unambiguous) vocabulary. While inferencing is an important aspect of Web, and all other, knowledge representation languages, the ability for terms to be linked is a critical difference between RDF-based languages and earlier KR languages.⁷

The terms in Semantic Web documents are, indeed, linked in many ways. Within an ontology, the terms can be linked to each other directly. Thus, where flickr, asked to find photos in Poland, will not include those labeled Lubusz, in an OWL ontology it is easy to assert that Lubusz is a voivodeship (or province) that is located in Poland. The links from Lubusz to Poland are made explicit, and thus the link space is there to be exploited. These links are also easily defined between documents. For example, if another document wants to assert that the two capitals of Lubusz are Gorzów Wielkopolski and Zielona Góra, those cities can be assigned their own URIs and linked to those in the earlier document about Poland.

This linking between ontologies, and between instances in documents that refer to terms in another ontology is where much of the latent value of the Semantic Web lies. The vocabularies, and particularly linked vocabularies using URIs, of the Semantic Web create a graph space with the ability to link any term to any other. As this link space grows with the use of RDF and OWL, Metcalfe's law will once again be exploited – the more terms to link to, and the more links created, the more value in creating more terms and linking them in.

Unfortunately, while the link space of the Semantic Web is large and growing, the social constructs to exploit these links have been slow in coming. Many of the first generation Semantic Web tools focus on developing ontology documents with little provision for linking, or provide inferencing capabilities only as long as all the terms are collected into a single triple store (preferably without too much instance data). New Semantic Web tools such as Tabulator⁸ and Zitgist⁹ are starting to change this by providing browsers that follow these links, making the graph space more explicit. To date these tools are comparatively simple, and the Semantic Web graph they browse is still fairly sparse. Applications to help create the links that the Semantic Web can exploit are still, unfortunately, few and far between.

Another problem for the Semantic Web is that, so far, applications have not largely caught on to exploiting the social mechanisms that are powering the Web 2.0 sites. All too often, Semantic Web researchers have been focused on trying to somehow utilize tagging and folksonomies in their current flat and ambiguous form, and have missed the

⁷ For a detailed discussion of the relationship between Semantic Web and other KR languages see Hendler and van Harmelen (forthcoming).

⁸ <http://www.w3.org/2005/ajar/tab>

⁹ <http://zitgist.com/>

point that this is precisely the space where semantics is needed and can most easily be exploited. Conversely, instead of exploiting the community contexts, interest groups and personal relationships that make sites such as flickr work, or the complex social dynamics of a Wikipedia, many Semantic Web applications focus solely on expert system-like applications with expressive semantics to the exclusion of all else. These systems make good use of the fact that OWL has become a standard, and therefore offers advantages in that respect, but they are not exploiting the Web nature of the Semantic Web.

A major exception to the above is the Friend of a Friend ontology¹⁰, which is without doubt one of the successes of the Semantic Web to date. FOAF was originally developed as a small ontology to describe people and to allow them to link to each other in a social network like way. FOAF was designed to be relatively lightweight and easy to use, rather than to push for an expressive representation of the properties of humans. A particular idiom, using RDF's seeAlso construct, was developed to allow FOAF files to link to each other and create a social network. Most FOAF files are now created automatically by other Web sites such as browsing or social networking sites, and thus the number of these files (and thus the value of the connections between them) grows rapidly (Golbeck 2008). There are tens of millions of FOAF profiles which, when reasoned over, add connections among the social networks produced from different websites (Golbeck, Rothstein, 2007). FOAF has largely been successful because of its modeling of the social networks it encodes, although the link space is still not as large as some Web 2.0 sites, and there is still a lot of effort going into working out how to create more linking of FOAF to other ontologies, and more instances, to increase the value the network effect brings.

Putting it together

A recent boom in Semantic Web technologies has been occurring in the so-called "Web 3.0" technologies. In these systems, an attempt is being made to exploit more of the link spaces inherent in RDF-based systems coupled with capturing some of the social dynamics of Web 2.0 applications. One difference between these and earlier AI systems is the attempt to figure out how to exploit the increased value of the network effect that can come from using Semantic Web technologies to provide links between diverse sets of content or users. Coupled with languages such as SPARQL, GRDDL and RDFa, which provide a technology base for making Semantic Web applications interoperate more smoothly with traditional Web applications, we see an increasing awareness in the importance of creating and exploiting Semantic Web links.

One example of an interesting Web 3.0 site is the RealTravel¹¹ site developed by Tom Gruber and described in his talk entitled "Where the Social Web meets the Semantic Web" at the 5th International Semantic Web Conference¹². RealTravel "seeds" a Web 2.0 travel site with the terms from a gazetteer ontology. This allows the coupling of place names and locations, linked together in an ontology structure, with the dynamic content

¹⁰ <http://www.foaf-project.org/>

¹¹ <http://realtravel.com>

¹² Available online at <http://tomgruber.org/writing/social-web-meets-semantic-web.pdf>.

and tagging of a Web 2.0 travel site. The primary user experience is of a site where travel logs (essentially blogs about trips), photos, travel tools and other travel-related materials are all linked together. Behind this, however, is the simple ontology that knows that Warsaw is a city in Poland, that Poland is a country in Europe, etc. Thus a photo taken in Warsaw is known to be a photo from Poland in a search, browsing can traverse links in the geolocation ontology, and other "fortuitous" links can be found. The social construct of the travel site, and communities of travelers with like interests, can be exploited by Web 2.0 technology, but it is given extra value by the simple semantics encoded in the travel ontology.

Sites like this are a good start, and show that coupling the social and semantic networks produces several layers of semantic and social linking that leads to increased value through the network effect, but we contend they are just a start. A much more powerful network effect will arise with the linking of different sites, containing different materials, based on common terms found in the persistent links of the Semantic Web. The use of common terms, or of OWL's inference power to make "sameAs" inferences, to link between these applications can be used to create Web spaces that will have far more links leading to the real power in the Web 2.0 to Semantic Web link.

Consider, for example, the new Web application Dopplr¹³ - a site in which users create a social network among their friends, and share travel itineraries. This allows users to find out when they have overlapping trips with others. Dopplr, like RealTravel, uses a simple location ontology to help manage its information. Developed separately, the places in Dopplr don't align one to one with those in the RealTravel, but both sites do have persistent URIs for places. This means that a relatively straightforward mashup of this information could be created (if both sites were willing) simply by creating a mapping between place names. Users from Dopplr could learn more about the places they intend to visit. Users of RealTravel could quickly find out if any of the places they are reading about have been visited, or plan to be visited in the future, by any of their friends.

This is a simple example where combining two sites could add value to both. Now consider linking to these all the photos of places in flickr using the same URI, or LiveJournal blog entries about the places visited, or any other site that uses geographic terms that can be reliably mapped to other sites (and creating such mapping ontologies is easy using owl:sameAs). Further, the people known to these networks, having FOAF files, can be linked to others in mySpace or Facebook, or to other sites that use FOAF and comply with the FOAF model of identity. Given a few simple ontologies of locations and the simple rules in FOAF, value could be added by the network effect emerging from the linking of these many different sites.

Going beyond locations (or better, coupling to them), we could also see similar linking in many other ontological areas. Currently, the Semantic Web contains a number of important resources that have large vocabularies of static URIs useful for creating these "mega" applications. For example, the National Cancer Institute ontology (Golbeck et al, 2003) could be used for coupling many different sites exploring different aspects of this

¹³ <http://dopplr.com>

major disease. The US National Library of Agriculture has released a large vocabulary (using SKOS) of useful agricultural terms¹⁴. Other ontologies already being developed, many of which are public, include vocabularies of science, medicine, common objects, projects, and hundreds of other useful areas.

In addition to the potential of linking terminologies between sites like these, there is also another dimension of sharing which is being made possible by the Semantic Web. Currently there are a number of projects focused on making high value datasets available in RDF to make them more available for applications to exploit. The simple semantics of these RDFized datasets make them easy to link to, and to describe using the more expressive constructs of RDFS, OWL and the emerging rule languages. For example, the BBC has released their programme catalog in an RDF compatible form. This makes 75 years of BBC programming available for linking to Semantic Web sites. Thus, for example, it would be easy for RealTravel to link to all the BBC shows taking place in, or reporting on, the known locations. This in turn, as above, would link to Dopplr, flickr, Wikipedia, mySpace, and so on. The potential network effect created by linking the URI space of Web resources, the social networks of current Web 2.0 applications, and the URIs in these vocabularies is huge: Metcalfe's law, exploiting the potential linkages of content in these many spaces, predicts a value that is truly staggering.

A research vision

The Web is an interesting place for browsing, but its real power derives from people finding what they need. Similarly, using Semantic Web technologies, social networks, and terminologies to label and link content will be powerful only when it enables people to do powerful things. Creating these links is a first necessary step, and the research challenges lie in understanding how to use them.

Building expressive Semantic Web ontologies is very difficult to do well, but once they are built a lot can be done by using the semantics of the links. Tagging, on the other hand, is very easy, but there is no structure and, as described above, many searches will miss relevant results that are not tagged with exactly the right term (e.g. dogs tagged with their name or breed will not show up when users search for "dog"). There is a balance that can be struck between these two extremes. For example, adding minimal structure to tags can bring a lot of advantages.

Some techniques have tried to add structure to tags using clustering methods. Though this can sometimes create sensible "hierarchies", the links between concepts do not indicate parenthood as we would normally expect. For example, one branch of a tag hierarchy generated from Del.icio.us in (Heymann, 2006) is software->mac->osx->apple->ipod. This kind of hierarchy will not significantly improve search and information structure as well as one that is human engineered. The first challenge, then, is how to build a structure around tags. In a social, collaborative web environment, communities are the logical group to be creating this structure, preferably delivered in machine-readable (Semantic Web) format with persistent URIs as discussed.

¹⁴ <http://agclass.nal.usda.gov/download.shtml>

Once that structure exists, we can begin to study the meaning of connections within the network of knowledge. Some of the mashup applications now available are indicative of what to expect in the first stages of integrating social networks, structured tags, and the annotated content; direct links can be exploited to bring data from one space into another (like showing photos tagged as depicting someone listed as a friend in a user's social network).

However, greater promise lies in exploiting the connections that extend out through the network. Analysis of paths that connect users in a social network can provide recommendations about how much they trust one another or how similar they are. The hierarchical structure of tags can be used to determine the relevance of matches to user queries. These networks can even be combined, where relationships are computed by combining social network and information profiles of users, and those relationships are used, in turn, for collecting and filtering information. There has been some research on computing relationships in social networks and using those relationships to filter content (Massa, Avesani, 2004) (Golbeck 2006). Those results show potential for how the integration of social and semantic networks can bring great improvement to how people see, and begin to trust, information on the web.

As the trend continues, the integration of social networks, semantics, and content has the potential to revolutionize web interaction. The creator's pages of data will no longer need to be the main vehicle for accessing content. Rather, resources can be aggregated, shared, and accessed from many different places, and users will be able to choose which has the most appropriate presentation and set of tools for the tasks they need to accomplish.

While we have primarily discussed technologies in this article, there are also important user interface challenges here that are possibly the most critical element for making the vision we present succeed. Tags work to a large extent because they are trivially easy to use. Butterfield (2004) puts it clearly: "I think the lack of hierarchy, synonym control and semantic precision are precisely why [tagging] works. Free typing loose associations is just a lot easier than making a decision about the degree of match to a pre-defined category (especially hierarchical ones). It's like 90% of the value of a proper taxonomy but 10 times simpler." (Mathes (2004) rightly says that the 90% value and 10 times simpler estimations are vastly overstated, but Butterfield captures the core point.) You get *something* from tags with very little effort, so additional effort will need to yield significant additional benefits. How to create user interfaces where people can easily label resources with tags from a pre-defined structured environment is an important line of this research.

Conclusion

Although there is great mythos about Web 2.0 and the Semantic Web, there is no real reason to believe they function significantly differently with respect to linking than other existing information systems, particularly the original "Web 1.0." Metcalfe's law makes it clear that the value of these systems, viewed as networks of communicating agents

(whether human or machine), arises from the many connections available between online resources. To exploit this space, however, there must be explicit linkages between the resources: when it comes to the network effect, if you don't have links, you don't get it.

Web 2.0 and Semantic Web applications currently are exploiting different sets of link spaces to different advantage. At a technical level, it is not the folksonomies of Web 2.0 per se where the strength derives, but from the social linkages that are enabled by the applications. For the Semantic Web, the linkages enabled by the URI-based languages provide a set of semantic linkages that applications are starting to take advantage of. Combining these two, and finding ways to combine (link) the social structures of the Web 2.0 applications with the semantic structures of the Semantic Web is a compelling way to bring together two different networking spaces, allowing the total value to increase enormously. Building these applications remains a challenge, and interface issues are still a limiting factor, but the potential value that can arise from the combined social and semantic networks is huge.

REFERENCES

BBC catalog available online from - <http://catalogue.bbc.co.uk/catalogue/infax>

Briscoe, Bob, Andrew Odlyzko, and Benjamin Tilly. 2006. Metcalfe's Law is Wrong., *IEEE Spectrum*, July 2006.

Brin, Sergey and Page, Lawrence. 1998. The anatomy of a large-scale hypertextual Web search engine. *Proceedings of the seventh international World Wide Web Conference*.

Butterfield, Stewart. 2004. "Sylloge." August 4, 2004 -- <http://www.sylloge.com/personal/2004/08/folksonomy-social-classification-great.html>

Golbeck, Jennifer. 2006. "Generating predictive movie recommendations from trust in social networks," in *Proceedings of the fourth international conference on trust management*.

Golbeck, Jennifer, Gilberto Fragoso, Frank Hartel, James Hendler, Bijan Parsia, Jim Oberthaler. 2003. The national cancer institute's thesaurus and ontology, *Journal of Web Semantics*, 1(1).

Golbeck, Jennifer. 2008. The Dynamics of Web-based Social Networks: Membership, Relationships, and Change, *Sunbelt XXVIII: International Sunbelt Social Network Conference*.

Golbeck, Jennifer Matthew Rothstein. 2007. Cross-network linkages with FOAF. *Technical Report TR-MSI292, MINDSWAP, University of Maryland*.

Hendler, James and van Harmelen, Frank. (Forthcoming) The Semantic Web and Knowledge Representation, in Lifshitz, Porter and van Harmelen (eds.) *Handbook of Knowledge Representation*, Elsevier (forthcoming).

Heymann P, H. Garcia-Molina. 2006. Collaborative Creation of Communal Hierarchical Taxonomies in Social Tagging Systems. *Stanford InfoLab Technical Report, No. 2006-10*.

Massa, P and P. Avesani. 2004. "Trust-aware collaborative filtering for recommender systems," in *Proceedings of the Federated International Conference On The Move to Meaningful Internet: CoopIS, DOA, ODBASE*.

Cameron Marlow, Mor Naaman, Danah Boyd, Marc Davis, 2006. HT06, tagging paper, taxonomy, Flickr, academic article, to read. *HYPERTEXT '06: Proceedings of the seventeenth conference on Hypertext and hypermedia*.

Mathes, Adam. 2004. Folksonomies - Cooperative Classification and Communication Through Shared Metadata -- <http://adammathes.com/academic/computer-mediated-communication/folksonomies.html>

O'Reilly, Tim. 2005. What Is Web 2.0? Design Patterns and Business Models for the Next Generation of Software. September 30, 2005. <http://www.oreillynet.com/pub/a/oreilly/tim/news/2005/09/30/what-is-web-20.html>

Marcus, Aaron and Angel Perez. 2007. "m-YouTube Mobile UI: Video Selection Based on Social Influence" *Proceedings of the 12th International Conference, HCI International*.

Reed, David P. 2001. The Law of the Pack. *Harvard Business Review*, February 2001, 23-24.

Shirky, Clay. 2003. The Semantic Web, Syllogism, and Worldview, -- http://www.shirky.com/writings/semantic_syllogism.html