

Contextual Authority Tagging: Cognitive Authority through Folksonomy

Terrell Russell

May 2005

SILS @ UNC-CH

INLS 302 - Gollop

Abstract

Contextual Authority Tagging is the use of folksonomies to discover and define cognitive authority through reputation within communities of users. Authority is granted by individual users to other individual users with regard to their perceived domains of knowledge via free text tags or labels. This allows discovery of at least two things, 1) which users in a group are authority figures on a certain topic area, and 2) what areas of expertise a particular user possesses. A basic proposal is laid out along with a few examples to foster communication and thought on this new distributed way to discover cognitive authority.

Contents

1 Introduction	4
2 Background	4
2.1 Folksonomies	4
2.1.1 Classification	5
2.1.2 Del.icio.us	5
2.1.3 Flickr.com	6
2.2 Traditional Authority	7
2.2.1 Chains of Authority	7
2.2.2 Perspective	8
2.2.3 Contextual Authority	8
2.3 Community Effect	10
3 Proposal for Contextual Authority Tagging	11
3.1 Related Tags	11
3.2 Degrees of Separation	12
3.3 Non-negative Authority	12
3.4 Option to Defer	13
3.4.1 Spreading Untruths	13
3.4.2 Privacy	14
3.5 Power Law Observation	14
4 Global and Personal Views	15
4.1 Type A Queries - User/Global	15
4.2 Type B Queries - User/Personal	15
4.3 Type C Queries - Topic/Global	16
4.4 Type D Queries - Topic/Personal	16

5 Usefulness of this Proposal	17
References	18

1 Introduction

We are social creatures. Nearly every interaction we have with the world today involves, at some level, other people. These interactions are constantly changing our opinion of the world and of the other people in it. Additionally, everyone else is going through the same thing. Our reputation is defined by others' opinion of us. Our social identity and place in the world is defined by others, but not in a way that can be measured very easily.

Contextual Authority Tagging brings measurement to this phenomenon. Each of us is seen in a variety of light by a variety of others. We have personal shared histories with every acquaintance. Each of these people only sees and knows a thin slice of who we are and what we know. And when asked to represent that slice, they would each paint a different picture. This contextual slice has always been represented, if at all, by a single value, by a single idea. Research in computer science and economics, trust networks and cooperation, have assigned this value in aggregate[32, 11, 21]. Scores have been a single number. Trust has been represented as a percentage, some fraction of "full trust"[16, 15]. I argue that this loses valuable information and can be better captured and represented through folksonomy.

2 Background

The concept of having a person tag another person with regards to their authority on certain topics merges two areas of research, folksonomies and authority. Folksonomies are new and exciting and relatively little has been written academically about them at this time. Authority and reputation, however, have an extensive history of literature both in psychology and sociology as well as business and economics. Some of the more relevant work is summarized below.

2.1 Folksonomies

Folksonomies have recently given us the ability to create messy, free-text, user-created metadata for existing artifacts (books, images, URLs, etc.). The tagging done by a user labels the "aboutness" of an object and later perhaps allows that object to be more easily found, sorted and used by that user or others. Coined in only August of 2004 by Thomas Vander Wal[27], folksonomy has quickly become a hotbed for in-situ applied

research and development in online communities and social networks. The two among these to receive the most credit for advancing folksonomies are the social bookmarking website del.icio.us[26] and photo sharing website flickr.com[2].

2.1.1 Classification

Classification has traditionally been handled by experts, by authority figures our culture has deemed knowledgeable enough to manage the task. These experts wield controlled vocabularies and authority files and synonyms and term lists. They have a relatively fixed set of knowledge about a domain that has been learned over the years and is useful in terms of how to store things for later retrieval. This system works very well and should not be supplanted from its role. However, it does have some weaknesses, not least among these that it does not scale well enough to handle the volume of information being created by individual users today[22]. It works fine for limited or finite amounts of information (journals, offices, libraries, newspapers), but it does not allow individuals a real voice in the matter, individuals who could be helping do some of the work. Folksonomies allow dynamic, decentralized, distributed, user-created metadata to aggregate and take most of the heavy lifting off the backs of the experts. In aggregate, users are pretty good at classification[28]. However, they are also susceptible to herd mentality, and this is one reason why folksonomies will not supplant traditional library science methods for “hard” or “life-and-death” classification. They will only augment and facilitate a scaling of effort.

2.1.2 Del.icio.us

Del.icio.us[26] is a web site where the community of users can “tag”, or freely associate text labels, with URLs. This allows a user to bookmark a site with the added incentive to “label it well” so it can easily be found later, by themselves. It is a selfish system that works because it is simple and easy to understand. It has also proven to be fairly valuable when compared to traditional bookmark lists that simply gathered the title and the URL of a web site and were sorted by date added (and usually not reverse chronologically, putting the most recent additions at the bottom of the screen, if not off the screen entirely). Additionally, the del.icio.us lists of bookmarks themselves are available on the web, as opposed to locally on the user’s machine, so all of a user’s bookmarks are held in a centralized place. No synchronization is necessary and they are fully searchable.

The more social aspect of del.icio.us occurs because everyone can see everyone else's tags, and search and sort based on those tags and related tags and related users. It all becomes very interesting very quickly. Many tools have been created since the inception of del.icio.us to manipulate and investigate the myriad patterns created by users, their objects, and the terms they use to describe them. These three artifacts make up the infrastructure of a tagging system. Each of these three artifacts can be used as a "pivot" point through which users may find other relevant listings of information[12]. Users can find other users who share the same bookmarks, the same language habits, the same interests. They can also find related topics and related links. Perhaps most elegantly, it scales to whatever speed and whatever shape the internet may take[22].

2.1.3 Flickr.com

Flickr[2] works on this same model, but with a slight difference. It too is a selfish system in that users tag their photographs, their images, concerning their aboutness. The difference is that users are tagging their own content instead of others' content (like within del.icio.us). Thomas Vander Wal has labeled this a "narrow folksonomy" as opposed to the del.icio.us "broad folksonomy" where users tag other users' content[29]. They are adding metadata to something that is intrinsically harder to search, or mine, for aboutness since it isn't text and cannot be analyzed with a dictionary or easily compared against other bodies of work. Tagging allows text to be associated with images and therefore brings some of the tools available in mining and latent semantic indexing to bear, but not much. Information retrieval has been historically limited to the text associated with the image rather than the images themselves[25]. Mostly, the tagging allows for contextual information to be made available for search and classification by the owner and other users. By associating free text labels with images already "tagged" with author, date, time, location, it is scalable and simple. Nothing more is required than a simple text box to allow users to annotate existing images. Future implementations might be more a part of our real world, more integrated to our life away from a computer monitor.

Recently, Flickr has given its users the option of allowing other users to tag their content (images), but this is only used by a very small percentage of users as it is not the default behavior of the software. This would bring Flickr more along the spectrum towards a broad folksonomy. Flickr.com is included here as a second example of tags being used in the "popular" sense. In this paper, the ideas put forth will be more in the model of del.icio.us in that tags will be applied to things (other users) *not* created by the user

doing the tagging.

2.2 Traditional Authority

“People understand authority on the basis of personal experience with particular sources or reputations.” (Rieh & Belkin, 1998)[24]

Authority, in the sense that it will be written about in this paper, is concerned with how people trust one another’s opinions and thoughts. As it will be modeled, the standard nomenclature comes from Kleinberg with regards to hubs and authorities in communities. Hubs have many outlinks and authorities are characterized by many inlinks[19, 20]. Rieh and Belkin[24] were most clear in summarizing the literature in this area of authority in terms of trust and reputation. Citing both Wilson[31] and De George[14], authority can be of two types, epistemic or deontic, meaning “an authority” or “in authority”[14]. Wilson[31] named them differently, “cognitive authority” and “administrative authority”. The first is a granting of a knowledgable reputation in a subject area while the second is more the recognizing of someone who can “tell others what to do”[31]. Cognitive authority is what tags can grant. Contextual Authority Tagging can only show who is “an authority”, not who is “in authority”. Research done in administrative authority has typically been centered on subordinates and their bosses in the workplace[4, 3, 6, 10].

Axelrod said in 1984 that trust allows us to give value to the “shadow of the future” today[5]. All our information flows have “historical residues” and should be considered with every decision we make[13]. The decision to grant cognitive authority to someone explicitly should not be taken lightly as it affects how other decisions are made later.

2.2.1 Chains of Authority

The way we grant authority and reputation on a human level has everything to do with who we know and what we perceive they know. A friend might know a lot about horses; at least, from one’s perspective, she does. She is “an authority”, in one’s eyes, on the topic of horses. Paradoxically, she may not consider herself an authority on horses; simply for the fact that because she knows a lot about them, she also knows how much she does *not* know about them.

While this may be common or not, this friend would probably consent to being labeled “an authority” on horses, by those who know her. While not being an ultimate authority

on horses in her own eyes, her friends and acquaintances perceive her as being knowledgeable in that area and she probably does not disagree with their assessment. If those who have labeled her an authority were asked by a third party some question in this area, they would probably defer, or direct the asker, to her. This referral system, pointing up the chain of granted cognitive authority, is practiced daily by all of us. We each know someone who would know someone who would know the answer.

These chains of authority are important in that they provide a deterministic framework for measuring cognitive authority. They allow us to calculate the relative cognitive authority of a person on the network. Those “higher” in the chain have more authority granted to them.

2.2.2 Perspective

A person has a view of the world from his perspective. While, objectively, a reality may exist, it may not seem to exist to the person who does not know about it. A tree falls in the woods. Perspective is what gives each of us our lens on the world. If one feels another has expertise on a topic, that is one’s perspective. A different person may adamantly disagree or simply not know that their acquaintance knows a great deal about a topic. And that’s okay. When a person has granted some cognitive authority on a topic to a friend, the friend, when writing or saying something on that topic, has credibility, as per the granted authority. The friend’s reputation is trusted in that domain. This is how all of us assess the information around us every day. Our incoming information, both news and gossip, are driven ultimately by who we trust and who we believe to be credible. This is Wilson’s epistemic or cognitive authority.

2.2.3 Contextual Authority

Cognitive authority is something that must be granted, not claimed. It can be given for different reasons (fear, trust, etc.) but for the purposes of this paper, it does not matter the reason. Cognitive authority is something that is bestowed on someone by someone else who feels they are sufficiently skilled or knowledgeable in a domain, or in the words of Wilson[31], someone “who knows what about what.”

In this sense, authority is not unlike aboutness. Aboutness is the ethereal quality that an object or person has whereby it defines what the object or person is about. Both authority and aboutness can be conveyed by other people in the form of free text tagging, metadata created by the very people whose perceptions matter. Hjørland[17]

posits that “‘subject’ and ‘aboutness’ should be considered real synonyms in information science.” While perfect synonyms may exist, non-perfect synonyms definitely exist, and it is the job of our graph theory and social network theory to provide that degree of synonymity, that correlation value, by which two things are related. An authority on one topic might not be explicitly defined in one’s network, however, there may be an authority in a closely related topic that would aid one in one’s search.

Additionally, Hjørland[17] writes that we can have differing opinions of aboutness when it comes to objects. That is how he explains the 40% discrepancy of labeling aboutness in the experiments of Bruza et al.[8]. Bruza wrote that there was a “core of agreement” where everyone could agree on aboutness, but this is too simple. The edge cases are where the action is and the idea that different individuals are seen as authorities in different areas only gives weight to this argument. This differing cognitive authority in different areas is “contextual authority” and can be conferred by anyone on any topic.

Contextual Authority Tagging should be very applicable when determining who your group of acquaintances, your “network”, sees as an authority figure on a topic. In the real world, each of us has authority figures for different domains. One gives authority to different individuals about different things, ala Hjørland. Some of those individuals may know about many things, and alternatively, many of those individuals may know about the same thing. The person whose opinion you trust most about basketball is probably not the same person you’d trust concerning issues of hazardous waste, although it could be. There is no single individual in your life who is an authority on all things. You seek different people to answer different questions, and since they know different things, this is normatively rational.

In fact, consider you have a question concerning submarines. You may not necessarily know anyone who knows about submarines directly. Or, rather, you do not think you know anyone who knows about submarines directly. You may very well know someone who has this expertise, but if you aren’t aware of that, you would not ask them for help or for their opinion. Their expertise in that regard is not apparent to you and so you haven’t granted them cognitive authority in that domain. Others, however, may know your friend served in the Navy for ten years and worked with sonar equipment. If you do not know this about him, you would never ask for his help regarding your question.

2.3 Community Effect

Muller[9] writes “authority relationships are mainly built through community members’ contribution to the work of the community (those contributions may consist either in the disclosure of pieces of knowledge or of information or in any other tasks aimed at enhancing the work of the community).” This is important since each person in a community has a reputation. Be it good or bad or related to one thing or another, they have one. In fact, I argue that they have more than one. Each member of a community has a contextual reputation within a certain topic area. They can be “trusted” to give good knowledgeable information regarding a certain finite list of topics. Other topic areas, they are not trusted to know about. They have not demonstrated to their network of acquaintances that they know about those other areas. In a sense, all the algorithms in computer science and social network theory that have been used to distill reputation and trust into a calculable value are really looking at an aggregate opinion across all topic areas. This is a lossy operation. There is information being lost that cannot be recovered after the aggregation of topic areas.

Research with trustmaps and cooperation and the prisoner’s dilemma use trust and reputation as a single variable. This allows for decisions to be made more easily but does not give context to this trust.

What actually happens in the real world is as follows. What Adam thinks about Bob, in general, has been defined largely as an average based on all the points of reference Adam has concerning Bob and their experiences together. Adam weighs his personal experiences, his coworkers’ shared opinions, Bob’s work, all of these, when looking for a single number or value to label as Bob’s reputation. Bob’s reputation, or authority, on how to fix the copier may not be in any capacity the same as his reputation for telling funny jokes in meetings. When aggregated, these subtleties are lost and Adam has to generalize about Bob if asked to characterize his opinion about Bob. If asked specifically, Adam would probably have very precise, and wildly different, opinions about Bob with regards to the copier and with regards to his punchlines. Cindy may have yet another set of opinions on Bob and his copier prowess and ability to tell a joke.

Contextual Authority Tagging allows for these subtleties to be assessed and measured. When enough community members are tagging one another, Adam’s ranked authority list with regards to “copiers” can be easily calculated and viewed. We can see, elegantly, who ranks highest in Adam’s opinion with regards to copiers. And then, just as easily, we can see who in Adam’s opinion tells the funniest jokes at work. If the preceding

example still holds, Bob will unlikely be at the top of both lists.

3 Proposal for Contextual Authority Tagging

The following is a proposal for a system that can be designed and overlaid on an existing system of users or community members. This system would allow the discovery and definition of those who are authorities. There are people in every community who are knowledgeable. This system could allow for those who are knowledgeable to be justly recognized.

There are two main differences in Contextual Authority Tagging and other uses of folksonomy to date. The first is the fact that only other users are tagged. Instead of having three artifacts in the system (users, tags, and objects being tagged), there are only two (users and tags). The users become the objects. The second main difference is that users are not tagging objects with “aboutness” or labels. They are not labeling what something is or what it is about. They are labeling what someone knows, what they are good at, what they are an authority on.

3.1 Related Tags

Amazon.com[1] has for some time now displayed to shoppers what other shoppers additionally bought when purchasing a certain item. This is done through clustering algorithms and purchase histories (data mining). Del.icio.us does this as well, except with tags applied to URLs instead of purchased items. URLs that have been tagged with a certain word have usually also been tagged with other words. These other words are therefore related and can be inferred to be close to the original tag in terms of aboutness. With enough users, this system elegantly begins to classify words together, not semantically, but mathematically, statistically. It seems as though the system knows something about the meaning of the tags themselves, but it does not. The users know the meanings, and through their usage patterns, the system can simply report what it counts. This same effect could be leveraged to apply heuristics to related tags for authority. If a user is searching for “submarines”, the words “Navy” and “sonar” may be related terms.

Relatedness can be measured either by correlation with other tags applied to the same person, or by a synonym database generated by the users in aggregate, or some combination thereof. I’m concerned that relatedness strictly determined by correlation

of inbound links to users tagged with a certain word will be too mutually exclusive to find “like terms”, in practice. I’m not sure the correlation is that of topicality rather than something else less helpful. A third party synonym database or tagged synonym database itself would alleviate this potential shortcoming. It would also allow for a system consisting of a smaller number of users to be more effective more quickly. The time required to reach a sufficient number of active tags would be much less. I think at this point, the most robust place for this system to work, all things being equal, is the Internet as a whole (like del.icio.us). Small groups may find the exercise most entertaining, but it may not aid in the discovery of really new information. More likely it will simply formalize some structure that the small group already knew about but had never quantified. Perhaps, still, this might be relevant by itself.

3.2 Degrees of Separation

With enough users and enough tags in a system, almost every topic imaginable will have an authority figure within your close network. This is related in part to the Small World phenomenon[30] and has been shown to have broad influence in many diverse areas of scientific research. An individual user may not have anyone that they know who is an authority on a particular topic, but some of their trusted friends in a set of related tags, does. And if not, then their trusted friends, with authority, know someone, etc. You may not know someone who knows about submarines, but you know someone who was in the Navy who probably does.

In regards to how much this should affect a ranking, a second order connection like this should not be worth as much as a first degree ‘hit’ for a topic area, but it should be worth more than a third or fourth degree ‘hit’. And likewise, the closer the word ‘Navy’ is associated with the word submarine, the higher the ‘score’ should be for my friend’s friend as an authority on submarines, from my perspective.

Generally, if exact words are used to describe someone’s authority and they’re in your network, they should percolate to the top of your list of whom to ask. As the words used to describe someone’s authority begin to drift away from the original query, they should descend the list of authorities in your network on that topic.

3.3 Non-negative Authority

Having first considered the possibility of having users assign a value to the authority they were bestowing on someone, this idea quickly passed. If authority ratings were to

be kept on a scale of -1 to 1, the community would become susceptible to wild swings of ratings, especially over contentious subjects or people. These wars of opinion would not help anyone figure out who is being looked to as an opinion leader and would hurt the commons terribly. A far more elegant solution is to simply have binary authority assigned on a topic for a first degree neighbor.

Removing the idea of negative authority allows perfect strangers to hold the same sway over a particular user's perspective of authority as someone they feel strongly should *not* be pointed to as an authority figure on a particular topic. A user can tag users he finds authoritative on a subject with that topic or word. Users who do not fit that description are not tagged. Cognitive authority, as a value that can be assigned in the system, can only be zero (0) or one (1). Positive values are standard when attributing value to a directed graph in social networks, and this is no different[23].

After a calculation is performed, the authority for a person is only a non-negative value with a range from zero (0) upwards. Zero (0) means a user has no connection to a topic from a particular user's perspective. The higher the score, the higher the amount of relative cognitive authority that user has.

3.4 Option to Defer

Deferral of cognitive authority is good for two reasons, it would allow flame wars to be preemptively disincentivized and privacy to be protected.

In terms of thinking through how this system might be open to gaming by those who want to affect their own rankings or the rankings of others, I have been led to consider the possibility of users having the option to defer specific authority tags. This creates the complexity of having bidirectional directed graphs (with reciprocal links), but I think makes the model that much more robust[7].

3.4.1 Spreading Untruths

If a user has been tagged by someone as an authority on something that is not flattering or simply not true or even just questionable in the eyes of the tagged, they should be allowed to defer that "authority". Since there are no negative values associated with tags (only 0 or 1), deferring would simply have the effect of not ever having been tagged, and therefore, lowering their own authority in that area of expertise. This would make a system more robust (and slow) as it would require both parties to play along

but would also require more work of the members to “police” their own reputation or authority tags. Of course, we do this everyday regardless of there being a system where it is mathematically deterministic and published for the world to see. Each of us constantly monitors what we publish about ourselves and manage our reputation amongst our peers. It is only natural to consider what others think of us and if nothing else, this allows for greater introspection and self-awareness.

3.4.2 Privacy

Another reason for considering bidirectionality of links (both users have to take an active role) is privacy. Without a tag being blessed by the receiver, it should not be made public. We only present a certain slice of ourselves to others. Perhaps we don’t want their opinion of us broadcast. They might know things that have intentionally been kept secret from others.

If a friend of mine has been tagged as an authority on abortion, yet she’s never been pregnant, doesn’t have a boyfriend, and should not really have any reason to be considered an authority on abortion, she may not wish that information to be broadcast to the network. She would wish to suppress (defer) the fact that two of her close girlfriends consider her an authority on abortion.

Perhaps early systems could be implemented both ways (with and without deferral) to see where the problems occur. I would guess that “Deferral/Blessing” would quickly emerge in a preferred default feature set.

3.5 Power Law Observation

The power law has been observed in tag profiles for artifacts (in del.icio.us and other systems) and this should hold for people or “users” as well[18]. There is no observed reason to expect otherwise. This power law distribution will attribute cognitive authority to certain individuals in certain areas. While someone might be known for a few different things to different people, the overall opinion of a user will be visible from the community and their true valued position in the world will be apparent.

4 Global and Personal Views

There are four types of query results that can be returned by this system. Type A and B queries are concerned with the topic areas that a particular user is an authority on. Type C and D queries would return the users who have been granted authority in this community on the topic of interest.

	Global	Personal
What is user X an authority on?	Type A	Type B
Who is an authority concerning topic Y?	Type C	Type D

Each type of search would return a ranked list based on the proposed algorithm in the next sections. Keep in mind that a system that decides to implement Contextual Authority Tagging could choose to expose any or all of these types of queries to the users. That would be a policy decision at the time of design and has no bearing on this discussion here. Each type of query is also affected by whether or not the system has implemented the Option to Defer (bidirectional links).

4.1 Type A Queries - User/Global

This type of query is very straightforward. The list of topics returned for a Type A query can be calculated by recursively (moving back down the chain of authority) looking at each tag that has been attributed to that user. Each inbound tag gets a score of 1 and all like tags are summed. Each tagging user's Type A score for the same topic is added to the total. The topics are then ranked and presented as results to the querying user. Related terms do not influence this list but can be generated as pivot links in an additional list. Since this is the global view, privacy is not an issue. If there are no bidirectional links, all links should be considered equal. If there are bidirectional links, then only they should be considered. Sort order for this list of returned topics should be from highest authority score to lowest authority score.

4.2 Type B Queries - User/Personal

If there are no bidirectional links, and privacy issues are inherent in the structure of the network, this calculation is the simplest of the four types. The list of returned topics

will be the exact list of tags the querying user has attached to the user in question. There are no network effects and no chains of authority. An associated list of related tags could be generated from the querying users' own statistical corpus of tags, but this may not be very helpful without a large number of links entered by the user. The only users who would have a non-zero authority score are those who are in the personal first-degree network of the querying user; those users who have been tagged directly. The sort order of this topic list is fairly irrelevant as the scores would all be 1.

If bidirectional links are in use, and tagged users have allowed links to be public, this type of query becomes more like a type A query. Public links can be followed up a chain of authority and accumulated. I would suggest this type of query be capped at second-degree links since beyond the "Friend-of-a-Friend" distance in a network, there is no common hub or shared friend. Users beyond the second-degree can be argued to be beyond the reach of a user's personal network. Sort order for this list of returned topics should be from highest authority score to lowest authority score.

4.3 Type C Queries - Topic/Global

This type of query is also recursive in nature. The user with the highest matching inbound tag score is listed first followed by the rest of the users who are tagged with the queried topic. An inbound tag score is calculated by summing all the matching inbound tags recursively down the chain of authority. For the matching topic, exactly, each inbound tag is given the value 1. Related topics have a correlation value that is less than 1, but greater than 0. The relatedness of these terms comes directly from whatever system was implemented to calculate relatedness. This could be any of the three listed earlier: clustering algorithms with only the tags in the system as the corpus, an external synonym database, or a hybrid of these two. Each inbound tag for a related topic is given the correlation value and summed recursively as well. Users who are designated with enough authority, regardless of which tag(s) contributed to their total authority score, are sorted and presented as the results of a Type C query. If bidirectional links are in use, then only they should be considered.

4.4 Type D Queries - Topic/Personal

If there are no bidirectional links, only first-degree links can be used to calculate who is an authority in the querying user's network. This would mimic most closely what we do in our heads in real life. We do not know what authority other users have given

one another, and only if we ask and they comply do we have access to that information. Therefore, the returned list of users for a Type D query will be all the users who have been tagged with the topic by the querying user. Next on the list will be the users tagged with the most related topic as determined by the user's personal corpus of tags or a third party synonym database, or a hybrid of the two. The entire system corpus is not available at the Personal level of query.

If bidirectional links are in use, then the "authority's authority" can play a role like in Type B. Authority scores can be accumulated up the chain of authority, but not past the second-degree. The authority score for each user returned should be the cumulative authority score for the queried topic and all related topics. The resulting list should be sorted with highest scores first.

5 Usefulness of this Proposal

Contextual Authority Tagging allows for the explicit discovery and definition of cognitive authority in social networks. Any organization could benefit from knowing where the expertise lies in their ranks. Those who know things can be more properly recognized for their expertise. If used widely, this system of cognitive authority recognition would allow an organization, or an entire society, to move closer to a true meritocracy. This system encourages a richer information economy.

References

- [1] Amazon.com, April 2005.
- [2] Flickr.com, April 2005.
- [3] Philippe Aghion and Jean Tirole. Formal and real authority in organizations. *The Journal of Political Economy*, 105(1):1, February 1997.
- [4] Nabil I. Al-Najjar. A reputational model of authority. *Journal of Economic Behavior & Organization*, 46(2):165–191, October 2001.
- [5] Robert Axelrod. *The Evolution of Cooperation*. Basic Books, New York, 1984.
- [6] George Baker, Robert Gibbons, and Kevin J. Murphy. Informal authority in organizations. *Journal of Law, Economics & Organization*, 15(1), 1999.
- [7] Lennart Björneborn and Peter Ingwersen. Toward a basic framework for webometrics. *Journal of the American Society for Information Science and Technology*, 55(14):1216–1227, 2004.
- [8] P. D. Bruza, D. W. Song, and K. F. Wong. Aboutness from a commonsense perspective. *Journal of the American Society for Information Science*, 51(12):1090–1105, 2000.
- [9] Danish Research Unit for Industrial Dynamics (DRUID). *The role of authority in the governance of knowledge communities*, Winter 2003.
- [10] Wouter Dessein. Authority and communication in organizations. *Review of Economic Studies*, 69(241):811–838, October 2002.
- [11] Chris Ding, Xiaofeng He, Parry Husbands, Hongyuan Zha, and Horst D. Simon. Pagerank, hits and a unified framework for link analysis. In *SIGIR '02: Proceedings of the 25th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 353–354, New York, NY, USA, 2002. ACM Press.
- [12] Bowen Dwelle. dwelle.org: How i use tags, January 2005.
- [13] Raymond Fisman and Tarun Khanna. Is trust a historical residue? information flows and trust levels. *Journal of Economic, Behavior & Organization*, 38:79–92, 1999.

- [14] Richard T. De George. *The Nature and Function of Epistemic Authority*, chapter Authority: A Philosophical Analysis, pages 76–93. University of Alabama Press, University, AL, 1976.
- [15] Jennifer Golbeck and James Hendler. Inferring reputation on the semantic web. In *Proceedings of the 13th International WWW Conference (WWW 2004)*, 2004.
- [16] Jennifer Golbeck and James Hendler. Reputation network analysis for email filtering, 2004.
- [17] Birger Hjørland. Towards a theory of aboutness, subject, topicality, theme, domain, field, content ... and relevance. *Journal of the American Society for Information Science and Technology*, 52(9):774–778, 2001.
- [18] Ben Hyde. Power-law distribution of popular url’s at del.icio.us, February 2005.
- [19] Jon M. Kleinberg. Authoritative sources in a hyperlinked environment. *J. ACM*, 46(5):604–632, 1999.
- [20] Jon M. Kleinberg. Hubs, authorities, and communities. *ACM Comput. Surv.*, 31(4es):5, 1999.
- [21] Michael W. Macy and John Skvoretz. The evolution of trust and cooperation between strangers: A computational model. *American Sociological Review*, 63:638–660, October 1998.
- [22] Adam Mathes. Folksonomies - cooperative classification and communication through shared metadata, December 2004.
- [23] Josep M. Pujol, Ramon Sangüesa, and Jordi Delgado. Extracting reputation in multi agent systems by means of social network topology. In *AAMAS '02: Proceedings of the first international joint conference on Autonomous agents and multiagent systems*, pages 467–474, New York, NY, USA, 2002. ACM Press.
- [24] Soo Young Rieh and Nicholas J. Belkin. Understanding judgement of information quality and cognitive authority in the www. In *Proceedings of the 61st ASIS Annual Meeting (Vol 35)*, pages 279–289. ASIS, 1998.
- [25] Y. Rui, T. Huang, and S. Chang. Image retrieval: current techniques, promising directions and open issues. *Journal of Visual Communication and Image Representation*, 10(4):39–62, April 1999.
- [26] Joshua Schachter. Del.icio.us, April 2005.

- [27] Gene Smith. Folksonomy: Social classification, April 2005.
- [28] James Surowiecki. *The Wisdom of Crowds: Why the Many Are Smarter Than the Few and How Collective Wisdom Shapes Business, Economies, Societies and Nations*. Doubleday, May 2004.
- [29] Thomas Vander Wal. Folksonomy explanations, January 2005.
- [30] Duncan J. Watts and Steven H. Strogatz. Collective dynamics of 'small-world' networks. *Nature*, 393(6684):440–442, June 1998.
- [31] Patrick Wilson. *Second-hand knowledge: An inquiry into cognitive authority*. Greenwood Press, Westport, CT, 1983.
- [32] Bin Yu and Munindar P. Singh. Detecting deception in reputation management. In *AAMAS '03: Proceedings of the second international joint conference on Autonomous agents and multiagent systems*, pages 73–80, New York, NY, USA, 2003. ACM Press.