Augmented Social Cognition

Ed H. Chi, Peter Pirolli, Bongwon Suh, Aniket Kittur, Bryan Pendleton, Todd Mytkowicz

Palo Alto Research Center
3333 Coyote Hill Road, Palo Alto, CA 94304
{echi, pirolli}@parc.com

Abstract

Research in Augmented Social Cognition is aimed at enhancing the ability of a group of people to remember, think, and reason; to augment their speed and capacity to acquire, produce, communicate, and use knowledge; and to advance collective and individual intelligence in socially mediated information environments. In this paper, we describe the emergence of this research endeavor, and summarize some results from the research. In particular, we have found that (1) analyses of conflicts and coordination in Wikipedia have shown us the scientific need to understand social sensemaking environments; and (2) information theoretic analyses of social tagging behavior in del.icio.us shows the need to understand human vocabulary systems.

Introduction

One enduring core value of research in Human-Computer Interaction (HCI) at PARC and elsewhere has been the development of technologies that *augment human intelligence*. This mission originates with Douglas Engelbart, who inspired researchers like Alan Kay at PARC in the development of the personal computer. The aim of augmented human cognition has remained a core value in the development of, for example, information visualizations, information foraging theory, personalized search, and information scent tools and technologies.

Over the last few years, we have realized that many of the information environments are gradually turning people into social foragers and sharers. People spend much time in communities, and they are using these communities to share information with others, to communicate, to commiserate, and to establish bonds. This is the "Social Web". While not all is new, this style of enhanced collaboration is having an impact on people's online lives.

Augmented Social Cognition research area at PARC has emerged from this background of activities aimed at understanding and developing technologies that enhance the intelligence of users, individually and in social collectives, through socially mediated information production and use. In part this is a natural evolution from our work around improving information seeking and sense making on the Web. In part this is also a natural expansion in our scientific efforts to understand and enhance the intelligence of the individual users coupled to information systems.

Definition of Augmented Social Cognition

A natural extension of augmenting human intelligence in the Social Web and Web2.0 world is the development of technologies that augment social intelligence. In this spirit, the meaning of "Augmented Social Cognition" builds on Engelbart's vision, and can be explained by deconstruction of the term:

- ***Cognition*** means the ability to remember, think, and reason; the faculty of knowing; to have functions associated with intelligent action such as perceiving, remembering, planning, deliberating, and learning (acquiring knowledge and experience).
- ***Social Cognition***[1] is the ability of a group of people to remember, think, and reason; the construction of knowledge structures by a group of people; socially mediated acquisition and use of knowledge.
- ***Augmented Social Cognition*** means the enhancement via technical systems of the ability of a group of people to remember, think and reason, acquire and use knowledge.

Our interest in this area obviously also arouse due to the emergence of Web2.0 and Social Web applications. In this short article, we will summarize examples of recent Augmented Social Cognition research on:

- How conflict and coordination have played out in Wikipedia?
- How social transparency might affect reader trust in Wikipedia?
- How to understand the efficiency of tagging vocabularies in social tagging systems using information theory?

Technology Trends

Web2.0 is a broad term used to mean a new wave of new technologies that is hitting the Web in full-force. What is different about this new Web 2.0 environment is that people are sharing information today in a fundamentally

[1] "Social cognition" has been used for years in psychology to designate the cognitive mechanisms people employ in social interactions. (See for example, Z. Kunda, *Social Cognition: Making Sense of People*, MIT Press, 1999.) Our definition here is intended to include this previous definition, as cognition around social interactions is often a component of the social construction of knowledge structures.

different way from how they are used to. One example is Wikipedia, which is a fascinating collaborative editing environment for creating an encyclopedia. The collaboration that happens here is very different from passing documents back and forth using traditional email, because you have (1) automatic versioning, and (2) you can always go back and find out who contributed what (transparency). Developments like this have taken a lot of the burden off of users. The features reduce the time it takes to collaborate with each other. This wave of new technologies is generated by a combination of new developments, including:

1. *Software as a Service or Web as Platform*. Web technologies have advanced to the point that the Web itself (and other connected networks) has become a computing platform for the delivery of novel features, tools, applications, and services. The computing platform involves a heterogeneous mix of technologies including REST, XML Web Services, RSS/Atom, and AJAX. The web platform provides the plumbing and necessary parts to support rich user interaction, mashups or remixing of Web Services, and the formation of social groups and interactions. *Mashups*: One consequence of the Web as platform is that it fosters innovative combinations of services, such as the connection of search engines or RSS feeds to Google Maps (a web service) to deliver results on geographical data to the end-user.
2. *Rich interaction*. New Web user interfaces no longer rely on the old paradigm of submitting results to the server and waiting for a new page to load. Instead, in its place, we have rich interactive applications that use asynchronous communication to servers to deliver fully interactive user experiences. With higher bandwidth, not only is there more "rich media" (e.g., video), but a richer variety of user-friendly ways to interact with content.
3. *Harnessing network effects of knowledge production and use*. Perhaps the most significant and exciting consequence of the evolution in technology is the emergence of novel architectures of participation that draw users to contribute value, and that gain value as more users cooperate. Novel systems support the creation and aggregation of knowledge through cooperative peer production (e.g., Wikis, blogs, social bookmarking), and others that augment intelligence through cooperative reasoning and judgment (e.g., prediction markets; voting).

Research trends

Researchers are also similarly seeing a surge of new research on Web2.0 technologies distributed in a wide variety of disciplines and associated conferences.

- At the light-end of collaboration spectrum, we have researchers trying to understand the micro-economics of voting systems, of individual and social information foraging behaviors, processes that govern information cascade, and wisdom-of-the-crowd effects. Economists are trying to understand peer production systems, new business models, and consumption and production markets based on intrinsic motivations.

- At the middle of the collaboration spectrum, researchers are building algorithms that mine new socially constructed knowledge structures and social networks. Here physicists and social scientists are using network theories and algorithms to model, mine, and understand these processes. Algorithms for identifying expertise and information brokers are being devised and tested by information scientists.

- At the heavy-end of the collaboration spectrum, the understanding of coordination and conflict costs are especially important for collaborative creation systems such as Wikis. Researchers had studied characteristics that enable groups of people to solve problems together or collaborate on scientific endeavors. Discoveries such as the identification of invisible colleges have shown that implicit coordination can be studied and characterized.

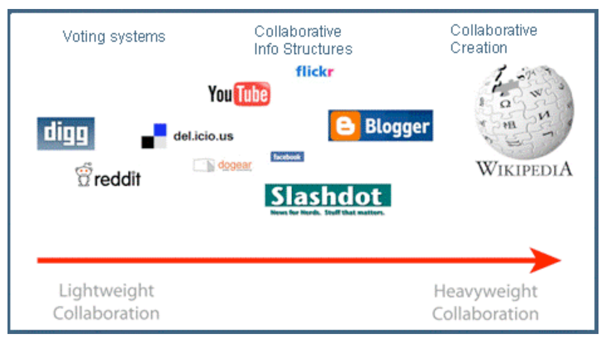

Figure 1: research spectrum in Augmented Social Cognition.

Also, modelers and scientists are trying to understand how to bring down the cost of social interactions, and understand the cost/reward structure for individuals. They are also building characterization models of what, how, and why people are behaving the way they do. Field studies, log file and content analysis, as well as cognitive task analysis are possible studies to conduct in this space.

Research: Conflict and Coordination in Wikipedia

As an example of building models and understanding how Web2.0 systems operate, we have been engaged in understanding how conflicts and coordination works in Wikipedia [Kittur07]. Wikipedia, a wiki-based encyclopedia, has become one of the most successful experiments in collaborative knowledge building on the Internet. As Wikipedia continues to grow, the potential for conflict and the need for coordination increase as well. Researchers have seen similar costs in other computer mediated communication (CMC) systems such as MOOs and MUDs [Curtis92, Dibbell93]. Even though researchers have documented the growth of Wikipedia [Voss05], the impact of coordination costs has largely been ignored. Conflict in online communities is a complex phenomenon. Though often viewed in a negative context, it can also lead to positive benefits such as resolving disagreements, establishing consensus, clarifying issues, and strengthening common values [Franco95].

Global Coordination

Here we try to understand the conflict and coordination costs through the concept of indirect work. Viewed from the goal of trying to create high quality content for a collaborative encyclopedia, we define "indirect work" or "conflict and coordination costs" as *excess work in the system that does not directly lead to new article content*. This allows us to develop quantitative measures of coordination costs, and also has broader implications for systems in which maintenance and consolidation occur.

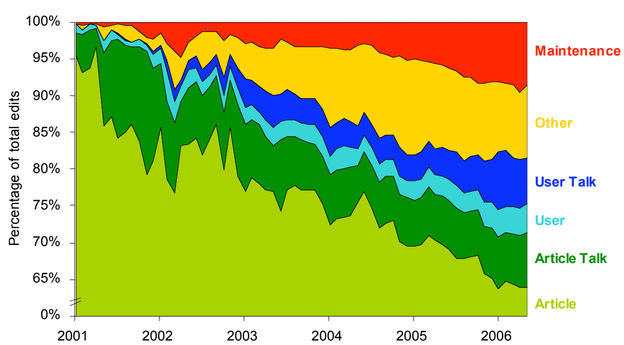

Figure 2. Changing percentage of edits over time showing that decreasing direct work (article) and increasing indirect work (article talk, user, user talk, other, and maintenance).

Overall, user, user talk, procedure, and other non-article pages have become a larger percentage of the total edits made in the system. These trends are summarized in Figure 2, which clearly shows the decreasing percentage of edits going to direct work (article edits) and the increasing percentage of edits going to indirect work across different page types.

Article Conflicts

We wanted to better understand and characterize article-level conflicts. Our goal was to develop an automated way to identify what properties make an article high in conflict using machine learning techniques and simple, efficiently computable metrics. We used the Support Vector Machine (SVM) algorithm to learn what page features predict article conflict scores.

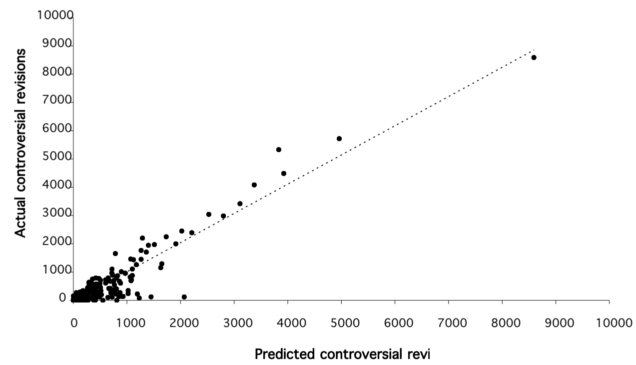

Figure 3. Model performance on articles tagged as controversial. $R^2 = 0.897$.

The machine learner provides insight to this in the weights it assigns to various page metrics. These weights are determined by the utility of a metric in predicting CRC scores, and are shown in order of importance in Table 1.

↑	1. Revisions (article talk)
↑	2. Minor edits (article talk)
↓	3. Unique editors (article talk)
↑	4. Revisions (article)
↓	5. Unique editors (article)
↑	6. Anonymous edits (article talk)
↓	7. Anonymous edits (article)

Table 1. Highly weighted metrics, rank ordered. Up arrows indicate positive correlation with conflict; down arrows indicate negative correlation with conflict

By far the most important metric to the model was the number of revisions made to an article talk page (#1 above). This is not unexpected, as article talk pages are intended as places to discuss and resolve conflicts and coordinate changes. Some of the metrics are more surprising; for example, one might expect that the more points of view are involved, the more likely conflicts will arise. However, the number of unique editors involved in an article *negatively* correlates with conflict (#5 above), suggesting that having more points of view can defuse conflict.

Another interesting finding is that while anonymous edits to the article talk page correlate with increased conflict (#6), they correlate with reduced conflict when made to the main article page (#7). This suggests that anonymous editors may be valuable contributors to Wikipedia on the article page where they are adding or refining article content. However, anonymity on the article talk page, where heated discussions often occur, seems to fan the flames. This suggests that anonymity may be a two-edged sword, useful in lowering participation costs for content but less so in conflict resolution situations.

User Conflicts

The characterization of conflicts between users is crucial to understanding the motivation of users and the sources of conflicts. The goals are to 1) identify users involved in conflicts; 2) characterize ongoing conflicts; and 3) develop a tool that can help in understanding the conflicts.

We built a tool called Revert Graph to visualize user conflict on a particular article. Revert Graph retrieves all users who have participated in reverts and visualizes a graph based on revert relationships between the users (Figure 4 and Figure 5).

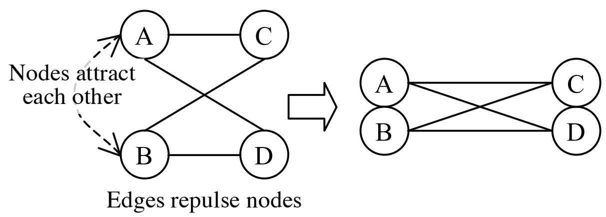

Figure 4. Force directed layout structure employed in Revert Graph. Users (represented as nodes) attract each other unless they have a revert relationship. A revert is represented as an edge. When there are reverts between users, they push against each other. Left figure: Nodes are evenly distributed as an initial layout. Right figure: When forces are deployed, nodes are rearranged in two user groups.

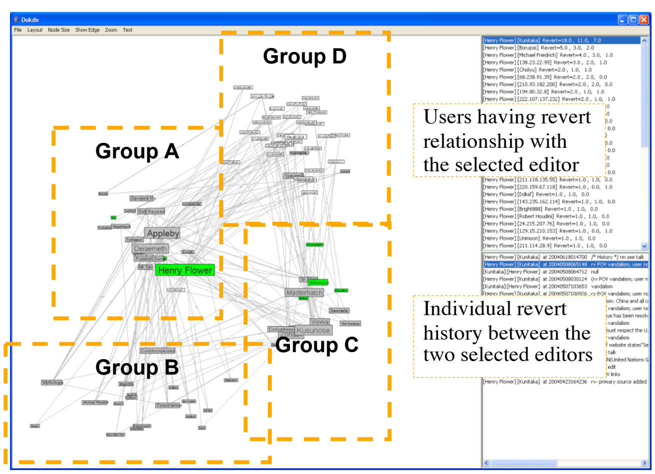

Figure 5. Revert Graph for the Wikipedia page on *Dokdo*. Revert Graph uses force directed layout to simulate revert relationship between users. The tool also allows users to drill down into revert relationships, which enables them to investigate the nature of the conflicts.

We can identify user clusters based on the assumption that a group of users have closer views on a topic the more they revert users in another user group.

The Wikipedia page on Dokdo (Figure 5) is one example where we were able to find interesting user clusters. Dokdo is a disputed islet in the Sea of Japan (East Sea) currently controlled by South Korea, but also claimed by Japan as Takeshima. Figure 5 shows user groups discovered on the Dokdo article. We manually labeled each user based on his/her position on the issue. The majority of users in Group A supports the Korean claims while users in Group C show the opposite pattern. Unlike Group A and C, users in Group D and B showed mixed opinion on the issue.

WikiDashboard

As the visualization above shows, accountability has been recognized as an important factor influencing trust in many online interactions and it plays an increasingly important role in collaborative knowledge systems such as wikis [Denning05]. Although users can access past revisions of every page, it is difficult and time-consuming even for dedicated users to make sense of the history of a page, because many page histories run into the thousands of edits. We are investigating how providing access to this type of accountability information, i.e. who edits how many revisions for an article, in a digestible form could affect users' trust and interpretation of an article. If so, the approach can result in reducing the risks many perceive as inherent to a system [Denning05] in which anyone can contribute or change anything.

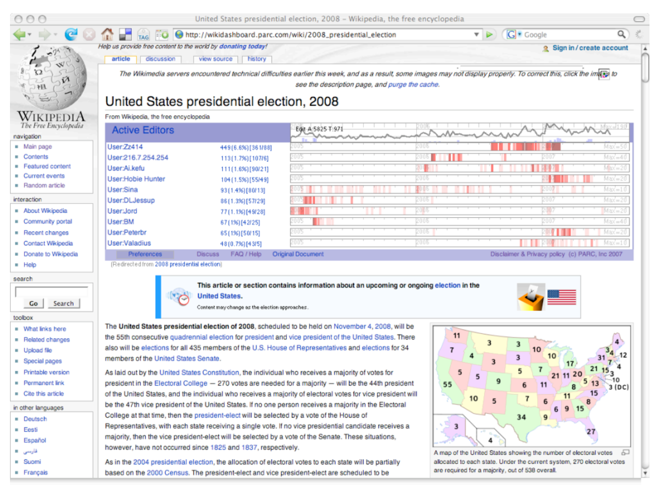

Figure 6. WikiDashboard is a visualization overlay for live Wikipedia pages. The dashboard provides a useful visual digest about who edits how many revisions on each Wikipedia page. It allows users to easily evaluate social activities and patterns around the page, which may be hard to detect otherwise. This figure shows an example of the tool applied to the Wikipedia article "United States presidential election, 2008"

To address this challenge, we designed WikiDashboard (http://wikidashboard.parc.com), a tool that helps users to identify interesting edit patterns in Wikipedia pages, patterns that may be very hard to detect otherwise [Suh07]. As shown in Figure 6, the site provides a dashboard for each page in Wikipedia, while proxying the rest of the content from Wikipedia. The dashboard provides a visualization overlay onto every live Wikipedia page, enabling users to be aware of social dynamics and context around the page they are about to read. The prototype can be used just as if users are on the Wikipedia site itself.

Each article has an associated article dashboard that displays an aggregate edit activity graph representing the weekly edit trend of the article, followed by a list of the active users who made edits on the page.

A user page is like a home page to display information relating to a user. In our system, each user page has a User Dashboard embedded, displaying the article contribution and editing patterns of that user (Figure 7).

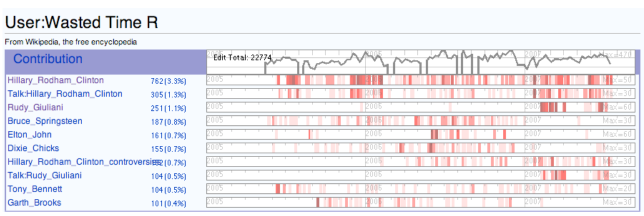

Figure 7. User Dashboard is embedded in each user page of Wikipedia. The dashboard displays weekly edit trend of an editor as well as the list of articles that the editor made revisions on. This example shows a user, "Wasted Time R" made significant edits on articles related to New York politicians and pop singers.

Theories of social translucence [Erickson02] state that three building blocks are necessary for effective communication and collaboration: making socially significant information visible and salient; supporting awareness of the rules and constraints governing the system; and supporting accountability for actions. The idea of social translucence suggests that WikiDashboard could benefit not only readers but also improve the effectiveness of active writers.

WikiDashboard has been available for around just a month with over 4,200 visits and 21,000 page views. Thus, we have already been able to capture a number of insightful feedbacks from various users:

"*WikiDashboard appears to be a valuable tool that can provide some good insights into individual edit patterns and edit conflicts on specific articles. As a means of learning about the tool I have found it useful to use it on articles that I have an intimate understanding of development in order to get a feel of how it can be used and interpreted*." [2]

"*This is very useful for getting a quick glance of the user's editing interests over time. … I actually think a tool like WikiDashboard presents significantly more utility, and is the beginning of an interesting trend of repurposing metadata to create a trust heuristic*." [3]

Research: Understanding Vocabulary Systems in Social Tagging with Information Theory

Given the rise in popularity of social tagging systems, it seems only natural to ask how efficient is the organically evolved tagging vocabulary in describing any underlying document objects? Does this distributed process really provide a way to circumnavigate the traditional categorization problem with ontologies? Shirky argues that since tagging systems does not use a controlled vocabulary, it can easily respond to changes in the consensus of how things should be classified [Shirky05].

Furnas mentioned that a potential cognitive process for explaining how social tagging works might arise out of an analysis of the "vocabulary problem" [Furnas06]. Specifically, Furnas mentioned that the process for generating a tag for an item that might be needed later appears to be the same process that is used to generate search keywords to retrieve a particular item in a search and retrieval engine.

Furnas' comment pointed to the usefulness of social tagging systems as a communication device that can bridge the gap between document collections and users' mental maps of those collections. Social navigation as enabled by social tagging systems can be studied by how well the tags form a vocabulary to describe the contents being tagged.

We analyzed a social tagging site, namely del.icio.us, with information theory in order to evaluate the efficiency of this social tagging site for encoding navigation paths to information sources [Chi07].

We show that entropy analysis from information theory provides a natural way to understand the descriptive encoding power of tags, which appears to be weaning. We found that users appear to have responded by increasing the number of tags they use to describe each item. This metric should be helpful in future analysis of social tagging systems.

As shown in Figure 8, one can see that the entropy of the document set, $H(D)$, continued to increase. We know that the number of documents in the system is increasing, contributing to this increase in entropy. This means that,

---

[2] The Wikipedia Review, http://www.wikipediareview.com

[3] Unit Structures, http://chimprawk.blogspot.com/

over time, users continue to introduce a wide variety of new documents into the system and that the diversity of documents is increasing over time.

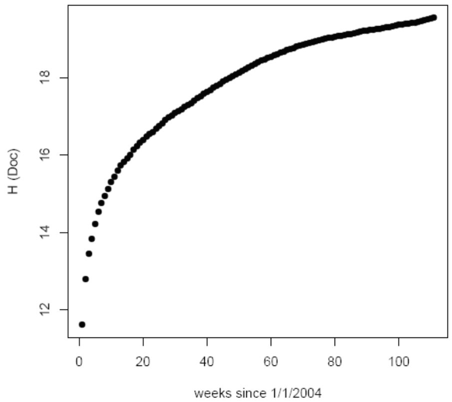

Figure 8. Entropy of documents H(D) is increasing over time.

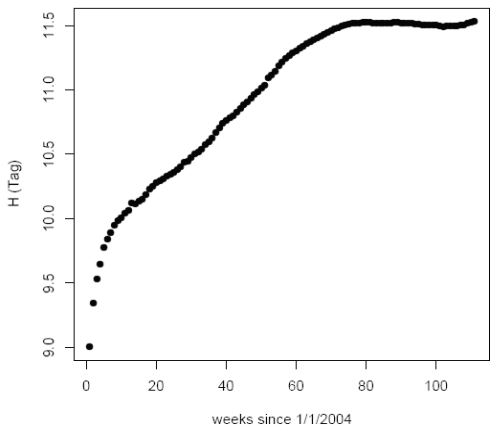

Figure 9. Entropy of tags *H(T)* is increasing at first, then started to plateau around Week 75 (mid-2005).

Figure 9 shows a marked increase in the entropy of the tag distribution *H(T)* up until week 75 (mid-2005) at which point the entropy measure hits a plateau. At the same time, the total number of tags is increasing, even during the plateau section. Since the total number of tags kept increasing, tag entropy can only stay constant in the plateau by having the tag probability distribution become less uniform. What this suggests is that eventually the tagging vocabulary saturated, and coming up with new keywords became difficult. That is to say, a user is more likely to add a tag that is already popular than to add a tag that is relatively obscure.

More importantly, the entropy of documents conditional on tags, *H(D|T)*, is increasing rapidly (see Figure 10). What this means is that, even after knowing completely the value of tags, the entropy of the document is still increasing. Conditional Entropy gives us a method for analyzing how useful a set of tags is at describing a document set. The fact that this curve is strictly increasing suggests that the specificity of any given tag is decreasing. That is to say, as a navigation aid, tags are becoming harder and harder to use. We are moving closer and closer to the proverbial "needle in a haystack" where any single tag references too many documents to be considered useful.

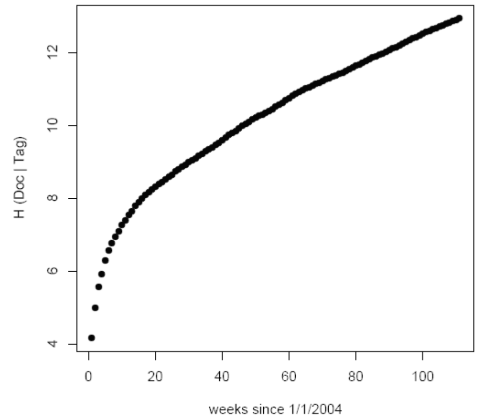

Figure 10. Entropy of Documents conditional on Tags *H(D|T)* increases over time.

Conclusions

Augmented Social Cognition is a new area to understand and develop engineering models for systems that enhance a group's ability to remember, think, and reason. While more enterprises contemplate the benefits of Web 2.0 social software (enhanced collaboration, innovation, knowledge sharing), the coordination and interaction costs that occur in social systems are often overlooked. In this article, we outlined our current research:

First, we are characterizing the various social web spaces in order to understand its collaboration and coordination models. Based on extensive studies of social systems such as del.icio.us and Wikipedia, we have started to identify multiple factors that must be managed to realize the full benefits of these systems within the enterprise.

Second, we are building new social web applications based on the concepts of social transparency and balancing interaction costs and participation levels.

Acknowledgements

This technical note contains research results from collaborations with current and past members of our research group. We thank them for great advice and comments.

References

1. Chi, E.H., Mytkowicz, T. Understanding the Efficiency of Social Tagging Systems using Information Theory. Manuscript Submitted for Publication. 2007.

2. Curtis, P. *Mudding: Social phenomena in text-based virtual realities.* Palo Alto Research Center, 1992.
3. Denning, P., Horning, J., Parnas, D., and Weinstein, L. “Wikipedia risks”, *Comm. ACM,* 48(12), ACM Press (2005), 152-152.
4. Dibbell, Julian. A Rape in Cyberspace. *Village Voice, 21* (1993), 26-32.
5. Erickson, T., Halverson, C., Kellogg, W.A., Laff, M., and Wolf, T. "Social translucence: designing social infrastructures that make collective activity visible", *Comm. ACM*, 45 (4), ACM Press (2002), 40-44.
6. Franco, V., Piirto, R., Hu, H. Y., Lewenstein, B. V., Underwood, R., and Vidal, N. K. Anatomy of a flame: conflict and community building on the Internet. *Tech. and Society Magazine, IEEE, 14* (1995) 12-21.
7. Furnas, G. W., Fake, C., von Ahn, L., Schachter, J., Golder, S., Fox, K., Davis, M., Marlow, C., and Naaman, M. 2006. Why do tagging systems work?. In *Proc. CHI '06 Extended Abstracts*. ACM Press, New York, NY, 36-39.
8. Kittur, A., Bongwon Suh, Bryan Pendleton, Ed H. Chi. He Says, She Says: Conflict and Coordination in Wikipedia. In *Proc. of ACM CHI 2007 Conference on Human Factors in Computing Systems*, pp. 453--462, April 2007. ACM Press. San Jose, CA.
9. Shirky, Clay. Ontology is Overrated: Categories, Links, and Tags. Blog entry. http://shirky.com/writings/ontology overrated.html (retrieved Sept 21, 2006).
10. Suh, B., Chi, E, Kittur, A., Pendleton, B. Lifting the Veil: Improving Accountability and Social Transparency in Wikipedia with WikiDashboard. To appear in Proc. CHI2008 Conference on Human Factors in Computing Systems. ACM Press, 2008.
11. Voss, J. Measuring Wikipedia. In *Proc. 10th Intl. Conf. Intl. Soc. for Scientometrics and Informetrics* (2005).