System Design for Social Translucence in Socially Mediating Technologies

David W. McDonald[1], Ivan Beschastnikh[2], Travis Kriplean[2], Alan Borning[2] and Mark Zachry[3]
dwmc@u.washington.edu, ivan, travis, borning@cs.washington.edu, zachry@u.washington.edu

[1]The Information School, University of Washington
[2]Computer Science & Engineering, University of Washington
[3]Technical Communication, University of Washington

1.0 Introduction

Socially mediating technologies (SMTs) like blogs, web forums, Slashdot and Wikipedia, have drastically changed the equation of mass social interaction online. As individuals participate and contribute content, many of their actions are visible to other participants. The work of making a contribution, supporting and maintaining the community are equally valuable to understanding what a community is about. Further, interpretation and understanding of those actions can influence the direction of public decision making (Borning et al. 2005), influence reputations (Resnick et al. 2000), notions of expertise (McDonald and Ackerman 1998) as well as other aspects of collaboration.

Wikipedia is an example SMT where large numbers of online actions can be observed. The Wikimedia Foundation provides regular releases of the complete contents of Wikipedia. These releases include all article, user, discussion, and administrative pages, each with the complete revision history. These regular releases have made it possible for us to consider future system architectures that can assist users making observations and interpretations of others activities.

Our approach takes social translucence (Erickson and Kellog 2000) as a conceptual framework to design a system infrastructure that supports a wide range of social activity inference. At a high level, a socially translucent system conveys socially salient information to participants in an online system, without inundating them with detailed low-level interaction data. Our system architecture allows community members to compose activities that they deem relevant into more general role type representations. Further, based on our approach, those compositions can be viewed and interpreted by other members of the community.

In the following we describe some research results that inform one specific dimension of social translucence related to work activity in Wikipedia. We consider how dimensions of Wikipedia work could be composed into viable representations of user reputation. We then describe our proposed system infrastructure and finally close with four issues that link our empirical social study of Wikipedia with our system development agenda.

2.0 Initial Interpretation and Relative Values of Online Activity

Wikipedians engage in a wide range of online activities that grow and maintain Wikipedia – both as an encyclopedia and as an online community. Understanding the range of work contributions and how Wikipedians come to understand that work is critical to the system we are developing.

In some recent research, we focused on what types of work Wikipedians acknowledge and what work they value (Kriplean et al. 2008).

Wikipedians have evolved a practice that acknowledges significant contributions of other Wikipedians. The practice involves giving little electronic certificates of achievement or recognition – much like the paper certificates one might earn at school or at work. They call these certificates *barnstars* (see Figure 1). Like any other part of Wikipedia, any editor can give a barnstar to any other editor for any reason – and they do. Yet many barnstars are awarded for specific work activity like editing an article, fixing broken links, or welcoming a newcomer.

We collected and analyzed barnstars, as one way to understand the dimensions of work that are valuable to the Wikipedia community. Every registered user on Wikipedia is given a user page in the user namespace, as well as a user talk page where others can leave messages. It is on these user and user talk pages that barnstars are typically given. We developed a parser to extract the text of these barnstars. Our extraction technique found over 14,500 unique barnstars.

Figure 1 – A sample Barnstar.

A random sample of 200 barnstars was analyzed using open coding to develop an initial codebook (Strauss and Corbin 1990). The codebook has seven broad dimensions of work. Each of these seven dimensions is comprised of from 4 to 10 different work type instances. For example, one of the dimensions of work commonly acknowledged by both Wikipedians and prior research is vandal fighting – which belongs to a dimension we call "Border Patrol." In this dimension we found eight different types of work activity, including vandal fighting on specific page types (article pages, user pages, administrative pages etc.), article deletion, spam detection and removal, and identifying copyright violations. We then applied the codebook to a second random sample of 200 barnstars and further refined the codebook.

A final random sample of 2400 barnstars was selected and evenly divided into six bins. Two coders independently coded each bin. One coder then reviewed codes for any discrepancies. The coders for each bin then met to resolve any discrepancies and come to agreement. Rather than force fitting a barnstar into a given category we allow multiple acknowledgements to be recognized in one bit of text. For example, in Figure 1, the giver of the barnstar acknowledges editing type work and a social "welcoming" type work. A set of consistency checks were then performed across correlated codes, outliers, small categories, keywords that might trigger codes as well as any suspected systematic inconsistency by coders.

Our results show that encyclopedic content work ("editing work") represents a minority of the actual work acknowledged by Wikipedians (see Table 1). In particular Wikipedians acknowledge a significant amount of social and community support work and a significant amount of work around detecting and repairing vandalism. If "work" in Wikipedia is an important component of reputation, then a reputation system for Wikipedia needs to have a richer notion of both the work that is valued and to what degree that work contributes to a Wikipedian reputation.

Our study of barnstars begins to illustrate in a concrete way that Wikipedians are able to observe others' activities and make judgments about what roles are being played in the community.

Table1 – Distribution of work acknowledged by Wikipedians through barnstars

Dimensions of Work Acknowledged	Frequency	Percentage
Editing Work	852	27.8%
Social and Community Support Actions	763	24.9%
Border Patrol	342	11.2%
Administrative	284	9.3%
Collaborative Actions and Disposition	244	8.0%
Meta-Content Work	128	4.2%
Undifferentiated Work	447	14.6%

These are not just discrete actions, but actions that illustrate a trace or trajectory of connectedness that imply more complex work and social behaviors.

An example might make this more concrete. Vandalism is a common problem in open online communities. There are (almost always) some individuals who would take advantage of the community in some way be it crass commercial advertisements or explicit delivery of misinformation or attempts to undermine social connectedness through trolling and baiting. In Wikipedia we characterize a wide range of vandal detection and repair as Border Patrol. A Wikipedian who is good at Border Patrol might have a number of observable features such as comments in an edit history indicating reverts of vandalism, reverts that persist (i.e., not reverted in return), and endorsements from other users who recognize other Border Patrol activity.

That people currently recognize and acknowledge these activities suggests that it might be possible for systems to compile the right kinds of activity data that would allow observers to more easily see the broader patterns of contributions by a wide range of participants. We know that this is possible; while not the specific role of the research, some visualization research has illustrated that these patterns can be made more observable. Motivated to understand how Wikipedians allocate their time, Wattenberg et al. (2007) developed a simple visualization technique that identified systemic editing patterns, which they applied to administrator edit summaries. Considering the visualizations of administrators, they noticed that administrators' work is often heterogeneous, but that they frequently have a current "focal task."

Our current research trajectory is to develop a system infrastructure that facilitates better understanding of the activities in an online community. We take the stance that the infrastructures that would do this should be motivated by social translucence. In the next section we present our current system architecture and follow that with some open questions.

3.0 A Socially Translucent Architecture

Our architecture is focused on supporting an open and reflective composition of user reputation in the community. We use reputation as a concrete example of understanding the roles that members play and contributions that members make to online communities. We plan to build a reputation system that allows members to compose their own notions of which activities are most salient. Our hypothesis is that systems that enable social introspection and understanding are more likely to be incorporated into practice. We describe our architecture and then outline some open issues that will inform our evolving system architecture.

3.1 A Reputation Framework with Social Translucence

An open composable reputation system requires components that collect information, language that supports composition, and interfaces that visualize the reflected reputation. Figure 2 gives our framework for a composable reputation system based on Wikipedia.

Our framework has five levels. The lower levels facilitate finding, mining, and exposing salient social activities. HCI researchers have contributed a number of tools to Wikipedia, such as SuggestBot (Cosley et al. 2007) and article highlighting (Adler & de Alfaro 2007) that commonly focus on editing activity. However, our barnstar study illustrates that there are many more activities on Wikipedia than just article editing. The upper levels of our framework will allow community members to compose and integrate representations of salient activities into useful tools.

At the base level, **datasets** provide an information source. This includes the full edit history on Wikipedia, as well as page view data (Priedhorsky et al. 2007). Raw data is often very difficult for users to understand or interpret. A **feature extractor** parses the data to extract potentially valuable tokens. One or more of these tokens comprise **primitive features** such as reverts, administrative actions like page locking, co-editing activities, barnstars, text persistence (e.g., Adler et al. 2007), and policy citations (Beschastnikh et al. 2008). Such features can be treated as indices of higher-level, socially salient activity.

A **semantic interpreter** transforms primitive features into socially salient **compound features**, such as roles based on a model of work, social network structures, and maps of conflicts amongst users and pages (e.g., Kittur et al. 2007). The semantic interpreter links the observable, extracted activities into social constructs. Such linking requires an understanding of how Wikipedians understand the observable actions of other Wikipedians.

All of the architectural components described so far help to find and represent social activity. The rest of the architecture is oriented toward exposing and supporting user generation of tools that enable greater social translucence.

The **reflexive composer** exposes representations of the compound and primitive features through a *reflexive composition language*. Community members and developers use the language to express (compose) their own meaningful combinations of relevant features. The language will need to be more complex than a markup language but somewhat less

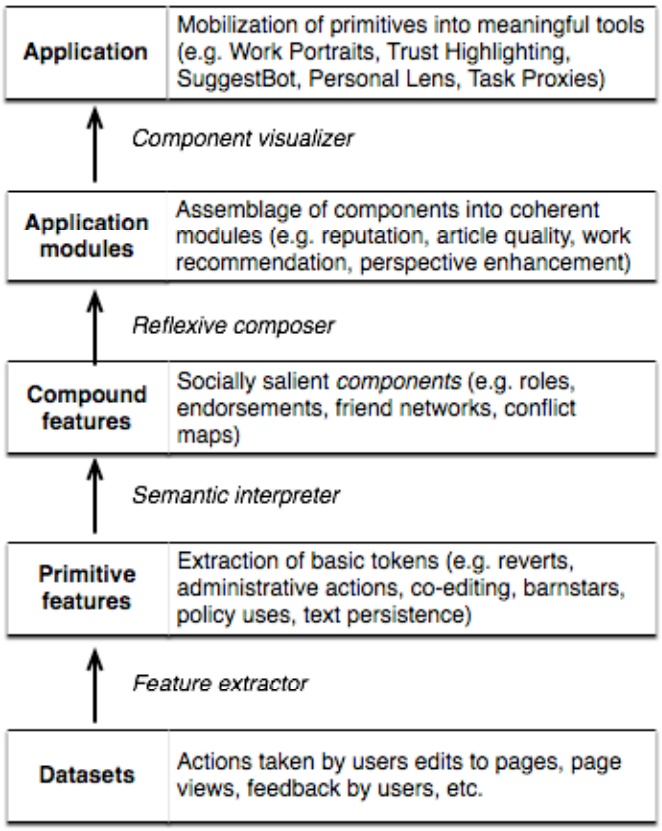

Figure 2 – Framework for supporting social translucence through reflections of activity.

complex than a traditional scripting language. The reflexive composer allows the specification of **application modules.** The reflexive composer is currently focused around expressions of reputation, but this could generalize to understanding other factors relating to involvements in the community.

Finally, the **component visualizer** links one or more application modules to a meaningful user presentation. The visualizer requires (1) a persistent representational model of the **application modules** as defined by the reflexive composition language and (2) a means to integrate these representations into working tools. Our efforts here are to develop a basic declarative, xml-based representation of user-composed reputations and useful visualizations of those representations (e.g., Social Proxies (Erickson et al. 2004)). The ultimate goal, however, is to build, and enable the community to build, applications for social translucence.

4.0 Issues for Socially Translucent Systems

Socially mediating technologies – as a function of their mediating influence – often make forms of interaction explicit. Both the participants and the systems that mediate the interactions can observe the interactions. As we have suggested here, this visibility of interactions can be extended to the design and construction of socially translucent systems. Social translucence assists members in an online community make sense of others activities while avoiding dumping low-level details. But there are still a number of critical issues that need to be addressed.

4.1 Interpretation of Mediated Activity

Mediated activity is sometimes very discrete and only through repeated observations overtime can a set of individual activities be understood as some form of behavior. We do not readily understand they way people perceive and interpret sets of activity as complex behaviors and how those in turn become characterized as more fully elaborated roles in a community. We fundamentally believe that the link between human interpretation and others activities can be understood through observation, engagement and the analysis of large community datasets such as Wikipedia.

4.2 Usable Reflexive Languages

Modern text parsing and data mining techniques allow effective means for identifying and selecting individual or groups of activities. Building upon these mining techniques – combined with heuristic models, machine learning, and user attributions – can allow a system to infer and interpret in some thin way what a set of actions by an individual might mean. Those interpretations are only useful when they can be combined and put to use. Our approach is to develop a type of scripting language. While we recognize that our language needs traditional flow control and other mechanisms, we also believe that there are other operators, like social conditionals (i.e., include a behavioral component only if the current user/viewer has expressed a similar behavioral component), which are likely to be necessary. Creating understandable and usable language constructs to express relationships among behaviors is a difficult problem. A social notation or social calculus and how people actually think about complex relationships might lead to usable language constructs.

4.3 Effective Visualizations

There are few examples of easily interpretable visualizations of social behavior. Some early work on Social Activity Indicators (Ackerman and Starr 1995) illustrated the effectiveness of social event notification that had some implications for maintaining awareness. Subsequent work on Social Proxies (Erickson et al. 2004) went further to use real-time animation to illustrate awareness of social process. However, each of these is far from representing a set of complex roles that individuals play in a community. Individuals fill more than one role at a time and visualizations of activity in a socially mediating technology should be capable of illustrating what actions constitute a given role and how different roles are filled by community members.

4.4 Evolving Socially Translucent Software Systems

While we have a software architecture design that facilitates social translucence, the current architecture facilitates only a certain range of flexibility. Socially mediating technologies cannot remain static as social practices change around them or they risk being abandoned, as they would no longer serve the needs of their users. Another open question is how to evolve socially translucent systems to account for new social sensibilities, new actions, and new interpretations of those actions.

5.0 Conclusion

Socially mediating technologies (SMTs) provide the opportunity to study mass social interaction. Our research is examining how to use social translucence to design infrastructure for SMTs that can help users make sense of the complex social activities and roles that others play in SMTs. We have a framework motivated by social translucence and have identified four critical issues that need further exploration and elaboration before our infrastructure can be realized.

6.0 Acknowledgements

We thank Overland Storage for providing a Snap Server with 8TB of storage to facilitate our data analysis. We acknowledge the support of Tom Anderson and Arvind Krishnamurthy. This work is partially funded by NSF grants IIS-0534094, IIS-0705898, and IIS-0811210.

7.0 References

1. Ackerman, M.S. and B. Starr. (1995). Social Activity Indicators: Interface Components for CSCW Systems. In *Proceedings of the ACM Symposium on User Interface Software and Technology* (UIST'95).
2. Adler, B.T., J. Benterou, K. Chatterjee, L. de Alfaro, I. Pye and V. Raman (2007). Assigning Trust To Wikipedia Content. *Technical Report UCSC-CRL-07-09*, School of Engineering, University of California, Santa Cruz, CA, USA.
3. Adler, B. T. and L. de Alfaro (2007). A Content-Driven Reputation System for the Wikipedia. *Proceedings of the 16th International Conference on World Wide Web (WWW'07)*.
4. Beschastnikh, I., T. Kriplean, and D.W. McDonald. Wikipedian Self-governance in Action: Motivating the Policy Lens. In *Proceedings of the International Conference on Weblogs and Social Media (ICWSM '08).*

5. Borning, A. B. Friedman, J. Davis, and P. Lin. (2005). Informing Public Deliberation: Value Sensitive Design of Indicators for a Large-scale Urban Simulation. In *Proceedings of the Ninth European Conference on Computer Supported Cooperative Work (ECSCW '05):* 449-468.
6. Cosley, D., D. Frankowski, L. Terveen, J. Riedl (2007). SuggestBot: using intelligent task routing to help people find work in Wikipedia. In *Proceedings of the 12th international Conference on intelligent User interface (IUI '07)*: 32-41.
7. Erickson, T. and W. Kellogg (2000). Social Translucence: An Approach to Designing Systems that Support Social Processes. *ACM Transactions on Computer-Human Interaction*. 7 (1). 59-83.
8. Erickson, T., W. Huang, C. Danis, and W.A. Kellogg (2004). A social proxy for distributed tasks: design and evaluation of a working prototype. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '04)*: 559-566.
9. Kittur, A., B. Suh, B.A. Pendleton, and E.H. Chi (2007). He says, she says: conflict and coordination in Wikipedia. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '07)*: 453-462.
10. Kriplean, T., I. Beschastnikh, D. W. McDonald and S. Golder (2007). Community, Consensus, Coercion, Control: CS*W or How Policy Mediates Mass Participation. *Proceedings of the 2007 International ACM Conference on Supporting Group Work*: 167-176.
11. Kriplean, T., I. Beschastnikh, and D. W. McDonald (2008). Articulations of Wikiwork: Uncovering Valued Work in Wikipedia through Barnstars. *Proceedings of the 2008 ACM Conference on Computer Supported Cooperative Work* (CSCW'08).
12. McDonald, D. W. and M. S. Ackerman (1998). Just Talk to Me: A Field Study of Expertise Location. *Proceedings of the 1998 ACM Conference on Computer-Supported Cooperative Work*: 315-324.
13. Priedhorsky, R., J. Chen, S. K. Lam, K. Panciera, L. Terveen and J. Riedl (2007). Creating, Destroying, and Restoring Value in Wikipedia. *Proceedings of the 2007 International ACM Conference on Supporting Group Work*: 259-268.
14. Resnick, P., R. Zeckhauser, E. Friedman, and K. Kuwabara (2000). Reputation Systems. *Communications of the ACM* 43 (12). 45-48.
15. Strauss, A. and Corbin, J. (1990). *Basics of Qualitative Research: Grounded Theory Procedures and Techniques*. Sage Publications.
16. Wattenberg, M., F. B. Viegas and K. Hollenbach (2007). Visualizing Activity on Wikipedia with Chromograms. *Human-Computer Interaction – INTERACT 2007*. Springer LNCS 4663: 272-287.