Collective Intelligence in the Social Web

John Riedl

GroupLens Research
University of Minnesota
Minneapolis, MN 55455 USA
`riedl@cs.umn.edu`

Many online communities are emerging that, like Wikipedia, bring people together to build valuable long-lasting artifacts in *peer production communities*. These communities exhibit several forms of collective intelligence, including effective decentralized cooperation within large groups of humans, partnerships between human and computer agents, and interaction effects between the two.

On the one hand, effective cooperation among large groups of humans should not surprise us, since such cooperation has been occurring since at least the time of the pyramids. On the other hand, it is unusual to see cooperation within such large groups without at least some members of the group holding whips! More seriously, the emergent behavior of millions of people working together with only the power structures they can evolve themselves is remarkable. We will seek to understand the causes of the successes and failures of these communities by looking at examples of both.

The individual cognitive processes of the human part of the equation are independently interesting as well. What is the nature of people's participation in peer production systems? What are their motives? In what ways is their behavior destructive instead of constructive? How can computer agents interact with the humans to encourage the positive contributions, and damp down the negative? We will explore these questions in the context of a wide variety of the most successful Social Web systems.

Partnerships between human and computer agents have been talked about for decades, but are beginning to see full-scale deployment on the Internet. Agents scour the Web seeking information of interest to humans, who react to that information in computer-readable ways, leading in some cases to a positive feedback cycle of increasingly frenzied responses by the agents. The most powerful of these feedback cycles – PageRank – has had powerful and unexpected affects on the structure of the Web itself. What can we predict from past experience about the future evolution of these human/computer partnerships?