Bias in Distributed Cognition as a Bridge between the Individual and the Network

Overview This is a high level sketch of the potential benefits of research into biases in distributed cognition. The nutshell summary is that distributed cognition may show biases similar to well documented and understood biases in individual cognition and that understanding these biases may help us bridge the gap between the actions of individuals in and the network properties of technology-mediated social environments. A second goal for this position paper is to consider one possible area of research that explicitly crosses traditional study domain boundaries (psychology and network science).

Biases in individual cognition Cognitive psychology has documented a number of properties of individual cognition that we might draw on as we develop our understanding of distributed cognition. Here I'll focus on the various forms of cognitive bias: biases in judgment and decision making that impact the way we process information in systematic and measureable ways. Attributional biases such as the fundamental attribution error (overestimating personality-based explanations for others' behavior), for example, describe the ways we skew information when assessing others. Many other cognitive biases directly address information processing. As one example, the availability heuristic reveals how people are more likely to incorporate new information based on information that is most readily accessed in memory. While these biases cover a broad spectrum, many of them assume a 'spreading activation' model of information in the mind. For example, the anchoring effect in which initial "anchor" points are overly weighted when making decisions, demonstrates how the initial cognitive activation caused by the anchor carries over to influence the final decision.

What do we mean by distributed cognition here? For the purposes of this position paper, distributed cognition is the "mind" of any crowd system, where a crowd system is any socio-technical system with multiple, dependent inputs on judgment and decision making or content creation (e.g., Digg, Q&A systems, ratings and reviews, aspects of Wikpedia). While "mind" is in quotes due to its fuzzy nature, the use here is intended to highlight the implication that distributed cognition may be an entity with properties that are systematic and measurable. I propose that this entity of distributed cognition occurs in socio-technical systems when individual cognition 1) happens under the influence of the actions of others in the system and 2) feeds back into the system. In other words, the requirement of independence of observations for wisdom of the crowds is violated in most cases, but the effects of this violation may yield systematic output patterns that allow us to connect individual behaviors to network phenomena.

Bias in distributed cognition? As an example of bias in distributed cognition, let's look at confirmation bias, the tendency to interpret information in a light that confirms what you already thought. There is plenty of evidence of this happening at the level of the individual. For example classic studies on confirmation bias [4] show how people will rate research on the death penalty as better or worse depending on whether the research is in line with their stance on capital punishment. What would confirmation bias look like in distributed cognition? A simple version likely happens regularly in sites like Digg where people are exposed to cumulative ratings of content before they can actually see the content themselves. That is, any given user may be biased to like articles with lots of diggs and not like articles with few diggs. Users see what they expect to see -a high or low quality article- and then strengthen this bias in distributed cognition when they themselves digg the article. Of course articles with more diggs are also likely to actually be better than those with few diggs, but the extent to which there is skew in the ratings beyond independent and objective ratings may be due to bias in distributed cognition.

If we assume carry over effects from one user to the next in these systems (spreading activation), then a whole host of cognitive biases in distributed cognition may come into play. For instance, suppose a person is reviewing a restaurant on Yelp. There happens to be an ad for Hawaiian vacations on the web page, which puts the person in a slightly positive state of mind and in turn influences their review in a positive direction. This is similar to research on biases due to use of the affective heuristic in which people's judgments are influenced simply by being primed with smiling and frowning faces [6]. To the extent there is carry over in distributed cognition, later users will be influenced by this review ("I have to agree with Diner72 that the dim lighting was romantic rather than dingy.")

As a third example, collaboratively constructed documents (e.g., Wikipedia articles) may be influenced by an anchoring bias (among other things): authors may be subtly biased by what they start working with and may make changes that are more in line with existing work than what they would have done had they started from scratch or would be different had they begun from a different starting point. Thus there is a bias that in part governs the rate of change in collaboratively constructed documents. This can be good. In fact, Wilkenson and Huberman [5] show that edits in Wikipedia on average correlate with article quality and Kittur and Kraut [2] detail how in cases of implicit coordination Wikipedia article quality will be higher if a small group of core contributors lays the groundwork for future authors to play supporting roles. To the extent that drastic change is worse than too little change, anchoring bias toward existing content in collaboratively constructed documents should be useful.

Bias in collaboration environments may not be limited to the active creation of content. For example, the initial WikiDashboard work [3] showed that articles were less trusted when they were shown to have more recent edits. This is might be thought of as a "recency effect" in distributed cognition: more recent information is weighted more heavily, often for good reason, but perhaps too heavily in some cases. As that research shows, there may be no good reason to weight the recent past more heavily yet it happens. This might play out on product review sites where more recent reviews, which are in fact more relevant, are weighted more heavily than they should be.

Bridging the individual and the network. Continuing the Digg example, each individual's liking of an article is presumed to be influenced by previous ratings. Thus, there is a conditional probability of any given individual's rating given existing ratings that is distinct from their probability of their rating given reading and rating the article in isolation. The difference between the two reflects bias in distributed cognition that may generate certain network dynamics. In this case, this type of bias may lead to an information cascade if it continues to influence users. The Yelp review example is similar: carry over from one user to the next (individual level effect) generates a bias in the collective cognition around the review and ultimately may lead to an information cascade (network dynamic).

Carry over between users in the first Wikipedia example suggests a different type of network model. When an editor comes to an article it is in a local minimum state that requires some amount of perturbation to be broken out of. In the best cases (the right type of coordination in Kittur and Kraut), the initial authors have moved the article to a local minimum that is close to the global minimum (perfect article) after which future authors continue to nudge the article toward that global minimum. In the worst cases, bias in anchoring around existing content increases the effort required to break the network out of an undesirable local minimum.

High level implications for research This was a proposal to consider the notion of bias in distributed cognition as an area that might allow us to bridge what we know about the psychology of individual users (cognitive bias) with what we know how network level dynamics and phenomena (e.g., information cascades). At a high level, the implications for research then are first to identify which biases exist in distributed cognition (if any) and to what extent. This could be done experimentally and should focus on cross-domain and context generalizability (i.e., are different biases more or less present in collaboration systems than in rating systems?). Second we would determine how these biases arise given system design decisions. This is an important step that is similar to efforts to leverage social psychology in system design [1]. Users in socio-technical systems are influenced by past behaviors of other users, so design efforts should focus on what information about those past behaviors should be surfaced (and when in the interaction process, how much of the information, etc.). Finally, we would explore how incorporating parameters for bias in distributed cognition can improve network-level statistical analyses and models. For example, if we can measure experimentally the degree of bias in a voting system like Digg, can we use that metric to inform statistical models that might for example be overweighting an item simply because it has a lot of diggs? This last step can also run the other direction. That is, statistical models of behaviors in of socio-technical systems might uncover biases on the part of individual users.

A brief note about implications for a research agenda and graduate training In part this was an exercise in sketching a research area that requires knowledge of socio-technical systems that runs the gamut from the individual to the network. The proposed study of bias in distributed cognition was an example of what happens in between these two. Explicitly this requires students to think about human interaction with socio-technical systems from a very broad and diverse perspective. The world appears headed toward increased specialization and we may be better off designing our research agenda to support collaboration between experts in the different areas. This however, may make it more difficult to synthesize across the different levels of analysis of socio-technical systems. We should consider a more holistic approach in which students are at least conversant in each of the various analysis levels one can bring to the study of socio-technical systems. This type of synthesis is important in industry research scenarios where you have very large socio-technical systems that require both large scale network analyses and “rubber meets the road” design recommendations that are focused ultimately on individual users.

References

1. Beenen, G., Ling, K., Wang, X., Chang, K., Frankowski, D., Resnick, P., Kraut, R. (2004). *Using Social Psychology to Motivate Contributions to Online Communities*. In Proc. CSCW 2004.
2. Kittur, A., & Kraut, R. (2008). *Harnessing the wisdom of the crowds in wikipedia: quality through coordination*. In Proc. CSCW 2008.
3. Kittur, A., Suh, B., & Chi, E. (2008). *Can you ever trust a wiki?: impacting perceived trustworthiness in Wikipedia*. In Proc. CSCW 2008.
4. Lord, C., Ross., L., & Lepper, M. (1979). Biased assimilation and attitude polarization: the effect of prior theories on subsequently considered evidence. *Journal of Personality and Social Psychology, 37(11)*, 2098-2109.
5. Wilkenson, D., & Huberman, B. (2007). *Cooperation and quality in Wikipedia*. In Proc. WikiSym 2007.
6. Winkielman, P., Zajonc, R., & Schwarz, N. (1997. Subliminal affective priming attributional interventions. *Cognition and Emotion, 11(4)*, 433-465.