

Establishing the Micro-to-Macro Link in Cognitive Engineering: Multilevel Models of Socio-Computer Interaction

Wai-Tat Fu and Peter Pirolli

Abstract

We argue that the time is ripe to develop multilevel models that provide explanations at both the individual and social levels to extend research in human-computer interaction to the emerging area called socio-computer interaction. We provide three examples of such multilevel models. We first describe the SNIF-ACT model, which was developed to explain individuals' information foraging behavior as they navigate through web pages to find target information, and to explain the general tendency for people to stay on any particular website. We then describe the social information Foraging theory, which characterizes how search efficiencies may be improved by utilizing social signals generated by other information foragers. Finally, we describe the semantic imitation model, which characterizes how individual users may generate social tags to index their own documents and could lead to aggregate indices that act as semantic structures that help other users to find relevant information. The models demonstrate the significance in capturing the social dynamics involved in systems that afford socio-computer interaction.

Key Words: socio-computer interaction, cognitive models, multilevel models, information foraging, social dynamics

Introduction

The modal view of the web as a place for individuals to search and browse massive amounts of information through a single dominant application—the web browser—has evolved into a more participatory and efflorescent social-technological system. It has become a platform for exploring information, learning, sharing information with others, and collaborating in the production of new content. This evolution of social information foraging, learning, sharing, and co-creation is a great opportunity for cognitive engineering because there is so much new territory to explore and explain, and so many new ideas about how to do so. It has also expanded the traditional research of human-computer interaction to include a broader range of intelligent activities that span multiple brains that work together, either

formally or informally, across a prolonged period of time through the mediation of socio-technological systems. To reflect the importance in capturing both the cognitive and social aspects in such a system, we call this emerging field of research “socio-computer interaction.”

A closer look at recent research on multi-user socio-technological systems has, however, revealed an unsettling feature that makes the study of socio-computer interaction challenging: It is fragmented into multiple disciplines, each focusing on a specific aspect and level of analysis, with a lack of integration and communication among them. Different levels of analysis are often found in disciplines that have non-overlapping publication outlets, and results from these disciplines are often presented in separate conferences.

One reason for the fragmentation can perhaps be attributed to the assumed efficiency achieved by division of labor among researchers in different disciplines, but another major reason can be attributed to the presupposition that analysis at a particular level can be isolated to events in that level (e.g., see also arguments by Anderson, 2002; Coleman, 1994; Hedstrom, 2006; Sun, 2006). This is reflected by the frequent practice, mostly in the fields of applied physics and information science, to develop models of social behavior that capture the high-level descriptive characteristics (e.g., power-law structures) but provide little information on lower-level behavior, such as how a person actually shares and distributes information, or how a person processes information represented by different interface representations. Similarly, research in sociology is typically concerned with aggregate outcomes, such as the strength of different social norms or properties, which are not definable for any single individual. These theories are therefore typically not concerned with how actions and interactions of individuals may lead to such social norms or properties. The lack of connection between emergent social behavioral patterns and activities at the individual level is particularly troubling for cognitive engineering researchers, as these models simply do not provide enough information to make sense of these social behavioral patterns, and therefore cannot provide any useful guidelines to inform engineering decisions on systems that are used by individual users. Conversely, in the area of cognitive science, computational models are developed to explain the fine-grained temporal structures of cognitive operations in individual behavior (e.g., how students learn to solve a physics problem), but they seldom consider how these operations could be shaped by social constraints, and therefore lack the capability to predict behavioral outcomes that require observations over an extended period of time in more natural social settings (e.g., educational outcomes).

In this chapter, we argue that the time is ripe to develop cognitive engineering models that span multiple levels of analysis, and these models are essential for designs of future socio-technological systems. We will first review recent “surprises” that occur when pushing macro-level social or network theory down to individual behavior, which will highlight the importance of an integrated theory that provides coherent explanations that span multiple levels of events and activities. We argue that multilevel models are important because many models of social behavior have ignored or made

over-simplistic assumptions about the actions and interactions of individuals. We will then describe examples of recent attempts to develop these models, and their implications for designing future socio-technological systems that afford more effective socio-computer interaction (Carayon & Hoonakker, this handbook).

The Importance of Establishing the Micro-to-Macro Link

The idea of establishing the micro-to-macro link is not new. One well-known example is the segregation model by Schelling (1971, 1978). In his model, Schelling showed that a small preference for one’s neighbors could lead to total segregation and provided an explanation of an important social observation (segregation) based on individual behavioral tendencies. Coleman (1987) provides an elaborated framework to study the micro-to-macro link. He argues that one reason why a theory that bridges micro- and macro-behavior is important is that, while the goal of studying a social system is to understand how it functions, observations are often done at the individual level. Being able to explain the function of a social system based on observation of individual behavior will provide a much stronger connection between data and theory. He uses the example of how panic occurs when a fire alarm is sounded in a crowded place—a macro-level relation. One explanation of this relation is to assume that the fire alarm created in each person a sense of fear of being trapped, which leads to running toward an exit, which, together with others, leads to blocking of the exits. A better explanation, Coleman argues, is to assume that each individual has a choice of running or walking, and that this decision depends on whether the person sees that other people are running or walking and can be explained using theories that explain individual choices in the prisoner’s dilemma. This explanation provides a richer set of predictions, as it does not predict that panic always occurs. In addition, it predicts that certain persons, particularly those at the center of attention (e.g., performers on the stage), may be less likely to run than others under certain assumptions of rational behavior at the individual level. The example not only shows how a theory that explains individual decisions can generate predictions at the social level (e.g., predicting collective behavior patterns of a socio-technological system based on theory and observation of individual users), but it also provides a richer set of predictions of macro-outcomes. In addition, given that the

cost of conducting experiments to study individual behavior is often much lower than large-scale social experiments, developing and testing individual theories and applying them to explain social-level outcomes will also be more cost-effective.

Recent research has begun to show that developing models at *only* the social or aggregate level is not sufficient for understanding how individuals interact with a social information system. Many have shown that the intuitive assumptions about how individuals interact with each other in many of these models are different from what were found by empirical studies (Hedstrom, 2006; Iribarren & Moro, 2009; Ku, Galinsky, & Murnighan, 2006; Leskovec, Backstrom, Kumar, & Tomkins, 2008; Liben-Nowell & Kleinberg, 2008; Salganik, Dodds, & Watts, 2006). While analysis within a single level often leads to descriptive measures that establish the relationship among some variables, analysis *across levels* requires specification of a *dynamic process* that provides the micro-to-macro link, which is critical for a deep understanding of a social phenomenon.

An example that demonstrates the importance of establishing this micro-to-macro link is the study by Salganik et al. (2006). Salganik et al. found that, as expected, preferences of individuals were influenced if they could see the preferences of others. On the one hand, their results were consistent with the common assumption that the more information people have regarding the preferences of others, the greater agreement they will display in their preferences. On the other hand, their study surprisingly showed that as social influence increased, the collective outcomes became more *unpredictable*, and the increased unpredictability was found to be caused by the dynamic process of social influence through the interactions among individuals, not simply by aggregating pre-existing individual preferences as previously assumed. The results highlighted the importance of the influence of individual processes on the final collective outcomes, and the inseparability of the two levels in the phenomena observed.

Another example is the research on information flow in social networks. Many models of information flow assume that information spreads in a similar way as an epidemic propagates on scale-free networks (Barabási & Albert, 1999). There are, however, differences between information flows in social networks and the spread of viruses. Individuals, unlike viruses, tend to organize both formally and informally into groups. Thus, information flow will be more selective than viruses, which tend to infect individuals indiscriminately and expand widely in a

short number of steps as in “small-world” networks. In the study by Wu et al. (Wu, Huberman, Adamic, & Tyler, 2004), they show that information flow indeed does not necessarily spread as quickly in a community as predicted by the scale-free network model. Instead, when similarities of individuals in the network are assumed to influence the spread of information, information flow tends to stop much quicker, suggesting that the epidemic propagation assumption is too simplistic to be useful for predicting the rate of information flow. Similarly, by tracing the flow of Internet chain letters, Liben-Nowell and Kleinberg (2008) show that information does not flow like viruses in the social network. Rather, information flows in a narrow tree-like pattern that continues to reach people several hundred levels deep. The finding reflects the highly clustered nature of social networks, in which information can take a widely varying amount of time to traverse from one cluster to another. These results demonstrate that although the small-world model is a simple and parsimonious characterization of most social networks, it is a model based on wrong assumptions at the individual level.

In sum, these studies not only show that there are inherent dynamics between actions at the individual level and aggregate patterns at the social level, they also demonstrate that it is important to establish this micro-to-macro link to seriously understand the dynamic process that connects actions and interactions among individuals with aggregate social behavioral patterns. Next, we will describe some of our initial efforts to characterize the micro-to-macro link.

Multilevel Models of Human-Information Interaction

Given the emerging perspectives on the dynamics between multiple levels of activities, there seems to be a need to develop multilevel models that are capable of providing an integrated explanation of behavior at both the individual and aggregate levels. Examples of the “surprises” described above also suggest that we need a better grounding for the aggregate processes. We will show that computational cognitive models of individual behavior can provide not only precise quantitative predictions on user-environment interactions but also a basis for incorporating realistic constraints, capabilities, and tendencies of individual users in terms of testable cognitive representations and processes (and perhaps also in terms of their physical and neurological underpinnings). Cognitive models, on the other

hand, are also in need of rigorous testing for their ability to explain phenomena at the aggregate and social level. Indeed, for cognitive models to be useful, they should also take into account how cognitive representations and processes are shaped by social and cultural constraints imposed by the environment with which the person is interacting.

SNIF-ACT: A Multilevel Cognitive Model of Information Foraging

Information foraging theory (Pirolli, 2005, 2007, 2009; Pirolli & Card, 1999) is an example of a multilevel theory. Information foraging theory assumes that people adapt their information-seeking behavior to maximize their rate of gaining useful information to meet their ongoing goals. It assumes that adaptive information systems evolve toward states that maximize gains of valuable information per unit cost. The theory draws upon computational models of cognition (Anderson et al., 2004), rational choice under uncertainty (McFadden, 1974), optimal foraging theory (Stephens & Krebs, 1986), and stochastic process theory (Lukose & Huberman, 1998). The crucial element in the information foraging theory is the measure of information scent, which is defined based on a Bayesian estimate of the relevance of a distal source of information conditional on the proximal cues.

A cognitive simulation model called SNIF-ACT (Fu & Pirolli, 2007; Pirolli & Fu, 2003), developed in an extension of ACT-R (Anderson et al., 2004), was fit to detailed moment-by-moment web surfing behavior of individuals studying in a controlled laboratory setting. The basic idea of this part of the SNIF-ACT model was identical to that of a cognitive model called the Bayesian satisficing model (Fu, 2007; Fu & Gray, 2006), which was developed to explain individual learning and choice behavior in repeated sequential decision situations. The Bayesian satisficing model is composed of a Bayesian learning mechanism and a local decision rule. When applied to web surfing, the model assumes that when users evaluate links on a web page, they will incrementally update their perceived relevance of the web page to the target information according to a Bayesian learning process. A local decision rule then decides when to stop evaluating the link: The evaluation of the next link continues until the perceived relevance is lower than the cost of evaluating for the next link. At that point, the best link encountered so far will be selected.

To illustrate the behavior of the model, we will focus on a case where the model is facing a single web

page with multiple links. There are three possible actions, each represented by a separate production: Attend-to-Link, Click-Link, and Backup-a-Page. Similar to the BSM model in the map-navigation task, these productions compete against each other according to the softmax equation; in other words, at any point in time, the model will attend to the next link on the page, click on a link on a page, or decide to leave the current page and return to the previous page. The utilities of the three productions are derived from the link likelihood equation, and they can be calculated as:

$$\text{Attend-to-Link: } U(n+1) = \frac{U(n) + IS(Link)}{1 + N(n)}$$

$$\text{Click-Link: } U(n+1) = \frac{U(n) + IS(BestLink)}{1 + k + N(n)}$$

$$\text{Backup-a-Page: } U(n + 1) = MIS(\text{Previous Pages}) - MIS(\text{Links 1 to } n) - GoBackCost \quad (1)$$

In (1), $U(n)$ represents the utility of the production at cycle n , $IS(Link)$ represents the information scent of the currently attended link, $N(n)$ represents the number of links already attended on the web page after cycle n (one link is attended per cycle), $IS(BestLink)$ is the link with the highest information scent on the web page, k is a scaling parameter, $MIS(Page)$ is the mean information scent of the links on the web page, and $GoBackCost$ is the cost of going back to the previous page. The values of k and $GoBackCost$ are estimated to fit the data.

Figure 34.1 shows an example of how the probabilities of selecting the three productions change (Figure 34.1b) as the model sequentially processes links on a page (Figure 34.1a). We can see that initially, the probability of choosing Attend-to-Link is high. This is based on the assumption that when a web page is first processed, there is a bias in learning the utility of links on the page before a decision is made. However, as more links are evaluated, the utilities of the productions decreases (as the denominator gets larger as $N(n)$ increases). Because the utility of Attend-to-Link decreases faster than that of Click-Link (since $IS(Best) = 10$, but $IS(Link)$ decreases from 10 to 2), the probability of choosing Attend-to-Link decreases, but that of Click-Link increases. The implicit assumption of the model is that since evaluation of links takes time, the more links that are evaluated, the more likely that the best link evaluated so far will be selected (otherwise the time cost may outweigh the benefits of finding a

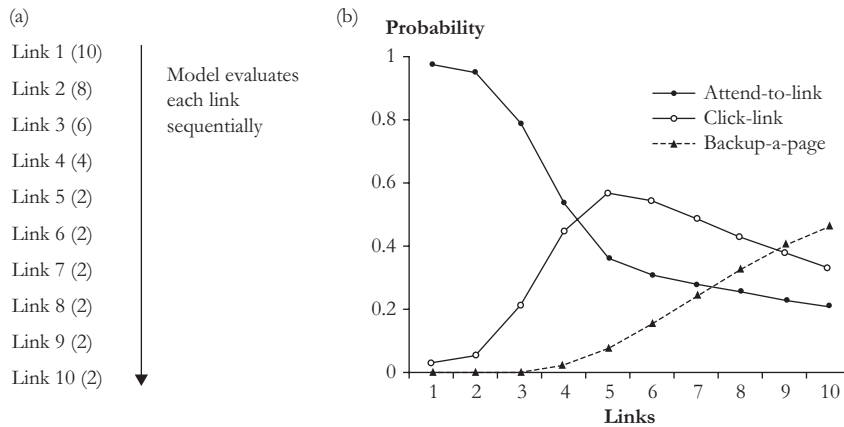


Figure 34.1 (a) A hypothetical web page in which the information scent values (number in parentheses) of links decreases linearly from 10 to 2. (b) The probabilities of choosing each of the competing productions when the model processes the page. The mean information scent of the previous page was assumed to be 10.

better link). As shown in the figure, after four links on the hypothetical web page have been evaluated, the probability of choosing Click-Link is larger than that of Attend-to-Link. At this point, if Click-Link is selected, the model will choose the best (in this case the first) link and will continue to process the next page. However, as the selection process is stochastic (because of the softmax equation), Attend-to-Link may still be selected. If this is the case, as more links are evaluated (i.e., as $N(n)$ increases), the probability of choosing Attend-to-Link and Click-Link decreases. On the other hand, the probability of choosing Backup-a-Page is low initially because of the high *GoBackCost*. The utility for Backup-a-Page is calculated based on a moving average of the information scent encountered in previous pages. However, as the mean information scent of the links evaluated (i.e., $MIS(\text{Links } 1 \text{ to } n)$) on the page decreases, the probability of choosing Backup-a-Page increases. This happens because the mean information scent of the current page is “perceived” to be dropping relative to the mean information scent of the previous page. In fact, after eight links are evaluated, the probability of choosing Backup-a-Page becomes higher than that of Attend-to-Link and Click-Link, and the probability of choosing Backup-a-Page keeps on increasing as more links are evaluated (as the mean information scent of the current page decreases). We can see how the competition between the productions can serve as a local decision rule that decides when to stop exploration.

Figure 34.2 shows the results for matching the link selection data from a group of 74 users searching using the Yahoo! website (Fu & Pirolli, 2007) across a range of information search tasks

(e.g., “finding the 2002 holiday schedule”). All links selected by both the model and the human subjects were extracted, and the total frequencies of visits for each of these links were plotted in Figure 34.2. We see that the model provided good fits to the data ($R^2 = 0.91$), suggesting that the dynamic selection mechanism in the Bayesian satisficing model describes the human link selection process well. The model is therefore very useful for predicting what links will be selected by users when they are engaged in different information search tasks. For example, the model can provide direct quantitative predictions on how likely users will find their information for different designs of Web pages (e.g., what link text should be used, their layouts, etc).

In addition to predicting behavior at the individual level, SNIF-ACT also has the capability to simulate aggregate behavior. One robust finding on how likely users will stay on the same page is the law of surfing by Huberman and his colleagues (B. A. Huberman, Pirolli, Pitkow, & Lukose, 1998). They show that, by assuming that users make a sequence of decisions to proceed to another page and they continue as long as the value of the current page exceeds some threshold, the probability distribution for the number of pages that a user visits within a given Web site follows an inverse Gaussian distribution. Although SNIF-ACT was derived based on cognitive mechanisms at the individual level, the link selection process assumes that the probability that a user will stay on a given Web page depends on the relevance of links to the user’s information goal. Thus, the model can also predict how likely multiple users will stay on a particular page. Figure 34.2 (right) shows the probability distribution derived

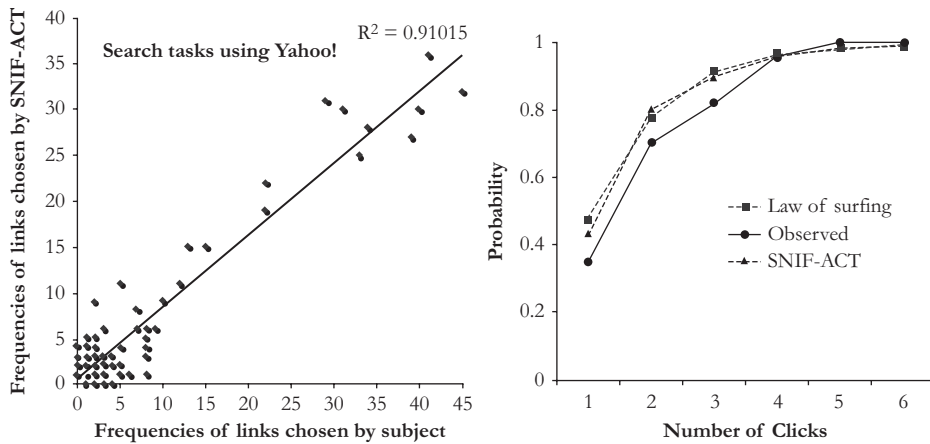


Figure 34.2 Left: The scatter plot for the frequencies of links chosen by the SNIF-ACT model and by human subjects when searching using the Yahoo! web interface. Right: The cumulative frequency distribution for the number of clicks on a web page as predicted by the law of surfing and SNIF-ACT. The estimated distribution from the empirical data (observed) was also shown for comparison.

from the SNIF-ACT model, the law of surfing, and the actual data collected from our experiment. As one can see, even though SNIF-ACT was derived based on completely different assumptions than the law of surfing (although the SNIF-ACT model is capable of predicting individual human behavior that the law of surfing cannot), it predicts the same aggregate behavior as the law of surfing. SNIF-ACT therefore serves as an example of a cognitive model that can make predictions that span across the individual and aggregate levels of analysis.

From Solo to Social: The Basic Social Information Foraging Model

The SNIF-ACT model, and the information foraging theory in general, is mainly focused on information seeking by a single user. Pirolli has extended the information foraging theory to explain cooperative information foraging theory (Pirolli, 2009). This multilevel model, called the basic social information foraging (basic *SIF*) model, is derived from the quantitative theories of cooperative problem solving (Clearwater, Hogg, & Huberman, 1992) and social foraging theory (Clark & Mangel, 1986; Giraldeau & Caraco, 2000). The model assumes that individual information foragers (e.g., users surfing the web) are engaged in a heuristic process of search for useful knowledge in a space in which useful information is encountered sporadically, in patches. It is assumed that a patch of information will yield some amount of utility. The information environment can be characterized by the expected number of steps required to find the next useful patch of information by random search

process, and search heuristics can be characterized by improvements to this baseline random search process (i.e., in terms of the proportion of search steps eliminated). It is assumed that heuristic hints are exchanged in cooperative information foraging regarding the likely location of useful information patches. For instance, hints are the social tags in systems such as delicious.com. Shared tags provide navigation paths and ontological organization to available content that subsequently improve individual information foraging. However, there are also various forms of interference effects in cooperative information foraging.

The basic SIF model assumes a heuristic process of search for useful knowledge in a space of discrete patches of information. It is assumed that a patch of information will yield some amount of utility for one or more foragers. The model assumes that heuristic hints are exchanged in cooperative information foraging regarding the likely location of useful information patches (for instance, as was observed to occur among analysts in the business intelligence agency discussed above). Another example of hints are the social tags in systems such as delicious.com. Hints may vary in the validity of the search information conveyed, in how they are interpreted by the information forager who receives them, and in effectiveness, depending on when they are exchanged in the search process. Based on the derivation by Huberman (Bernardo A. Huberman, 1990), the probability density distribution for finding a valuable information patch can be cast as a lognormal distributive function of the sample of hints H contributed by collaborators. Figure 34.3a shows the

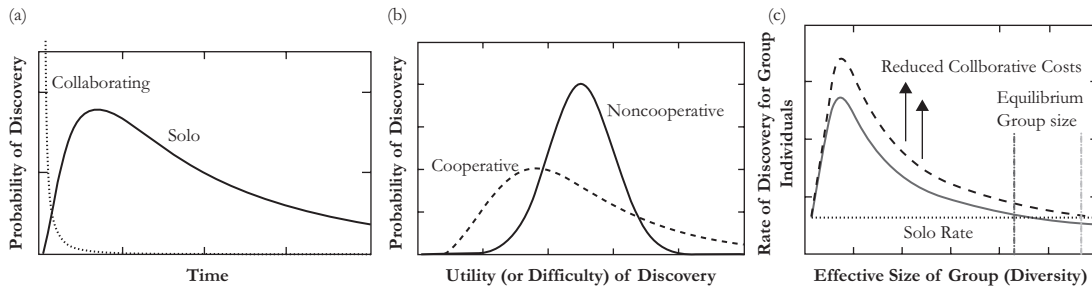


Figure 34.3 Collaboration with a diverse group improves the rate of return to the individual information forager. (b) Cooperation among information foragers improves the probability of making difficult discoveries. (c) The combination of expected benefits and interference costs to an individual determines the effective size of a group. People typically join a group only if the benefits (to the individual) outweigh the costs of cooperation (Clark & Mangel, 1986). A reduction in the costs of collaboration is predicted to increase the effective size of the collaborating group.

difference in search efficiency as H increases from 1 (solo) to 10 (collaborating).

As discussed in greater detail in Huberman (Huberman, 1990), the lognormal distribution of performance times makes interesting predictions about productivity of a cooperative group with respect to high-utility search results (see Figure 34.3b). If one assumes that the various states of a search space have a binomial distribution of utilities, then, assuming a mildly effective search heuristic, the search performed by noncooperating searchers will return a distribution of result values shown in Figure 34.3b. Increasing cooperation (e.g., by sharing knowledge as a group when searching for information, such as where one can find what information) will shift that distribution to a lognormal and will especially increase the likelihood of search results at the higher end of the utility spectrum. In a sense, this accounts for the “standing-on-the-shoulders-of-giants” effect that is frequently observed where an individual with an average amount of smarts benefits from the cooperation of others in making a better-than-average discovery.

Many species, besides humans, forage in groups. Although there may be positive effects of foraging in a group, foraging groups do not become arbitrarily large, suggesting that there may be some form of interference costs (e.g., intra-group competition) that at some point outweigh the advantages of further increments in the size of groups. It has been found empirically (Hassell & Varley, 1969) and in computational modeling (Anil, 2002) that there is often a power-law relationship between the number of foragers in a patch and the rate of consumption intake by each forager. Based on this assumption, it can be shown that the rate of discovery for individual and group foragers can be characterized by the curves shown in Figure 34.3c.

Figure 34.3c can also be used to discuss why the equilibrium group size may be greater than the optimum group size, \hat{n} , as discussed by Clark and Mangel (1986). Suppose solitary foragers have joined a group until it has the optimum size n^* . Solitary foragers should continue to join the group so long as the rate of return for group foraging is still above the rate of return for solitary foraging, as shown in Figure 34.3c. Members of the group may see their individual rates of return diminish from the optimum as new members join the group, but remaining in the group is still better than solo foraging. Consequently, individuals will join the group until the addition of new members makes the individual rate of return less than solitary foraging. Consequently, when the rate of return for group foraging has peaked, as in Figure 34.3c, we may expect the equilibrium size to be larger than the optimal size (i.e., $n^* > \hat{n}$). Another implication of the model presented graphically in Figure 34.3c is that any changes to the technology or policies that reduce the impact of having to deal with others should increase the overall participation rates. Reducing the costs of cooperation extends the tail of the rate of returns curve, which also extends the point at which it crosses the solo foraging threshold. Consequently, the equilibrium group size is predicted to increase.

We have shown how the information foraging model can be generalized to predict cooperative information foraging behavior. This multilevel predictive model serves as a tool to reason generally about several aspects of the power of cooperation and the social capital that is relevant to finding information. The model suggests that so long as the diversity of agents increases with group size, then the size of a group increases the overall power of cooperative discovery. As individual foragers increase the

diversity of their cooperating contacts, they will improve in performance. This provides a mathematical rationale for the idea that brokerage positions in social networks provide social capital. In the example shown in Figure 34.3, social capital can be quantified, based on the mathematical model, as the marginal advantage in performance when an individual is cooperating with a diverse group of agents, who tend to share non-overlapping information that helps the group to achieve better performance. The model also provides a rationale for the observed lognormal distribution of innovative discoveries. A variety of technologies have emerged to exploit or enhance social information foraging. Web, blogs, email, internet groups, collaborative tagging, wikis, recommender systems, and other technologies are all aimed at supporting cooperative information sharing, and their success implies their effectiveness. Given the increased ease with which it is possible to study social networks and information flow in the electronic world, it is likely that there will be more studies of the effects of technologies on social structure and social capital, hence a need for a suitable theoretical framework.

The Semantic Imitation Model: A Multilevel Model of Exploratory Search and Sensemaking

As described in the basic SIF, exchange of information cues is an important feature that allows more efficient human-information interaction. One increasingly popular form of this exchange of information cues is the use of social tags for indexing web resources. Social tagging systems, such as del.icio.us (<http://del.icio.us>) and CiteUlike (<http://citeulike.org>), allow users to annotate, categorize, and share their web content using short textual labels called tags. The popularity of tagging arises from its benefits for supporting personal online search, ability to browse content using tags, and organization and sharing of tagged contents. In contrast to the traditional keyword annotation system, a social tagging system provides users with an unstructured mechanism for organizing and managing content. In fact, the major characteristics of tags are an open vocabulary and nonhierarchical nature, and they are created by users of the information documents rather than by professional annotators. The relative unstructuredness of tagged content is also suggested as its potential weakness, as the openness of the systems may result in a large number of tags that are not meaningful to other users (Furnas, Landauer, Gomez, & Dumais, 1987). However, recent research

has shown that social tagging systems do exhibit stable patterns over time, which prompted researchers to investigate how these structural patterns emerge from multiple users in the system.

In our previous work (Fu, 2008; Fu & Kannampallil, 2009; Fu, Kannampallil, & Kang, 2009, 2010; Fu, Kannampallil, Kang, & He, in press), we argued that, rather than imitating other users at the word level, one possible explanation for this kind of social cohesion could be grounded in the natural tendency for people to process tags at the semantic level, and it was at this level of processing that most imitation occurred. This explanation was supported by research in the area of reading comprehension (Kintsch, 1998), which showed that people tended to be influenced by the meanings of words, rather than the words themselves, during comprehension. Assuming that the background knowledge of people in the same culture tends to have shared structures (e.g., using similar vocabularies and their corresponding meanings in order to conform and communicate with each), users of the same social tagging system may also share similar semantic representations of words and concepts, even when the use of tags may vary across individuals at the word level. In other words, we argued that part of the reason for the stability of social tagging systems can be attributed to the shared semantic representations among the users, such that users may have relatively stable and coherent interpretations of information contents and tags as they interact with the system. Based on this assumption, we developed the semantic imitation model (Fu et al., 2009; Fu et al., 2010; Fu et al., in press), which predicts how different semantic representations may lead to differences in individual tag choices and eventually to different emergent properties at the aggregate behavioral level. The model also predicts that the folksonomies (i.e., knowledge structures) in the system reflect the shared semantic representations of the users.

The model assumes that when a user is navigating in a social tagging system, existing tags associated with documents will invoke a tag-based topic inference process, such that the user can infer the topics contained in the documents based on the semantic interpretation of these tags. Based on the inferred topics, the user chooses tags to represent these topics. In other words, background knowledge structures of multiple users are dynamically “connected” through the interpretation and creation of tags in the social tagging system. The major components of the model are shown in Figure 34.4 and will be elaborated below.

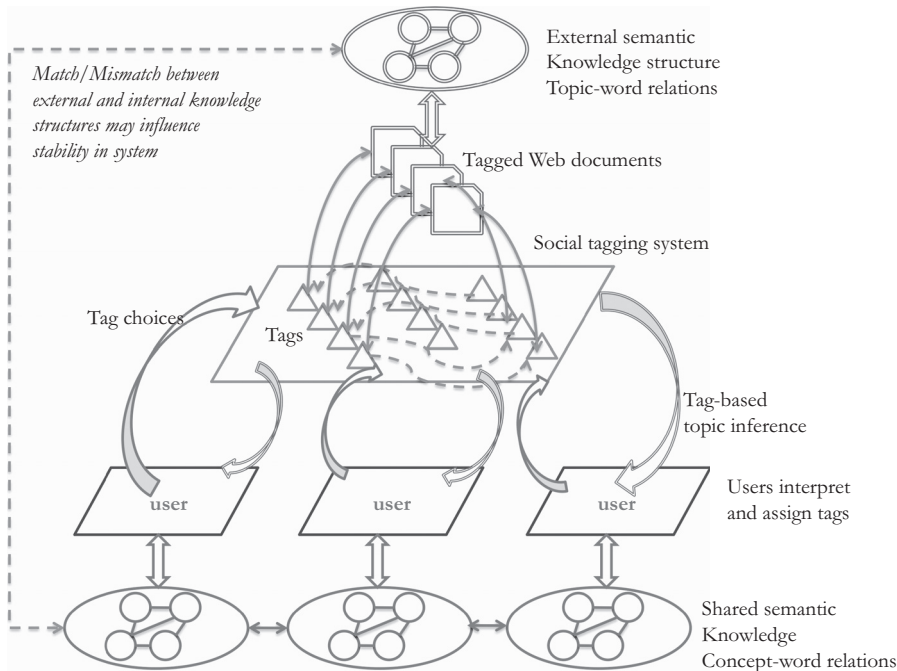


Figure 34.4 The semantic imitation model: Users interpret tags based on their background knowledge and infer topics in a document (tag-based topic inference). Based on these inferred topics, users choose tags to represent the latent semantics in the document (tag choices). It is assumed that there is a shared (but not identical) semantic knowledge background among users. The model shows that the match between the internal knowledge structures of users and the external knowledge structures contained in the documents will influence stability of the folksonomies formed in a social tagging system.

There are three main components in the model: words, semantic concepts (or topics) underlying the words, and documents that contain both the words and the topics. It is assumed that concepts can be represented as a probability distribution of words. For example, the concept “health” can instantiate many associated words such as hospital, doctor, surgery, etc. If c represents the set of available concepts and w represents the set of words, then $p(w|c)$ is the probability distribution of words, given a set of concepts. Using Bayes’ theorem, we can then calculate the set of concepts given a set of words $p(c|w)$.

As the user navigates through the system, existing tags of a resource act as retrieval cues for inferring the related concepts. As users browse through a set of tags that were previously created for a resource, semantic representations for those tags (words) are activated. This tag-based topic inference can be represented as a measure that estimates the probability that a set of tags will instantiate a specific set of topics (or concepts), $p(c|t)$, where c is the set of concepts and t is the tags that are associated with a resource. This tag-based topic inference process will facilitate the evaluation of the relevance of the document, as well as the comprehension of the document if a document is selected. Tag creation is based on

concept-word and word-concept relations. In order to do this, the probability of selecting a new tag given the set of existing tags and the words in the document is computed as in Equation 2, in which w_{new} is the new tag created by the user and d represents the aggregate of all words in the document and the existing tags. The utility of assigning the word w to the document d can then be calculated as U_w .

$$U_w = p(w_{new} | d) + \sigma = \sum_k p(w_{new} | c_k) p(c_k | d) + \sigma \quad (2)$$

The tag assignment process is based on the assumption that a *rational tagger* will choose tags so as to best represent the concepts contained in a document. The degree of representativeness of a tag (i.e., $p(w|d)$) provides such a measure, as it reflects the extent to which a tag represents the semantic contents of a document. We use a random variable, σ , to incorporate a degree of uncertainty into the model. Similar to the SNIF-ACT model, we used the random utility model (RUM) to represent the stochastic nature of the tag selection process.

As opposed to other models of social tagging, the semantic imitation model was based on a

cognitively plausible tag choice mechanism that was coupled with the formal representations of semantic knowledge that exist in both external documents and internal knowledge structures of the users. The model provides an explanation on generally accepted concepts of tag convergence and stabilization based on a cognitive model of individual taggers (users). Additionally, it has a mechanism for generating testable predictions about behavioral patterns in systems generated by different user populations (e.g., experts vs. novices).

To illustrate the importance of the role of semantic representations in the overall stability of the system, we created two sets of simulated users who differed in their background knowledge structures. First, simulated users who had word-concept distributions that matched perfectly with the word-topic distributions in the documents were created. These simulated users could represent users who

have strong domain expertise and therefore have developed highly structured knowledge that is well adapted to knowledge represented in various documents. In contrast, novices in a domain are unlikely to have a well-structured knowledge representation (similar to that in the document). We simulated novices by changing the spread of the prior distributions of words over concepts in their background knowledge. The wider spread in prior distributions implied that words were less accurate in predicting any particular concept (thus less effective topic inference), and, given a particular concept, there was a higher variance in the choice of words (tags) to represent the concept.

Results shown in the top left panel of Figure 35.5 were obtained from the simulated experts, and those in the top right panel were obtained from the simulated novices. The results show that experts reached stability much faster than novices.

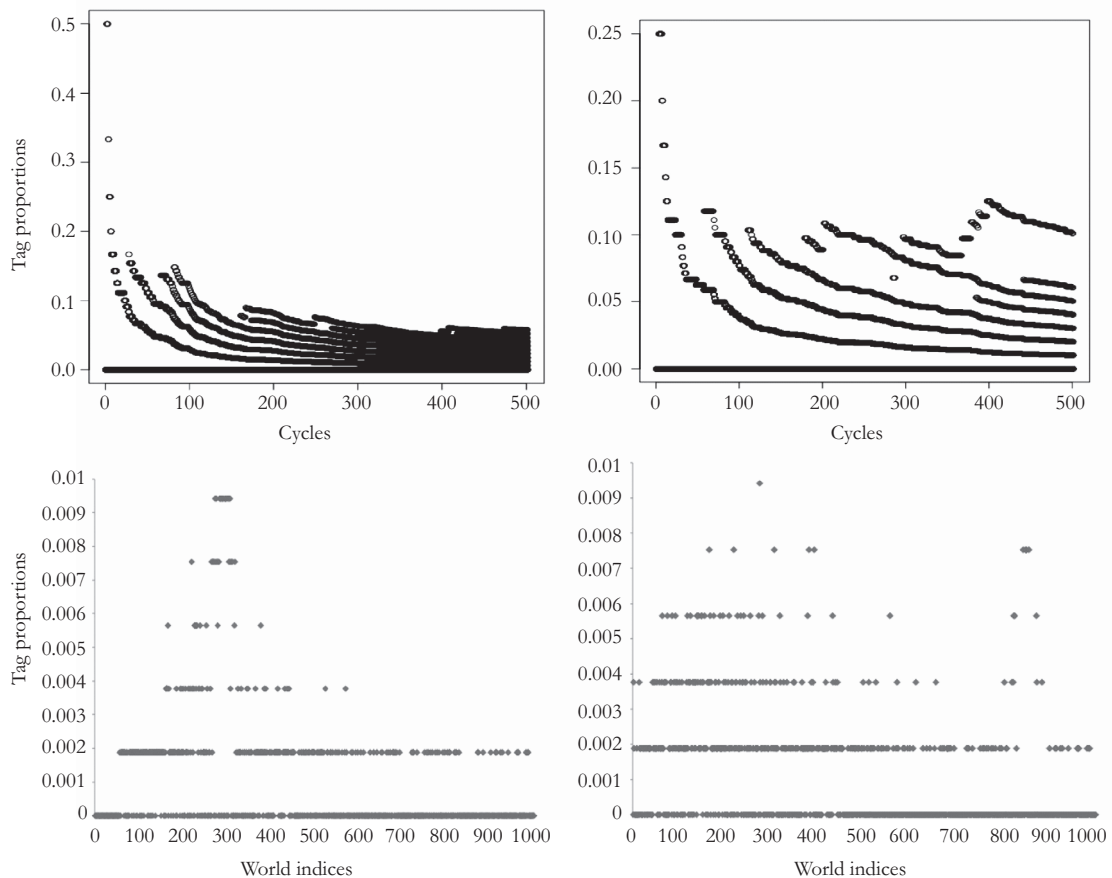


Figure 34.5 Top: Scatter plots of tag proportions against tag choice cycles by the semantic imitation model when there was high (left) and low (right) level of match between the word-topic distributions in the document and word-concept distributions in the background knowledge of the simulated user. Bottom: Scatter plots of choice proportions of each tag assigned to a single-topic document simulated by the semantic imitation model when there was high (left) and low (right) level of match between the word-topic distributions in the document and the word-concept distributions in the background knowledge of the simulated user.

The faster convergence in the case of experts could be explained by the fact that tags assigned to each document were more predictive of the topics contained in the document, and that the experts were much better at extracting the correct concepts based on “high quality” tags created by other experts. On the other hand, novices create tags that are less representative of the concepts. Novices were also less effective in extracting the topics (in the Bayesian sense) from the documents (McCarley & Benjamin, this handbook), and their choice of words resulted in more diverse concepts. Tags created were therefore more diverse, and thus convergence was much slower than experts.

The bottom panels show the scatter plots of the relative tag frequencies of one special document that we created to illustrate this difference. This special document contained a single topic, with the mean of the prior distribution of words over this single topic at word 300. As expected, for both experts and novices, tag proportions were highest around the most representative words. However, experts clearly had a much more focused vocabulary than novices, as shown by the wider spread of tag choices. In addition, novices seemed to have “misinterpreted” the topic and chose tags around word 800 (the initial choice of this tag was due to random noise) to represent the wrong topic, which led others to follow (the initial choice led to a second cluster of words around the “wrong” topic). Both the wrong interpretation of topics and the higher variance in word choices contribute to the slower convergence for novices than experts. The model therefore predicts that systems that are often used by domain experts (e.g., by academic researchers, as in CiteUlike or CiteSeer) will likely converge faster and have more high-quality tags than those that are designed for general users (e.g., Del.icio.us).

The simulation results in Figure 34.5 show that the semantic imitation model was successful in explaining the same stability of tag proportions as found by others. The major contribution of the current model is that the prediction was based on a cognitively plausible tag choice mechanism that was coupled with the formal representations of semantic knowledge that exist in both external documents and internal knowledge structures of the users. The results show that not only was the model capable of offering a sophisticated explanation of the stabilization of tag proportions based on a cognitive model of individual users, but also that it can generate testable predictions of emergent social behavioral patterns in systems used by different user populations

(e.g., experts vs. novices). The results demonstrate how this multilevel modeling approach can explain the impact of different user profiles on social behavior. It also has the potential to include even lower-level models (such as how information presented on the interface may influence cognitive processing of information) and eventually influence social behavior. It can also be incorporated into higher-level network models to explain characteristics exhibited by different social networks.

Conclusions and Future Directions

An apparent trend for future information technologies is to increase the connectedness among individuals to encourage flow of information. This implies that individual actions are no longer independent of each other. Rather, actions of any individual in the connected information environment will likely impose externalities on other individuals, and thus change the structure of the environment confronting them. This is, in spirit, similar to the idea in the segregation model of Schelling (1971, 1978), who showed that the preference of individuals may lead to changes in the environment that in turn influence how other individuals may behave, and the effect may be magnified as the dynamics grow in a neighborhood, leading to segregation. Cognitive engineering theories should therefore aim to capture the characteristics of the micro-to-macro link, instead of focusing on either the individual behaviors by isolating them from the actual social context or the statistical relationship of variables derived from aggregated behavioral data. We argue that, as we transition from the information age to the network age, it is no longer sufficient to develop good cognitive engineering principles based on isolated levels of analysis. Instead, good theories should be equipped with the explanatory power to make the transitions between micro- and macro-levels successfully. We show that computational models, combined with testing that aims to reveal dynamic relations among multiple levels of behavioral analysis, are building blocks for achieving such research goals.

There are a few challenges for future socio-computer interfaces. First, with the rapid growth of networked systems, the amount of information generated from the interactions among human operators as well as between human and machines is exponentially increasing. How to extract meaningful patterns from this massive amount of information to facilitate performance at both the individual and social levels is challenging. Second, while more

“connectedness” is often considered useful for the spread of information, more studies have shown that “structural holes”—or in general certain forms of “disconnectedness”—are desirable for certain socio-technical systems (e.g., Burt, 2004). How specific network structures may facilitate performance in different socio-technical systems is therefore an important area of research that is relatively underdeveloped. Third, as the world becomes more connected, social and cultural differences will likely moderate how effectively people collaborate, and how technology can be designed to mitigate the potential negative effects induced by these differences will be an important challenge. Last but not the least, research has only begun to harness the rich networks of knowledge embedded in socio-technical systems, and there is still so much that we do not know about how these networks of knowledge can facilitate learning, knowledge acquisition, and performance (“Learning and Retention,” Ritter et al., this handbook).

References

- Anderson, J. R. (2002). Spanning seven orders of magnitude: A challenge for cognitive modeling. *Cognitive Science*, 26(1), 85–112.
- Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y. (2004). An integrated theory of the mind. *Psychological Review*, 111, 1036–1060.
- Anil, K. S. (2002). Modeling group foraging: Individual sub-optimality, interference, and a kind of matching. *Adaptive Behavior: Animals, Animats, Software Agents, Robots, Adaptive Systems*, 9(2), 67–89.
- Barabási, A.-L., & Albert, R. (1999). Emergence of scaling in random networks. *Science*, 286, 509–512.
- Burt, R. (2004). Structural holes and good ideas. *The American Journal of Sociology*, 110(2), 349–399.
- Clark, C. W., & Mangel, M. (1986). The evolutionary advantages of group foraging. *Theoretical Population Biology*, 30(1), 45–75.
- Clearwater, S. H., Hogg, T., & Huberman, B. A. (1992). Cooperative problem solving. In B. A. Huberman (Ed.), *Computation: The micro and macro view* (pp. 33–70). Singapore: World Scientific.
- Coleman, J. S. (1987). Microfoundations and macrosocial behavior. In J. C. Alexander, B. Giesen, R. Munch, & N. J. Smelser (Eds.), *The micro-macro link* (pp. 153–173). Berkeley, CA: University of California Press.
- Coleman, J. S. (1994). *Foundations of social theory*. Cambridge, MA: Harvard University Press.
- Fu, W.-T. (2007). Adaptive tradeoffs between exploration and exploitation: A rational-ecological approach. In W. Gray (Ed.), *Integrated models of cognitive systems*. New York, NY: Oxford University Press.
- Fu, W.-T. (2008). The microstructures of social tagging: A rational model. In *Proceedings of the ACM 2008 Conference on Computer Supported Cooperative work* (pp. 229–238). San Diego, CA.
- Fu, W.-T., & Gray, W. (2006). Suboptimal tradeoffs in information seeking. *Cognitive Psychology*, 52(3), 195–242.
- Fu, W.-T., & Kannampallil, T. G. (2009). Harnessing Web 2.0 for context-aware learning: The impact of social tagging system on knowledge adaption. In N. Lambropoulos & R. Margarida (Eds.), *Educational social software for context-aware learning: Collaborative methods and human interaction* (pp. 98–113). Hershey, PA: IGI Global.
- Fu, W.-T., Kannampallil, T. G., & Kang, R. (2009). A semantic imitation model of social tag choices. In *Proceedings of the IEEE Conference on Social Computing* (pp. 66–72). Vancouver, BC.
- Fu, W.-T., Kannampallil, T. G., & Kang, R. (2010). Facilitating exploratory search by model-based navigational cues. In *Proceedings of the 2010 International Conference on Intelligent User Interfaces*. Hong Kong, China.
- Fu, W.-T., Kannampallil, T. G., Kang, R., & He, J. (in press). Semantic imitation in social tagging. *ACM Transactions on Computer-Human Interaction*.
- Fu, W.-T., & Pirolli, P. (2007). SNIF-ACT: A cognitive model of user navigation on the World Wide Web. *Human-Computer Interaction*, 22, 355–412.
- Furnas, G. W., Landauer, T. K., Gomez, L. M., & Dumais, S. T. (1987). The vocabulary problem in human-system communication. *Communications of the ACM*, 30(11), 964–971.
- Giraldeau, L.-A., & Caraco, T. (2000). *Social foraging theory*. Princeton, NJ: Princeton University Press.
- Hassell, M. P., & Varley, G. C. (1969). New inductive population model for insect parasites and its bearing on biological control. *Nature*, 223(5211), 1133–1137.
- Hedstrom, P. (2006). Sociology: Experimental macro sociology: Predicting the next best seller. *Science*, 311(5762), 786–787.
- Huberman, B. A. (1990). The performance of cooperative processes. *Physica D: Nonlinear Phenomena*, 42(1–3), 38–47.
- Huberman, B. A., Pirolli, P., Pitkow, J., & Lukose, R. M. (1998). Strong regularities in World Wide Web surfing. *Science*, 280, 95–97.
- Iribarren, J. L., & Moro, E. (2009). Impact of human activity patterns on the dynamics of information diffusion. *Physical Review Letters*, 103(3), 038702.
- Kintsch, W. (1998). *Comprehension: A paradigm for cognition*. New York, NY: Cambridge University Press.
- Ku, G., Galinsky, A. D., & Murnighan, J. K. (2006). Starting low but ending high: A reversal of the anchoring effect in auctions. *Journal of Personality and Social Psychology*, 90(6), 975–986.
- Leskovec, J., Backstrom, L., Kumar, R., & Tomkins, A. (2008). Microscopic evolution of social networks. In.
- Liben-Nowell, D., & Kleinberg, J. (2008). Tracing information flow on a global scale using Internet chain-letter data. *Proceedings of the National Academy of Sciences*, 105(12), 4633–4638.
- Lukose, R. M., & Huberman, B. A. (1998). Surfing as a real option. In *Proceedings of the first International Conference on Information and Computation Economics* (pp. 45–51). Charleston, NC.
- McFadden, D. (1974). Conditional logit analysis of qualitative choice behavior. In P. Zarembka (Ed.), *Frontiers of econometrics* (pp. 105–142). New York, NY: Academic Press.
- Pirolli, P. (2005). Rational analyses of information foraging on the web. *Cognitive Science*, 29, 343–373.
- Pirolli, P. (2007). *Information foraging theory: Adaptive interaction with information*. New York, NY: Oxford University Press.
- Pirolli, P. (2009). *An elementary social information foraging model*. Paper presented at the Proceedings of the 27th International

- Conference on Human Factors in Computing Systems CHI '09, Boston, MA.
- Pirolli, P., & Card, S. (1999). Information foraging. *Psychological Review*, 106(4), 643–675.
- Pirolli, P., & Fu, W.-T. (2003). SNIF-ACT: A model of information foraging on the World Wide Web. *Proceedings of the Ninth International Conference on User Modeling*, 2702, 45–54.
- Salganik, M. J., Dodds, P. S., & Watts, D. J. (2006). Experimental study of inequality and unpredictability in an artificial cultural market. *Science*, 311(5762), 854–856.
- Schelling, T. C. (1971). Dynamic models of segregation. *Journal of Mathematical Sociology*, 1, 143–186.
- Schelling, T. C. (1978). *Micro motives and macro behavior*. New York, NY: Norton.
- Stephens, D. W., & Krebs, J. R. (1986). *Foraging theory*. Princeton, NJ: Princeton University Press.
- Sun, R. (Ed.). (2006). *Cognition and multi-agent interaction: From cognitive modeling to social simulation*. Cambridge University Press.
- Wu, F., Huberman, B. A., Adamic, L. A., & Tyler, J. R. (2004). Information flow in social groups. *Physica A: Statistical and Theoretical Physics*, 337(1–2), 327–335.

