Social Navigation:

Techniques for Building More Usable Systems

A. Dieberger, P. Dourish, K. Höök, P. Resnick, and A. Wexelblat

The term "navigation" conjures images of maps, compasses, and guidebooks. These may be tools we use to get around from time to time, but are they how we usually find our way? Imagine walking down a street in your hometown, trying to decide what to do. You notice a crowd outside your favorite café. Knowing that the café often has live music, you can guess that a special event must be happening tonight. You might decide that you're in the mood for a lively evening and join the line, or you might decide that you prefer a quiet night and look for a different café. Or imagine you're in a library, looking for a book about interface design. One of the books on the shelf is much more worn and dog-eared than the other, suggesting that lots of people have read it. You may decide it's a better place to start learning than the pristine books beside it on the shelf. In both cases, you didn't rely on maps or guides; instead, you used information from other people to help make your decision. This is a different sort of "finding your way." We call it "social navigation," a topic we discussed on a panel at CHI'99 in Pittsburgh.

Social navigation ideas, often taken from the physical world, can be valuable elements of design in the digital domain. Imagine logging into an unfamiliar chat system. Beside the chat topics, you see representations of recent activity in each chat room. The system shows you how many people are chatting, along with icons that change as users in the chat rooms change topics. Maybe the chat system also indicates movement of people between rooms. At a glance, you can see hot topics, crowd formation, and the general tenor of the system. You may find a room to your liking and go chat there, or move on. Again, information about what others are doing has guided and informed (but not controlled) your decision. This is social navigation in the online world.

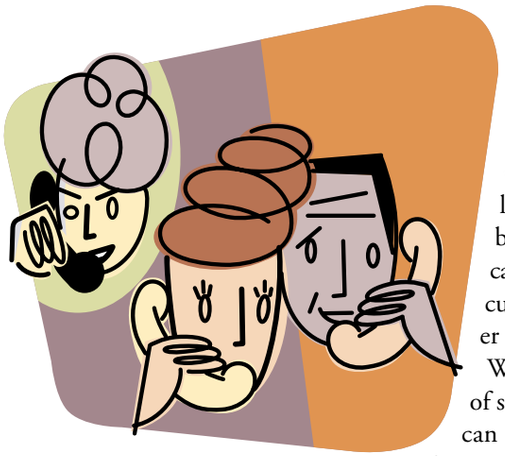

Social navigation can also work through information traces left by previous users for current users. Just like the dog-eared book, these traces can show us shortcuts that take us closer to our destinations. With the right kind of software support, we can allow people to leave useful traces with digital information as well. The reader reviews, comments, and message boards popular on websites today all can be used for this purpose. In addition, observations that "people who liked X also liked Y" provide another kind of useful trace for social navigation.

These scenarios have one thing in common: they describe forms of navigation—in a general sense—through which decisions are informed by the behavior of other people. This behavior can be directly visible, such as when people move from one place to another, or it can be aggregated and hidden in the interaction history of a physical object or place, such as in the library. This information can be anonymous or it can be directly associated with a person, such as when a reviewer signs a recommendation for a book and even provides background information about him- or herself.

Social navigation is not a new idea: it's something that we use every day. However, although other ideas from real-world navigation have been incorporated into interface designs, such as maps, guides, and short cuts, social forms of navigation are only slowly being adopted in software systems. Most digital systems don't help us navigate socially. On most areas of the Web, for example, users are given the illusion of being the only person present. The only indication they may get of other people's using the system is a slower-than-usual response time. Systems actually showing other people directly or indirectly are in the minority, with virtual communities, chat rooms, and so on being more the exception than the rule. Even on sites that provide these community features, it is usually not possible to take advantage of the work done by earlier visitors to solve problems, retrieve information, and so on. Vannevar Bush's famous 1939 article "As We May Think," which is often cited as an early source of hypertext ideas, proposed not only the idea of links between information, but that people might share the "trails" they create through information space. Now that we live in a hypertext-based world, perhaps it is time to realize the other elements of Bush's vision.

What would this mean for building information systems? We do not propose putting a chat box on every Web page. Instead, when we are building our systems, we think of facilities that make us aware of other people's activities and select ones that seem appropriate for the task. This must be done without infringing on people's privacy and without interrupting work, except if explicitly desired.

An important point here is that although we rarely make direct or conscious use of information related to social navigation, in many situations we want to be peripherally aware of others, particularly within a group of friends and colleagues. In many cases within the Web or other software systems, users are not solving novel problems. Previous problem-solvers could leave traces of their work, and those traces could be picked up, used, improved on, and left again for future need. If

we can make available the results of activities of previous users, we believe that solving repeated problems could be made easier.

Key Social Navigation Properties

Taking advantage of information created by other people can occur in many ways, but not all of them capture the sense of social navigation as we mean it here. Two additional properties are needed to describe the phenomena we aim to capture: *personalization* and *dynamism*. Two examples borrowed from Svensson [15] illustrate the importance of these ideas:

1. Walking down a path in a forest is social navigation, but walking down a road in a city is not.
2. Talking to a person at an airport help desk who explains how to find the baggage claim is social navigation, but reading a sign with more or less the same message is not.

In these examples both methods seem to involve the same navigational advice; the difference lies in how advice is given to the navigator. In the first example, the navigator chooses to follow a path based on the fact that other people have walked that way. Conversely, walking down a street is not driven by the fact that other people have walked the same street. The street is an intrinsic part of the space. One way to think about this is that social navigation traces are not preplanned aspects of a space, but rather are "grown"—or created *dynamically*—in a more organic, or bottom-up, fashion. In this way, social navigation is a closer reflection of what people actually do than it is a result of what designers think people should be doing.

In the second example the navigator gets the impression that the navigational advice is *personalized* to her and the situation allows her to ask for additional information. Also, the advice ceases to exist when the communication between the navigator and advice provider ends. The person at the help desk may have to use different terms, or even speak a different language, to convey the same message to each particular customer. The help-desk worker can also recognize a repeat visitor, and modify the presentation of information according to knowledge that a past attempt has failed.

Another significant distinction between social navigation and general navigation is how navigational advice is mediated. Social navigation has a strong temporal and dynamic aspect. A person chooses to follow a particular path in the forest because she makes the assumption that people have walked it earlier. Forest paths are transient features in the environment; if they are not used they vanish. Their state (how well-worn they are) can indicate how frequently or recently they have been used, which is typically not possible with a road.

Social navigation is a closer reflection of what people actually do than it is a result of what designers think people should be doing.

We see therefore that social navigation relies on the way that people occupy and transform spaces, leaving their marks upon them—turning a "space" into a "place" in the terminology of Harrison and Dourish [6]. In time, the social cues people leave behind can be-come sedimented and formalized, transformed into social practices (such as letting people get off the train before you get on), rules and regulations (such as those governing driving), or artifacts (such as signs and landmarks). Social navigation, in the sense of our individual actions being designed around collective social behavior, is not just something that is "layered on top of" a space, but comes to transform both the space and the ways that people act within it. To design with such ideas is to leave yourself open to the possibility that users will

render your system unrecognizable by you and your co-designers.

Styles of Social Navigation Systems

Social navigation has been explored in various research systems and is beginning to find its way into commercial offerings in a variety of ways.

A popular example of traces on the Web is recommendations, such as those encountered at online retailers like Amazon.com: "people who bought this book also bought…" Customers purchasing goods at Amazon.com leave an anonymous trace of their activities in the system. These traces can assist other people in their browsing and decision making, even at this coarse level of aggregation. Amazon.com has displayed this sort of social navigation for some time and recently incorporated other features such as customer reviews and "purchase circles," which shows that a certain book is popular with, say, people from Boston, or that people from the apple.com domain buy a given list of books.

Social navigation relies on the way that people occupy and transform spaces.

Services like these are based on *recommender systems* [2], which help people make selections by looking at what other people with similar interests have done. They collect data (usually ratings, but sometimes comments) from decision makers and then analyze the data to find patterns that suggest similar sets of interests. The simplest recommendation systems produce ratings or rankings that are the same for all users, much like best-seller lists. More sophisticated systems, such as GroupLens [9], group the data according to similarity metrics. Some systems exploit profile information to make recommendations. For example, if I own a Volvo, I might get a better recommendation from data submitted by other Volvo owners than from data submitted by car owners in general. Other recommendation systems find patterns in the data to discover similarities and improve their advice without needing to know more about users.

Although recommendation systems have existed for a number of years, important challenges to their successful deployment remain. The first is that they often need a lot of rating information on which to base their recommendations, which can make it hard to "bootstrap" such systems. This is also called the "early rater" problem, because it tends to penalize people who enter ratings early; their work benefits people who ask for recommendations later, but they don't get good recommendations themselves until sufficient ratings have been entered. Joe Konstan, John Riedl, and their colleagues in the GroupLens project at the University of Minnesota have explored approaches that can help relieve this problem. A second challenge is that it's difficult to derive good metrics for similarity in different domains, which often leads to people being given low quality ratings until the system has learned their profile. This problem is compounded by the difficulty of explaining why a rating was generated, because ratings are based on broad statistical trends. A third challenge is that recommender systems rely on disclosure of information by people about their actions (in the form of ratings or purchase histories) that they may not wish to disclose. Systems that track user actions are subject to a host of privacy problems, even when the information is used to enrich the interaction. As designers, we need to understand not only how to manage privacy sensitively, but also how to make sure that people understand what information they are disclosing and how it is used—a balance of visibility, awareness, and accountability that Erickson and Kellogg [4] title "social translucence."

In addition to the commercial recommender systems used today, a number of

social navigation tools are being explored in research settings. For example, the EFOL system developed by Kristina Höök and her colleagues in the PERSONA project is used for shopping for food over the Internet. Cooks select a set of recipes, and the necessary ingredients are added to the users' shopping chart. Unlike the recommender approach, EFOL uses the idea of a *populated space*—an information space in which other people can be encountered.

Figure 1: The interface of EFOL, an online store.

The EFOL food store has been enriched with a number of different functions that improve social navigation using a toolkit named the Social Navigator [15]. First, the recipes themselves are ordered by collaborative filtering methods. Choosing a recipe means that the ingredients will be sent to the user's doorstep, so picking a recipe is a strong vote for that recipe. In addition to recommending individual recipes, recipes are grouped into *recipe clubs*, which are places with special themes, such as "vegetarian food." Users can move between clubs to get different recommendations. Users also have a virtual presence in the shop through icons (avatars) representing them in an overview map. As recipe seekers move from one club to another, their avatars are shown moving from one location to another in the map (see Figure 1). The system provides a chat function, allowing cooks to talk with each other in the same recipe club. Finally, the system provides social annotations in more anonymous ways: each recipe bears signs of who put it there (the author) and how many times it has been downloaded.

Another approach that has been explored in a number of systems is to augment information with cues about the history of previous activities over that information. Andreas Dieberger explored this idea of "history-enriched" navigation on the Web in the social navigation system CoWeb. The system tracks how often pages on a collaborative Web server have been accessed and when they were last modified. It then annotates links to these pages with markers indicating the amount of recent traffic on that page, whether the page hasn't been accessed for a long time, or if that page was recently modified. The annotation indicates a history of the page itself, not of the link. It also aggregates history over all users so that these markers can be used as an indirect and anonymous social navigation tool.

Another example of a history system was built based on IBM's WBI toolkit [13]. It observes people's paths through the Web and looks for recurring paths. For example, I might not remember the URL of Kia's home page, but I know I can get there from my own home page by going to "friends" and then to "Kia Höök." If I follow these steps repeatedly the history mechanism will insert a shortcut

to Kia's home page at the top of my home page. This system is based on "intermediary computation." It does not actually modify the page; rather, it creates a personalized version of the page for me, based on my own browsing history, and "inserts" it between my browser and the website. Such a system can be used also by a group of people—shortcuts then represent commonly observed browsing behavior within the group.

Alan Wexelblat's Ph.D. system, Footprints [16], took a different approach to history-enriched information. Footprints used a client proxy server to connect users anywhere to a common database of history information. By keeping the data separate from both the Web browser and Web server, Footprints could be used on any existing Web page with no prior modification. Footprints presented different visualizations of history information as maps, trails, and annotations allowing users to see where within a page activity had taken place. Wexelblat's experiments with Footprints showed that within the context of trying to solve a particular information-finding problem, history information was not useful for naïve users, but was useful for people who had some familiarity with the type of problem. In effect, it seems that social navigation can help people understand that they're "on the right path" to a solution, but it's unclear what else needs to be done to help general users find such paths.

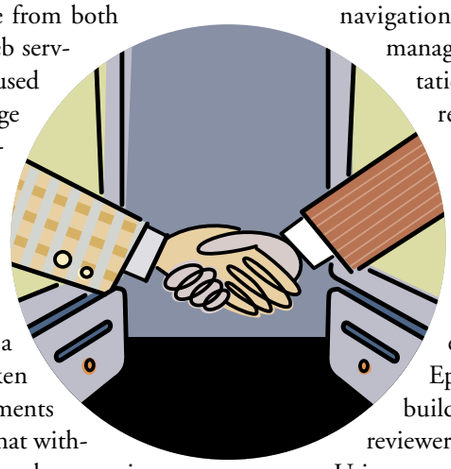

Open Problems and Interesting Questions

Social navigation is still emerging as an approach to information system design. As we apply it to new information tasks and new information communities, we encounter new problems and solutions.

Recommendations and Reputations

Store-based recommenders like that at Amazon.com provide feedback largely by what you purchased. Missing in many recommendations is feedback on whether the item you bought, read, or visited really met your needs and whether you enjoyed it. Reviews can provide some of this feedback, but reviews can vary widely in quality. The advice site Epinions.com tries to improve on this by rating reviewers themselves so that you can see right away whether a review was written by an "expert reviewer." An "expert reviewer" is not the same as a "domain expert." Ideally we would want highly rated reviews by domain experts who can write well on such a system and who know what your personal preferences are.

An active area of investigation in social navigation systems is reputation management. A person's reputation lets you evaluate their recommendations and determine how much trust you might want to put in them. Auction sites such as eBay incorporate information on the reputations of buyers and sellers, whereas Epinions allows people to build up reputations as expert reviewers on one topic or another. Using reputations solves some of the problems in recommender systems that we outlined earlier, such as the difficulty of interpreting a recommendation. At the same time, though, it requires that all recommendations be personally associated with an individual, eliminating the anonymity that comes from statistical measures such as "people who bought this book." Reputation systems are also subject to various forms of "spoofing," in which reputations are artificially inflated to appear more reliable than they really are. The use of reputations needs to be balanced with privacy and security concerns when deciding what approach to take to system design.

Forms of Social Navigation

Another area of investigation is the precise form that social navigation information will take. We can distinguish between different forms of social navigation:

- Whether other users co-exist or if it is

the aggregated history of previous, nonconcurrent, usage that is displayed in the interface;

- Whether other users can directly contact one another (direct social navigation) or are only in anonymous, indirect contact (indirect social navigation).

An agreed-on theory does not exist of which forms of social navigation are most relevant in the design process and how to best design for social navigation. For example, Paul Dourish and Paul Resnick believe that the recommender function should be part of every aspect of computer systems—it should be a general process running in the background of all applications, so that all aspects of interaction are suffused with information about other people's actions. Such an approach would allow users to get recommendations on anything from how to set their network mask to what query syntax to use for search engines. In different situations, recommendations would be based on different subsets of other users. The network mask recommendation, for instance, might be based on the user's co-workers' settings, whereas the Internet search might be based on data from a topical mailing list to which the user subscribes.

In contrast, Kristina Höök holds a different view; she claims that social navigation will be more or less useful in different domains and situations and should not be built in by default. She is also more concerned that a social navigation function might fail, leading users astray.

Andreas Dieberger and Alan Wexelblat have followed a third approach. As well as incorporating social navigation directly into different aspects of a system, they look toward spatial metaphors as the basis for enabling social navigation, drawing on work in architecture and urban design. These disciplines have a long tradition of observing how people navigate socially in the everyday spatial environment (see [1, 12]). From this work, they draw out models and metaphors that incorporate similar combinations of sociality and spatiality into information systems (see, for example, [3]).

How Does Social Navigation Help?

How might the presence of social navigation capabilities affect users' behavior? There is a difference between concluding that social navigation happens in the world no matter what we do and deciding that it is a good idea to design systems based on this phenomenon. How will our perceptions of our systems change? Incorporating social navigation might not, for example, change the paths we might take through the Web, but we think it can profoundly influence how we experience our systems and will significantly change interactions we have with other people mediated through our systems. Simply put, we believe social navigation can strongly influence users' experience of a system that is used by many people. Although we still lack the empirical grounds needed to know which kinds of social navigation forms are most relevant in different circumstances and domains, there are some early suggestions.

Filtering

History-enriched environments will help users find the most relevant information [11, 16]. Users guided by history markers will have help choosing what bears a closer look. Recommender systems have similar effects. They help users pick out a reasonable set of items from a huge space. This depends in part on the nature of the domain; some domains more easily lend themselves to recommendations because they allow users to judge the content before choosing it. In other domains one might not discover that the item chosen (for example, a food recipe) is bad until after choosing it. History must include feedback mechanisms in order to be useful in these situations. Some domains depend more heavily on expert recommendations; to function properly in these domains, history-enriched environments may need to include information on users' credentials and expertise.

Quality

History-enriched environments will aid users in finding good *quality* information, in the sense that the information is interesting, that it is valid, and that the author or artist or producer who created or provided it did a good

job. In some situations, certain people's use of a particular information source is more relevant than the existence of another information source that provides more reliable information, as Harper [5] points out in a study on information workers at the International Monetary Fund. Harper uses this example to introduce a "framework of relevance." He shows that a piece of information might very well be valid and important and still completely uninteresting because the people with power are not reading and acting on it. Thus, quality is a more complex concept than simply "contains the right keywords."

Social Affordance

History-enriched environments will make users more aware of each other and contribute to a social experience of the information space. When entering a new application, a new space, a new service, we often need help learning. Visible actions of other users can inform us what is appropriate behavior—what can or cannot be done—and provide *social affordance.* At the same time, this awareness of others and their actions make us feel that the space is alive and might make it more inviting. Here we are not really interested in whether users navigate more efficiently, or find exactly what they need more quickly; instead, we want to make them stay longer in the space, feel more relaxed, and perhaps be inspired to try out a new function or to pick up new products and new information items or to try out new services that they would not have considered otherwise.

Study of the EFOL system showed that awareness of others and how they moved around the food recipes influenced users to explore more of the functions of the system [7].

Use Reshapes Experience

Social navigation design will alter the organization of the space. In Amazon.com, the structure of the space experienced by visitors is changed: one can follow the recommendations instead of navigating by the search-for-terms structure. Social navigation thus could be a first step toward empowering users to, in a natural subtle way, make the functions and structure "drift" and make our information spaces more "fluid" [10]. In current systems, only the designer of the system can influence the structure of the space. Through social navigation we can allow aggregated behaviors or expert behaviors to take part in shaping the system.

Digital information systems would be improved if their designers considered how one user within the system could help another.

Conclusions

As a design approach, social navigation is still in its infancy. Few systems have been deployed outside laboratories and test user groups, and much more work needs to be done to evaluate the concrete benefits we expect to come from this style of system-building. Each step we take reveals a wealth of important issues that need to be explored. Additionally, as systems move into real use, new challenges will arise.

For example, privacy concerns about use information are still inadequately understood. Solutions involving anonymity, pseudonymity, and attribution will need to be explored. Although we are used to being observed in public spaces (the supermarket, the corner pub), we might not wish to be similarly observed in electronic spaces. Privacy concerns must also interact with trust concerns; how can users be encouraged to trust the hints given by social navigation systems?

We must also address concerns that arise from using these systems over longer periods of time. For example, should history information "fade"? If so, at what rate? How can history information be usefully aggregated so that users will not be overwhelmed by the

wealth of social navigation suggestions or clues? How can accumulated history information be shown without overwhelming or obscuring whatever users were looking for in the first place?

Despite these challenges, we believe that social navigation is today mature enough to be introduced as a choice for system builders and interaction designers. We believe that many, if not most, digital information systems would be improved if their designers considered how one user within the system could help another. Such thoughts could turn the lonely, socially void information spaces we have now into more humane environments, and maybe into real places.

References

1. Brand, S. *How Buildings Learn: What Happens After They're Built*. Penguin, New York, NY, USA, 1995.

2. Communications of the ACM. Special issue on recommender systems. *Communications of the ACM 40*, 3 (1997).

3. Dieberger, A. Social connotations of space in the design for virtual communities and social navigation. In Munro A., Höök , K., and Benyon, D., *Social Navigation of Information Space*, Springer, London, 1999, pp. 35–54.

4. Erickson, T. and Kellogg, W.A. Social translucence: An approach to designing systems that support social processes. *ACM Transactions on Human Computer Interaction 7*, 1 (2000), pp. 59–83.

5. Harper, R. Information that counts: A sociological view of information navigation. In Munro A., Höök , K., and Benyon, D., S*ocial Navigation of Information Space*, Springer, London, 1999, pp. 80–89.

6. Harrison, S. and Dourish, P. Re-Placeing Space: The Roles of Space and Place in Collaborative Systems. ACM Conference on Computer-Supported Cooperative Work (CSCW'96), Boston, 1996, pp. 67–76.

7. Höök, K., Rudström, Å., and Waern, A. Edited Adaptive Hypermedia: Combining Human and Machine Intelligence to Achieve Filtered Information. Workshop "Flexible Hypertext", Hypertext'97, Southampton, U.K., 1997.

8. Höök, K., Laaksolahti, J., et al. Individual Differences in Social Navigation. Workshop on "Social Navigation: A Design Approach?" ACM Conference on Computer–Human Intereraction (CHI'2000), The Hague, 2000. (available from SICS at www.sics.se/)

9. Konstan, J., Miller, B., et al. GroupLens: Applying collaborative filtering to Usenet news. Communications of the ACM 40, 3 (1997), pp. 77–87.

10. Levy, D.M. Fixed or Fluid? Document Stability and New Media. European Conference on Hypertext (ECHT'94), Edinburgh, 1994, pp. 24–31.

11. Lönnqvist, P., Dieberger, A., et al. Usability Studies of a Socially Enhanced Web Server. Workshop on "Social Navigation: A Design Approach?" ACM Conference on Cimputer-Human Interaction (CHI'2000), The Hague, 2000. (available from SICS at www.sics.se/)

12. Lynch, K. The Image of the City. MIT Press, Cambridge, MA, 1960.

13. Maglio, P. and Barrett, R. Intermediaries personalize information streams. Communications of the ACM 43, 8 (2000), pp. 96–101.

14. Munro, A.J., Höök, K., and Benyon, D.R. Social Navigation of Information Space. Springer Verlag, London, 1999.

15. Svensson, M. Defining and Designing Social Navigation. Bachelor's thesis. University of Stockholm, 2000.

16. Wexelblat, A., and Maes, P. Footprints: History-Rich Tools for Information Foraging. ACM Conference on Computer-Human Interaction (CHI'99), Pittsburgh, 1999, pp. 270–277.

17. Dieberger, A. Where did all the people go? A Collaborative Web space with social navigation information. Poster at WWW9. The Hague, The Netherlands, May 2000.