Social Navigation: Modeling, Simulation, and Experimentation

Mark O. Riedl and Robert St. Amant
Computer Science Department
North Carolina State University
Raleigh NC 27695 USA
+1 919 515 7938
moriedl@eos.ncsu.edu, stamant@csc.ncsu.edu

ABSTRACT

The term *social navigation* refers to the process of seeking social interaction as a source of navigational support. In this paper we present a computational model of social navigation as an extension to an existing conceptual, non-computational model of general navigation. We argue that such models are important in designing effective shared environments for information navigation. We describe support for the model with a simulation environment for social navigation and an analytical model that addresses some of the quantitative implications of social interaction on the process.

Categories and Subject Descriptors

H.5 [**Information Interfaces and Presentation**]: Hypertext/ Hypermedia - *navigation;* User Interfaces - *theory and methods*.

General Terms

Design, Human Factors.

Keywords

Social navigation, simulation

1. INTRODUCTION

Almost all real-world activities at some point involve interaction with others or are influenced by the presence and opinions of others in a social setting [7]. Navigation, while studied extensively as an individualistic activity, is no exception. Navigation can be understood as situated action where an agent is embedded in the surrounding environment [12]. The agent can sense navigation-related information in the environment or can act upon knowledge stored in a cognitive map, where a cognitive map is a spatial description of the environment acquired through observations of the environment over time; it is used to find routes to the agent's goals [13, 11]. These two sources of information, the environment and stored knowledge, are complementary, with perception giving the agent access to information concerning immediate action, and cognitive maps allowing the agent to look ahead at possible future states rather than act only on local information [3].

AAMAS '03, July 14-18, 2003, Melbourne, Australia.

In accounts of human navigation in real-world environments, a more complex picture of the process emerges. A study of route finding through unfamiliar city roads shows that the vast majority of people choose to interact socially with others in order to acquire route-finding information [9, 10]. Other studies support the concept of social communication as a preferred form of knowledge acquisition in the face of uncertainty [16]. Navigation is a social, and often times collaborative, task [11].

The term *social navigation* refers to the process of seeking social interaction as a source of navigational support [6, 8, 4, 5]. While there is some debate whether social navigation is distinct from navigation in the abstract, social navigation can be distinguished by the practical tools used to solve the navigation problem [1, 18]. Social navigation is characterized by the use of other people's experiences in order to acquire knowledge for navigation [14] in addition to affordances for action that can be perceived in the environment. These experiences are acquired by interaction with others through a variety of social media, including computer-mediated communication systems, and are integrated into one's own cognitive map. Social navigation characterizes task-oriented activities in physical and information spaces that are populated by other people and agents. Practical examples of social navigation include finding a location in a city and finding products that meet desirability requirements in a shopping center or e-commerce Web site.

This paper presents a model of social navigation as an extension to an existing model of general navigation and describes the model in computational terms. This description is followed by applications of the model and a simulation of navigation that addresses some of the quantitative implications of social interaction on the process. This work extends and refines an earlier, preliminary model based on the same concepts [14].

2. A MODEL OF SOCIAL NAVIGATION

A recent conceptual model of navigation, due to Spence [17], treats navigation as the creation and interpretation of an internal mental model. The model contains four stages of processing, as shown in the central cycle of Figure 1: browsing, modeling, interpretation, and revision of the browsing strategy. An agent begins with the navigational goal of visiting a specific state. In the domain of navigation, the current state is the local environment and operations move the agent from one locale to the next. During the *browsing stage*, the navigational agent senses the surrounding environment and registers what is referred to as the *environmental content*. The environmental content is what the environment has to offer the agent in terms of the navigational task. Once the environmental content has been registered, the *modeling stage*

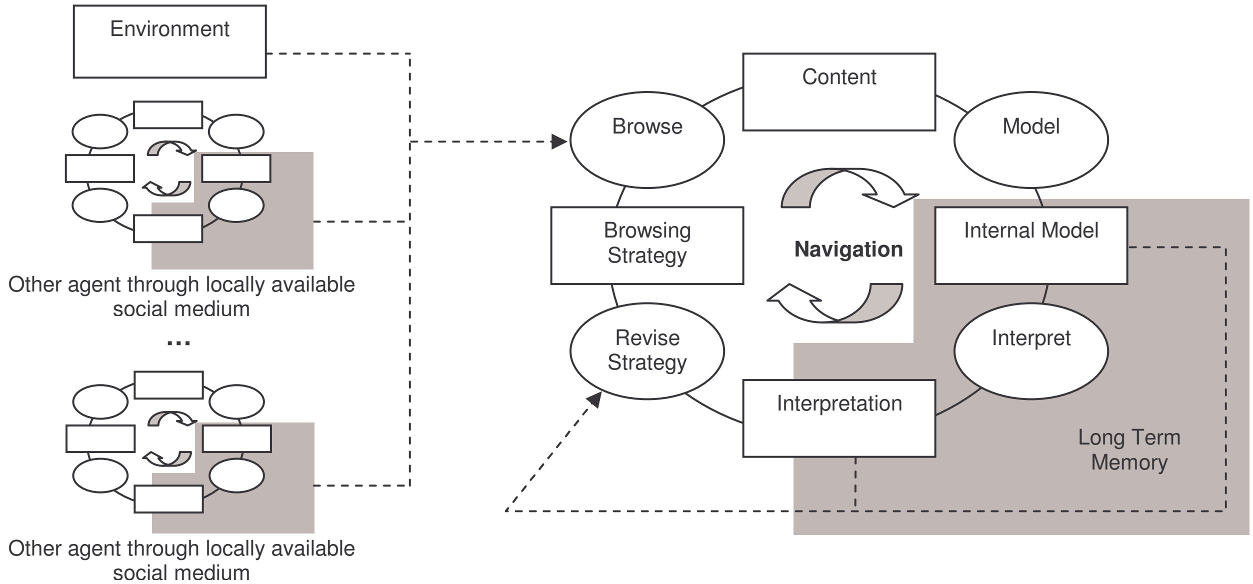

Figure 1. Decision-making in Spence's framework for general navigation

takes place in which the content is used to build an internalized model of the local surroundings and to add to the internalized global model of the environments that have been visited and will be visited. During the *interpretation stage*, the internalized models, both local and global, are analyzed to determine whether the current locale meets the criteria for successful navigation and, if not, how much farther the goal is thought to be. Finally, during the *revision stage*, the internalized model of the environment and the interpretation are used to determine a *browsing strategy* – what the next best step, or series of next best steps, will be for the navigation task. The browsing strategy is a task-oriented plan for interacting with the environment in order to achieve progress towards the navigational goal. Executing the browsing strategy causes movement, delivering the agent to another locality where it can begin the iterative process again by browsing the environmental content.

Spence's model, as it stands, is not sufficient to deal with the complexities of situating a navigational agent in a social environment. While at the most general level, navigating in social environments is still fundamentally a task of sensing, planning, and execution, the possibility of social interaction as a source of navigational information introduces the complexities of inter-agent communication. Traditionally, navigation is viewed as the interaction between the agent and the environment; all navigational decisions stem either from observations of the environment or knowledge in the agent's memory. Social environments yield a third possible source of navigational knowledge: other agents in the environment who have previous experience. The agent can choose to interact with other agents socially in order to learn from their experiences and thus delegate the decision-making process to another agent instead of relying strictly on the environment or personal memory. Social navigation is no different from navigation in a social vacuum, in the abstract, but what changes when an agent performs in a social environment is the richness of sources of decision-making knowledge that can be drawn from.

In any given state, the agent is tasked with updating its browsing strategy during the revision stage. This is tantamount to decision-making using knowledge about the local and global environment. We extend Spence's navigation model by re-formalizing the way in which environmental content is encoded into the internal model and how that model is interpreted in order for the agent to make a decision about appropriate browsing strategy. Our extension to Spence's model is as follows. The three sources of decision-making information—memory, the environment, and other agents—are collected into the internal model of the environment. It should be noted that information from other agents is stored in those agents' long-term memories and is accessed through communicative actions. Since the information from other agents is elicited through communication, those agents do not need to be locally situated, as long as the social medium (e.g. a phone booth or computer terminal) is situated in the local environment. Once formed, the internal model and its interpretation are used to revise the browsing strategy. Information from the agent's memory can be interpreted as a pre-existing model of the navigational task held by the agent before the navigational task begins. External knowledge, both from the environment and from others, is acquired during the browse stage and appended to the internal model. Therefore, when the strategy revision stage comes, the choice of action is a matter of choosing the best response, based on heuristic pattern matching with the internal model stored in memory. Figure 1 shows how the stages of Spence's model interact with long-term memory to form the decision-making process of navigation. The gray region is long-term memory. The internal model of the environment, which is stored in long-term memory, is interpreted with the use of heuristics. The internal model plus the interpretation form the basis for decision making.

The decisions involved in formulating and revising a browsing strategy are related to those faced by an agent that performs interleaved planning and execution. Knowledge about the current state of the agent is considered and a decision about the next operations to be performed is made. However, due to the nature of the navigation task and uncertainty in the environment, the agent often only has partial knowledge of the solution path. During the revision stage, the agent must decide whether it is better to stay and attempt to refine its strategy or to proceed with its partial results. In the latter case, a significant motivation for execution before the reasoning process is complete is to move the agent into a state where information is available and further processing is possible. This decision is encountered in a number of different contexts; it is sometimes referred to as an exploration/exploitation tradeoff.

The agent can make a rational choice in this situation by comparing the relative values (or cost) of its options. In real-world environments that support social navigation, revising a browsing strategy is almost entirely dependent on information gathering and interpretation. The value of reaching a given state is thus based on the navigational information that it provides – information relevant to reaching the agent's navigational goal. Equivalently, the information gained in a state is evaluated by the extent to which it reduces future navigational cost. This information is obtained from the local environment, from memory, or from others' memories. In any state, the agent can compare the cost of continuing to revise its strategy, based on locally available information, against the cost of moving to another state for further information. (Dean et al. [3] refer to the latter type of actions as point-to-point traversal tasks.) In other words, navigation is called for when the strategy revision process has completed or when the current state does not yield enough information for the agent to generate further profitable revision.

There are two sources of uncertainty in the agent's evaluation, arising from the correctness of the information gained in any state, and the potential difficulty of reaching that state to obtain the information. For example, I may know that a colleague has had information in the past about where to retrieve documents that I need, but I also know that my colleague is somewhat forgetful, rendering his current information suspect, and that he is rarely in his office, which means that at any given time he may or may not be reachable for this information. In dealing with this uncertainty, an agent can rely on estimates of information quality and the probability of reaching specific states, in order to compute expected costs for navigation actions.

To summarize: Operating under the navigation model, the agent builds a representation of the environment and stores it in long-term memory. This representation includes what the agent knows about the environment through which it is navigating, what the agent believes it can further learn from the environment, and what sources of social media are available. The nature of the agent's internal representation is such that cost estimates of revising the browsing strategy through communication through social media are quickly determined. Along with the cost estimates are also estimates of message uncertainty. Space limitations preclude a more thorough description of how the internal representation is used to determine these estimates (but see [15] for more detail.)

These value and uncertainty estimates quantitatively comprise utility values of different strategies for revising the browsing strategy. Comparing the utility for revising the strategy to the utility of executing the partial plan for navigation provides the agent with the foundation for rational navigational behavior in the presence of social media. We assume that the agent always has a partial plan of at least one step. This is the next locale the agent should navigate based solely on local environmental cues. This step may or may not be optimal. The agent may also contain a more substantial partial plan, which is often desirable. We also assume that the estimates the agent makes of cost and uncertainty with regards to both plan revision and plan execution are not guaranteed to be accurate and will reflect internal biases. Our model of social navigation can be expressed in decision-theoretic terms and is equivalent to one specific case of congregating multi-agent systems [15]. A congregation [1] is a confederation of agents that co-locate in order to benefit from the abilities other agents provide. In the case of social navigation, congregations tend to be small and short-lived, forming where the environment facilitates co-location or provides social media through which to communicate.

In the next section we discuss a simulation based on this model, to explore its implications for social navigation in specific situations.

3. SIMULATION AND ANALYSIS

To evaluate our model of social navigation, we have developed an abstract simulation to test how social awareness affects navigation. The simulation should provide a reasonable approximation of agent behavior when confronted by a variety of environments and given various preferences for social media usage. By analyzing the simulation results, we will be able to determine whether the model is capable of generating reasonable behaviors in navigational agents.

The simulation does not model specific agents or specific environments, but instead operates on *classes* of environments and agents, defined by the parameterization of each. The abstraction in this design allows us to analyze the broad patterns of behavior in a large variety of environments without getting bogged down by the details of implementing and testing real agents and the design and construction of real environments. A more complete study can be performed by developing actual social navigation agents, as supported by the environment described in Section 4.

This simulation is intended to provide insight into behavioral patterns of agents that use the model of social navigation proposed above. Through the simulation, we can vary the environment through which a simulated agent would navigate and aspects of the simulated agent itself to determine how different combinations of parameters will affect overall performance. We expect to see patterns of behavior that are plausibly explained by our model and by anecdotal evidence in the social navigation literature. Additionally we would like to know what features of the environment and of the social media used in social navigation have effects on agent performance when navigating through unfamiliar spaces? Specifically, under what conditions are social navigational strategies appropriate and when should more traditional navigation strategies be adopted?

3.1 Simulation Design

The simulation is a way of quantifying the ways in which social interaction affects navigation. The simulation engine assumes an arbitrary environment in which navigation can be represented as following links through an undirected graph. At each node in the graph, one or more out-link is optimal in that it is part of the path that will take the agent from the current node to the goal in the shortest number of steps. There can be several out-links that are part of sub-optimal, but valid, paths to the goal. Yet other out-links, referred to as "dead ends," take the agent to nodes that are not part of any path to the goal unless backtracking is performed. The complexity and uncertainty of the environment can be parameterized by six variables, independent of the connectivity of the graph. For each of these three classes of out-links, there is an associated cost of following an out-link and a probability that that out-link class will be chosen. In a perfectly ambiguous environment (e.g., some mazes), every out-link has the same probability of being chosen. If a maze is composed of four-way intersections, then the probability of choosing the optimal out-link is 1/3 (we do not consider the link from which we arrived at the

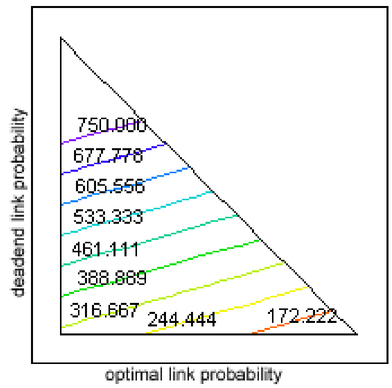

Figure 2. Constant D_{final} contours by $P_{optlink}$ and $P_{deadlink}$.

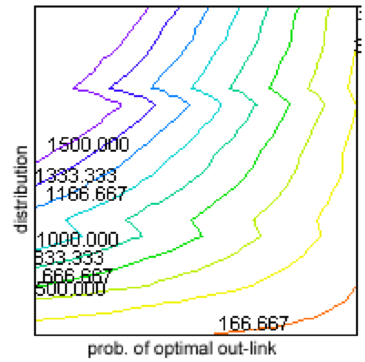

Figure 3. Constant D_{final} contours by $P_{optlink}$ and D_{social}.

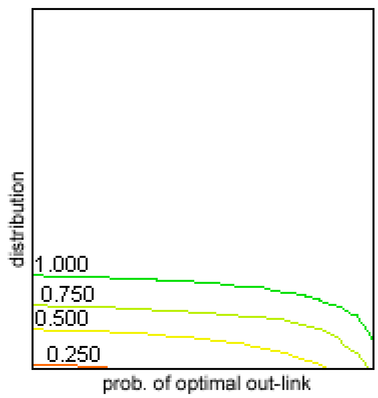

Figure 4. Constant R_D contours by $P_{optlink}$ and D_{social}.

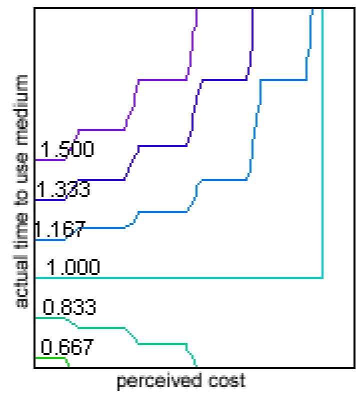

Figure 5. Constant R_D contours by C_{social}, and T_{social}.

current node). In a less ambiguous environment, this probability increases. Given the probabilities of choosing a class of out-links at any given node in the graph and the optimal length from start state to goal state, we compute the expected number of nodes visited using a binomial distribution. This expected number of nodes visited gives us a baseline for which to compare results using social navigation. In summary, the navigation environment is parameterized by six values:

- D_{opt}, the optimal distance from starting point to goal,
- $P_{optlink}$, the probability of choosing an optimal out-link in any state,
- $P_{deadlink}$, the probability of choosing a dead-end,
- $C_{optlink}$, the cost of following an optimal out-link,
- $C_{sublink}$, the cost of following a sub-optimal link, and
- $C_{deadlink}$, the cost of following a dead-end link,

The probability of choosing a sub-optimal link is left implicit, since this value plus the sum of $P_{optlink}$ and $P_{deadlink}$ must be one. $P_{optlink}$, $P_{sublink}$, and $P_{deadlink}$ implicitly capture the ease through which an agent can navigate a network of nodes without specifying why it is easy or hard. For example, $P_{deadlink}$ may be low when the environment contains many navigational cues, such as landmarks. A more specific method of analyzing the impact of environmental attributes is described in Section 4.

Socialization during navigation is parameterized by four additional variables:

- D_{social}, the distribution of social media, as measured by the average number of nodes the agent must visit before encountering another instance of a social medium,
- C_{social}, the perceived cost of using social media,
- T_{social}, the actual time it takes to use the social medium, and
- L, the number of steps generated through illocution.

Social media can be other people, phones, email, or any communication technology through which illocution can occur. On a university campus, the distribution of other people through the environment might be quite high. In a city environment, telephones are distributed such that one can be found every few miles. Alternatively, email might have a very low distribution because publicly accessible computer terminals are rare.

Following the model of social navigation, the simulated agent decides at each stage whether to use social navigation or to reactively search for the goal. If reactive search is chosen, the simulation engine computes the expected number of nodes visited until the agent must make another decision. While the model calls for a decision to be made at every node, in practicality the decision will not vary until something in the environment has changed significantly, which is based on the distribution of social media in the environment. If social navigation is chosen, the simulation engine computes the expected number of nodes visited before reaching an instance of the social medium. This is exploratory navigation that is not directly goal-related. The number of plan steps received through illocution is then subtracted off the total distance to the goal as goal-directed navigation; we assume that once directions are received, no error in navigation is made while the plan is being executed. The cycle of planning and execution is repeated iteratively until the goal is reached.

The simulation engine allows for only one type of source of social medium at a time, although the model is not limited in this way. In more complex environments, different remote sources of information might not overlap in the portions of the navigation environment they are knowledgeable about, which means that not all social information sources are equivalent. For the purposes of the simulation we assume that remote transactions cost the same and that one always exists that can assist the navigational agent as long as some instance of a social medium is locally available.

The simulation engine computes five values as results:

- D_{final}, the distance the agent is expected to travel,
- T_{final}, the elapsed time the agent is expected to arrive at the goal,
- R_D, the percentage of the baseline (without social navigation) distance traveled when social navigation is used,
- R_T, the percentage of the baseline elapsed time when social navigation is used,
- and N_{social}, the number of times the agent chooses to interact socially.

Expected distance and expected time traveled are two commonly measured variables in navigation research from which we can deduce the level of difficulty navigation in the parameterized environment will pose the agent. The percentage measurements, R_D and R_T, tell us whether social navigation is more efficient in certain circumstances, how much more efficient social navigation is, and, as other parameters change, how much slower the distance and time traveled using social navigation changes with respect to distance and time traveled without social navigation. N_{social} is redundant because it grows proportionally with D_{final}, but turns out to be important when analyzing certain anomalies observed in the other measured values. It should be noted that while the behavior of the simulation is expressed in terms of turn-based agent decision-making, the output values are actually computed using parameterized mathematical formulae.

3.2 Simulation Procedure

The simulation procedure was purely exploratory. Of the ten parameters to the simulation engine, two were chosen at a time and varied across a range of values while all other parameters are held constant. For example, D_{opt} was varied from 100 to 400 by a step size of 20; $P_{optlink}$ and $P_{deadlink}$ were varied from 0.1 to 1.0 by 0.1; $C_{optlink}$, $C_{sublink}$, and $C_{deadlink}$ took on ratios of 1:2:4, 1:3:10, and 1:5:15; and so forth, all pairs varying independently. The dataset generated by such a pairing can be examined for any interactions between the two parameters. There are 45 possible M-by-N analyses that can be made, although not all combinations prove to be useful. Of the M-by-N analyses that were chosen, each one was run several times with different sets of constants in order to assess whether there are any additional interactions. Different sets of constants were used and the data set was regenerated in order to assess whether there were any 3-way interactions. Data sets were graphed in various ways and observed for interesting patterns. Initially, the simulation engine was run without any of the social navigation parameters. This enabled us to ensure that there were no unexpected patterns that arose from various combinations of environmental parameterizations as well as provided baseline patterns that we could contrast to patterns involving social navigation.

3.3 Simulation Results

Running the simulation engine without the four social navigation parameters gives predictable results.[1] As the probability of choosing the optimal out-link decreases, the expected distance and time to reach the goal increases proportionally. The probability of choosing a sub-optimal out-link and the probability of choosing a dead-end link determines which of the respective cost parameters dominates the growth. In Figure 2, each contour represents the (linear) relationship between these two probabilities, $P_{optlink}$ and $P_{deadlink}$, for a constant level of distance traveled, D_{final}. As described above, D_{final} increases in a northwesterly direction over the contour plot, as $P_{optlink}$ decreases.

With this baseline established, we can look at the effects that social interaction has on navigation. Social navigation decreases the expected distance the agent must travel to reach the goal, but only in certain circumstances. When conditions are right, the use of social interaction *increases* the expected travel distance, as indicated by $R_D > 1$. The remainder of this section is broken into a discussion of the parameters that cause monotonic increases in the four dependent variables and the discussion of the parameters that cause non-monotonic increases in the dependent variables.

Once we compensate for the integer settings of N_{social}, the monotonic patterns can be summarized as follows.

- The distribution of social media throughout the environment, D_{social}, has a significant effect on the distance the agent will travel, D_{final}, as shown in Figures 3.[2] As the distance the agent must travel to reach other social media increases, the overall distance to travel to reach the goal will increase. In fact, D_{social} does not need to grow very large at all before social navigation is less efficient than blind search. The only tempering factor is the uncertainty of the environment. When $P_{optlink}$ is close to 1.0 (i.e., when the uncertainty of the environment is very low), on the right hand side of the plot, the distribution of social media has a reduced impact on D_{final} because fewer errors are made while the agent is trying to congregate.

  As the environment becomes less uncertain, however, the usefulness of congregation is reduced (unless the usefulness of congregation, L, increases proportionally, which is not addressed here.) The basic pattern is shown in Figure 4. When uncertainty in the environment is high, on the left hand side of the plot, low distribution of social media throughout the environment can result in the agent following

[1] A more extensive discussion of all the results in this section can be found in a technical report [15].

[2] The irregularities in this graph are due to N_{social}, which increases in a step pattern as the number of decision points increases, only taking on integer values.

a lengthy path. However, in an environment with low uncertainty, on the right hand side of the plot, any navigation is extraneous because reactive search can reveal the path to the goal just as easily or even more easily if the distribution of social media is low.

- The perceived cost of social interaction, C_{social}, has little effect on distance unless C_{social} is very large. When the cost of social interaction is very large, then the number of social interactions, N_{social}, quickly drops to zero. It appears that perceived cost is primarily a factor for choosing between the best source of interaction and best social medium through which to conduct the interaction. In this simulation, which only looks at the possibility of one social source at a time (assuming the given source is the best), perceived cost has little effect.
- The actual time to use an instance of the social medium, T_{social}, predictably causes an increase in the overall time to reach the goal, T_{final}. The actual time required to complete illocution is not factored into the utility of planning because the agent does not know this value. The only defense the agent has against high and unwieldy actual transaction times is to estimate the time required and incorporate that into its perceived cost of the social medium. Ideally, we would like to see perceived cost have some effect on T_{final} as T_{social} is varied. In fact, we do find that, in general, if actual time T_{social} is low, the agent benefits from a low cost to social media usage C_{social}, as shown on the left hand side of the plot in Figure 5, which shows contours for R_D, the gain/reduction in distance traveled when social navigation is used. Conversely, it is also apparent that if actual time is high, the agent benefits from a high value of C_{social}. Only when T_{social} and C_{social} are inversely related, in the lower right part of the plot, do we see poor performance from the simulated social navigation agent.

The non-monotonic observations, which we will summarize without plots, show complex interactions between two or more parameters that would not be immediately obvious without generating large data sets from the simulation. The most interesting interaction occurs between L, D_{social}, and C_{social}. The first thing that is noticed is that if $D_{social} > L$, then $R_D > 1$; i.e., when the number steps between instances of social media is greater than the number of steps generated by illocution, then distance traveled with social navigation is greater than without.

This observation makes sense, because if D_{social} is large and L is small, then the agent is going to expend more energy getting assistance than it gets back from obtaining assistance. The value L can be thought of as the quality of responses through illocution. A small value of L means that the help the agent is receiving is not advancing it greatly towards its goal. The value of L can be attributed to the knowledgeableness of social guides in the environment and can also be attributed to the expressiveness of the communication mediums available. Aside from the quality of social contact, we assume that social contacts do not lie. The non-monotonic relationship is evident; as L increases the distribution of social media has a decreasing effect on D_{final} and R_D. This lessening of effect occurs more rapidly than linear and, at some point, actually reverses itself. The non-linear decrease and reversal is non-intuitive, since one would guess that a large L could only be a benefit to a navigation agent. However, upon closer inspection, we observe that when the distance the agent must travel to use an instance of a social medium is greater than the number of steps acquired through illocution, social navigation is less efficient than reactive search; the agent expends more energy reaching the social medium than it receives as reward for illocution. When social navigation is beneficial, increasing L is beneficial up to a certain point – where L is slightly smaller than D_{opt}. After this point, social navigation can actually be harmful to the agent's performance unless C_{social} is great enough to discourage the agent from using social navigation when the goal is within reaching distance; a reactive strategy will be more beneficial. If fact, other simulation results show that increasing C_{social} does diminish the non-linearity because the agent will be less likely to engage in social navigation when close to the goal.

In general, these results are consistent with independent observations that human navigators, when faced with navigation in an unfamiliar environment, will adjust their goals to seek out advance information through social interaction, even when such goals took them off of the direct path [9]. Human navigators prefer social navigation to map reading because social instructions narrow the space and provide contextually relevant information that cannot be acquired easily without prior experience [10]. Other investigators [4, 8, 9, 10] cite phenomena where human navigators perform non-goal-related navigation in order to ease the cognitive demands of an unknown environment.

4. MUNE

The goal of our work is to understand real environments for social navigation, in which people communicate with one another for information. Our model and simulation are a step toward this goal, in that they can help us evaluate our general intuitions about the nature of social navigation environments. Nevertheless there remains a gap between theory and practice. We need an environment in which we can evaluate detailed design decisions that may influence social navigation. For example, the model and simulation say little about how to position instances of social media of different types throughout the environment (e.g., public telephones and information desks scattered through a shopping mall.) That is, to fully test our model of social navigation, we must be able to implement specific environments rather than rely on a simulation of abstract environments. If one understands how synchronicity, directness, and social presence [14] affect the perceived cost of using social media for a given population, one can select the social media with the most desirable traits to place in the environment. Similarly, if social media do not exist that already have the most desirable traits, new social systems can be designed that have the most appropriate values of synchronicity, directness, and social presence. The evaluation of specific environments can indicate whether social navigation is sufficiently supported and, if not, how the environment can be redesigned to better support user tasks and preferences.

To evaluate a social environment for navigability, we need a better way to simulate an arbitrary navigation environment that reflects relevant details about the information in the domain. The environment to be navigated with or without social navigational aid could be semantic, such as a database or the World Wide Web, or spatial, such as a MUD, CVE, or a real city. In order to simulate as wide a range of environments as possible, we have created MUNE, a Multi-User Navigation Environment based on the concept of the MUD. A MUD is a textual virtual environment that users can navigate through and interact with other users. Where a MUD emphasizes identity and interpersonal relationships among human users, MUNE is designed to be a navigational test bed for software agents. Instead of natural-language descriptions of "rooms," MUNE uses an expressive communication protocol that is easily parsed by software agents. Beyond navigation and the basic face-to-face interactions that can occur in a MUD, MUNE allows objects – referred to as "social media" – to be scripted into the world that can facilitate communication such as phones, email, etc.

MUNE, unlike most other MUDs, does not assume a geographical world; its "rooms" can be designed to look like nodes in a semantic environment, such as the World Wide Web or a relational database. Instead of links between rooms being described as "east" or "north," the links can be described as hyperlinks or other such mechanisms for navigation. MUNE's flexibility stems from its simplicity and its extensibility. MUNE only knows about room descriptions and links between rooms and user locations but allows for additional world descriptors. Additional world descriptors can be included in the environmental description which MUNE does not attempt to understand. Instead additional descriptors are passed to the software clients who are responsible for their interpretation. For example, in the real world, a person standing on top of a hill is able to see for miles around. While it is possible for the room representing the top of the hill to be described in human-readable form in such a way as to give the appearance of being able to see for miles around, a software agent would need a more formal description. Such formal descriptions can be coded into the hilltop room as a world descriptor extension. Figure 6 show a conventional, spatial world based on geographic terrain in MUNE. The LINK descriptor is a basic field that tells MUNE which actions are legal in a given room; the LOOKAHEAD descriptor is an extension allows

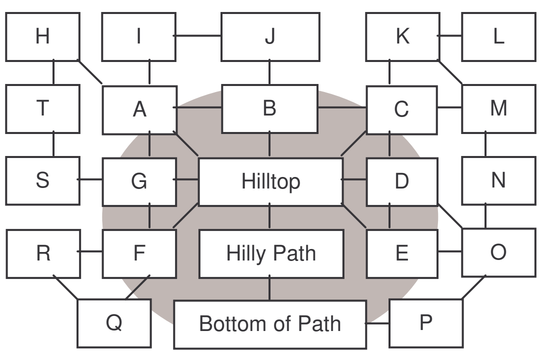


```
LOCATION
   ID Hilltop END-ID
   DESCRIPTION
      You are on a hilltop…
   END-DESCRIPTION
   LINK
      COMMAND south END-COMMAND
      DESTINATION Hilly Path
         END-DESTINATION
   END-LINK
   …
   LOOKAHEAD Bottom of Path …
      END-LOOKAHEAD
   LOOKAHEAD Hilly Path …
      END-LOOKAHEAD
   LOOKAHEAD D … END-LOOKAHEAD
   LOOKAHEAD E … END-LOOKAHEAD
   …
END-LOCATION
```

Figure 6. A spatial MUNE world with description

software agents to see one or more rooms ahead. The gray circle designates what an agent can see from the hilltop room.

Inside the MUNE world, software agents can be given goals to navigate to. MUNE is agnostic with respect to agent design; we have tested agents in MUNE that employ depth-first search, reactive search, and search using social navigation. MUNE is also flexible enough that a large variety of environments can be simulated. Social navigation agents can be tweaked to behave according to certain preferences and preconceived notions about social interaction that real users of the simulated environment might have. The navigability, both in terms of general navigation strategies as well as social navigation strategies, of the environment can be evaluated by measuring and comparing the performance of social navigation software agents, conventional navigation software agents, and human users in the simulated MUD environment. Human users can interact with the simulated MUD environment through client software that converts MUNE world specification format into human readable room descriptions, although human users may find the textual nature of MUDs more cumbersome to interact with than a visual spatial environment. Because the leanness of textual descriptions is limiting to human users and not to software agents, it may not be possible to directly compare the performances of human users and software agents operating within MUNE, unless the conditions are well controlled for or more expressive client programs are developed. Once performance measures have been collected, the navigation environment can be easily adjusted and social media can easily be re-distributed until a desirable level of performance is reached.

Although we have implemented some simple environments in MUNE, we have not yet carried out an extensive evaluation. Nevertheless our preliminary work leads us to believe that it can play a complementary role to the simulation described in Section 3. As is generally the case in building interactive systems, automated tools can inform and improve the preliminary design of a software environment, but eventually users must become involved in order to evaluate the details. MUNE is designed for this purpose, to support a smooth development path from abstract navigation agents moving through an abstract simulation, through implemented software agents moving through a more detailed environment, further through real human users interacting in the same detailed environment, and finally to a deployed social navigation system that incorporates the lessons learned at each of the previous stages.

5. CONCLUSIONS

The simulation of navigation in social environments has shown to produce reasonable results and shows that, in most circumstances, social navigation results in superior navigation performance over strictly reactive approaches. One would assume from the commonality of social navigation in human behavior [8] that social navigation is strictly more efficient than asocial navigational practices, such as reactive search. The simulation, however, revealed the following situations in which using a social navigation strategy actually detracts from navigation performance:

- The environment is so uncertain that reaching the social medium incurs large penalties.
- The reward for pursuing social navigation does not dominate the cost of reaching the nearest social medium.

- The perceived cost of using a social navigation strategy is so low that the agent pursues social navigation when a reactive approach would serve better.

Social navigation assumes helpfulness and is bound to fail when this assumption does not hold because of environmental conditions.

We have also seen that an agent that can reasonably estimate the cost of using social navigation will perform much better than an agent that cannot. The estimate, however, does not have to be accurate for performance to be good, only that the agent estimates the cost to be high when the time to use the social medium is in fact high or that the agent estimates the cost to be low when the time to use the social medium is in fact low. It is fortunate that the agent only requires a reasonable estimate and not an accurate estimate because human use of heuristics to make decisions relies on approximate situation assessment.

Finally we have seen that social navigation is more beneficial in environments that are highly uncertain than otherwise. Within highly uncertain environments, benefit from social navigation comes only when instances of social media are widely available. In an uncertain environment, reactive search will cause the agent to make many mistakes that will be costly in terms of the distance the agent travels. Social navigation can reduce the number of mistakes, but only if the agent can reach social media without too much difficulty.

While the simulation has shown that our model of social navigation can produce plausible results, it remains to be evaluated in more detail. A more detailed analysis of social navigation is not within the scope of this research, but the framework for a more detailed experiment is already in place. The MUNE system can be used to extract more detailed behavior patterns from agents that use social navigation. Due to the nature of MUNE, these agents can be software agents, implementing the model of social navigation described above, or they can be human agents. We expect any detailed analysis to fall within the patterns observed through our more limited simulation. Further work is needed in order to determine if the cost framework for choosing between planning and execution can sufficiently integrate all the factors that an agent might consider when choosing between planning – using the environment, memory, and others' experiences – and execution.

6. ACKNOWLEDGMENTS

This effort was supported by the National Science Foundation under award 0083281. The information in this paper does not necessarily reflect the position or policies of the U.S. government, and no official endorsement should be inferred.

7. REFERENCES

1. Brooks, C.H., Durfee, E.H., Armstrong, A. (2000). An introduction to congregating in multiagent systems. In *Proceedings of the Fourth International Conference on Multiagent Systems*.
2. Dahlbäck, N. (1998). On spaces and navigation in and out of the computer. In N. Dahlbäck (Ed.) *Exploring Navigation; Towards a Framework for Design and Evaulation of Navigation in Electronic Spaces*. Technical Report T98:01. Kista, Sweden: Swedish Institute of Computer Science, pp. 15-29.
3. Dean, T., Basye, K., Chekaluk, R., Hyun, S., Lejter, M., & Randazza, M. (1990). Coping with uncertainty in a control system for navigation and exploration. *Proceedings of the Eighth National Conference on Artificial Intelligence*, Boston MA, July 1990.
4. Dieberger, A. (1997). Supporting social navigation on the World Wide Web. *International Journal of Human-Computer Studies*, 46(6), 815-825.
5. Dieberger, A. (1999). Social connotations of space in the design for virtual communities and social navigation. In A.J. Munro, K. Höök, & D. Benyon (Eds.), *Social Navigation of Information Space* (pp. 15-32). London: Springer-Verlag.
6. Dourish, P. & Chalmers, M. (1994). Running out of space: models of information navigation. *Proceedings of HCI'94*, Glasgow, August 1994.
7. Dourish, P., (1999). Where the footprints lead: tracking down other roles for social navigation. In A.J. Munro, K. Höök, & D. Benyon (Eds.), *Social Navigation of Information Space* (pp. 15-32). London: Springer-Verlag.
8. Erickson, T. (1996). The World-Wide-Web as social hypertext. *Communications of the ACM*, 39(1), 15-17.
9. Gordon, D.A. & Wood, H.C. (1970). How drivers locate unfamiliar addresses – an experiment in route finding, Public Roads, 36(2), 44-47.
10. Höök, K. (1998). Route guidance issues; verbal versus map instructions, and route choices. In N. Dahlbäck (Ed.) *Exploring Navigation; Towards a Framework for Design and Evaulation of Navigation in Electronic Spaces*. Technical Report T98:01, Swedish Institute of Computer Science, Kista, Sweden, 1998.
11. Hutchins, E. (1995). *Cognition in the Wild*. Cambridge, MA: MIT Press.
12. Jul, S., & Furnas, G. (1997). Navigation in electronic worlds. *SIGCHI Bulletin*, 29(4).
13. Kuipers, B. (1982). The "map in the head" metaphor. *Environment and Behavior*, 14(2), 202-220.
14. Riedl, M.O. (2001a). A computational model and classification framework for social navigation. *Proceedings of the International Conference on Intelligent User Interfaces*, Santa Fe, NM, January 2001.
15. Riedl, M.O. (2001b). A computational model of navigation in social environments. Master's Thesis, Department of Computer Science, North Carolina State University, Raleigh, NC, June 2001.
16. Rosen, D.L. & Olshavsky, R.W. (1987). The dual role of informational social influence: implications for marketing management. *Journal of Business Research*, 15(2), 123-144.
17. Spence, R. (1999). A framework for navigation. *International Journal of Human-Computer Studies*, 51, 919-945.
18. Svensson, M. (1998). Social navigation. In N. Dahlbäck (Ed.) *Exploring Navigation; Towards a Framework for Design and Evaulation of Navigation in Electronic Spaces*. Technical Report T98:01. Kista, Sweden: Swedish Institute of Computer Science, pp. 73-8.