

The Impact of Communication and Memory in Hive-based Foraging Agents

Paul Schermerhorn and Matthias Scheutz
Human-Robot Interaction Laboratory
Cognitive Science Program
Indiana University
Bloomington, IN 47406, USA
{pscherme,mscheutz}@indiana.edu

Abstract—Many hive-based agents, such as some bees, rely on memory and communication to aid in foraging. The benefits of these abilities seem obvious, but it is unlikely that they are beneficial in every environment. In this paper, we present the results of experiments examining the effect that environmental structure can have on the utility of communication and memory for hive-based agents, finding that there are some environments in which they do not contribute substantially to the agents’ ability to survive.

I. INTRODUCTION

The wide presence of communication in the animal kingdom implies that it must be beneficial to creatures that use it. In particular, many colony and hive-based insects (bees [5], [7] and, some would argue, ants [8] are examples¹) use communication to aid in locating and gathering food resources. However, it is unlikely that communication is always beneficial to them, in particular if environmental conditions do not provide enough structure to make the information shared among agents through communication valuable enough to offset the costs of communicating. That is, if the food resources are sufficiently scattered, or resources are easy to find, communication is unlikely to provide substantial benefit. Indeed, experimental evidence indicates that, for honey bees in environments where resource clustering is low (e.g., in temperate climates), the amount of nectar gathered is not affected when the bees are deprived of their ability to reliably share information about resource locations. In contrast, though, in environments where food resources are clustered (e.g., in tropical climates), inhibiting communication does reduce the amount of nectar gathered [3].

To the extent that hive-based agents are explicitly cooperative, they should be ideal candidates for demonstrating the benefit of communication; because they maintain a community store of energy (e.g., for hive sustenance and procreation), there should be an advantage to sharing information about resource locations rather than keeping them private, since all hive members should be working together. Agents that do not share a centralized resource may well reduce their own fitness by sharing information about food sources [12]. Structuring the task so that agents are explicitly cooperative may “stack the deck” in favor of communication

being beneficial, but, as we show below, that is not sufficient to guarantee that communication is profitable.

Any analysis of the benefit of communication is incomplete without some consideration of its cost. This includes not only the cost of transmitting and receiving, but also the cost of whatever infrastructure is required in the agent architecture to support communication. For example, a hive-based agent must have mechanisms in place to store the locations of recently-encountered food sources in order to share them with colony-mates. This memory component is likely to be a substantial part of the cost of communication, unless it is beneficial in its own right.

This paper presents our investigation of the utility of communication and memory for hive-based colonies and the role environmental conditions play in determining it. Specifically, we explore the utility of information storing and sharing (i.e., memory and communication) in hive-based colonies by varying two dimensions of the environment: the distribution and quality of food sources. We examine two distributions: *random*, in which food items are placed randomly in the environment with no structure (similar to the distribution of resources in temperate climates), and *clustered*, in which food items are placed together in clusters at random locations in the environment (analogous to food distribution in tropical environments). Surviving in the random environment is easier than in the structured environment because hive agents are likely to find food regardless of the direction in which they set out. In clustered environments, there are large areas in which no food sources can be found. This should increase the value of the food locations stored in memory and transmitted during communication in cluster environments. The results of our simulations confirm that the value of food locations in random environments is likely not sufficient to justify the cost of memory and in fact, communication leads to a performance *decrease* in these environments, while in cluster environments both memory and communication are found to provide substantial benefits.

The following section describes the agent model and environment in detail and provides an overview of the experimental design used in this study. Section III presents the results of the simulation experiments. Some implications of these results are discussed in Section IV. In Section V, we provide a brief overview of some related work. Concluding remarks are presented in Section VI.

¹The study described here does not consider simple indexical signaling (“I see food here”), nor stigmergy (e.g., pheromone trails), but rather more complex communication requiring representations of food source locations.

II. METHODOLOGY

This paper is an examination of communication and memory in hive-based colonies in the biological setting, focusing on *the effect environmental factors have on the utility of communication and memory*. Our approach is to use agent-based modeling to study how the behavior of simple agents is impacted by the distribution of food resources in the environment, contrasting the performance of agents with and without memory and communication mechanisms. In this section, we describe the agents and environments used in the simulations.

A. Agent Model

The agent model employed is highly idealized, compared to biological hive agents such as bees (a more realistic model would include separation of duties, etc. [14]). In particular, simplified models of memory and communication are used, as the particular details of how information is stored and transmitted do not affect the utility of either operation. The details of determining, storing, and maintaining food source locations are abstracted away, with agents simply storing the heading and distance relative to the hive. Similarly, no particular mode of communication (e.g., dancing) is specified; when agents return to the hive to contribute food, the location of the most recently perceived food source is communicated (by virtue of the angle to some reference point, such as the sun, and a distance from the hive, similar in spirit to the way bees communicate locations) directly with any other agents at the hive, which can then choose whether to target that food source. This is not a pheromone-based information transmission scheme; communication is explicit and only possible when both transmitter and receiver are at the hive. Clearly, real hive-based agents will have to come up with more realistic implementations of these mechanisms.

The hive agents are very simple reactive agents that always target the nearest perceived food source and move directly toward it when foraging. They gather energy by consuming food and use energy as they move throughout the environment; energy use (i.e., food consumption) is comprised of two components, *maintenance cost* $c_m = 1$ and *movement cost* $c_v = 4$. These costs are deducted from the agent's food stores every cycle (c_m) and when the agent is in motion (c_v), respectively. When an agent has collected a sufficient surplus of food, it returns to the hive to contribute the surplus to the hive's food store, consuming resources from the food store on the way back. Agents can only collect a fixed amount of food per cycle ($C = 50$) and are limited in the total they can carry ($T = 20000$); when an agent's collected food level E reaches T , it returns to the hive. Each agent recalls the location of the food source it most recently visited, and if there is no food source within sensor range ($R = 100$), it instead targets the memorized food source and moves toward it. If there is no memorized food source (e.g., when the agent returns to a memorized location and the food source is gone), the agent begins foraging. Foraging behavior consists of moving in a random direction for a fixed number of cycles ($W = 200$). If no food source has been perceived after W cycles, the agent

makes a random turn (1-45 degrees in either direction) and begins again.

Formally, the agent model can be specified with three sets of rules corresponding to the three (broadly defined) functional roles that agents assume: *foraging and consumption*, *return*, and *home*. Foraging and consumption rules apply between the time an agent leaves the hive until it decides to return, during which time agents are searching for and gathering food:

- *Rule F1*: if no unoccupied food source is perceived or memorized and energy $E \geq T/4$, duty $D = \text{forage}$
- *Rule F2*: if no unoccupied food source is perceived or memorized and energy $E < T/4$, $D = \text{return}$
- *Rule F3*: if at least one unoccupied food source is perceived and $E \leq T$, move to nearest and consume it
- *Rule F4*: if no unoccupied food source is perceived but one is memorized and $E \leq T$, go to it
- *Rule F5*: if $E > T$, $D = \text{return}$

The hive's energy stores (H) are maintained by hive members bringing food back to the hive. The stores are used for two purposes: food and reproduction. When agents forage without success, their own energy stores are depleted. When an agent's stores drop below a critical value, it will return to the hive (per rule *F2*) to acquire energy there (if available). It can then recommence foraging and consumption. When an agent runs out of energy, it dies.

Return rules apply from the point at which an agent decides to return to the hive until it arrives there. If the agent is returning to contribute a surplus, it ignores any food it encounters, otherwise it will revert to foraging when food is detected:

- *Rule R1*: if $E \geq T/2$, continue to hive, ignoring food
- *Rule R2*: if $E < T/2$ and no unoccupied food perceived, continue to hive
- *Rule R3*: if $E < T/2$ and unoccupied food perceived, $D = \text{forage}$

Agents can only reproduce when they are at the hive and the hive's energy stores are sufficiently high (procreation energy $P = 48000$). If an agent is at the hive (either to contribute or consume food) and the hive's energy threshold for reproduction is met, that agent can reproduce. Reproduction is asexual, and no mutation is employed; all agents are identical. Hive rules:

- *Rule H1*: if $E > T/2$, contribute $E - T/2$ to hive store H
- *Rule H2*: if $E < T/4$, withdraw $\min(T/2 - E, H)$
- *Rule H3*: if $H > P$, procreate

Colonies implementing the above agent model demonstrate certain behavioral tendencies: at any given time some members will be foraging, wandering aimlessly throughout the environment looking for food. Others will be steadily moving between food sources and the hive and back, gathering food and bringing it back to contribute to the community store.

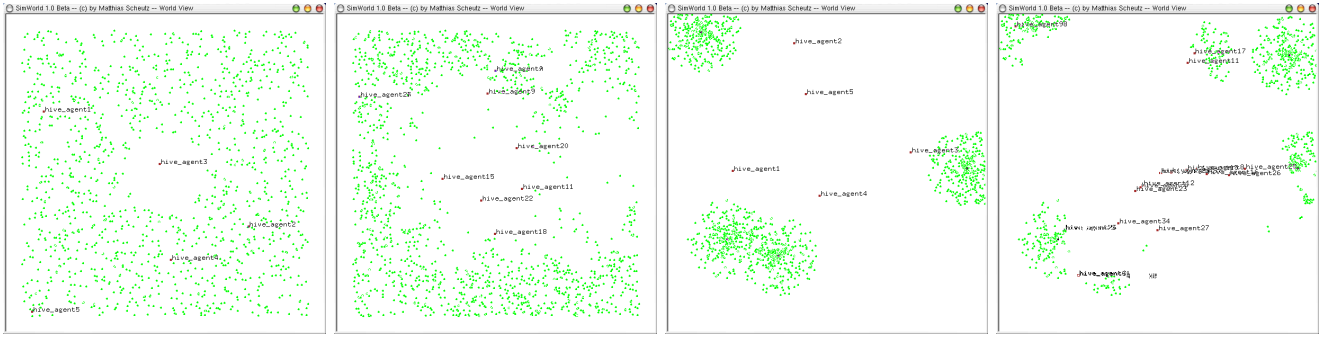


Fig. 1. Random environment at cycle 1 (far left) and cycle 10000 (middle left), cluster environment at cycle 1 (middle right) and cycle 10000 (far right). Hive members are labeled (hive_agent N), unlabeled dots represent food sources. The hive is located at the center of the environment.

B. Communication

Communicating agents add to the basic model the ability to share with other agents the location of memorized food stores. Communication occurs only near the hive. When an agent returns to the hive with food, it communicates the location where the food was found to whoever else is near the hive. Those other agents can then choose to target the new source (if it is closer than their own target, or if they are foraging and have no target).

Additional hive rules for communicating agents:

- *Rule H4*: if memorized food source M , broadcast its location
- *Rule H5*: if broadcast location B nearer than M or $M = \emptyset$, $B \rightarrow M$

The behavior of communicating agents is very similar to non-communicating agents. Because communication takes place only near the hive, there is no wide-scale change in the action patterns. At times an agent will not return to the food source it just left, instead choosing a closer source given by another agent. More often, foraging agents (i.e., those that have returned to the hive to withdraw food) take on a remote food source as a target. The communication mechanism meshes well with the agents' default behaviors: when agents are low on energy, they return to the hive. But they are low on energy just when they cannot find food. Thus, they are returning to the hive, and potential sources of information regarding the location of food, exactly when they need to. Note also that communicating agents are not charged extra cost for their added capability. Section IV discusses the role of cost in determining the net benefit of communication (and memory).

C. Memory

A third type of agent is introduced via a minor architectural tweak: the elimination of memory. These *memoryless agents* use the same ruleset as the non-communicating agents, effectively finding no food source location stored whenever a rule checks memory. Hence, foraging and consumption rules $F1$ and $F2$ will fire whenever no food source is perceived and the energy level requirements are met. Similarly, rule $F4$ never fires, because there is never a memorized food source location. Communicating memoryless agents are not modeled, as their behavior would be identical to non-communicating memoryless agents; rule $H4$ would never

fire, because no agent would ever return to the hive with a memorized location, and no agent would be broadcasting the location of a food source to trigger rule $H5$ (not to mention there would be no mechanism to store the information).

Memoryless agents behave similarly to (normal, with-memory) non-communicating agents, except for the memory-specific aspects. A non-communicating hive agent will return to its most recent known food source when leaving the hive or foraging in a region with no perceived food source. The memoryless agent begins each trip from the hive afresh, and must forage from scratch each time. In some environments, this could be a serious handicap. Finally, as in the case of communication, there is no cost levied for having and using memory.

D. Environment Models

There are two environments in which we test the utility of communication: *random* and *clustered* (see Figure 1). In the random environments, food sources are generated at random locations on average every 16 cycles. This creates a steady influx of energy into the environment without introducing structure. In the clustered environments, on the other hand, food sources are generated much less frequently (on average every 2048 cycles), and are placed at random locations in clusters of 128. The clusters of food sources are circular, 100 units in radius, and food sources are randomly placed within the cluster. Clustered environments introduce energy into the environment in spurts, with some structure. However, in all cases the amount of energy created on average is the same. Also, in all cases, food is generated outside of the agent's sensor range if located in the hive (i.e., agents cannot perceive food sources from the hive; they must remember or forage for it). The hive is located in the center of the environment.

E. Experimental Setup

The simulations reported below were performed in SWAGES, an artificial life simulation environment under development in our lab [13]. The simulation environment is a continuous 2D world which is limited to a 3200 by 3200 square region. Each experiment consists of 100 experimental runs in different randomly generated initial conditions. The same set of 100 initial conditions was used for each experiment in the same environment type, allowing us to




Fig. 2. Comparison of hive agent performance (average number of survivors at the end of an experimental run) in both environments.

compare directly between agent types. The results reported are averages over the 100 experimental runs that make up an experiment.

III. EXPERIMENTS AND RESULTS

We conducted extensive simulation experiments to explore the role of environment on the utility of communication. There are six basic configurations: memoryless agents in random environments, non-communicating agents in random environments, communicating agents in random environments, memoryless agents in cluster environments, non-communicating agents in clustered environments, and communicating agents in clustered environments. Each of these configurations was tested in environments where the food quality was varied from 200 to 3000 units of energy per food source, in steps of 200, for a total of 9000 experimental conditions, each of which consisted of 100 experimental runs (with different initial placements of the agents and food sources). The performance measure used here is the number of agents surviving at the end of a 10000 cycle experimental run (the reported values are the averages over the 100 runs).

The results of the experiments are presented in Figure 2. As predicted, the average number of survivors is greater in random environments than in clustered environments with the same net amount of energy. While the density of the food is lower, agents are more likely to come across food sources in random environments than in clustered ones, and hence less likely to die. Performance of communicating agents in random environments is similar to non-communicating agents and memoryless agents. A two-way 3×15 ANOVA was conducted for *agent type* (memoryless, non-communicating, and communicating) and *food quality* (200 to 3000) as randomized variables and *average survivors* as dependent variable for random environments. There were highly significant main effects for both *agent type* ($F(2, 4455) = 214.10, p < .001$) and *food quality* ($F(14, 4455) = 20625.22, p < .001$). The latter indicates, as expected, that as food quality increases the average number of survivors also increases. The former indicates that there is a difference between memoryless, communicating, and

non-communicating agents in random environments: non-communicating agents perform best, followed by memoryless agents, then communicating agents. Additionally, there is a highly significant interaction between *agent type* and *food quality* ($F(28, 4455) = 8.77, p < .001$). Tukey’s HSD test confirms that, despite the presence of the interaction, each of the three agent types performs differently than the others. The interaction points to an effect of the foraging strategy dependent on the amount of energy in the environment. Sharing information about the locations of food sources has the effect of focusing a subset of agents on a particular geographic region; when food quality is low, that region gets cleared of resources quickly, making it more difficult for communicating agents to survive and return to the hive with food contributions. However, having no memory does not allow the agent to return to food sources that may be near to the hive (e.g., when it happens that more agents have foraged in one direction than another).

Figure 3 shows the average net contribution per agent (i.e., the amount of energy contributed to the hive minus the energy withdrawn from the hive over the agent’s lifetime) for random environments. The three agent types contribute similarly, as might be expected from their similar survival performance. Net contribution is almost perfectly correlated with survival rate (Pearson’s $r = .997, p < .001$). When agents are able to contribute more to the hive, reproduction will occur more frequently, producing more hive agents to gather food resources for the hive.

The distance an agent has to travel to find a food source also can affect survival, as a longer return trip to the hive will lead to smaller amounts of energy to be contributed. Direct comparison of the average distance of a food source gathered by an agent is made difficult by the fact that, when there are more agents in an environment, food sources nearer to the hive are more likely to have been gathered, so an agent will need to travel further to find food. This can be seen in Figure 4, which shows the counterintuitive result that memoryless and non-communicating agents find their food sources further away from the hive than communicating agents, despite the worse performance of communicating

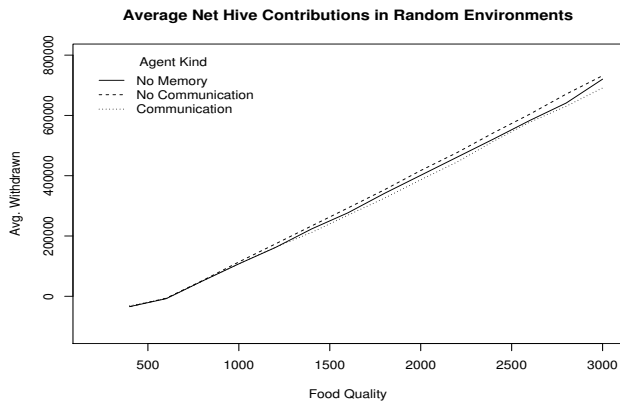


Fig. 3. Comparison of net hive agent contributions (average energy contributed less energy withdrawn) in random environments.

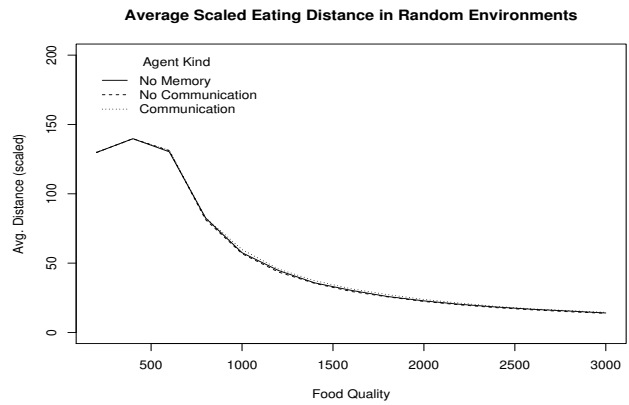


Fig. 5. Comparison of the scaled average distance to hive agent food sources in random environments.

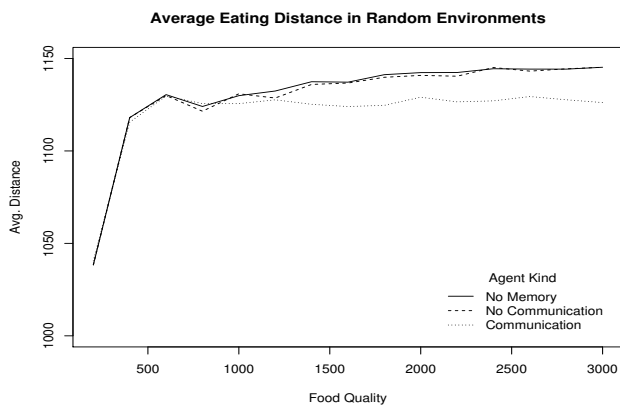


Fig. 4. Comparison of the average distance to hive agent food sources in random environments.

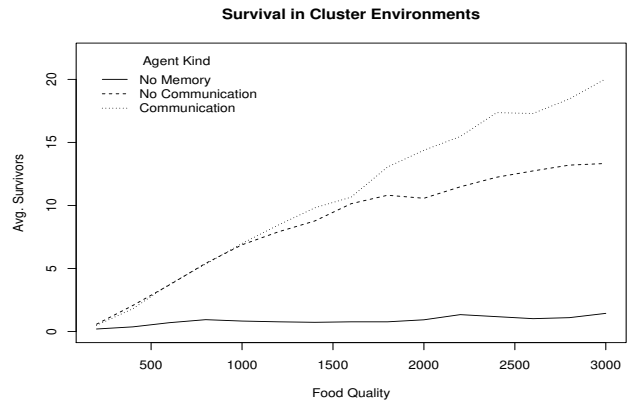


Fig. 6. Comparison of communicating and non-communicating agent performance in clustered environments.

agents. The scaled distance to gathered food sources, with the distance from Figure 4 divided by the *total* number of agents in the environment over the experimental run (not just the surviving agents) is shown in Figure 5. This shows that there is really very little difference in how far the three agent types must travel to find food sources once differences in the number of agents carried by the environment are taken into account. With the exception of the lowest food qualities, the curve in Figure 5 has the form of a power law. The values for the lowest food qualities exhibit different characteristics due to the interaction between the food quality and the agents' capacity for carrying food; when food quality is low, agents must spend more cycles foraging and collecting before reaching the threshold needed to trigger the return behavior.

Figure 6 zooms in on hive agent survival in clustered environments. The change is dramatic, with substantial performance differences between agent types. In these environments, an analogous two-way 3×15 ANOVA was conducted for *agent type*, *food quality*, and average number of *survivors*. Again, there were highly significant main effects for *agent type* ($F(2, 4455) = 573.05, p < .001$) and *food quality* ($F(14, 4455) = 52.33, p < .001$), and a highly significant interaction between *agent type* and *food*

quality ($F(28, 4455) = 13.12, p < .001$). Tukey's HSD shows that the performance of the three types is significantly different from each of the others. All three agent types start out with similar performance with very low food quality, reflecting the difficulty of surviving in these conditions at all, regardless of strategy. As food quality increases, both memory and communication provide an increasing advantage, with communicating agents beginning to outperform non-communicating (with-memory) agents as food quality surpasses 1000. Clusters are difficult to come by at random, so the extra information is very beneficial. When memoryless agents have to begin the foraging process over again after each return trip to the hive, it is difficult for them to survive individually, much less contribute back to the hive.

Figure 7 bears this out, showing the average net contribution of each agent type in cluster environments. Here we see that memoryless agents fail to break even on their contributions, remaining negative for all values of food quality. This indicates that they repeatedly returned to the hive and withdrew energy resources, while very seldom returning with energy reserves to contribute, having found a cluster on that trip. As in the random environments, average net contribution very closely matches agent performance

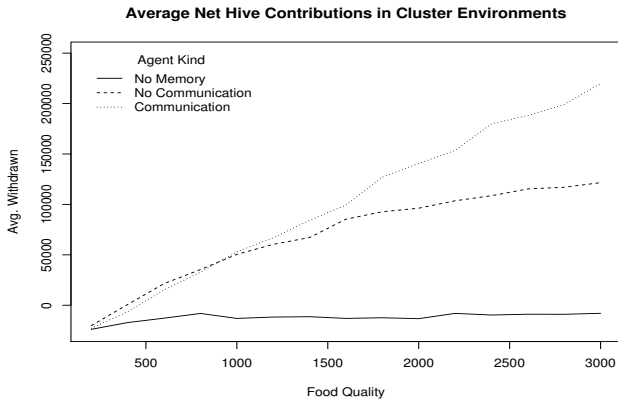


Fig. 7. Comparison of net hive agent contributions (average energy contributed less energy withdrawn) in clustered environments.

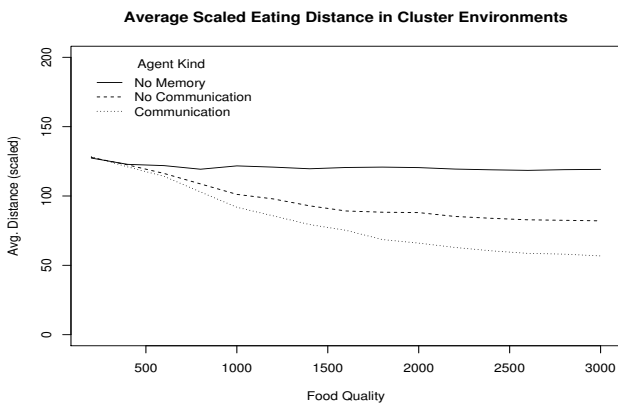


Fig. 8. Comparison of the scaled average distance to hive agent food sources in clustered environments.

(Pearson’s $r = .923, p < .001$).

One real benefit of memory and communication can be seen in Figure 8, which depicts the scaled eating distance (as explained before) for the three types in clustered environments. Here, where the likelihood of locating food is much lower than in the random environments, the handicap of not having any idea where to look for food leads to agents traveling unnecessarily far to find energy. As expected, memory provides a clear advantage, as non-communicating (with-memory) agents will return to the most recent food source seen. In contrast, memoryless agents must forage, and (when they are lucky enough to find a cluster) may not approach it directly from the hive. Similarly, communicating hive agents are more often able to choose the nearest cluster, thereby reducing their trip back to the hive. Again, the disadvantage to longer distance is that agents expend more energy on the way back to the hive and therefore have less energy to contribute once there.

An interesting effect here is the apparent divergence of communicating and non-communicating agents in clustered environments. It may seem that the advantage of communication should be constant, or to change slowly as in the random environments. However, the difficulty of for-

aging does not decrease as food quality increases; non-communicating agents still need to come across food in their wanderings. Increasing the value of the food once found does increase non-communicating agents’ performance, but not as much as communicating agents’. In other words, increasing the quality of food sources not only boosts the benefit of finding a food source, it also increases the benefit of sharing information about food sources. Thus, communicating agents improve performance at a faster rate.

IV. DISCUSSION

The results of these hive experiments confirm the conclusions of some previous work: that communication is beneficial under only particular assumptions [9], [12]. We find here that, even for agents that are similar to biological hive-based agents, communication will not prove beneficial to performance in environments where food sources are easily found. In fact, communication *harms* agent performance in these environments, due to the effect of agents tending to group together and visit one particular region, thereby “clearing out” that region faster and increasing the distance agents must travel to return to the hive with energy contributions. One potential mechanism to attack this problem would be to have communicating agents ignore the communicated information with some degree of probability, potentially allowing them to break the cycle of group foraging. However, this is unlikely to resolve the problem, as the benefit of communication is dependent on the value of the information being communicated. In random environments, heading off in one direction versus another does not make too much of a difference (unless everybody is moving in the same direction), hence, the value of a food source location is very small (and possibly negative overall) in these environments. Moreover, the performance values reported here ignore an important aspect of agent performance: the cost of developing and using the communication mechanisms.

In order for an architectural improvement I to be a net benefit to the agent, the performance of the agent with I must be better than its performance without I by a sufficient margin to justify the cost of I . If the performance with I is *worse* than the performance without I , then it is certainly not a net benefit to the agent, since each architectural improvement must be accompanied by additional cost—there is no free lunch. If performance is better, then whether I is a net benefit depends on the extent of the performance increase. The smaller the increase, the less likely it is to outweigh the added cost of the architecture. When the performance increase is large, the new component is likely to be a net benefit (for more detailed discussion of cost-benefit tradeoffs, see [10], [11]). Hence, the cost of communication must be (at least) offset by added performance benefits, or communication will be a net loss. Given the relative ease with which agents find food in random environments, it is unlikely that communication can be “fixed up” enough to make it profitable.

The same holds true for memory in random environments. For example, in the food quality 3000 condition,

the average number of survivors in the with-memory, non-communicating condition (82.6) was significantly greater than the number of survivors in the no-memory condition (80.87). This comes out to a performance increase of just slightly more than 2%. From an evolutionary perspective, the added benefit of memory will only be worthwhile if its costs are very small, a proposition that seems unlikely given these environmental conditions. Once again, the value of the information of a food source location, compared to foraging without such knowledge, is probably too small to justify the additional cost of the mechanisms required to support memory.

The situation is much different in the clustered environments. Here we see memoryless agents at a severe disadvantage, as expected, due to the crucial importance of information regarding food sources. Unlike the random environments, in which a hive agent is likely to find food by setting out in any direction from the hive, successful foraging is likely in only a limited proportion of headings from the hive. Heading out to the “border” of the environment and returning empty-handed is very costly, as not only does the agent fail to contribute to the hive’s energy stores, it must then withdraw the energy it expended searching from the hive’s stores. Memory alone, of a single food source location, is sufficient to provide an enormous improvement in performance over memoryless agents; looking again at food quality 3000 conditions, the average number of survivors increases from 1.44 to 13.33, an increase of 825%! At lower levels of food quality, the performance improvement may not be sufficient to overcome the cost of memory mechanisms, but in some cases it is likely to be profitable when food quality (and, therefore, the value of information regarding its location) increases. Cluster environments are so difficult to survive in that memoryless agents could likely improve their performance significantly by remaining in a cluster once it is found, instead of returning to the hive and contributing; although no reproduction would occur, it is likely that a greater proportion of the original 8 agents would survive to the end of the simulation run.

Communication also fares better in clustered environments than in random environments, again due to the increased value of the information communicated. While there is no performance benefit to communication when food quality is low (less than 1000), once it reaches a critical value, communication pays off, progressively improving survival rates. By the time food quality reaches 3000, the benefit of communication is slightly over 50% (increasing the number of surviving agents from 13.33 to 20.05). This margin seems likely to be sufficient to offset the cost of communication in some organisms, particularly given that the cost of memory has already been offset by the benefits of memory, as argued above. Memory is critical for survival in sparse environments like this, and communication builds on that mechanism to further improve performance.

Hence, the structure of the environment plays a crucial role in the benefits of memory and communication. Moreover, the structure of the task itself helps determine whether such

mechanisms are worth their costs. We demonstrated previously that communication is unlikely to provide sufficient benefit to make it a plausible strategy in similar random and clustered environments where the task involved only locating the food sources, but not returning to a central location periodically [12]. In these cases, drawing the attention of other agents to food sources decreases the probability that the communicating agent will be able to consume the food itself, and has the further effect of clustering agents in regions that quickly become “stripped” of resources while leaving other regions with ungathered food. Increasing the value of the information shared (e.g., by introducing clustered environments) improves the performance of communicating agents, but not sufficiently to overcome non-communicating agents, particularly when the cost of communication is factored in. One substantial advantage of the hive-based task is that agents have to communicate only over short distances (i.e., within the hive), albeit at the cost of having to return periodically. The cost of biologically plausible communication increases at least quadratically with the distance between the communicating agents (this is the energy cost for signals based on distance), while information shared with agents nearby is of less value, as they are more likely to already know of the locations communicated than is they were further away. The centralized hive allows agents to communicate at a very short distance (thereby keeping the cost of communication low—when the distance is short, there is no need to shout louder!), while simultaneously allowing agents to learn of food sources far distant from the regions in which they had previously been foraging (thereby increasing the benefit of communication). The central meeting place imposed by the task has a substantial impact of the benefit of communication.

At the same time, the nature of the task virtually ensures that at least some of the supporting infrastructure required for effective communication of food sources (e.g., cognitive representations of locations) will be in place for communication to build upon. For example, in an evolutionary scenario, the results above strongly suggest that memory is greatly beneficial in itself in clustered environments, and would likely evolve in hive-based agents there. This means that the benefits of communication only need to offset the “marginal” cost of adding communication mechanisms to an architecture that already includes memory mechanisms, rather than having to offset the whole cost of the underlying infrastructure, as is likely to be the case in the “decentralized” task described in [12]. This makes the evolution of communication much more likely for tasks like the one used in the present paper. Note, however, that there may be other, lower-cost, mechanisms that could be nearly as effective. For example, the returning agent could simply indicate that it knows where there is food, and then rely on others to follow it back to the food source. In that case, representational memory is not required, as some form of path integration is probably sufficient. Moreover, communication is not a given, even among agents that have memory. Some bees, for example, do not dance, but do well using only memory

even in conditions that seem to favor communication, based on the results above. These possibilities are worthy of further study, but are beyond the scope of the current paper.

V. RELATED WORK

In contrast to other simulation studies that investigate the utility of communication for artificial tasks [16], [6], our investigation utilizes a *generational survival task* to measure the benefit of communication in a biological setting. Colony members forage for food, consuming energy in the process, and return to a central “hive” to contribute any surplus. The resources of the hive determine when reproduction is possible, so it is in the best interest of the population to maximize the efficiency of foraging. A related approach is taken in [2], which found that “silent” (i.e., non-communicating) agents will invade a population if they are allowed to consume the entirety of any food source they locate, but not if they are only allowed to consume one half of those food sources. This result hints at the results described here, although the approach is somewhat different (i.e., a GA-style experiment in which a fixed number of agents was chosen to reproduce at fixed intervals) and the mode of communication appears to be unlimited broadcast. Follow-up work indicates that behavioral diversity resulting from miscommunication can increase population fitness [1], for reasons similar to those discussed in [12].

Other groups have used agent-based simulations to examine communication in hive-based agents. Schmickl et al study the division of labor in hives [14], [15], an aspect that we have ignored for the purposes of this paper. They identify two channels of communication related to the transfer of nectar when foragers return to the hive. However, the kinds of communication channels that they have in mind use only deictic signals ([12]) rather than representational signals (e.g., to denote locations). Dornhaus et al examine the relationship between environmental structure and communication, following up on their previous work with honey bees in the field [4]. Their findings are similar to those found here with regard to communication, but they do not look at the utility of memory as a potential intermediary step in the evolution of communication.

VI. CONCLUSION

In this paper, we investigated the effect of the environment on the utility of communication and memory in hive-based agents. Little is known about the conditions that must be met in order for these mechanisms to be worthwhile, yet we know that they have evolved in many colony insects. In order to better understand the food configurations most conducive to the evolution of memory and communication, we created simplified but biologically plausible hive-based agents and tested them in a variety of environments, comparing the performance of agents with and without communication, and without memory. We found that when resources are clustered throughout the environment instead of randomly distributed, communicating agents enjoy a significant advantage over non-communicating agents if the quality of the food sources

is sufficiently high. As the quality increases, so does the communicating agents’ advantage. The benefit of memory is even greater in clustered environments, demonstrating the value of information about food locations in sparse environments. However, in environments with randomly-placed (i.e., unclustered) communication actually *decreases* hive agent performance, while memory increases performance only slightly, and likely not enough to justify its added costs. This demonstrates that environmental conditions play an important role in determining the utility of communication and memory in hive-based colonies.

REFERENCES

- [1] Jin Akaishi and Takaya Arita. Misperception, communication and diversity. In *Artificial Life VIII: Proceedings of the Eighth International Conference on Artificial Life*, pages 350–357, 2002.
- [2] Takaya Arita and Yuhji Koyama. Evolution of linguistic diversity in a simple communication system. *Artificial Life*, 4:109–124, 1998.
- [3] Anna Dornhaus and Lars Chittka. Why do honey bees dance? *Behavioral Ecology and Sociobiology*, 55(4):396–401, 2004.
- [4] Anna Dornhaus, Franziska Klügl, Christoph Oechslein, Frank Puppe, and Lars Chittka. Benefits of recruitment in honey bees: Effects of ecology and colony size in an individual-based model. *Behavioral Ecology*, 17:336–344, 2006.
- [5] Fred C. Dyer. The biology of the dance language. *Annual Review of Entomology*, 47:917–949, 2002.
- [6] Maja J. Mataric. Using communication to reduce locality in multi-robot learning. In *AAAI/IAAI*, pages 643–648, 1997.
- [7] Randolf Menzel, Karl Geiger, Jaskan Joerges, Uli Müller, and Lars Chittka. Bees travel novel homeward routes by integrating separately acquired vector memories. *Animal Behavior*, 55(1):139–152, 1998.
- [8] D. J. Nicholson, S. P. D. Judd, B. A. Cartwright, and T. S. Collett. Learning walks and landmark guidance in wood ants (*Formica rufa*). *The Journal of Experimental Biology*, 202(13):1831–1838, 1999.
- [9] Paul Schermerhorn and Matthias Scheutz. Investigating the adaptiveness of communication in multi-agent behavior coordination. *Adaptive Behavior*, 15(4):423–445, 2007.
- [10] Paul Schermerhorn and Matthias Scheutz. Social, physical, and computational tradeoffs in collaborative multi-agent territory exploration tasks. In *Proceedings of the First IEEE Symposium on Artificial Life*, pages 295–302, April 2007.
- [11] Matthias Scheutz and Paul Schermerhorn. Predicting population dynamics and evolutionary trajectories based on performance evaluations in alife simulations. In *Proceedings of GECCO 2005*, pages 35–42. ACM Press, June 2005.
- [12] Matthias Scheutz and Paul Schermerhorn. The limited utility of communication in simple organisms. In *Proceedings of Artificial Life XI*, pages 521–528, August 2008.
- [13] Matthias Scheutz, Paul Schermerhorn, Ryan Connaughton, and Aaron Dingler. SWAGES—an extendable parallel grid experimentation system for large-scale agent-based alife simulations. In *Proceedings of Artificial Life X*, pages 412–418, June 2006.
- [14] T. Schmickl and K. Crailsheim. Analysing honeybees’ division of labour in broodcare by a multi-agent model. In *Proceedings of Artificial Life XI*, pages 529–536, August 2008.
- [15] R. Thenius, T. Schmickl, and K. Crailsheim. How to know without having been there? — investigating communication channels in the nectar collecting system of a honeybee colony (abstract). In *Proceedings of Artificial Life XI*, page 807, August 2008.
- [16] V. Trianni, T. H. Labella, and M. Dorigo. Evolution of direct communication for a swarm-bot performing hole avoidance. In *Proceedings of the 4th Intl. Workshop on Ant Colony Optimization and Swarm Intelligence*, pages 131–142, 2004.