The Limited Utility of Communication in Simple Organisms

Matthias Scheutz and Paul Schermerhorn
Cognitive Science Program and Complex Systems Group
College of Arts and Science and School of Informatics
Indiana University, Bloomington, IN 47406, USA
{mscheutz,pscherme}@indiana.edu

Abstract

Many forms of communication have evolved in the animal kingdom for different purposes. In this paper we investigate the limits of communication for simple reactive organisms and show that communication has only limited benefits in biologically inspired foraging tasks and can even have detrimental effects in certain environments. Based on these results, we argue that simple agents with simple architectures need very special environmental conditions for communication to benefit them and thus to evolve.

Introduction

Various forms of communication have evolved in the animal kingdom, ranging from broadcasting simple signals, to the complex linguistic exchanges of humans. Much work in ALife has attempted to demonstrate *that*, *when* and *how* communication can evolve, but paid little attention to cases where communication *did not* or *will not* evolve. Yet, we believe that a full appreciation of the utility of communication for natural and artificial agents is not possible without understanding both its potential *and* limitations.

In this paper we attempt to delineate the kinds of circumstances that would limit the evolution of communication for biologically inspired tasks. We start with a few methodological points about the notion of "communication" and lay out the argument structure we will use here for investigating the limitations of communication. Then we define a biologically inspired task called *t*-MATES (for "timed Multi-Agent Territory Exploration Task with Satiation") and introduce various agent models for that task. Simulation results will paint a surprising picture, showing that communication is of very little utility for *t*-MATES tasks. We discuss implications of the results for the evolution of communication and relate them to previous findings in the literature, concluding with a brief summary and suggestions for future work.

Communication and Mechanism

Biological agents have evolved different forms of communication for different purposes, ranging from signaling danger (e.g., danger caws of lookout crows), to indicating readiness for mating (e.g., mating calls of frogs), to reporting locations of food (e.g., food dances of honey bees), to initiating joint action (e.g., dogs' bows to initiate social play), and to sharing mental states (e.g., human reports of their beliefs).

These different forms of communication require different functional capabilities of the agents' architectures. Processing simple signals emitted from another agent that only indicate the agent's presence in some location does not require much more than a perceptual system that can pick up those signals as such and determine the direction from which they originated (e.g., female frogs can determine the signal strength and direction of male callers in a swamp; similarly, ants can sense the gradient of pheromones left in the environment by other ants). In fact, simple signals indicating a particular state of affairs as perceived by an agent can be construed as *indexicals* (in the Peircian sense), i.e., a food call effectively communicates indexical information of the form "I see food here now", see also Perconti (2002). Note that this message containing three indexicals is different from the message "Agent A sees food at location X at time t" even if the content of the message is the same (i.e., the variables A, X, and t are replaced by the respective names so that agent, location and time agree with the utterance of the indexical message). For messages of the latter sort can realistically not be encoded in simple indexical signals (unless one has a large number of distinct signals for all occasions of interest at hand, which is practically almost never feasible). Hence, representational devices are needed to represent agents, places, and times in the second case. Those representational devices, in turn, require a systematic encoding (i.e., representations with formal rules defining well-formed expressions) and mechanisms that can encode and decode information (i.e., parsers). Moreover, to determine times, locations, and agents (as in the above case), agreed-upon scales (e.g., clocks, maps, and naming conventions) are required together with "measuring devices" (i.e., algorithms and possibly tools) to determine that determine where instances fall on the scale (i.e., what time it is, where in the map an item is located, and who the speaker is). All of this, in turn, requires much more sophisticated functional capabilities in agent architectures that allow agents to determine what to communicate and how to use the communicated information. Processing expressions that can encode and thus communicate mental states like beliefs, for example, might require representational capabilities such as those used in modal or (fragments of) first order logic (e.g., to represent the belief that at least one member of a group has already had dinner).

More complex architectures that can handle the added sophistication of more complex messages (such as their syntax and semantics) comes at a price, however: the cost of building/growing and maintaining it.[1] In contrast, simple signals like "I see food here now" might not require much additional processing at all: the receiving agent could just move towards the perceived signal if it needs food, or ignore it otherwise.

Aside from the computational/architectural costs, the costs of communicating can also be substantive. An agent that continuously sends broadcast signals (like alarm, food, or mate calls) might use up a significant portion of its energy, possibly without any benefit if no other agent can hear the signal (the calls of male treefrogs, for example, are much more expensive than navigation, limiting them to participate in the calling chorus during mating season for only a few days out of several weeks, e.g., see Fellers (1979)).

From all of the above it is clear then that claims about the evolvability of communication or about the likelihood of communication evolving need to be very specific with respect to the form of communication they target, as all the above differences (with respect to communication schemes, functional, representational and computational capabilities of agent architectures, and the various costs) are typically subsumed under the general term "communication".

For example, only social insects seem to have evolved different ways to communicate information about food sources in partly non-indexical ways among their respective groups (from annotating the environment by leaving marks like ants do with pheromone trails, to using intricate dances like those of honey bees that encode direction and distance to food sources as well as food quality), despite the large number of different species of insects. For most other forms of insects, indexical mate signaling is the most that has been described (but see Cocroft (2005) for an example of food signaling in treehoppers).

We will in the following investigate two forms of communication about food: the simplest form of indexicals signaling "I see food here now", and a more complex form of "I see food in location *X* now", which removes one of the three indexicals and replaces it with an explicit value (namely the location of the food item). Since we are interested in determining the limits of communication, we will have to follow a different strategy from most work on the evolution of communication, which we will describe next.

Method

Most research on the evolution of communication (see the Related Work section below) attempts to demonstrate that communication is *beneficial* and thus *can* evolve in some agent in some given task. The logical form of these "evolvability claims" is typically an existence claim: for a given agent type and task there exists an initial distribution and an evolutionary trajectory from this distribution that leads to communication in those agents. The existential quantifiers here are often the result of a common strategy to establish claims about the evolvability of communication by examining the outcomes of runs of genetic algorithms or similar evolutionary computational tools. It is, however, important to note that the existential quantifiers critically limit the scope of the claim: it only says that for some initial conditions there are trajectories that lead to communication being beneficial. This leaves open whether communication could have or would have evolved for all trajectories, or the vast majority of trajectories, and thus whether it was *likely* for communication to evolve. For the likelihood of a property evolving in a task, we need to determine the conditional probability of communication evolving *given* as set of initial conditions. The conditional probability formulation can then be used both to confirm and disconfirm that a property $\mathcal{P}$ such as communication is likely to evolve in a set of agents by comparing the performance of agents with $\mathcal{P}$ and without $\mathcal{P}$ for each initial condition. If there is no *absolute performance* difference, then there will likely be no evolutionary trajectory resulting in agents with $\mathcal{P}$, for having and using $\mathcal{P}$ would at best incur an additional cost without yielding any gain in task performance. If there is no *relative performance* difference between agents with and without $\mathcal{P}$ in the given task (i.e., when the cost of having and using $\mathcal{P}$ is taken into account in the performance measure), then the answer to the question whether there is an evolutionary trajectory leading to $\mathcal{P}$ will depend on additional information about intermediary stages of the trajectories, e.g., what evolutionary operations are used and how frequently they are employed, whether these operations can produce viable architectures at any point along the trajectory, etc. Typically, it is difficult (if not infeasible) to obtain this kind of information.

Hence, we will aim at establishing that there is no *absolute performance* difference between agents with and without communication. While it is impossible to do this exhaustively for the sheer size of the set of initial conditions (even in our limited experimental setup), it is possible for a small, but representative subset of initial conditions randomly drawn from the set of all initial conditions. Statistical significance tests can then be used to reject the null hypothesis that there is an absolute performance difference between communicating and non-communicating agents. And the *p*-

---

[1] In humans, for example, the brain consumes up to 25% of the body's energy, in infants even up to 75%, e.g., see Cunnane (2006).

value of the significance test can be taken as an upper bound on the conditional probability that communication evolves in environments of the given type (a more detailed exposition of the employed experimental methodology can be found in Scheutz and Schermerhorn (2005)).

Task and Agent Models

To be true to the question about when and why communication evolved in nature, we define a generic biologically plausible territory exploration task that is intended to measure the efficiency with which agents can negotiate their environment (e.g., how they determine where to go in their environment based on their survival goals such as finding food).

Definition t-MATES: A *timed multi-agent territory exploration task with satiation* (t-MATES) $T(t,C,A,R,D,S)$ requires a group of identical agents A each with sensory range R to visit as many checkpoints in C as possible in a 2D environment within the allotted time t, where agents and checkpoints are placed according to a probability distribution D and each agent can visit up to S checkpoints (the "satiation level").

D is typically unknown to the agents, hence it cannot be *a priori* exploited by them. Agents neither know their own locations in the environment nor those of the checkpoints. Rather they can only detect relative locations of checkpoints based on their perceptions (e.g., the location of a checkpoint relative to the agent's heading). All checkpoints are marked so agents can perceive them when they are within sensory range. Whenever a checkpoint is visited by an agent, the agent removes the mark, thus effectively removing the checkpoint from the environment.

One way to conceptualize t-MATES tasks is to think of them as "foraging episodes" (of duration t) taken from an ongoing evolution of populations of biological agents: checkpoints can be viewed as food sources, and visiting can be taken to be consuming them, with the satiation level determining the maximum food intake an agent can have within the foraging period t. Performance of different agent types during t reflects the agent types' *foraging efficiency* (i.e., the efficiency with which agents can find food), which in turn is a fitness measure of their performance in the larger evolutionary context of survival and procreation. That is, if an agent kind K_1 has a higher foraging efficiency than another agent kind K_2 as measured in the t-MATES task, where foraging efficiency is given in terms of "average number of visited items per time unit", then one would expect K_1, on average, to perform better than K_2 in t-MATES tasks in an evolutionary setting.[2]

Next, we define a simple reactive, yet biologically plausible non-communicating base agent model (e.g., at the level of insect behavior) that meets the minimum requirements for the t-MATES task of being able to move about the environment, detect a checkpoint within the given sensory range R and move towards it. For simplicity's sake, we do not employ a particular sensory model (e.g., sonar or visual sensors), which would introduce complicating perceptual effects such as interference or visual occlusions, but rather assume that agents can detect any number of checkpoints within the circular $Area$ of radius R around them. Given that checkpoints have an extension in space (1 square unit), the maximum number of detectable non-overlapping checkpoints is limited by $Area = R^2 \cdot \pi$.

The behavior of a *non-communicating agent* is then determined solely by its sensory information (which is limited to checkpoints, other agents are not perceived) based on the following three rules:

Rule 1: if no checkpoint is sensed, perform a random walk $RW(rwd,\beta)$ (i.e., move in the direction of the current heading θ for *rwd* cycles, then change heading randomly to some value in $[\theta-\beta,\theta+\beta]$)

Rule 2: if some checkpoints are sensed and are not within visiting range (i.e., they are not within the extension of the agent's body of 8 units), go directly towards the closest checkpoint (the direction is given by α such that $\min_d\{\langle d,\alpha\rangle | \langle d,\alpha\rangle$ is within sensory range)

Rule 3: if some checkpoints are sensed, at least one checkpoint C is within visiting range, and the agent's count of checkpoints visited c is less than its satiation level S, remove the mark(s) of up to $S-c$ of the checkpoints (if it is/they are still there)

When an agent achieves satiation, it continues to execute the rules above (i.e., it will search for another checkpoint and move to it, but upon arrival will simply remain there until the checkpoint is removed by another agent).

Note that the basic model is parameterized by $RW(rwd,\beta)$ and R, hence its performance will critically depend on those parameters. In a sense, $RW(rwd,\beta)$ is an agent-internal parameter that should be chosen so as to maximize an agent's performance with a given sensory range in a given environment if we want to investigate the utility of communication. However, to be able to choose the best values for $RW(rwd,\beta)$, we need to understand how the random walk interacts with other parameters such as the agent's sensory range, the number of participating agents in the task, and the structure of the checkpoints in the environment (e.g., a random distribution). Hence, we conducted a large set of *calibration experiments* to determine the best random walk distance (*rwd*) for base agents for each *sensory range* $R \in \{25\cdot n | 1 \le n < 10\} \cup \{300+50\cdot n | 0 \le n \le 6\}$ and *group*

---

[2] The qualifier "on average" is critical here as there can always be special circumstances that punish normally fitter agents and can even lead to their extinction.

size $|A| \in \{2,3,4,5\}$, in both random and clustered environments.[3]

We extend the base agent model in two different ways to allow for two kinds of communication: purely indexical communication – call it *signaling agent* – and partly "representational" communication – call it *messaging agent*. We also allow for two reactions to signals: a typical *approach* behavior (e.g., like those exhibited by Toque Macaques when they hear a food signal Dittus (1984)) and, for contrast, an *avoidance behavior* that will cause agents to walk away from the direction of the food signal. The effect of the avoid behavior should contrast to potential clustering effects engendered by the approach behavior, potentially leading to better agent distribution, particularly in random environments. Thus, we will define four different types of communicating agents.

For the signaling agent we add the following two rules:

Sending: whenever a checkpoint is sensed, the agent turns on its "checkpoint" signal

Receiving: whenever no checkpoint is sensed but one or more checkpoint signals are sensed, the agent either approaches or moves away from the direction of the closest signal.

Similarly, we add two rules for the messaging agents:

Sending: whenever a checkpoint is sensed, the location of the checkpoint is communicated as the $\langle d, \alpha \rangle$ of distances $d \in [0,R]$ and angles $\alpha \in [0,2\pi]$ relative to the sending agent's position

Receiving: whenever no checkpoint is sensed but one or more checkpoint messages are received, the agent either approaches or moves away from the closest checkpoint[4]

Note that messaging agents will at any given time know the locations of all checkpoints that are perceived by any agent within communication range, while signaling agents will only know the locations of checkpoints they themselves perceive, even though they will know where other agents are that perceive checkpoints. Satiated agents will continue to send and receive according to the communication rules for their agent type.

Because we are interested in determining the limitations of communication, we will consider two different communication ranges: an (unrealistic) unlimited communication range (as a control condition) and a biologically plausible limited communication range that is the same as the agent's sensory range.

We thus arrive at nine different agents, which we will label from 0 to 8 for ease of presentation. Agents of type 0 do not use communication, while odd-numbered agents use messaging and even-numbered agents use signaling. Agents types 1 through 4 use *unlimited* communication, while agents types 5 through 8 use *limited* communication. Finally, agents types 1, 2, 5, and 6 use *approach* behavior, while agents types 3, 4, 7, and 8 use *avoidance* behavior.

Experiments and Results

All simulation experiments were conducted in the artificial life simulator SWAGES, which is a distributed agent-based artificial life simulation environment that consists of the parallelizable SIMWORLD simulator and an experiment grid-server used to schedule experiments on heterogeneous clusters of computers, automatically parallelize and distribute simulations over multiple hosts, collect statistics, and perform preliminary data analysis (Scheutz and Schermerhorn, 2006; Scheutz et al., 2006).

One simulation experiment consists of 100 experimental runs, each using different randomly generated initial conditions from a given distribution D (of initial conditions) in a continuous 2D world, which is limited to an 800 by 800 square region (in comparison, each agent occupies a circular region of diameter 8).[5] Two different distributions of checkpoints are used: *random* and *cluster*. In the random distribution, checkpoints are placed at random locations within the whole environment, while in the cluster distribution all checkpoints are placed according to a Gaussian distribution (with a radius of 150 units and a standard deviation of 75 units) centered in one 200 by 200 quadrant (with all checkpoints contained within the quadrant). We consider two different numbers of checkpoints, $|C| = 10$ and $|C| = 40$, and four group sizes of agents, $|A| = 2$ to $|A| = 5$ to investigate the possible effects of food density and group size on the utility of communication. Moreover, we fix the agents' satiation thresholds at $S = 10$, but vary their sensory ranges from 25 to 600. The same set of 100 initial conditions of checkpoint and agent placements is used for all variations of group size and sensory/communication range for a given number of checkpoints and checkpoint distribution to allow for a

---

[3] For space reasons we omit a detailed description of the results.

[4] The details of exactly how agents extract the exact location of a checkpoint relative to an agent's own heading from another agent's message are not straightforward; they usually involve additional communicated parameters such as heading of agents relative to each other or relative to a fixed coordinate system (e.g., as measured by a compass, etc.). Here we simply assume that the agent can compute the angle and distance to the communicated checkpoints based on where the message came from, and that they somehow have access to the source location. For it will turn out that messaging agents do in general not have better absolute performance than signaling agents, hence the details of the control architecture and the buried complexities and costs do not have to be considered explicitly (as would typically be the case for conditions where messaging agents performed absolutely better than signaling agents).

[5] Whenever an agent reaches the boundary of the environment, it will "bounce" off (similar to a billiard ball bouncing off the cushion) with some very small random error (this is to make sure that agents will not be able to leave the area within which checkpoints are located).

	$C = 10$		$C = 40$	
	Random	Cluster	Random	Cluster
0	6.93 (2.62)	8.07 (3.13)	20.80 (7.55)	20.69 (10.74)
1	6.65 (2.71)	8.41 (2.99)	20.32 (7.95)	25.22 (10.48)
2	6.62 (2.70)	8.37 (2.99)	20.29 (7.93)	24.89 (10.44)
3	6.60 (2.70)	7.82 (3.13)	19.98 (7.99)	18.57 (10.90)
4	6.60 (2.71)	7.81 (3.20)	19.99 (8.03)	18.58 (10.90)
5	6.90 (2.63)	8.08 (3.14)	20.74 (7.58)	21.11 (10.87)
6	6.89 (2.62)	8.07 (3.13)	20.72 (7.58)	21.08 (10.86)
7	6.95 (2.62)	8.05 (3.13)	20.85 (7.55)	20.36 (10.61)
8	6.95 (2.62)	8.05 (3.13)	20.86 (7.55)	20.36 (10.61)

Table 1: Average performance of all nine agent types across all sensory ranges (from 25 to 600) in both types of environments (random and cluster) with both numbers of checkpoints (10 and 40) for all four group sizes (from 2 to 5).

		10 Checkpoints			
		Random		Cluster	
	Df	F	p	F	p
Type	8	59.00	$< .001$	55.07	$< .001$
Range	16	5167.12	$< .001$	3390.00	$< .001$
Type:Range	128	3.93	$< .001$	3.24	$< .001$
Error	61047				

		40 Checkpoints			
		Random		Cluster	
	Df	F	p	F	p
Type	8	25.17	$< .001$	499.00	$< .001$
Range	16	2610.12	$< .001$	1756.89	$< .001$
Type:Range	128	3.82	$< .001$	15.13	$< .001$
Error	61047				

Table 2: Two-way 9×16 Analysis of variance in performance with the independent variables *agent type* and *sensory range* and the dependent variable *checkpoints visited.*

		10 Checkpoints							
		Unlimited Communication Range				Limited Communication Range			
		Approach		Avoid		Approach		Avoid	
	0	1	2	3	4	5	6	7	8
0		+++	+++	+++	+++				
1	- - -					- - -	- - -	- - -	- - -
2	- - -					- - -	- - -	- - -	- - -
3	- - -					- - -	- - -	- - -	- - -
4	- - -					- - -	- - -	- - -	- - -
5		+++	+++	+++	+++				
6		+++	+++	+++	+++				
7		+++	+++	+++	+++				
8		+++	+++	+++	+++				
0		- - -	- - -	+++	+++				
1	+++			+++	+++	+++	+++	+++	+++
2	+++			+++	+++	+++	+++	+++	+++
3	- - -	- - -	- - -			- - -	- - -	- - -	- - -
4	- - -	- - -	- - -			- - -	- - -	- - -	- - -
5		- - -	- - -	+++	+++				
6		- - -	- - -	+++	+++				
7		- - -	- - -	+++	+++				
8		- - -	- - -	+++	+++				

		40 Checkpoints							
		Unlimited Communication Range				Limited Communication Range			
		Approach		Avoid		Approach		Avoid	
	0	1	2	3	4	5	6	7	8
0		+++	+++	+++	+++				
1	- - -			+	+	- -	- -	- - -	- - -
2	- - -					- - -	- - -	- - -	- - -
3	- - -	-				- - -	- - -	- - -	- - -
4	- - -	-				- - -	- - -	- - -	- - -
5		++	+++	+++	+++				
6		++	+++	+++	+++				
7		+++	+++	+++	+++				
8		+++	+++	+++	+++				
0		- - -	- - -	+++	+++				
1	+++			+++	+++	+++	+++	+++	+++
2	+++			+++	+++	+++	+++	+++	+++
3	- - -	- - -	- - -			- - -	- - -	- - -	- - -
4	- - -	- - -	- - -			- - -	- - -	- - -	- - -
5		- - -	- - -	+++	+++				
6		- - -	- - -	+++	+++				
7		- - -	- - -	+++	+++				
8		- - -	- - -	+++	+++				

Table 3: Comparison of the nine agent kinds in 10 checkpoint (top two tables) and 40 checkpoint (bottom two tables). Within each checkpoint condition, the upper table is in random and the lower table is in cluster environments. “+” and “-” denote significant performance differences (between mean performance of the agent type in the row minus mean performance of the agent type in the column), where the number of symbols indicates the significance level based on Tukey’s honestly significant difference (HSD) multicomparison post-hoc test: one symbol for $p < .05$, two symbols for $p < .01$, and three for $p < .001$.

direct performance comparison among the different agent kinds and parameters. We use *the number of checkpoints visited within* t as performance measure and fix $t = 500$, which turned out to be long enough to highlight foraging differences among agent types.

The overall performance results for the nine agent types in the four environmental conditions averaged over all sensory ranges and group sizes are shown in Table 1, the results of ANOVAs for each environmental condition are shown in Table 2, and the statistically significant performance differences are shown in Table 3. The results in Table 2 show that the differences in average performance between agent types are significant (as is the effect of sensory range on performance, unsurprisingly). The interaction between agent type and sensory range is due to performance differences between types found in medium sensory ranges; when sensory range is very low, agents have a very hard time locating checkpoints about which to communicate, whereas when sensory range is high, shared information is seldom novel.

The results in Table 3 demonstrate that in random environments, regardless of the food density, communication does not pay off, not even in the simplest form of signaling (as demonstrated by the lack of minus symbols in the row with agent type 0). Quite to the contrary, unlimited communication can significantly hurt agent performance (see the plus symbols in the 0 agent row). In cluster environments, there is some benefit to communication: agents with unlimited communication range using approach behavior perform better than non-communicating agents (see the minus symbols in the first two columns of the 0 agent row), but not if they use avoid behavior, as expected (see the plus symbols in the third and fourth columns of the 0 agent row). Note that there was no performance difference between the two forms of communication. Messaging agents with unlimited communication range using approach behavior in random high density environments do, however, have a slight advantage over those using avoid behavior (see the single plus symbol in columns 3 and 4 of the 1 agent row). Overall, there is no statistically significant performance difference in any of the four environmental conditions between non-communicating and communicating agents with limited communication range (as evidenced by the lack of any symbols in columns 5 through 8 in the 0 agent row).

Discussion

The above results make a strong case for the limited utility of communication for simple insect-like agents in t-MATES tasks, especially since there was no statistically significant performance difference between non-communicating and communicating agents with limited communication range. With unlimited communication range, the question about the utility of communication becomes surprisingly dependent on the type of environment: in random environments, performance actually *decreases* due to agents wasting cycles pursuing checkpoints that will likely have been visited by other agents before them, while in cluster environments performance increases due to agents being attracted to the cluster quickly as soon as one agent has discovered it. The performance difference between communicating and non-communicating agents is particularly pronounced in the high density condition (of 40 checkpoints), where the satiation level limits agents to 10 visits (thus 4 agents are required to visit all checkpoints in the cluster; in the non-communication conditions this means that the cluster needs to be discovered independently at least four times, which can take a while). The performance improvement is less pronounced in the 10 checkpoint cluster (given that one agent could visit them all). In the random condition, the trend is in the opposite direction: the performance decrease is higher in the low density condition than in the high density condition, again for the reason that agents will chase checkpoints that other agents are likely to get first.

Note that the above results are based on absolute performance differences, as communicating agents are not charged any penalties for their communication mechanism (including processing and representational resources and computation time, additional sensors and effectors, etc.). The costs involved in communication, especially the cost for sensitive sensors with large sensory ranges (as is required for communication to be beneficial) can be quite expensive (e.g., see (Schermerhorn and Scheutz, 2006, 2007b) for comparison of the various tradeoffs). Hence, whether communication based on large communication ranges will evolve for high density cluster environments is an open question, but we can already say that if it evolves then it will use signaling and not messaging, given that there was no performance difference between signaling and messaging, but messaging requires and incurs much greater costs.

It is curious, then, that a small number of insect species – the social insects – did evolve messaging communication to communicate the locations of resources to their peers. This could be because these agents depart from and return to a common location which makes a difference in their foraging patterns that could favor communication. Moreover, honey bees (Capaldi and Dyer, 1999; Menzel et al., 1998) could not use the signaling mechanism for food employed in this study when they are at the hive. Interestingly, Dornhaus and Chittka (2004) provide evidence that bees can survive just fine *without communication* (i.e., when their ability to communicate is suppressed) depending on the food distribution and food quality in their environment. Hence, we would expect to observe this contingency of communication being beneficial depending on the distribution of food in the environment in modified t-MATES tasks if the additional constraint of always having to return to a common “base checkpoint” after visiting a “field checkpoint” is taken into account; and in fact Schermerhorn and Scheutz (2005) provides preliminary evidence from a related task that suggests that this might indeed be the case.

Related Work

Several authors have investigated the utility of communication or signaling in various tasks. There are, for example, purely game-theoretic studies that explore the role of communication in *coordination games* and show that non-binding, pre-play communication can can increase the probability of playing the Pareto-dominant strategy (e.g., Cooper et al. (1992)). Arkin et al. find that communication can aid coordination in robotic retrieval tasks (Arkin and Hobbs, 1992; Wagner and Arkin, 2004). Conversely, Werger et al. (Werger and Mataric, 2001) and Quinn et al. (2003) found communication to be unnecessary to achieve task formation in a system which uses behavior-based mechanisms to generate cooperative behaviors. However, the employed tasks are sometimes very different from t-MATES tasks making it difficult to compare the outcomes.

MacLennan found that communication will evolve in a task requiring coordinated behavior when agents are rewarded for agreeing on the meaning of a signal (MacLennan, 2002). However, this rewards agents directly for communication rather than demonstrating that communication can be beneficial to performance of a separate task. Similarly, Levin (1995), using a genetic algorithm approach with a fitness function that explicitly favors the evolution of communication, found it possible to progressively improve the ability of agents to correctly interpret other agents’ communications. Noble and Cliff (1996) extend MacLennan’s work to show that a structured language will evolve based on the benefits of communication.

Quinn (2001) describes experiments in which artificial agents evolve a signaling mechanism in the absence of predetermined communication channels. Pairs of simulated robotic agents starting within sensor range of one another are given the task of moving in their environment while staying within sensor range. Here, a signaling system evolved which was not part of the fitness function, but rather measured absolute task performance and behavior coordination.

Marocco et al. (2003) and Cangelosi et al. (2004) describe experiments with simulated robots which are required to recognize a sphere and a cube in order to maximize contact with the sphere and minimize contact with the cube. Once agents identify an object, they can communicate that information to other agents, allowing them, for example, to avoid contact with the cube without using first-hand proprioceptive information. Communication between parents and offspring was found to evolve.

Ackley and Littman (1991) note that models in which the speaker as well as the listener benefits from communication produces an unrealistic environment in which many observed phenomena related to communication do not make sense. In their model, agents can share information about nearby food and predators. They found that, in some conditions, communication can improve performance on the survival task (i.e., locating food and avoiding predators).

Noble (1999) examines various communication games to determine under what circumstances communication will evolve. Agents have the opportunity to communicate during encounters between a signaler and a receiver, and they are rewarded when the receiver responds appropriately to the signal. In this study, communication was found to evolve when the signaler receives a net reward. However, when signalers are not rewarded for receivers' successes, communication did not evolve.

Grim et al. (2002) examine the benefit of communication in a survival task requiring agents to consume food when present and hide from predators. Agents can share information about food and predators with neighboring agents. They find that communication will evolve in the absence of a cost for signalling, but that adding such a cost, even just to the level of 2% of the benefit of eating or the cost of predation, affects the viability of communication.

Reggia et al. (2001) found similar results to those presented above with regard to the effect of checkpoint (food) distribution. Their study examines only indexical signalling, and they do not examine the effect of sensory range or communication range. However, different from our study, their model includes predators, an important risk factor that should further decrease the utility of communication.

Conclusion

We investigated the limitation of communication for improving the performance of simple agents in timed multi-agent territory exploration tasks with satiation. Different from most work on the evolution of communication, our results paint a nuanced picture of the utility of communication. In environments with no structure communicating agents with limited communication range do not perform better than non-communicating agents, and unlimited communication range can result in a significant performance drop. In cluster environments, only communication with unlimited communication range (that covers the whole environment) leads to better performance (which may or may not be biologically plausible depending on the type of environment and type of sensory modality). More importantly, there was no significant performance difference between signalling and messaging agents suggesting that if communication were to evolve, then it would be of the simplest possible form using only indexical broadcast signals rather than non-indexical messages.

While the above results might seem largely negative from the perspective of someone arguing for the utility of communication and the likelihood of it evolving for biologically plausible foraging tasks, the results about the utility and high level of performance of simple agents with at best simple means of signalling in *t*-MATES tasks is highly relevant for and might find direct applications in a variety of engineering tasks where low-cost solutions or solutions with a high (if not the best) relative performance are of major interest (for example, in energy-efficient mobile rovers that explore the surfaces of planets, expandable autonomous mine-sweeping robots that search an area for mines and explode with the mine by driving over it, or unmanned surveillance vehicles that need to check various locations in an environment as they are dynamically reported as quickly as possible).

In addition to the already mentioned constraint of imposing a hive-like home base for foragers, we see at least two promising directions for further investigating the benefits and limits of communication in *t*-MATES tasks. The first concerns the idea of "structure in the environment" and the degree to which communication can benefit from it. Specifically, it would be interesting to define a measure of "structure" (ideally information-theoretic) for environments that gets at the kinds of distributions of checkpoints that would favor communication.

Another direction concerns the coordination of agent behavior, in which communication could play an facilitatory role. The idea here is to impose additional task constraints on the *t*-MATES task such as requiring multiple agents to visit the same checkpoint at the same time (as would be required for mating) to isolate scenarios where coordination can be significantly improved via communication (Schermerhorn and Scheutz (2007a) already started an exploration of these tradeoffs in a related task).

Acknowledgments

The authors would like to thank Colin Allen and three anonymous reviewers for their helpful suggestions for improving the paper.

References

Ackley, D. and Littman, M. (1991). Altruism in the evolution of communication. In Brooks, R. and Maes, P., editors, *Artificial Life IV: Proceedings of the Fourth Workshop on Artificial Life*, pages 40–48, Cambridge, MA. MIT Press.

Arkin, R. C. and Hobbs, J. D. (1992). Dimensions of communication and social organization in multi-agent robotic systems. In *Proceedings of the Second International Conference on Simulation of Adaptive Behavior*, pages 486–493.

Cangelosi, A., Riga, T., Giolito, B., and Marocco, D. (2004). Language emergence and grounding in sensorimotor agents and robots. In *First International Workshop on Emergence and Evolution of Linguistic Communication*, Kanazawa, Japan.

Capaldi, E. A. and Dyer, F. C. (1999). The role of orientation flights on homing performance in honeybees. *The Journal of Experimental Biology*, 202:1655–1666.

Cocroft, R. B. (2005). Vibrational communication facilitates cooperative foraging in a phloem-feeding insect. *Proceedings of Biological Science*, 272(1567):1023–1029.

Cooper, R., DeJong, D. V., Forsythe, R., and Ross, T. W. (1992). Communication in coordination games. *The Quarterly Journal of Economics*, 107(2):739–771.

Cunnane, S. C. (2006). *Survival of the Fattest: The Key to Human Brain Evolution*. World Scientific.

Dittus, W. P. (1984). Toque macaque food calls: Semantic communication concerning food distribution in the environment. *Animal Behaviour*, 32(2):470–477.

Dornhaus, A. and Chittka, L. (2004). Why do honey bees dance? *Behavioral Ecology and Sociobiology*, 55(4):396–401.

Fellers, G. M. (1979). Mate selection in the gray treefrog, Hyla versicolor. *Copeia*, 1979(2):286–290.

Grim, P., Kokalis, T., Tafti, A., and Kilb, N. (2002). Evolution of communication with a spatialized genetic algorithm. *Evolution of Communication*, 3(2):105–134.

Levin, M. (1995). The evolution of understanding: A genetic algorithm model of the evolution of communication. *BioSystems*, 35:167–178.

MacLennan, B. (2002). Synthetic ethology: A new tool for investigating animal cognition. In Bekoff, M., Allen, C., and Burghardt, G. M., editors, *The Cognitive Animal: Empirical and Theoretical Perspectives on Animal Cognition*, pages 151–156. MIT Press.

Marocco, D., Cangelosi, A., and Nolfi, S. (2003). The emergence of communication in evolutionary robots. *Philosophical Transactions: Mathematical, Physical and Engineering Sciences*, 361(1811):2397–2421.

Menzel, R., Geiger, K., Joerges, J., Müller, U., and Chittka, L. (1998). Bees travel novel homeward routes by integrating separately acquired vector memories. *Animal Behavior*, 55(1):139–152.

Noble, J. (1999). Cooperation, conflict and the evolution of communication. *Adaptive Behavior*, 7(3-4):349–370.

Noble, J. and Cliff, D. (1996). On simulating the evolution of communication. In Maes, P., editor, *Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*, pages 608–617, Cambridge, MA. MIT Press.

Perconti, P. (2002). Context-dependence in human and animal communication. *Foundations of Science*, 7:341–362.

Quinn, M. (2001). Evolving communication without dedicated communicatin channels. In *Proceedings of ECAL 2001*, pages 357–366.

Quinn, M., Smith, L., Mayley, G., and Husbands, P. (2003). Evolving controllers for a homogeneous system of physical robots: Structured cooperation with minimal sensors. *Philosophical Transactions of the Royal Society of London, Series A: Mathematical, Physical and Engineering Sciences*, 361:2321–2344.

Reggia, J. A., Schultz, R., Wilkinson, G. S., and Uriagereka, J. (2001). Conditions enabling the evolution of inter-agent signaling in an artificial world. *Artificial Life*, 7(1):3–32.

Schermerhorn, P. and Scheutz, M. (2005). The effect of environmental structure on the utility of communication in hive-based swarms. In *IEEE Swarm Intelligence Symposium 2005*, pages 440–443. IEEE Computer Society Press.

Schermerhorn, P. and Scheutz, M. (2006). Social coordination without communication in multi-agent territory exploration tasks. In *Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-06)*, pages 654–661, Hakodate, Japan.

Schermerhorn, P. and Scheutz, M. (2007a). Investigating the adaptiveness of communication in multi-agent behavior coordination. *Adaptive Behavior*, 15(4):423–445.

Schermerhorn, P. and Scheutz, M. (2007b). 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.

Scheutz, M. and Schermerhorn, P. (2005). Predicting population dynamics and evolutionary trajectories based on performance evaluations in alife simulations. In *Proceedings of GECCO 2005*, pages 35–42. ACM Press.

Scheutz, M. and Schermerhorn, P. (2006). Adaptive algorithms for the dynamic distribution and parallel execution of agent-based models. *Journal of Parallel and Distributed Computing*, 66(8):1037–1051.

Scheutz, M., Schermerhorn, P., Connaughton, R., and Dingler, A. (2006). Swages–an extendable parallel grid experimentation system for large-scale agent-based alife simulations. In *Proceedings of Artificial Life X*, pages 412–418.

Wagner, A. and Arkin, R. (2004). Multi-robot communication-sensitive reconnaissance. In *Proceedings of the 2004 IEEE International Conference on Robotics and Automation*.

Werger, B. B. and Mataric, M. J. (2001). From insect to internet: Situated control for networked robot teams. *Annals of Mathematics and Artificial Intelligence*, 31(1–4):173–198.