Evolving Referential Communication in Embodied Dynamical Agents

Paul L. Williams[1], Randall D. Beer[1,2,3] and Michael Gasser[1,2]

[1]Cognitive Science Program, [2]Dept. of Computer Science, [3]Dept. of Informatics
Indiana University, Bloomington, IN 47406 USA
plw@indiana.edu

Abstract

This paper presents results from three experiments which investigate the evolution of referential communication in embodied dynamical agents. Agents, interacting with only simple sensors and motors, are evolved in a task which requires one agent to communicate the locations of spatially distant targets to another agent. The results from these experiments demonstrate a variety of successful communication strategies, providing a first step towards understanding the emergence of referential communication in terms of coordinated behavioral interactions.

Introduction

Communication is traditionally viewed as the use of signals to transmit information (Hauser, 1997; Seyfarth and Cheney, 2003; Smith and Harper, 1995). We refer to this view of communication as the IT view, for *information transmission*. Numerous models of emergent communication adopt this view as their starting point (Cangelosi and Parisi, 1998; MacLennan and Burghardt, 1993; Steels, 2003). Agents are provided with signalling mechanisms and informational content (or "meanings"), and through some adaptive process they establish shared associations between signals and meanings. However, the IT view of communication is not uncontroversial (Di Paolo, 1997, 1998), providing motivation to study emergent communication without preconceived notions of signals and information transmission. Moreover, even if the IT view is accepted, models that begin with established signalling mechanisms cannot be used to address important questions of how signals may arise from initially non-communicative behaviors.

An alternative view of communication comes from autopoetic theory (Maturana, 1978; Maturana and Varela, 1980), with similar ideas expressed by researchers in cybernetics, psychology, and a wide range of other disciplines (see Di Paolo (1997) for an extended discussion). On this view, communication occurs whenever the behavior of one agent shapes the future behavior of another agent. Thus, communication is taken to refer to all kinds of socially coordinated behaviors. We refer to this view of communication as the CB view, for *coordinated behavior*. Importantly, the CB view does not assume the existence of signals or information transmission as fundamental aspects of communication. Rather, if anything, these ideas are left to the analysis of communication by scientific observers. The essential elements of communication are the structured interactions that take place between agents in a shared domain.

Several models have explored the emergence of communicaton from a CB perspective (Baldassarre et al., 2003; Di Paolo, 2000; Iizuka and Ikegami, 2003; Nolfi, 2005). In these models, agents are typically equipped with dedicated channels to use for communication. Through some adaptive process, the agents develop the ability to use these channels to signal each other, resulting in the improved coordination of their behaviors. Thus, since these models begin without pre-specified signals, they provide compelling demonstrations of how initially non-communicative behaviors can adapt to serve communicative functions. This is particularly true when agents are equipped with only sensors and actuators, and without dedicated communication channels (Quinn, 2001; Quinn et al., 2003). In this case, simulations can provide additional insights into the interplay between communicative and non-communicative behaviors, and explore how signals may emerge from behaviors that initially evolved for other purposes.

Models that study communication from the CB perspective have typically focused on certain kinds of tasks. For instance, common tasks are those which require agents to develop signals for the dynamic assignment of roles (e.g. "leader" and "follower") in some situation. In contrast, models of communication from an IT perspective have often studied tasks that are referential in nature, where agents must develop signals that refer or "point to" states of affairs that are removed in space and/or time. Referential tasks are certainly of principal importance for understanding the evolution of communication, but to our knowledge no such tasks have been addressed within a CB framework. Accordingly, referential communication provide an important challenge for models of communication based on a CB perspective.

In this paper, we present results from a series of experiments which explore the evolution of referential communication. In these experiments, agents interact with only simple sensors and motors, and hence without any pre-specified signalling mechanisms. These experiments thus provide an initial exploration of the evolution of referential communication through the lens of a CB perspective. The results demonstrate the successful evolution of referential communication using this approach; moreover, they provide insights into the kinds of subtle communication systems that are possible through the coupled interactions of embodied dynamical agents.

The rest of this paper is organized as follows. In the next section, we expand upon the notion of referential communication and consider how it might fit with the CB view. Then we present results from a series of three experiments which explore the evolution of referential communication under various conditions. Finally, we conclude with a general discussion of the experimental results, and outline some directions for future work.

Referential Communication

One of the most widely studied examples of referential communication in a nonhuman species is the waggle dance of honeybees (Crist, 2004; Dyer, 2002). When a forager bee discovers a lucrative food source, she will often return to her hive and congregate with other hive mates. The returned forager then performs a "dance", consisting of repeated runs in a small figure-eight pattern, amidst tightly packed adjacent bees. As was first identified by Karl von Frisch (von Frisch, 1950, 1967), and later elaborated by many others (Dornhaus and Chittka, 2004; Michelsen, 2003; Riley et al., 2005), various aspects of this dance correlate with the distance and direction to the previously identified food source. Having observed the dance, other bees are then able to successfully navigate to the new food source. Thus, the waggle dance provides an excellent example of referential communication, with signals used to communicate about states of affairs that are distant in both space and time.

In what follows, we present results from a series of three experiments which were inspired by the waggle dance of honeybees. Before proceeding, it is first necessary to establish an operational notion of referential communication in terms of behavioral coordination. The difficulty is that a standard conception of referential communication fits most naturally with an IT perspective. Intuitively, signals are "about" something, and that which they are about is the information that they convey. Identifying referential communication from a CB perspective, however, is less obvious.

Communication from a CB perspective is understood in terms of the effect that interactions between agents have on the future behavioral trajectories of those agents. Here we will consider only asymmetric interactions between pairs of agents, a *sender* and a *receiver*, in which case the primary concern is the effect of interactions on the behavior of the receiver. Intuitively then, we consider an interaction to be communicative when the future behavior of the receiver is sufficiently constrained as a result of its interaction with the sender. To demonstrate this idea, consider the waggle dance example. In absence of the dance interaction, the behavioral trajectories of would-be recruits are effectively unconstrained and, in principle, the recruits may travel to any location outside of the hive. Thus, the specific effect of the dance, and what makes it communicative, is that it serves to constrain the behavior of recruits to the subset of behavioral trajectories which result in their arrival at the food source.

The additional component necessary for an account of *referential* communication is the dependence of these constrained behavioral trajectories on an object of reference. That is, the receiver's behavior resulting from a communicative interaction should change as properties of the referent change. In the waggle dance, for example, the future behavioral trajectories of the receiver vary directly and reliably with properties of the referent, namely, the distance and direction of the food source.

Another important aspect of referential communication from a CB perspective is the nature of the interaction between the agents. Specifically, in order for an interaction to be considered referential there should be a degree of separation, spatial and/or temporal, between the communicative interaction and the object of reference. To demonstrate this idea, consider again the waggle dance. Rather than perform a dance, the forager could instead gather recruits and fly with them to the food source. This would result in the same behavioral outcome for the receiver, but in this case the interaction between the bees would persist until the recruits had reached the food source. Thus, we would not consider such an interaction to be referential. In contrast, the waggle dance is referential because the communicative interaction is spatially removed from the food source.

To summarize, we propose an operational notion of referential communication from a CB perspective based on the following considerations. Firstly, the future behavior of a receiver should be constrained by its interaction with a sender. Secondly, the nature of the receiver's constrained behavior should vary based on properties of the referent. Finally, the communicative interaction should have a degree of separation from the referent.

Methods

In all of the following simulations, two agents coexist in a one-dimensional circular environment (Figure 1). Agents are able to move around the circle with a maximum angular velocity of $\frac{\pi}{8}$ in either direction, with agents free to move past each other unimpeded. Each agent is equipped with two angular sensors, one each in the clockwise and counterclockwise directions, with each sensor having a maximum range of $\frac{\pi}{8}$. An angular sensor responds with a value inversely proportional to the distance at which it intersects the other agent, with sensor values $\in [0, 1]$. In this way, the

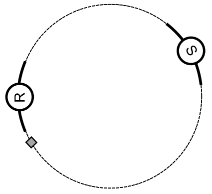

Figure 1: The agents and environment. Two agents, a sender (S) and a receiver (R), interact in a one-dimensional circular environment. On each trial, a target angle (gray diamond) is selected, and the sender must communicate the location of the target to the receiver.

angular sensors provide each agent with local information regarding the relative position of the other agent.

In addition to a pair of angular sensors, each agent is also equipped with two "bearing" sensors. The bearing sensors provide each agent with information about a certain angle ψ, with the two bearing sensors taking on the values $\frac{sin(\psi)+1}{2}$ and $\frac{cos(\psi)+1}{2}$, respectively. For one of the agents, henceforth the *receiver*, the angle ψ_R indicated by the bearing sensors is the receiver's current angular position. For the other agent, henceforth the *sender*, the angle ψ_S indicated by the bearing sensors is the separation between the sender's current angular position and a certain target angle in the environment. As will be elaborated later, the task for the agents is to get the receiver to the target angle, whose location is only available to the sender through its bearing sensors. As a result, the sender and the receiver must structure their interactions, which take place exclusively through the angular sensors, so that the target angle is successfully conveyed from the sender to the receiver.

The behavior of each agent is controlled by a continuous-time recurrent neural network (Beer, 1995) with the following state equation:

$$\tau_i \dot{s_i} = -s_i + \sum_{j=1}^{N} w_{ji}\sigma(s_j + \theta_j) + I_i \quad i = 1, \ldots, N$$

where s is the state of each neuron, τ is the time constant, w_{ji} is the strength of the connection from the j^{th} to the i^{th} neuron, θ is a bias term, $\sigma(x) = \frac{1}{1+e^{-x}}$ is the standard logistic activation function, and I represents an external input. The output of a neuron is $o_i = \sigma(s_i + \theta_i)$. In each agent, the two angular sensors and two bearing sensors are fully connected to a layer of five interneurons. The interneurons are fully interconnected and project fully to two motor neurons. The angular velocity of an agent is proportional to the difference between the outputs of the two motor neurons.

Neural parameters are evolved using a real-valued genetic algorithm with rank based selection. Each genome encodes two separate neural controllers for a sender and a receiver. Thus, rather than using a co-evolutionary procedure to evolve senders and receivers separately, we evolve a population of sender/receiver pairs. The following neural parameters, with corresponding ranges, are evolved: time constants $\in [0, 30]$, biases $\in [-16, 16]$, and connection weights (from sensors to neurons and between neurons) $\in [-16, 16]$. Successive generations are formed by first applying random Gaussian mutations to each parent genome with a mutation variance of 7 (see (Beer, 1996) for details). In addition, one-point modular crossover is applied with 5% probability, using two modules corresponding to the sender and the receiver neural controllers. A child replaces its parent if its performance is greater than or equal to that of the parent; otherwise the parent is retained.

Sender/receiver pairs are evolved for the ability of the receiver to successfully reach a number of specified target locations. On any given trial, a certain angle is designated as the target and the corresponding bearing angle inputs are given to the sender. The agents then interact for a fixed period of time, after which the receiver's final separation from the target angle is recorded. Since the information about the target angle is only available to the sender, success in this task requires that the sender and the receiver evolve a system for accurately communicating the target locations.

The performance of a sender/receiver pair is determined based on a number of evaluation trials. Each trial proceeds as follows. First, the neural states of both sender and receiver are initialized to 0. The sender is given an initial angular location of 0 and the receiver is positioned with an initial offset relative to the sender $\in \{-\frac{\pi}{32}, -\frac{2\pi}{32}, -\frac{3\pi}{32}, \frac{\pi}{32}, \frac{2\pi}{32}, \frac{3\pi}{32}\}$, such that the sender and the receiver are initially within sensory range of each other. Next, one of a set of target angles is chosen as the current target, with different sets of targets used for each experimental condition. The sender's bearing angle inputs are set to reflect the current target location and the agents are allowed to interact for an initial period of 80 time units, which is enough time for an agent moving at its maximum velocity to traverse the circular environment five times over. The duration of this initial period was chosen to minimize time pressure on the evolved communication systems. After the initial period has elapsed, the angular separation d_1 between the receiver and the target angle is recorded and normalized to run between 0 and π. The simulation is then continued for an additional 10 time units and the receiver's distance from the target is again recorded as d_2. This second value is incorporated into the fitness evaluation in order to ensure that the receiver remains at the target location. The score that a sender/receiver pair receives on a given trial is:

$$1 - \frac{d_1 + d_2}{2\pi}$$

and overall fitness is determined by averaging trial scores for every possible combination of target angles and initial receiver positions.

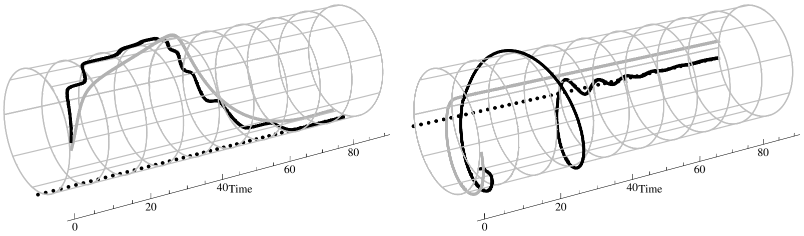

(a) (b)

Figure 2: Communication in the unconstrained condition. In each plot, the trajectories of the sender (gray line) and receiver (black line) are shown for an individual trial. The black dotted line indicates the target location. The shepherding strategy is shown in (a): the sender guides the receiver to the target location through a series of "push" or "pull" interactions. The "sit and wait" strategy is shown in (b): the sender sits a fixed distance away from the target location and waits for the receiver. Note that all plots of this kind in the rest of the paper follow the same format.

Unconstrained interactions

In the first set of simulations, agents were evaluated with a set of ten target locations uniformly distributed around the circle. Twenty evolutionary runs were performed, with a population of 400 sender/receiver pairs evolved for 2000 generations in each. The best sender/receiver pair in each run attained a fitness of at least 99%. Moreover, it was found that all of the best sender/receiver pairs had developed communicative strategies that could readily generalize to previously unseen target locations. When tested with 10,000 random trials, using target locations drawn uniformly from $[0, 2\pi]$ and initial receiver positions $\in [-\frac{3\pi}{32}, \frac{3\pi}{32}]$, the best pair from each run attained a score of at least 97%.

This first condition places no restrictions on the kinds of interactions that are available to agents. Any interaction which results in the receiver reaching the target location is acceptable, and the only aspect that matters for fitness is the receiver's final separation from the target location. Thus, in one sense, these simulations can be seen as an initial proof of concept, verifying that agents in this environment are capable of accomplishing the task. However, the results from these simulations also prove to be interesting in their own right, providing some initial insights into the kinds of interactions that agents may use for communication in this task. A preliminary inspection revealed two qualitatively distinct strategies, each of which we describe next.

The first strategy, used by 12 of the 20 best agent pairs, we refer to as "shepherding". An example of this strategy is displayed in Figure 2(a). Agents employing this strategy typically remain within or nearly within sensory contact for the entire duration of a trial. Thus, the trajectories of agents using this strategy are closely coupled. Over the course of a trial, the sender typically moves in a sustained direction towards the target location. In addition, the sender's movement towards the target is accompanied by a series of "push" or "pull" interactions with the receiver. That is, as the sender moves towards the target location, it brings the receiver along with it by closely governing its motion through repeated interactions. The sender then stops near the target location, causing the receiver to stop at or near the target. Thus, the sender's actions serve to effectively guide the receiver to the target.

The second strategy, used by 8 of the 20 best pairs, can be described as a "sit and wait" strategy. Figure 2(b) shows a characteristic interaction from a pair of agents using this strategy. At the beginning of a trial, the sender and receiver start off traveling together clockwise around the circle. The receiver continues traveling in this direction, making a full pass around the circle, while the sender stops and takes up a fixed location. Crucially, in all trials the position at which the sender stops is always the same fixed distance away from the target. The receiver continues traveling around the circle until it again reaches the position of the sender. At this point, the receiver begins moving in the opposite direction and oscillates back and forth before settling at the target location. In different versions of this strategy, the sender halts its motion at different distances away from the target and the specific details of the interaction vary, but the same general characteristics hold. To summarize, this strategy can be glossed as follows: (1) the sender moves to a fixed distance away from the target location and stops; (2) the receiver travels until it finds the sender; (3) the receiver positions itself the same fixed distance away from the sender, thus coinciding with the target location.

A common feature of the behavioral strategies evolved in this condition is that the communicative interactions continue until the receiver reaches the target location. In the "shepherding" strategy, the sender accompanies the receiver to the target location, while in the "sit and wait" strategy the sender indicates the location of the target directly using its own position. Consequently, these strategies are not considered referential, since there is no separation between the communicative interaction and the object of reference. The next two experiments address this issue by specifically

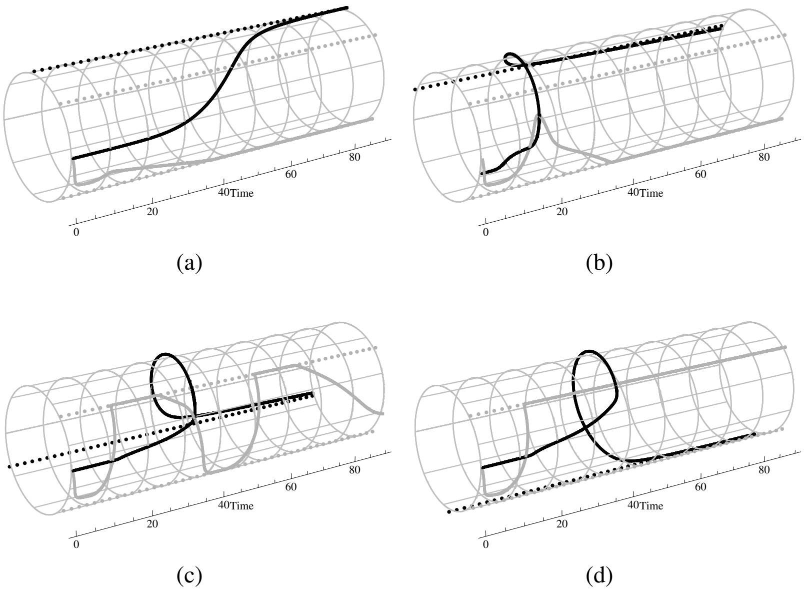

Figure 3: Communication strategy with discrete targets. Behavioral trajectories from the best sender/receiver pair are shown for each of the four target locations. The gray dotted lines indicate the boundaries of the sender constrained region. In this condition, the sender uses qualitatively distinct behavioral patterns to communicate each of the four target locations.

evolving for referential interactions.

Constrained interactions with discrete targets

There are several approaches that we could use to select for referential interactions in our simulation, which have interesting parallels with potential accounts for the evolution of referential communication in nature. One possibility would be to associate an energy cost with the behavior of the sender and to select for behaviors that minimize this cost. That is, we could select for interactions in which the sender does less while producing the same behavioral outcome for the receiver. In natural evolution, energy minimization presumably provides a strong selective advantage for referential communication. An alternative, though related, possibility is to impose spatial restrictions on the interactions between communicating agents. In considering the waggle dance, for example, it may be the case that an important selective pressure was the restriction to interactions that take place within the hive.

In the second experimental condition, we adopt the latter strategy of imposing a strict spatial constraint on the interactions between the sender and the receiver. Specifically, the sender is constrained to move within the region between $\frac{\pi}{4}$ and $-\frac{\pi}{4}$ (Figure 4). Agents in this condition are evaluated with a set of four target locations, uniformly distributed between $\frac{\pi}{2}$ and $\frac{3\pi}{2}$. Thus, given the constraint on the sender and the target locations, it is impossible for the agents to remain within sensory contact as the receiver goes to a target.

Twenty evolutionary runs were performed with a population size of 400 in each. Successful strategies proved to be more difficult to evolve in this condition, so populations were evolved for 10,000 generations as opposed to the 2,000 used in the first experiment. In all of the runs, the best pair attained a fitness measure of at least 92%, with the top 5 pairs achieving over 99%. Again we found a variety of behavioral strategies employed by different pairs of agents. We next outline one such strategy, which comes from the best sender/receiver pair.

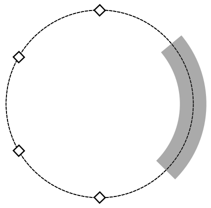

Figure 4: The constrained environment with discrete targets. The sender's motion is constrained between $-\frac{\pi}{4}$ and $\frac{\pi}{4}$ (gray region). Diamonds indicate the locations of the four targets.

Sample trajectories from the best pair of agents are displayed in Figure 3. Note the qualitatively different behavior exhibited by the sender for each of the four target locations. When the target is at $\frac{\pi}{2}$ (Figure 3(a)), the sender moves immediately to the lower boundary of the constrained region. As a result, the two agents interact for only the initial portion of the trial, after which time their trajectories quickly diverge. For the next target, proceeding counterclockwise

(Figure 3(b)), the sender motions back and forth at the beginning of the trial, before again moving to the lower boundary of the constrained region. In this case, the sender and receiver cross paths multiple times over the course of the trial. For the third target, the sender makes broad oscillatory motions between the two boundaries of the constrained region, resulting again in several intersections with the path of the receiver. Finally, for the last target location, the sender sweeps across the constrained region once before settling on the upper boundary of the constrained region, which results in the agents crossing paths relatively late in the trial.

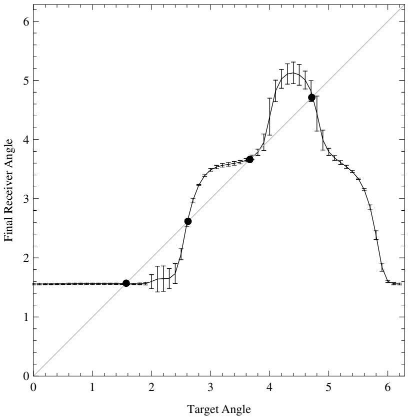

Figure 5: Generalization of the constrained agent with discrete targets. The receiver's final position across the range of target locations is shown. Black dots indicate the four target locations on which the agents were evolved.

The qualitatively distinct behaviors of the sender turn out to be a general feature of the behavioral strategies that evolved in this condition. For each of the four targets, the sender exhibits a different pattern of behavior which serves to distinguish the locations and appropriately guide the receiver. When tested on intermediate target locations, the sender makes abrupt shifts between its different behavioral regimes. A reasonable prediction then is that the behavior of the receiver will be similarly distinguished. Figure 5 shows the final position of the receiver across the range of target locations, and verifies this prediction. The receiver typically moves to one of the four standard target locations and exhibits sharp transitions between these locations for intermediate values. Thus, these results suggest that the evolved communication systems are, in a certain sense, symbolic: discrete and categorical "signals" are associated in an arbitrary fashion with a small set of possible locations.

Constrained interactions with a continuum of targets

The communicative strategies evolved in the previous experiment can be viewed as analogous to a simple form of words. Senders use a small set of essentially arbitrary signals to distinguish between a discrete number of targets. An alternative communicative strategy would involve the use of a continuum of signals to indicate a similarly continous range of possible targets. Such is presumably the case with the waggle dance, where a continuous range of dance maneuvers can be used to communicate a range of distances and directions. Other examples of this kind of communication system are the deictic indicators used by humans, such as finger pointing and eye gaze. The third experimental condition explores the evolution of this kind of referential communication.

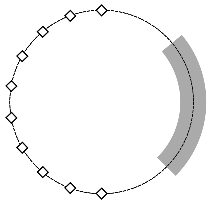

Figure 6: The constrained environment with a continuum of targets. The sender was again constrained (gray region), and the agents were evaluated with 10 target locations (diamonds).

In this experiment, agents were evolved with a set of ten target locations, uniformly distributed between $\frac{\pi}{2}$ and $\frac{3\pi}{2}$ (Figure 6). The larger number of targets was used to pressure agents into developing communicative strategies that can generalize to a continuum of locations, as opposed to using distinct signals for a small number of discrete targets. The sender was again constrained to move within the region bounded by $\frac{\pi}{4}$ and $-\frac{\pi}{4}$.

We again performed twenty evolutionary runs with this condition, with a population of 400 agents evolved for 10,000 generations in each. The best sender/receiver pair in each run achieved a fitness of at least 94%, with the top seven pairs attaining fitness scores over 99%. In order to verify that the evolved strategies could generalize to a continuum of targets, we evaluated each of the best pairs on 10,000 trials with target locations drawn uniformly from $[\frac{\pi}{2}, \frac{3\pi}{2}]$ and initial receiver offsets $\in [-\frac{3\pi}{32}, \frac{3\pi}{32}]$. The best pair from each run achieved a score of at least 90%, with the top seven pairs all scoring in excess of 98%. Figure 8 shows the generalization performance of the best sender/receiver pair. We next describe the behavioral strategy used by this pair of agents.

Sample trajectories from the best sender/receiver pair are displayed in Figure 7. In each case, the interaction between the agents consists primarily of two path crossings. At the beginning of the trial, the sender moves counterclockwise while the receiver makes a full pass around the circle in the clockwise direction. The paths of the agents then cross for the first time, after which the receiver turns back in the counterclockwise direction. The agents then cross each other for the second time, followed by the receiver continuing around

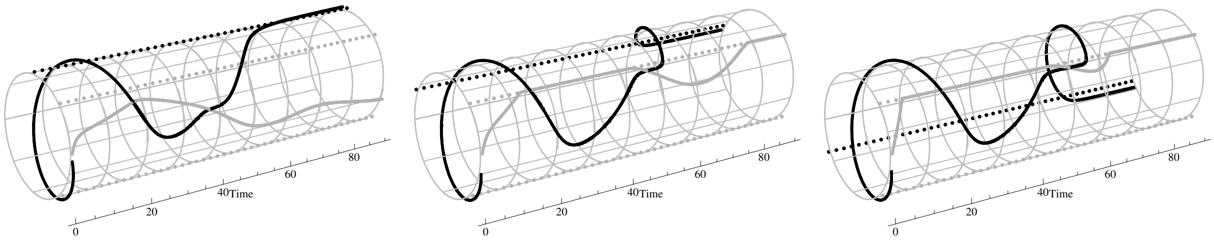

(a) (b) (c)

Figure 7: Communication strategy with a continuum of targets. The gray dotted lines indicate the boundaries of the sender constrained region. For targets further from the sender in the clockwise direction, the sender's trajectory becomes increasingly straightened against the upper boundary of the constrained region.

the circle in the clockwise direction before stopping at the target location. Note that, for targets further from the sender in the counterclockwise direction, the sender's trajectory becomes increasingly straightened, effectively flattening out against the upper boundary of the constrained region. Importantly, this change in the sender's trajectory affects the timing and location of the path crossings between the two agents. As the sender's path becomes straighter, the first intersection between the paths occurs earlier and the second intersection occurs later. Figure 9 demonstrates this more clearly, where trajectories are shown for a range of target locations. The interactions between the agents vary smoothly with the location of the target. As a result, the agents are able to systematically communicate the location of targets anywhere along the continuum.

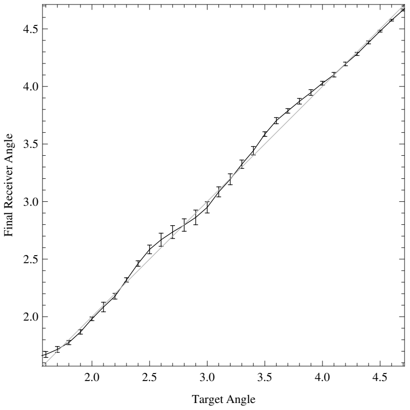

Figure 8: Generalization of the constrained agent with a continuum of targets.

Discussion

Previous models of emergent referential communication are based on the view of communication as information transmission. In this paper, we presented results from a set of experiments which explored the evolution of referential communication from the perspective of coordinated behaviors. Embodied dynamical agents, interacting with only basic sensory and motor capabilities, evolved strategies to communicate the locations of spatially distant targets.

In the first experimental condition, agents were evolved with no restrictions placed on their potential interactions. In this case, we found that agents developed two distinct strategies that were successful in guiding the receiver to the targets. These results provided an initial proof of concept, as well as giving some insights into the kind and variety of communicative strategies that are possible.

The second experiment selected specifically for referential interactions by placing spatial restrictions on the motion of the sender. With a small set of discrete targets we found that senders evolved a number of distinct behavioral patterns to communicate the different target locations. Thus, in a loose sense, the communication systems evolved in this condition can be viewed as similar to words.

Finally, in the third condition, we explored the possibility of evolving referential communication which could generalize to a continuum of target locations. The "signals" evolved in this condition were found to vary smoothly with the range of targets, resulting in successful generalization. This ability to indicate a continuum of locations is analogous to the various deictic indicators used in human communication, such as finger pointing and eye gaze.

Future directions

The next step for future work is to perform detailed analyses of the evolved communication systems. Using the mathematical tools of dynamical systems theory, we will explore how the underlying dynamical structure of individual agents forms the basis for their joint interactions. Additionally, we plan to apply analytical techniques from information theory to examine the possibility of developing a mathematically rigorous notion of information transmission between interacting agents.

In other work, we plan to extend the approach used here to a two-dimensional environment. With the added dimensionality, we could explore tasks in which agents have to

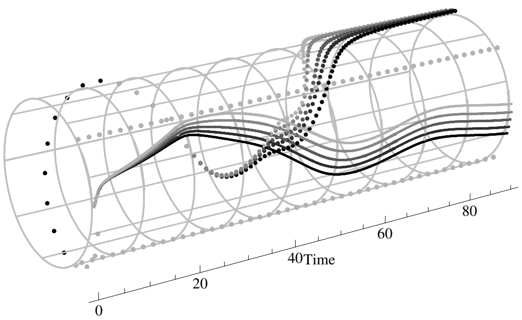

Figure 9: Signalling systematicity of the constrained agent with a continuum of targets. Trajectories of the sender (solid lines) and receiver (dotted lines) are shown for a set of equally spaced target locations, with color used to indicate trajectories from the same trial. The behavioral interactions of the agents vary smoothly with the location of the target.

communicate multiple kinds of information simultaneously. For example, we could bring the task closer to its waggle dance inspiration by evolving agents to communicate both the distance and direction to target locations.

Finally, we intend to investigate the evolutionary trajectories of the evolved communication strategies. To do so, we can track the individual lineages of sender/receiver pairs and study their behavioral interactions at different times during the course of evolution. Such an investigation should provide unique insights into the process by which various non-communicative behaviors adapt to serve communicative functions.

Acknowledgements

This investigation was supported by the National Institute of Health and National Research Service Award HD007475-13.

References

Baldassarre, G., Nolfi, S., and Parisi, D. (2003). Evolving mobile robots able to display collective behaviors. *Artificial Life*, 9(3):255–267.

Beer, R. D. (1995). On the dynamics of small continuous-time recurrent neural networks. *Adaptive Behavior*, 3(4):469–509.

Beer, R. D. (1996). Toward the evolution of dynamical neural networks for minimally cognitive behavior. *From Animals to Animats 4*, pages 421–429.

Cangelosi, A. and Parisi, D. (1998). The emergence of a 'language' in an evolving population of neural networks. *Connection Science*, 10(2):83–97.

Crist, E. (2004). Can an insect speak?: The case of the honeybee dance language. *Social Studies of Science*, 34(1):7–43.

Di Paolo, E. A. (1997). Social coordination and spatial organization: Steps towards the evolution of communication. *Fourth European Conference on Artificial Life*, pages 464–473.

Di Paolo, E. A. (1998). An investigation into the evolution of communication. *Adaptive Behavior*, 6(2):285–324.

Di Paolo, E. A. (2000). Behavioral coordination, structural congruence and entrainment in a simulation of acoustically coupled agents. *Adaptive Behavior*, 8(1):25–46.

Dornhaus, A. and Chittka, L. (2004). Why do honey bees dance? *Behavioral Ecology and Sociobiology*, 55(4):395–401.

Dyer, F. C. (2002). The biology of the dance language. *Annual Review of Entomology*, 47(1):917–949.

Hauser, M. D. (1997). *The evolution of communication*. MIT Press.

Iizuka, H. and Ikegami, T. (2003). Adaptive coupling and intersubjectivity in simulated turn-taking behaviours. *Advances in Artificial Life: 7th European Conference*.

MacLennan, B. J. and Burghardt, G. M. (1993). Synthetic ethology and the evolution of cooperative communication. *Adaptive Behavior*, 2(2):161–188.

Maturana, H. R. (1978). Biology of language: The epistemology of reality. *Psychology and biology of language and thought: Essays in honor of Eric Lenneberg*, pages 27–63.

Maturana, H. R. and Varela, F. J. (1980). *Autopoiesis and cognition: The realization of the living*. D. Reidel.

Michelsen, A. (2003). Signals and flexibility in the dance communication of honeybees. *Journal of Comparative Physiology A: Sensory, Neural, and Behavioral Physiology*, 189(3):165–174.

Nolfi, S. (2005). Emergence of communication in embodied agents: Co-adapting communicative and non-communicative behaviours. *Connection Science*, 17(3):231–248.

Quinn, M. (2001). Evolving communication without dedicated communication channels. *Advances in Artificial Life: 6th European Conference*, 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: Mathematical, Physical and Engineering Sciences*, 361(1811):2321–2343.

Riley, J. R., Greggers, U., Smith, A. D., Reynolds, D. R., and Menzel, R. (2005). The flight paths of honeybees recruited by the waggle dance. *Nature (London)*, 435(7039):205–207.

Seyfarth, R. M. and Cheney, D. L. (2003). Signalers and receivers in animal communication. *Annual Reviews in Psychology*, 54(1):145–173.

Smith, M. J. and Harper, D. G. (1995). Animal signals: models and terminology. *Journal of Theoretical Biology*, 177(3):305–311.

Steels, L. (2003). Evolving grounded communication for robots. *Trends in Cognitive Sciences*, 7(7):308–312.

von Frisch, K. (1950). *Bees: their vision, chemical senses, and language*. Cornell University Press.

von Frisch, K. (1967). *The dance language and orientation of bees*. Harvard University Press.