

Directional Communication in Evolved Multiagent Teams

In: *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 2014)*. New York, NY: ACM

Justin K. Pugh
Department of EECS
University of Central Florida
Orlando, FL 32816 USA
jpugh@eecs.ucf.edu

Skyler Goodell
Department of EECS
University of Central Florida
Orlando, FL 32816 USA
goodsky@eecs.ucf.edu

Kenneth O. Stanley
Department of EECS
University of Central Florida
Orlando, FL 32816 USA
kstanley@eecs.ucf.edu

ABSTRACT

The question of how to best design a communication architecture is becoming increasingly important for evolving autonomous multiagent systems. Directional reception of signals, a design feature of communication that appears in most animals, is present in only some existing artificial communication systems. This paper hypothesizes that such directional reception benefits the evolution of communicating autonomous agents because it simplifies the language required to express positional information, which is critical to solving many group coordination tasks. This hypothesis is tested by comparing the evolutionary performance of several alternative communication architectures (both directional and non-directional) in a multiagent foraging domain designed to require a basic “come here” type of signal for the optimal solution. Results confirm that directional reception is a key ingredient in the evolutionary tractability of effective communication. Furthermore, the real world viability of directional reception is demonstrated through the successful transfer of the best evolved controllers to real robots. The conclusion is that directional reception is important to consider when designing communication architectures for more complicated tasks in the future.

Categories and Subject Descriptors

I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence—*intelligent agents, multiagent systems, coherence and coordination*; I.2.6 [Artificial Intelligence]: Learning—*connectionism and neural nets*

Keywords

Multirobot Teams; Coordination; Communication; Artificial Neural Networks; Evolutionary Algorithms; HyperNEAT

1. INTRODUCTION

Because there exists a substantial class of tasks that require more than one agent to solve [28], evolving multiagent robot

teams has attracted significant interest. Due to the limitations inherent in centralized control systems [18], a popular alternative is to evolve controllers for multiple autonomous agents that confront the challenge of group coordination through distributed interactions. While some researchers have built controllers for non-communicating autonomous agents [4, 5], communication can potentially improve group coordination through sharing information about the sensed environment [17] or agent states [10]. The question then becomes how to design an effective communication scheme. While there are many important design features of communication systems to consider [14], this paper focuses on one of the most basic features: *directional reception* of signals. When this feature is present, agents can perceive the direction from which incoming signals originate.

Directional reception in the design of communication systems for autonomous robots is not a new idea [7, 8, 19, 33]. However, many systems without directional reception have also been proposed [3, 17, 22, 31] and despite the interest in both communication schemes with and without directional reception, few empirical studies have investigated the importance of directionality. The hypothesis in this paper, suggested by the prevalence of directional reception in natural communication systems, is that directional reception is beneficial to the evolution of communicating autonomous agents that must coordinate in a spatial environment. This hypothesis is important to confirm in particular to provide guidance to new researchers entering the field who may assume from prior such experiments that lacked directionality that such directionality is not necessary.

Consider the following thought experiment: Imagine that you are trying to help a friend who is on the other side of a crowded room to come to your location. If you can only communicate through text messages, the message you send may be similar to “come to the southwest corner of the room, near the grey statue,” or, “turn 120 degrees to your left and walk towards me.” However, by shouting across the room, you can simply say “come here.” The reason the message complexity reduces is that the positional information that is explicit in the longer messages is implicit in the compact message because humans benefit from directional reception of auditory signals. Similarly, agents equipped with directional reception in a spatial environment should be able to communicate with less complex messages than those without directional reception, which benefits evolution because simpler signals are easier to evolve.

This hypothesis is validated by an experiment comparing the evolutionary performance of directional and non-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
GECCO'14, July 12–16, 2014, Vancouver, BC, Canada.
Copyright 2014 ACM 978-1-4503-2662-9/14/07 ...\$15.00.
<http://dx.doi.org/10.1145/2576768.2598299>.

directional communication schemes in a multiagent foraging domain designed to require a “come here” type of signal. Controllers in the experiment are evolved with the HyperNEAT algorithm for evolving large-scale artificial neural networks (ANNs) [11, 24], which has been successfully applied to multi-agent domains in the past [3–5]. Unlike many neuroevolution methods, HyperNEAT can be informed by geometric information in the domain, making it an ideal platform for testing directional communication. Experimental results support the hypothesis that directional communication schemes are more evolutionarily tractable than their non-directional counterparts, highlighting that directional reception is important for enabling the evolution of group coordination tasks because it redirects evolutionary effort from learning a complex language to learning task-critical behaviors. Finally, the practical value of the evolved controllers is demonstrated by transferring the best teams to real-world Khepera III robots.

It is important to emphasize that while the main hypothesis may be intuitive, many robot coordination experiments in the past nevertheless have lacked such directional reception [3, 17, 22, 31]. The need for directional communication has in effect rarely if ever been addressed explicitly. Thus this paper makes a practical contribution by serving as a reminder that including directionality is often in practice more important than allowing complex communication when contemplating the setup of future experiments in evolved coordination.

2. BACKGROUND

This section reviews prior work in training communicating agents and summarizes the HyperNEAT neuroevolution algorithm featured in the experiments presented in this paper.

2.1 Communication in Multiagent Teams

Within the field of evolutionary computation, communication among agents has been studied from several perspectives. One of the most popular questions is how natural communication may have emerged under evolutionary pressure; Wagner et al. [30] provide a detailed review of past work in this area on the emergence of communication. One aspect of communication that is often present in simulated models but is rarely the subject of empirical testing is *directional reception*, which is the ability of an agent, upon hearing a message, to sense the direction of the message’s source. Directional reception is one of the 13 design features of animal communication identified by Hockett [14] and is present in even the least advanced communicating mammals. Perhaps the reason its importance is not often tested is that when tracing the evolution of humans back to land mammals, directional reception co-occurs with communication as an immediate implication simply of having ears.

Communication is also often important for agents in a cooperative multiagent team. One situation in which agents benefit from communication is when each agent senses only a fraction of the observable state of the environment. In such tasks, communication facilitates sharing information among team members to increase the amount of environmental context available to each agent [17]. Communication may also be necessary in tasks in which agents have limited information about the state of other agents on the team, such as when an agent requires the assistance of other agents to solve a subgoal [10].

Realizing the importance of communication in solving cooperative tasks, researchers developing evolved or learned

controllers for cooperative teams of robots have experimented with different forms of communication, many of which include a form of directional reception. For example, Yong and Mikkulainen [33] developed neural controllers that received as input the relative positions of all team members. Marocco and Nolfi [19] developed simulated robots with a single communication output that produces signals whose intensity and direction is perceived by team members through directional pie-slice sensors. Di Paolo [7] developed robots that could “speak” using simulated sound waves and who could thereby infer the relative position of a speaker through a set of ears positioned on opposite sides of the body. Floreano et al. [8] evolved physical robot controllers that communicate their positions through light.

Considerable effort has also been devoted to developing communicating agents that are not aware of the relative position of the speaker. In pioneering work on the evolution of communication, Werner and Dyer [31] evolved agents that communicate via three-bit binary signals to solve a task in which stationary “female” agents must guide blind “male” agents to their position on a discretized two-dimensional grid. Jim and Giles [17] evolved teams of agents to solve a discrete predator-prey task; each agent in the domain reads and writes bit strings to a shared message board. In a recent work, D’Ambrosio et al. [3] solved a multiagent synchronization task by evolving a neural controller that features direct neural connections between agents. In another recent work, Rawal et al. [22] evolved agents for a discrete predator-prey task in which blind predator agents receive real-valued communication signals from stationary predator agents. In each of Werner and Dyer [31], Jim and Giles [17], D’Ambrosio et al. [3], and Rawal et al. [22], positional information is critically important to solving the task even though it is not explicitly included in the communication scheme. The binary-string languages that subsequently emerged in Werner and Dyer [31] and Jim and Giles [17] and the real-valued languages from Rawal et al. [22] attach positional information to certain “words” in the language. D’Ambrosio et al. [3] in contrast achieved a solution by extracting positional information from assumptions about the starting positions and orientations of the agents.

This paper advances the hypothesis that knowing the relative location of a speaker is often critically important in solving group coordination tasks and that it is beneficial to include such information implicitly in the communication scheme so that evolutionary effort need not be expended incorporating it into the emergent language. A recent neuroevolution technique called HyperNEAT [11, 24] easily facilitates the evolution of agents with a position-aware communication scheme because HyperNEAT can take into account geometric information about the domain (such as the spatial relationship between an agent’s vision or audition sensors), making it a good platform for testing this hypothesis. The following sections review the HyperNEAT approach, which is applied in the experiments in this paper.

2.2 Neuroevolution of Augmenting Topologies

The HyperNEAT approach is itself an extension of the original NEAT (Neuroevolution of Augmenting Topologies) algorithm that evolves increasingly large ANNs [25, 27]. NEAT starts with a population of simple networks that then *increase in complexity* over generations by adding new nodes and connections through mutations. By evolving ANNs in this way,

the topology of the network does not need to be known a priori; NEAT searches through increasingly complex networks to find a suitable level of complexity. Because it starts simply and gradually adds complexity, it tends to find a solution network close to the minimal necessary size. However, as explained in the next section, the direct representation of nodes and connections in the NEAT genome cannot scale up to larger networks that can take advantage of domain geometry. For a complete overview of NEAT, see Stanley and Miikkulainen [25] or Stanley and Miikkulainen [27].

2.3 HyperNEAT

Like NEAT, many neuroevolution methods are *directly encoded*, which means each component of the phenotype is encoded by a single gene, making the discovery of repeating motifs expensive and improbable [32]. Therefore, indirect encodings [1, 2, 15, 20, 26] have become a growing area of interest in evolutionary computation.

One such indirect encoding designed explicitly for neural networks is in Hypercube-based NEAT (HyperNEAT) [11, 24], which is itself an indirect extension of the directly-encoded NEAT approach [25, 27] reviewed in the previous section. This section briefly reviews HyperNEAT; a complete introduction can be found in Stanley et al. [24] and Gauci and Stanley [11]. Rather than expressing connection weights as independent parameters in the genome, HyperNEAT allows them to vary across the phenotype in a regular pattern through an indirect encoding called a *compositional pattern producing network* (CPPN; [23]), which is like an ANN, but with specially-chosen activation functions.

CPPNs in HyperNEAT *encode* the connectivity patterns of ANNs as a *function of geometry*. That is, if an ANN’s nodes are embedded in a geometry, i.e. assigned coordinates within a space, then it is possible to represent its connectivity as a single evolved function of such coordinates. In effect the CPPN paints a pattern of weights across the geometry of a neural network. Because the CPPN encoding is itself a network, it is evolved in HyperNEAT by the NEAT algorithm, which is designed to evolve networks of increasing complexity. To understand why this approach is promising, consider that a natural organism’s brain is physically embedded within a three-dimensional geometric space, and that such embedding heavily constrains and influences the brain’s connectivity. Topographic maps (i.e. ordered projections of sensory or effector systems such as the retina or musculature) in natural brains preserve geometric relationships between high-dimensional sensor and effector fields [16, 29]. In other words, there is important information *implicit* in geometry that can only be exploited by an encoding informed by such geometry.

In particular, geometric *regularities* such as symmetry or repetition are pervasive throughout the connectivity of natural brains. To similarly achieve such regularities, CPPNs exploit activation functions that induce regularities in HyperNEAT networks. The general idea is that a CPPN takes as input the geometric coordinates of two nodes embedded in the *substrate*, i.e. an ANN situated in a particular geometry, and outputs the weight of the connection between those two nodes (figure 1). In this way, a Gaussian activation function by virtue of its symmetry can induce symmetric connectivity and a sine function can induce networks with repeated elements. Note that because the size of the CPPN is decoupled from the size of the substrate, HyperNEAT can compactly

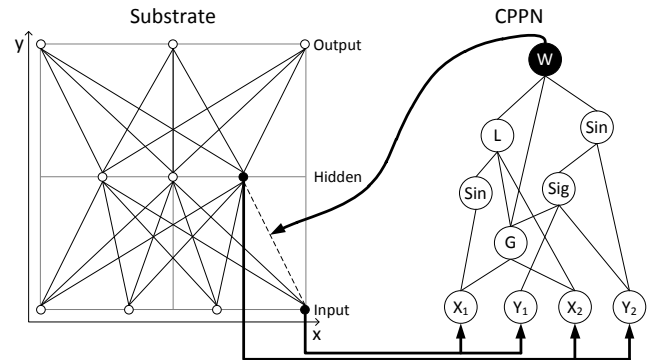


Figure 1: HyperNEAT example. An example substrate (left) for a simple ANN contains ten neurons that have been assigned (x, y) coordinates. The weight of every connection specified in the substrate is determined by the evolved CPPN (right): (1) The coordinates of the source (x_1, y_1) and target (x_2, y_2) neurons are input into the CPPN, (2) the CPPN is activated, and (3) the weight w of the connection being queried is set to the CPPN’s output. CPPN activation functions in this paper can be *sigmoid* (*Sig*), *Gaussian* (*G*), *linear* (*L*), or *sine* (*Sin*).

encode the connectivity of an arbitrarily large substrate with a single CPPN.

Additionally, HyperNEAT can evolve controllers for *teams* of agents. This multiagent HyperNEAT algorithm was first introduced by D’Ambrosio and Stanley [5], D’Ambrosio et al. [4], and D’Ambrosio and Stanley [6]. It can be used to evolve both homogeneous and heterogeneous teams; however, the experiments in this paper only necessitate the homogeneous case. Controllers for homogeneous teams are created by evolving a single controller that is duplicated for each agent on the team.

3. EXPERIMENT

Each of the communication schemes reviewed in Section 2.1 that succeed without directional reception share a common feature: the domain in which they are applied is simplified. For example, the simulated worlds in Werner and Dyer [31], Jim and Giles [17], and Rawal et al. [22] consist of discrete two-dimensional grids, and while the world in D’Ambrosio et al. [3] is continuous, agents are restricted to only forward-backward movement, effectively narrowing actions to one dimension. To test the hypothesis that non-directional communication schemes require more evolutionary effort than their directional counterparts, a more challenging domain is proposed that takes place in a continuous two-dimensional world.

A team of five agents is placed in a large rectangular room bounded on all four sides by walls but otherwise free of obstructions. Agents must collect the greatest number of food items possible from the room within a time limit. Only one food item is present in the room at a given time; whenever one is collected, a new one spawns randomly somewhere in the room. The key mechanism encouraging communication is that food items are not collected until they are touched by three agents simultaneously. While it is possible for all three agents to arrive at a food item through random wandering (without communication), the second and third agents can

find the food much faster if the first agent to find it signals its location. The type of communication that is required for the optimal behavior is thus relatively unsophisticated: agents must produce a *come here* signal upon discovering a food item. However, this type of signal is significantly more difficult to produce when directional reception is not embedded in the communication architecture because extra “words” must be discovered to describe the location of “here” (a concept that is implicit within directional reception). While this domain is relatively simple, the underlying dynamic of rallying a group of cooperating agents to salient locations in real time is fundamental to a range of important real world tasks such as rescue, patrol, and retrieval tasks. In this way, this experiment helps to highlight the potential importance of directional reception across a range of real-world problems.

In addition to a control scheme with no communication (**NoCom**), three different communication schemes are tested to determine the extent to which directional reception is important in solving this group foraging task. In the first scheme, **DirCom**, agents can emit signals on one channel (which can range from 0.0 to 1.0) and can hear signals coming from one of ten directions (equally spaced around the agent) via an array of ten pie-slice sensors. In the second communication scheme, **OneBit**, agents can also emit signals on only one channel, but cannot hear signals directionally. Instead, agents have five communication inputs for hearing signals from each of the five agents on the team. The input for sensing an agent’s own signals is disabled for consistency with the DirCom scheme, in which agents cannot hear themselves either. The final communication scheme, **FiveBit**, is the same as OneBit except that agents can emit signals on five channels and have an array of 25 communication inputs (one for each channel on each agent). The extra “bits” can potentially make possible a more complex language capable of expressing the positional information necessary to compensate for a lack of directional reception. However, whether such a language is evolutionarily tractable will be determined experimentally.

Individual agents are equipped with several arrays of pie-slice sensors for the various senses that serve as inputs into a HyperNEAT-evolved ANN. The HyperNEAT substrate that serves as the basis for all the variant setups is shown in figure 2. This kind of *multi-spatial substrate* is shown effective for tasks with multiple modalities in Pugh and Stanley [21]. Agents sense the location of food items within a maximum radius of 100 units with a set of five equal-size pie-slice sensors that span the frontal 180 degrees of vision. Similarly, agents detect walls with a set of five 100-unit rangefinders arranged across the front of each agent in 36 degree intervals. Each agent also detects the location of other agents with a set of ten unlimited-range¹ pie-slice sensors that surround the agent. Figure 3 depicts the different sensor types. Each agent has the ability to move forward and to turn left or right (a set of three output neurons control the movement: left, forward, and right, where the direction and magnitude of each turn is decided by the difference between the left and right turn outputs).

DirCom, OneBit, and FiveBit agents are identical in every way except for their communication scheme, which differ as follows. DirCom (figure 4a) and OneBit (figure 4b) agents each have a single output neuron for sending simple commu-

¹Activation of the friend sensors that detect other agents is floored at 20% if sensing an agent more than 500 units away.

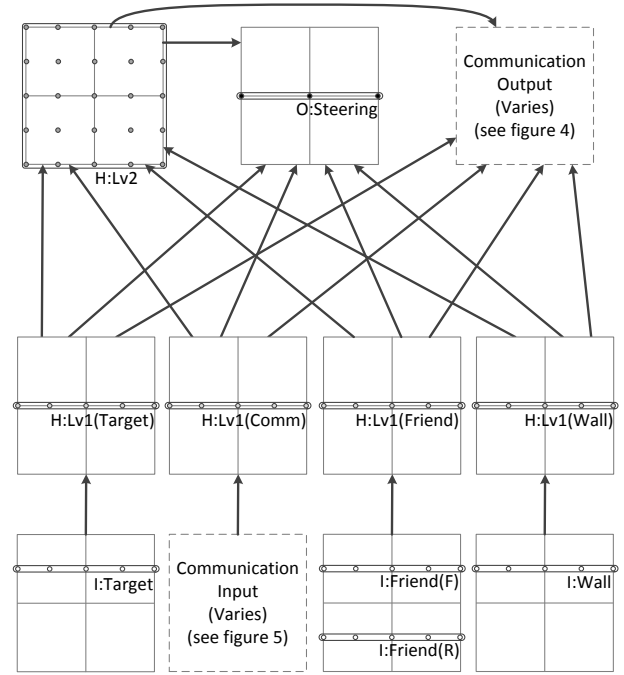


Figure 2: HyperNEAT substrate (all communication schemes). Each input and output modality and each hidden layer is placed on its own plane within the substrate, following the multi-spatial substrate configuration described in Pugh and Stanley [21]. The substrate is strictly feedforward: Each input modality feeds into a dedicated “level 1” hidden layer, all level 1 hidden layers feed into a common “level 2” hidden layer, and all hidden layers (level 1 as well as level 2) feed into all outputs. Individual neural connections are omitted for clarity. Instead, arrows indicate the existence of neural connections between two planes. Planes shown as connected are potentially fully connected (all neurons on the first plane are queried by HyperNEAT for connections to all neurons on the second plane). Communication input and output planes vary depending on communication scheme; see figure 4 for output planes and figure 5 for input planes.

nication signals, while FiveBit (figure 4c) has a set of five output neurons for sending more complex signals. DirCom agents have a set of ten input neurons for receiving communication signals (figure 5a), one for each pie-slice in their 360 degree non-overlapping array. OneBit agents have a set of five input neurons for receiving communication signals (figure 5b), one for each agent on the team (one of which is never activated because agents cannot hear themselves). FiveBit agents have a set of 25 input neurons for receiving communication signals (figure 5c), five for each of five agents on the team (five of which are never activated). Recall that because HyperNEAT is an indirect encoding, the dimensionality of the inputs is not an obstacle to effective learning [24].

The room is 1,000 by 900 units, and the five agents are initially arranged in a horizontal line in the center of the room spaced 100 units apart. Food items spawn randomly around the room, not closer than 40 units from a wall. During evolution, team performance is averaged over 20 trials to mitigate the evaluation noise caused by randomized food

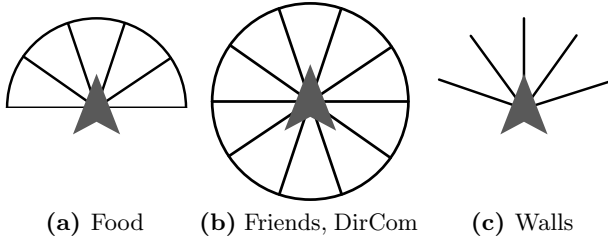


Figure 3: Agent sensor types. Agents have pie slice sensors for detecting food (a) and friends (b), and a set of rangefinders for detecting walls (c). DirCom robots also have a set of pie slice sensors for sensing communication signals (b). The activation level of all sensor types is based on distance (i.e. it decreases linearly the farther the distance) except for DirCom communication sensor activation, which is based on the strongest signal received from each direction.

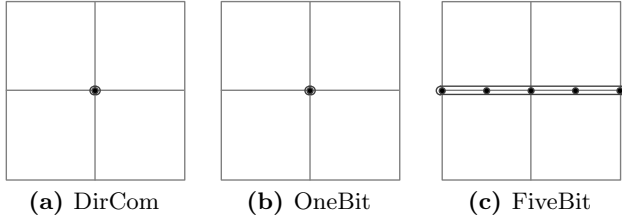


Figure 4: Substrate: communication output planes. The number of neurons on the communication output plane corresponds to the size of the signal vector sent to all other agents on the team on each tick: one for DirCom (a) and OneBit (b), and five for FiveBit (c).

item spawn points. Each trial, teams are given 2,000 ticks of simulation time to collect as many food items as possible, up to a maximum of ten food items. Team fitness on a particular trial is determined by the number of food items seen, the number collected, and the time at which each was collected. More specifically, each food item is worth a maximum of 100 points, 10 of which are awarded when a single agent comes within range (the food is “seen”), 40 of which are awarded when three agents come within range (thus collecting the food), and 50 of which are time dependent (50 points are awarded if the food is collected on the first tick of simulation, diminishing to 0 points awarded if the food is collected on the last tick). The time component is included to provide a smoother fitness gradient for evolution to follow.

Evolution is run for 1,000 generations, by which time fitness inevitably stops improving. The training phase consists of 20 runs of evolution for each of the three communication schemes, after which the champions are tested according to a more stable metric: Testing performance is determined by the raw number of food items collected in 5,000 time ticks, averaged over 10,000 trials. During testing, there is no artificial limit of 10 food items, although there are practical limits due to the time it takes agents to travel across the room (agents move at a maximum rate of 5 units per tick and have a maximum turn rate of 36 degrees per tick).

Finally, to highlight the practical value of effective communication, the best performing teams are implemented on

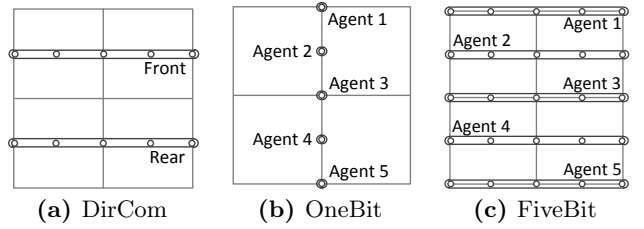


Figure 5: Substrate: communication input planes. Communication input for DirCom (a) consists of two sets of five neurons corresponding to the front and rear pie-slice sensors. Each neuron in the communication input plane for OneBit (b) corresponds to an agent on the team and is activated when detecting signals from that agent. The communication input plane for FiveBit (c) consists of a 5×5 grid arranged as five horizontal groups. Each horizontal group corresponds to an agent on the team, and each neuron within a group corresponds to a different communication channel.

real-world Khepera III robots. In the real world implementation, several sensor types are implemented by a central station that reads each agent’s position and orientation information on every tick via odometry on their wheel encoders. This information is synthesized to yield communication sensors, friend vision, and target vision. Hardware IR sensors serve as the wall sensors. While in practice positional information in the real world may not always be possible to compute through a central computing node, it is important to note that this setup approximates other real world setups that would convey similar information, such as true hardware-based emitters [8] or decentralized mutual localization techniques [9].

It is also important to note that in the real world the IR sensors also detect the presence of other robots as well as the target points (plastic cups), causing those entities to be treated like walls. This ambiguity is contrary to the case during evolution where wall sensors only activate when in range of a wall. To partially mitigate this potential discrepancy, walls in the real world are covered with retroreflective tape while robots and target points are not. Thus agents perceive a strong activation when seeing walls and a comparatively weak activation when seeing other solid objects in the room. The evolved policies are robust enough to work when transferred to the real world despite any remaining differences.

Because HyperNEAT differs from original NEAT only in its set of activation functions, it uses the same parameters [25]. The experiment was run with a modified version of the public domain SharpNEAT package [12]. The size of the population was 500 with 20% elitism. Sexual offspring (50%) did not undergo mutation. Asexual offspring (50%) had 0.96 probability of link weight mutation, 0.03 chance of link addition, and 0.01 chance of node addition. The coefficients for determining species similarity were 1.0 for nodes and connections and 0.1 for weights. The available CPPN activation functions were sigmoid, Gaussian, linear, and sine, all with equal probability of being added to the CPPN. Parameter settings are based on standard SharpNEAT defaults and prior reported settings for NEAT [25, 27]. They were

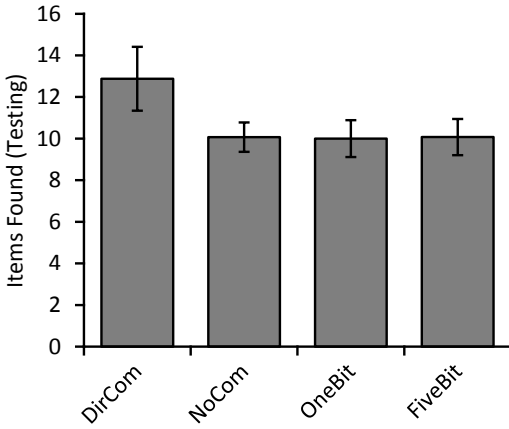


Figure 6: Test performance. The testing performance of a single run is determined by the average number of food items collected over 10,000 trials by the team with the highest training performance. Results for each communication scheme are averaged over 20 runs. Error bars depict a 95% confidence interval. The DirCom scheme performs significantly better than the other schemes ($p < 0.05$; Student’s t-test).

found to be robust to moderate variation through preliminary experimentation.

4. RESULTS

For each communication scheme, each of the best performing genomes found after 1,000 generations of evolution (one for each of 20 runs) is evaluated according to the testing metric described in the previous section. The test performance for each communication architecture, averaged over the 20 champions, is presented in figure 6. The directional communication scheme significantly outperforms all others ($p < 0.05$; Student’s t-test), while there is no significant difference between the OneBit, FiveBit, and NoCom architectures. The best performing NoCom champion collects 12.3 food on average during testing. Nine out of the 20 DirCom champions pass a success threshold of collecting more than 105% of the food collected by the best NoCom champion, while the number of such successes for OneBit and FiveBit are zero and one, respectively. It is important to note that the differences between DirCom and the other variants in figure 6 represent qualitatively very different behaviors.

For example, the best performing NoCom teams exhibit surprisingly effective exploratory behavior. The very best consists of a single agent closely following the walls and corners of the arena while the remaining four agents lag behind seeking the leading agent and at the same time checking the interior of the room. If food is discovered in the interior, the swarm of agents quickly collects it, and if food is discovered near a wall, the wall-exploring agent stops at it and the swarm of agents is able to catch up and thus collect the food in a reasonable time. However, this highly coordinated behavior is rare; most NoCom teams achieve scores closer to 10.0 through random wandering. The best OneBit teams exhibit a similar style of coordination as the best NoCom team, with scores of 12.7 and 12.6. All of the OneBit teams either send chaotic signals at all times, usually with no discernible purpose (signals changed little if at all when agents

were in range of a target point or wall), or else do not send signals at all. In contrast, the best FiveBit team produces a rudimentary “come here” signal that allows it to achieve good results, with a score of 14.4. Agents on this team spread out and explore the map individually. When an agent finds the food, it sends a signal that causes all other agents to change behavior and begin to seek nearby agents. This inevitably causes the agents to capture the food, although it is clear that they do not know the location of the agent sending the signal because often all four of the remaining agents will join up before moving to the fifth (and thus the food). While this strategy is effective, that it was only achieved in one of the FiveBit runs suggests the difficulty of discovering it. From the behaviors of teams without directional communication, it is evident that they fail to express position even implicitly within their communication schemes because once the food is found by one agent, no teams approach it right away. Furthermore, while a more sophisticated communication strategy is conceivable whereby agents e.g. near the wall indicate that they are near a wall when they see food, no such strategy actually evolved with FiveBit, highlighting the difficulty of fully exploiting the increased bandwidth without a directional component.

In contrast to the non-directional schemes, half of the DirCom teams learned to produce a “come here” signal upon discovering the food. The testing performance of these teams varied (from 13.0 to 19.3) according to the effectiveness of non-communicative behaviors such as exploration and the reliability of turning towards a detected food item without moving past it (and thus out of range). The best DirCom team, with a test score of 19.3, combines “come here” signaling with efficient exploratory behaviors (agents tend to spread out as much as possible).

The five best performing DirCom teams were transferred to real Khepera III robots and placed in an open arena containing a single food item. Transferred teams demon-

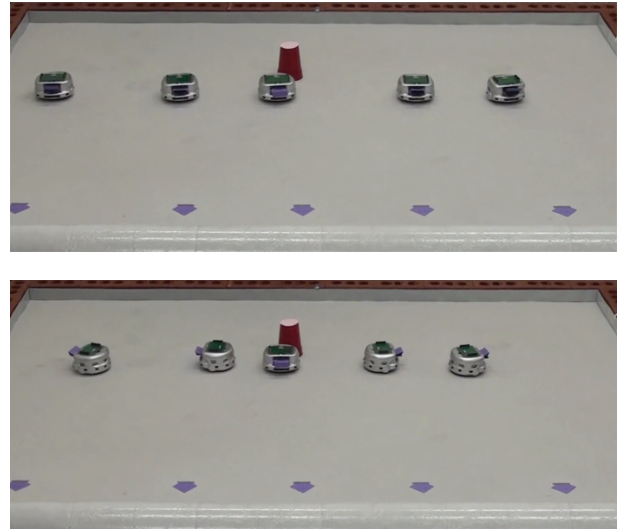


Figure 7: Real world implementation. In this example, the robots explore the arena by moving forward from their starting positions. The third robot finds the target point (top) and signals its location to the other robots who then begin moving towards it (bottom).

strated effective group coordination (figure 7), despite significant differences between the real world and the simulated world (e.g. robots can collide with each other in the real world, while they cannot do so in simulation). Real world teams are tested with varying team sizes of five, four, and three robots without losing the ability to solve the task. Videos of some of the transferred teams can be found at <http://tinyurl.com/DirComVideos>.

5. DISCUSSION

The results indicate that the communication scheme that includes directional reception is able to solve the task more efficiently than those schemes without directional reception. In fact, only one out of 40 teams across both non-directional communication schemes achieves a level of performance unreachable by non-communicating teams. By comparison, almost half of the teams with directional reception achieve such performance. This result demonstrates that even though it is possible to evolve effective communication without directional reception, doing so is less evolutionarily tractable than when directional reception is present in the system.

Messages in a system without directional reception must express directionality through a complex system of “words” that each contain different positional information. These words must be invented independently, and they may have little or no benefit until several have already been invented. This problem is difficult for evolutionary algorithms because there may exist no path of increasingly fit stepping stones to the discovery of the final working language. Directional reception alleviates this problem by incorporating positional information into the communication architecture itself so that agents automatically learn the entire vocabulary of positional information as soon as a single word is invented. In this way, effective signaling is more evolutionarily tractable because directionality can be discovered in a single step.

Directional reception can be implemented in the real world in a number of ways. For example, it is possible to perceive the direction of signals sent via sound or light [8]. However, naïve implementations of such approaches may not be robust to environmental factors such as bright rooms or poor acoustics. More promisingly, directional communication can also be simulated if robots have an accurate means of localization. The real world experiment in this paper demonstrates a centralized localization approach based on wheel encoder odometry that facilitates the real world implementation. However, stronger methods exist, such as the decentralized localization proposed by Franchi et al. [9], or the system for localized communication proposed by Gutiérrez et al. [13] that uses frequency modulation to overcome environmental IR-sensor noise.

The significant advantage demonstrated in the experiment by the communication scheme with directional reception over the non-directional approaches points to the importance of considering including directional reception when formulating multiagent communication architectures for future robot coordination problems. Furthermore, communication features other than directional reception likely have a significant impact on evolutionary performance, but have yet to be empirically tested. One such feature is directional projection - the ability of agents to project communication signals only in a specific direction. Overall, with realistic real world options for directional sensing and signal projection available, such

schemes merit serious consideration for complex multiagent tasks.

6. CONCLUSIONS

This paper presented an empirical study of directional reception for the evolution of neural controllers applied to a group coordination task. A communication scheme with directional reception was compared to two non-directional schemes with different potentials for message complexity, as well as to a scheme in which communication is disabled. The results indicate that directional communication is more evolutionarily tractable than non-directional approaches because it eliminates the need to discover a complex language to express the positional information that is crucial to optimally solving the task. The real world viability of the directional communication scheme was demonstrated by successfully transferring the best evolved neural controllers to Khepera III robots. Thus this paper serves as a reminder of the importance of directional reception and suggests that it is an important language feature to consider when designing communication architectures for more complicated tasks in the future.

7. ACKNOWLEDGMENTS

This work was supported through a grant from the US Army Research Office (Award No. W911NF-11-1-0489). This paper does not necessarily reflect the position or policy of the government, and no official endorsement should be inferred.

References

- [1] P. J. Bentley and S. Kumar. Three ways to grow designs: A comparison of embryogenies for an evolutionary design problem. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 1999)*, pages 35–43, 1999.
- [2] J. C. Bongard and R. Pfeifer. Evolving complete agents using artificial ontogeny. *Morpho-functional Machines: The New Species (Designing Embodied Intelligence)*, pages 237–258, 2003.
- [3] D. B. D’Ambrosio, S. Goodell, J. Lehman, S. Risi, and K. O. Stanley. Multirobot behavior synchronization through direct neural network communication. In *Proceedings of the 5th International Conference on Intelligent Robotics and Applications (ICIRA 2012)*, New York, NY, 2012. Springer-Verlag.
- [4] D. B. D’Ambrosio, J. Lehman, S. Risi, and K. O. Stanley. Evolving policy geometry for scalable multiagent learning. In *Proceedings of the Ninth International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, pages 731–738. International Foundation for Autonomous Agents and Multiagent Systems, 2010.
- [5] D. B. D’Ambrosio and K. O. Stanley. Generative encoding for multiagent learning. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 2008)*, New York, NY, 2008. ACM Press.
- [6] D. B. D’Ambrosio and K. O. Stanley. Scalable multiagent learning through indirect encoding of policy geometry. *Evolutionary Intelligence*, 2013. To appear.
- [7] E. Di Paolo. Behavioral coordination, structural congruence and entrainment in a simulation of acoustically coupled agents. *Adaptive Behavior*, 8(1):27–48, 2000.

- [8] D. Floreano, S. Mitri, S. Magnenat, and L. Keller. Evolutionary conditions for the emergence of communication in robots. *Current Biology*, 17(6):514–519, 2007.
- [9] A. Franchi, G. Oriolo, and P. Stegagno. Mutual localization in a multi-robot system with anonymous relative position measures. In *Proceedings of the International Conference on Intelligent Robots and Systems (IROS 2009)*, pages 3974–3980. IEEE, 2009.
- [10] A. Garland and R. Alterman. Multiagent learning through collective memory. In *Adaptation, Coevolution and Learning in Multiagent Systems: Papers from the 1996 AAAI Spring Symposium*, pages 33–38, 1996.
- [11] J. Gauci and K. O. Stanley. Autonomous evolution of topographic regularities in artificial neural networks. *Neural Computation*, 22(7):1860–1898, 2010.
- [12] C. Green. SharpNEAT homepage. <http://sharpneat.sourceforge.net/>, 2003–2006.
- [13] A. Gutiérrez, A. Campo, M. Dorigo, D. Amor, L. Magdalena, and F. Monasterio-Huelin. An open localization and local communication embodied sensor. *Sensors*, 8(11):7545–7563, 2008.
- [14] C. F. Hockett. The origin of speech. *Scientific American*, 203:88–111, 1960.
- [15] G. S. Hornby and J. B. Pollack. Creating high-level components with a generative representation for body-brain evolution. *Artificial Life*, 8(3), 2002.
- [16] D. Hubel and T. Wiesel. Receptive fields, binocular interaction and functional architecture in the cat’s visual cortex. *The Journal of Physiology*, 160(1):106, 1962.
- [17] K.-C. Jim and C. L. Giles. How communication can improve the performance of multi-agent systems. In *Proceedings of the Fifth International Conference on Autonomous Agents (ICAA 2001)*, pages 584–591, New York, NY, USA, 2001. ACM.
- [18] J. L. King. Centralized versus decentralized computing: organizational considerations and management options. *ACM Computing Surveys*, 15(4):319–349, 1983.
- [19] D. Marocco and S. Nolfi. Emergence of communication in embodied agents evolved for the ability to solve a collective navigation problem. *Connection Science*, 19(1):53–74, 2007.
- [20] J. Miller. Evolving a self-repairing, self-regulating, french flag organism. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 2004)*, pages 129–139. Springer, 2004.
- [21] J. K. Pugh and K. O. Stanley. Evolving multimodal controllers with HyperNEAT. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 2013)*, New York, NY, 2013. ACM Press. To appear.
- [22] A. Rawal, P. Rajagopalan, R. Miikkulainen, and K. Holecamp. Evolution of a communication code in cooperative tasks. In *Artificial Life*, volume 13, pages 243–250, 2012.
- [23] K. O. Stanley. Compositional pattern producing networks: A novel abstraction of development. *Genetic Programming and Evolvable Machines Special Issue on Developmental Systems*, 8(2):131–162, 2007.
- [24] K. O. Stanley, D. B. D’Ambrosio, and J. Gauci. A hypercube-based indirect encoding for evolving large-scale neural networks. *Artificial Life*, 15(2):185–212, 2009.
- [25] K. O. Stanley and R. Miikkulainen. Evolving neural networks through augmenting topologies. *Evolutionary Computation*, 10(2):99–127, 2002.
- [26] K. O. Stanley and R. Miikkulainen. A taxonomy for artificial embryogeny. *Artificial Life*, 9(2):93–130, 2003.
- [27] K. O. Stanley and R. Miikkulainen. Competitive coevolution through evolutionary complexification. *Journal of Artificial Intelligence Research*, 21:63–100, 2004.
- [28] P. Stone and M. Veloso. Multiagent systems: A survey from a machine learning perspective. *Autonomous Robots*, 8(3):345–383, 2000.
- [29] S. Udin and J. Fawcett. Formation of topographic maps. *Annual Review of Neuroscience*, 11(1):289–327, 1988.
- [30] K. Wagner, J. A. Reggia, J. Uriageraka, and G. S. Wilkinson. Progress in the simulation of emergent communication and language. *Adaptive Behavior*, 11(1):37–69, 2003.
- [31] G. Werner and M. Dyer. Evolution of communication in artificial organisms. In *Proceedings of the Second International Conference of Artificial Life (ALIFE 1991)*, pages 659–687, Redwood City, CA, 1992. Addison-Wesley.
- [32] X. Yao. Evolving artificial neural networks. *Proceedings of the IEEE*, 87(9):1423–1447, 1999.
- [33] C. Yong and R. Miikkulainen. Coevolution of role-based cooperation in multiagent systems. *IEEE Transactions on Autonomous Mental Development*, 1(3):170–186, 2009.