The Evolution of Communication in Adaptive Agents

Gregory M. Saunders and Jordan B. Pollack

Laboratory for Artificial Intelligence Research
Department of Computer and Information Science
The Ohio State University
Columbus, Ohio 43210
USA

saunders@cis.ohio-state.edu
pollack@cis.ohio-state.edu

Technical Report GS-94-EVCOMM

Abstract

This paper explores how communication can be understood as an adaptation by agents to their environment. We model agents as recurrent neural networks. After arguing against systems which use discrete symbols to evolve communication, we supply our agents with a number of continuous communications channels. The agents use these channels to initiate real-valued signals which propagate through the environment, decaying over distance, perhaps being perturbed by environmental noise. Initially, the agents' signals appear random; over time, a structure emerges as the agents learn to communicate task-specific information about their environment. We demonstrate how different communication schemes can evolve for a task, and then discover a commonality between the schemes in terms of information passed between agents. From this we discuss what it means to communicate, and describe how a semantics emerges in the agents' signals relative to their task domain.

Key Words: communication; evolutionary algorithms; autonomous agents; neural networks

Running Head: "The Evolution of Communication"

1 Introduction

The field of adaptive behavior holds that higher-level cognitive skills arise from the more primitive ability of an agent to adapt to its environment (Beer, 1990). Although many behaviors have been studied in this bottom-up fashion – perception (e.g., Arkin, Schweighofer, and Thach, 1994; Harvey, Husbands, and Cliff, 1994), action selection (e.g., Maes, 1991, 1992; Tyrell, 1993), navigation (e.g., Brooks, 1986, 1991a), food collection (Arkin, 1992; Deneubourg, et al., 1991), planning (e.g., Agre & Chapman, 1986; Donnart and Meyer, 1994), predator avoidance (Grefenstette, 1992; Schmajuk, 1994), locomotion (Beer, 1990; Beer and Gallagher, 1992), social activity (Mataric, 1993, 1994), etc.[1] – relatively few people have studied *communication* as adaptive behavior. In this paper, we explore how communication can be understood as an adaptation by agents to their environment.

The role of communication in multi-agent systems remains one of the most important open issues in multi-agent system design (Brooks, 1991b; Arkin and Hobbs, 1993). Some have approached this problem top-down, first incorporating perception, navigation, communication, competition, and cooperation within a single, rich environment (Steels, 1994; McFarland, 1994), although at this stage in the research the agents are still designed.

Others have approached the problem of communication bottom-up, turning towards adaptation and evolution as a source of insight (Yanco and Stein, 1993; Werner and Dyer, 1992; Robbins, 1994; MacLennan, 1992; MacLennan and Burghardt, 1993; Collins and Jefferson, 1991, 1992; de Bourcier and Wheeler, 1994), but their work all shares an emphasis on *discrete* communication.

Yanco and Stein (1993) investigate a “follow-the-leader” task in which one agent, the leader, receives a command which must be followed by a group of agents. The leader chooses one of *n*

---

[1] See Meyer and Guillot (1991, 1994) for extensive reviews.

symbols to represent the command, broadcasts the symbol to the other agents, and the subordinates respond. A reinforcement algorithm governs both the encodings of the leader and responses of the subordinates; over time, a consensus emerges between the two.

Werner and Dyer (1992) describe a more complex environment in which simulated animals must communicate to find mates. Females, while stationary, can sense potential mates within a limited range and "call out" to them by emitting a signal. Males, wandering around the environment, lack the capacity to produce signals or see the females directly, but they can sense the females' signals and respond by moving toward them. Using a neural network representation for agents and a genetic algorithm for search, Werner and Dyer show that the sexes can agree on a common language.[2] Robbins (1994) shows that introducing artificial parasites can create more robust languages in this task.

MacLennan (1992) adopts a higher-level view of language by defining an abstract task in which a group of agents must learn to communicate. Each agent possesses local information in terms of one of *n* symbols; it chooses a second symbol (from a set of *n*) to convey that information, and other agents must respond appropriately. Using finite state machines to represent agents and a genetic algorithm, MacLennan shows how the group of agents evolve a common symbol-symbol mapping.

Collins and Jefferson (1991, 1992) study AntFarm, a simulated ant colony in which agents must learn to communicate the presence of food. At each time step, an agent drops between 0 and 64 units of pheromone, which then diffuses throughout the environment as a signal to other ants.

---

2. Werner and Dyer (1993) propose a very interesting model "BioLand" which supports the evolution of communication as well, but the results focus on herding behavior rather than the evolved communication scheme, and it is unclear how the signals generated by the agents affect their behavior.

Although they have yet to evolve cooperative foraging, the work sheds some light on representational issues, in particular, on the use of neural networks as an agent representation.

de Bourcier and Wheeler (1994) adopt a slightly different focus: rather than studying cooperation among agents, they investigate how competition can lead to dishonest signalling schemes. The basic interpretation of communication is fixed, however – an increase in perceived signal strength *always* translates into an increased probability that an agent will move away from the source.

Two sorts of problems run throughout all this work, both engineering and theoretical. From an engineering perspective, adopting discrete symbols for communication implicitly make a strong assumption about the number of messages needed for a given task; namely, that this number is finite. Furthermore, because each symbol is topologically isolated from the others, these systems suffer from what we call the “lookup table syndrome.” If noise perturbs a symbol during communication (say from 0 to 1), there is not necessarily any relationship between how the two symbols are perceived by the recipient. Thus, barring the evolution of a parity bit or even more complex redundancy scheme (e.g., Shannon and Weaver, 1948), discrete-communication systems will be highly sensitive to noise.

From a theoretical perspective, adopting discrete symbols for communication implicitly solves the signal detection problem. When an agent perceives “001,” for instance, it knows with certainty that another agent is trying to communication. In contrast, when an agent in the real world perceives a sound, it might be another agent, or a nearby stream, a car, etc. By assuming discrete symbols, the above systems in effect assume well-defined units of communication. This lends no insight into how words arise from a continuous sound wave, or how words are then combined into variable-length sentences by the agents. Finally, some studies make an architectural distinction between the agent sending the message and the recipient (Yanco and Stein, 1993; Werner and Dyer, 1992; and to some extend MacLennan, 1992, in the sense that at any given time, there is a privileged agent attempting to convey its local information to the others).

2 Communication with Continuous Symbols

Our approach to understanding multi-agent communication differs significantly from all the work described above. Saunders, Kolen, and Pollack (1994) showed the importance of a continuous substrate for evolving adaptive agents. Subsequently, rather than assume the transmission of discrete signals between agents, we provide our agents with continuous channels capable of supporting a wide variety of communication schemes. Furthermore, we make no architectural distinctions between transmitter and receiver.

As will be described below, we model agents as connectionist networks. We supply each agent with a number of communications channels implemented by the addition of both input and output units for each channel. The output units initiate environmental signals which are perturbed by environmental noise and whose amplitude decays over distance. An agent does not receive input from other individuals, rather the agent's input reflects the summation of all other agents' output signals along that channel. Because we use real-valued activations, the agents communicate using real-valued vectors. Under an evolutionary program, the agents coevolve a communication scheme over continuous channels which in order to be successful conveys task-specific information.

This section describes our experiments. First we briefly describe GNARL, the algorithm we use to evolve our agents. Then we introduce an extension of the Tracker task (Jefferson et al., 1992), which will serve as a substrate for our experiments. Next, we describe the method of communication our agents employ. Experimental results will be discussed in Section 3.

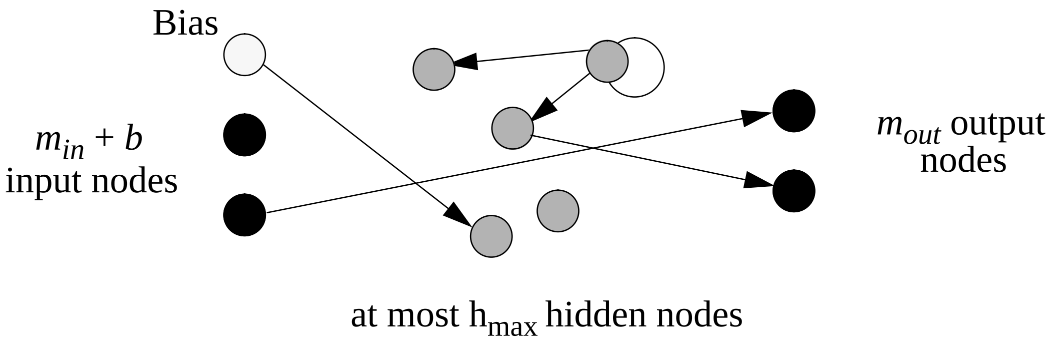

Figure 1: Sample initial network. The number of input nodes (m_{in}) and number of output nodes (m_{out}) is fixed for a given task. The presence of a bias node (b = 0 or 1) as well as the maximum number of hidden units (h_{max}) is set by the user. The initial connectivity is chosen randomly (see text). The disconnected hidden node does not affect this particular network's computation, but is available as a resource for structural mutations.

2.1 GNARL

GNARL (Saunders, Angeline, and Pollack, 1994; Angeline, Saunders, and Pollack, 1994) is an algorithm based on evolutionary programming (Fogel, 1992) that induces recurrent neural networks. It provides a mechanism for the simultaneous acquisition of network structure and weight values. GNARL employs a population of networks and uses a fitness function's unsupervised feedback to guide search through network space.

The algorithm is described in detail elsewhere (Angeline, Saunders, and Pollack, 1994). Briefly, it begins with a population of *n* random individuals; a sample network *N* is shown in Figure 1. The number of input nodes (m_{in}) and number of output nodes (m_{out}) are fixed for a given task; the number of hidden nodes as well as the connections among them are free to vary from 0 to a user-supplied maximum h_{max}. Links use real-valued weights, but their presence is free to vary, so that GNARL networks may have no connections, sparse connections, or full connectivity.

In each epoch of search, the networks are ranked by a user-supplied fitness function $f\colon \mathbf{N} \longrightarrow \mathbb{R}$, where $\mathbf{N}$ represents the space of networks, and $\mathbb{R}$ represents the reals. Reproduction of the best

$^{n}/_{2}$ individuals entails modifying both the weights and structure of each parent network *N*. First, the temperature *T(N)* is calculated:

$$T(N) = 1 - \frac{f(N)}{f_{max}} \tag{1}$$

where f_{max} (provided by the user) is the maximum possible fitness for a given task. This measure of *N*'s performance is used to anneal the *structural* and *parametric* (Barto, 1990) similarity between parent and offspring, so that networks with a high temperature are mutated severely, and those with a low temperature are mutated only slightly. This allows a coarse-grained search initially, and a finer-grained search as a network approaches a solution (cf. Kirkpatrick, Gelatt, and Vecchi, 1983).

More concretely, parametric mutations are accomplished by perturbing each weight with gaussian noise, whose variance is $T(N)^2$:

$$w \leftarrow w + \text{Normal}\,(0; T(N)), \qquad \forall w \in N \tag{2}$$

Structural mutations are accomplished by:

- adding k_1 hidden nodes with probability $p_{add\text{-}node}$
- deleting k_2 hidden nodes with probability $p_{delete\text{-}node}$
- adding k_3 links with probability $p_{add\text{-}link}$
- deleting k_4 links with probability $p_{delete\text{-}link}$

where each k_i is selected uniformly from a user-defined range, again annealed by *T(N)*. When a node is added, it is initialized without connections; when a node is deleted, all its incident links are removed. All new links are initialized to 0.

GNARL has been applied to several different problems (Angeline, Saunders, and Pollack, 1994). In particular, we have applied GNARL to the Tracker task (Jefferson et al., 1992) in which

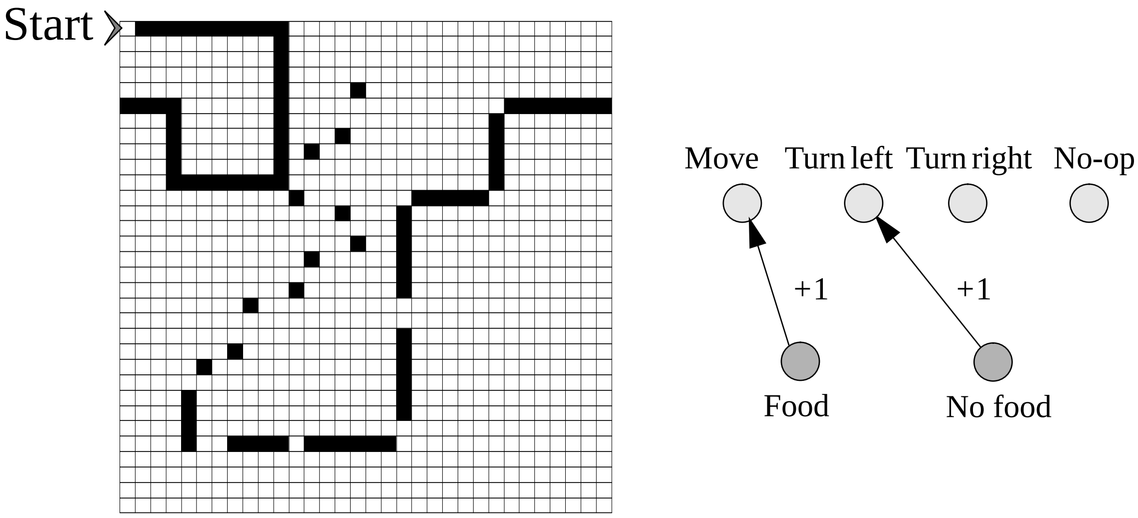

(a) **(b)**

Figure 2: The Tracker task. (a) The trail is connected initially, but becomes progressively more difficult to follow. The underlying 2-d grid is toroidal; (b) The semantics of the I/O units for the ant network. The first input node denotes the presence of food in the square directly in front of the ant; the second denotes the absence of food in this same square. No-op, from Jefferson et al., allows the network to stay in one position while activation flows through recurrent links.

a simulated ant must learn to follow a broken trail of food (Figure 2a). Each ant receives two inputs: one indicating the presence of food in the square directly before the agent; and another detecting the absence of food in that same square. Jefferson, et al., allowed four primitive actions: move-forward (and implicitly eat food if present), turn left, turn right, and no-op (Figure 2b). Under these conditions GNARL evolved several different networks as a solution, one of which closely approximates the finite-state automaton shown in Figure 3.[3]

2.2 The Tracker Task, Revisited

To study the evolution of communication in groups of agents, we extend the Tracker task in three ways (Figure 4):

---

[3] Note however that the network's behavior is not precisely captured by the FSA. Kolen (1994a, 1994b) shows that, in general, FSAs approximate networks only poorly. Another network induced by GNARL makes this point empirically. (See Saunders, Angeline, and Pollack, 1994).

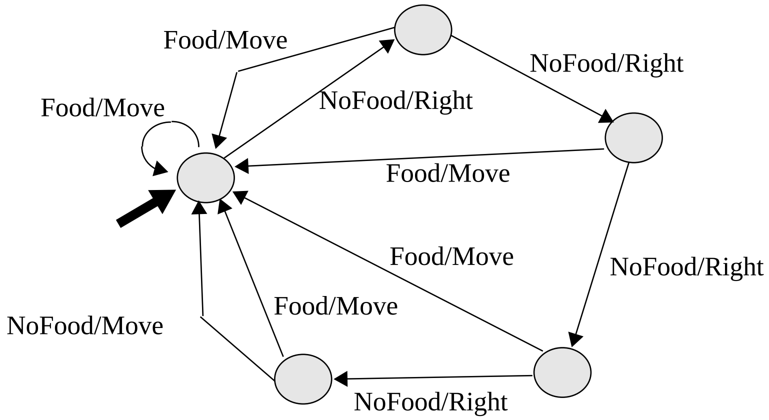

Figure 3: FSA hand-crafted for the Tracker task, from Jefferson, et al., 1992. The large arrow indicates the initial state. This simple system implements the strategy “move forward if there is food in front of you, otherwise turn right four times, looking for food. If food is found while turning, pursue it, otherwise, move forward one step and repeat.”

- increasing the number of agents
- increasing the size of the grid to accommodate these agents
- moving all the food to a small area in the center of the environment

We assume that these modifications will shift the emphasis of the task from evolution of local internal state to evolution of distributed external state, i.e., communication. We concentrate the food within one area so that when an agent finds it and communicates, some food remains by the time other agents arrive. The size of the environment and the amount of food it contains far exceed the capabilities of a single ant: in the limited time available an ant can neither search the entire space nor consume all the food therein. Thus (we assume) the task design ensures that the only method of complete success necessarily involves communication among the agents.

2.3 An Architecture for Communication

When faced with a task requiring communication, the architecture of Jefferson, et al., (Figure 2b) will certainly fail; namely, because it in no way supports communication. To remedy this shortcoming, we add *n* additional input and output units to the network of Figure 2b, representing *n* channels

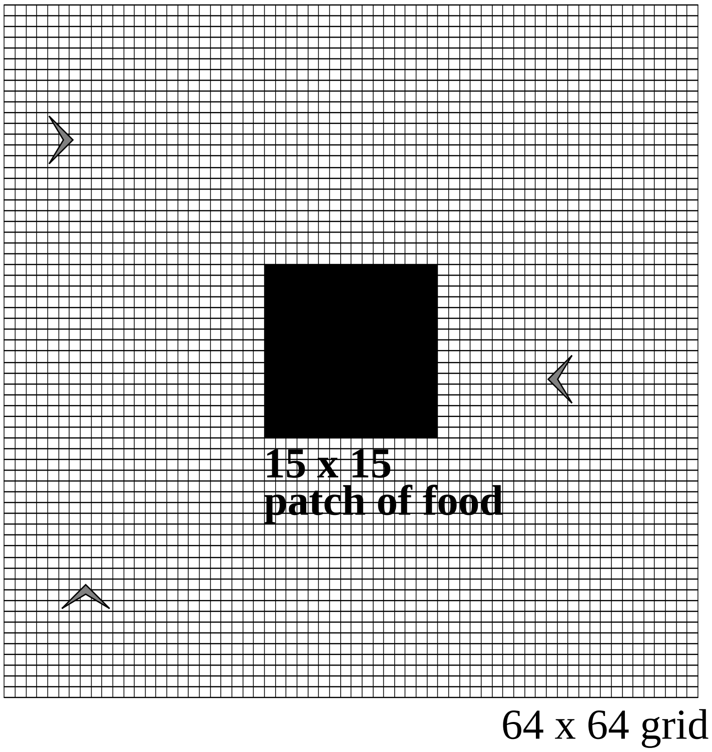

Figure 4: Environment for the modified Tracker task.

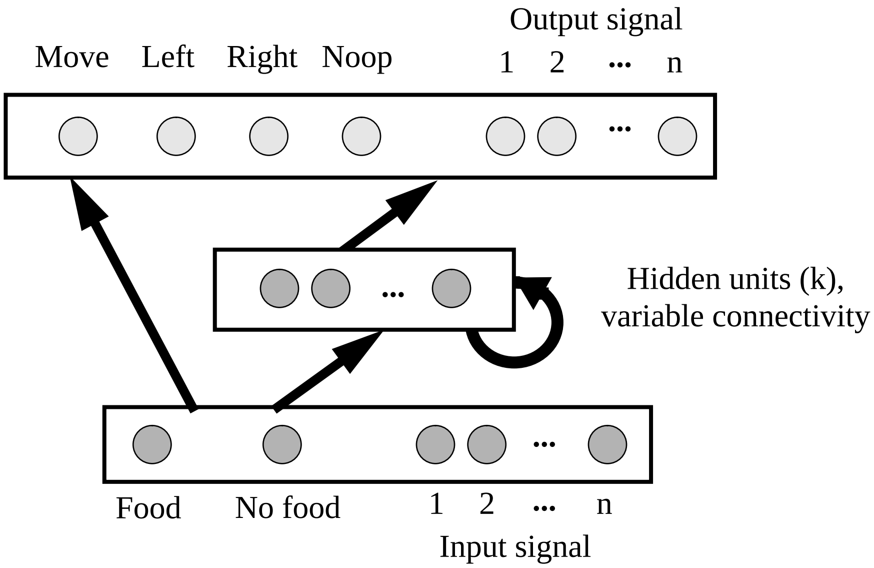

Figure 5: The semantics of the I/O units for evolving communication. The "food/nofood" inputs and "move/left/right/noop" outputs are from the Tracker task. The additional nodes, described in the text, give the agent the ability to perceive and generate signals.

of communication (Figure 5.) We maintain, from the original study, an implicit winner-take-all network on the (non-signal) outputs.

Output signals propagate throughout the environment, decaying in inverse proportion to squared distance.[4] Perception of these signals is governed by Equation 3. The input signal to agent a along the i^{th} channel, $s_{IN}(a, i)$, is a summation of the signals of all other agents along this channel. A is the set of agents, $s_{out}(b, i)$ is the i^{th} output signal of agent b. The noise in the channel, $U[-u_i, u_i]$ is a uniform random number with range specific to the channel, and σ is a linear threshold function, which bounds the signals in all channels to a user-specified range $[s_{min}, s_{max}]$. In the experiments below, $s_{min} = 0$ and $s_{max} = 40$.

$$s_{IN}(a, i) = \sum_{\begin{bmatrix} b \in A \\ b \neq a \end{bmatrix}} \frac{\sigma(s_{OUT}(b, i) + U[-u_i, u_i])}{distance^2(a, b)} \qquad \text{(Eqn 3)}$$

Effectively, this equation creates a sound-like model of signal propagation, an example of which is shown in Figure 6.

For the studies reported in this paper, all activations are continuous; only the hidden activation is squashed (with the standard sigmoid function). Fitness is measured by simply observing the total amount of food eaten by the group. Because our interest is in communication rather than structural variation, we use a version of the GNARL algorithm in which the number of hidden units is fixed. The number of links, though, varies as described above. Finally, all agents in an environment are homogeneous in that they share not only the architecture of Figure 5, but also common weights. As show below, however, their behaviors will be quite different depending upon each agent's perspective of its world, creating a heterogenous group at the behavioral level.

---

4. We assume that the signals propagate much faster than the agents react (as would a sound wave), so that effectively, at each discrete time step, an agent's output signals establish a wave front whose strength decays over distance.

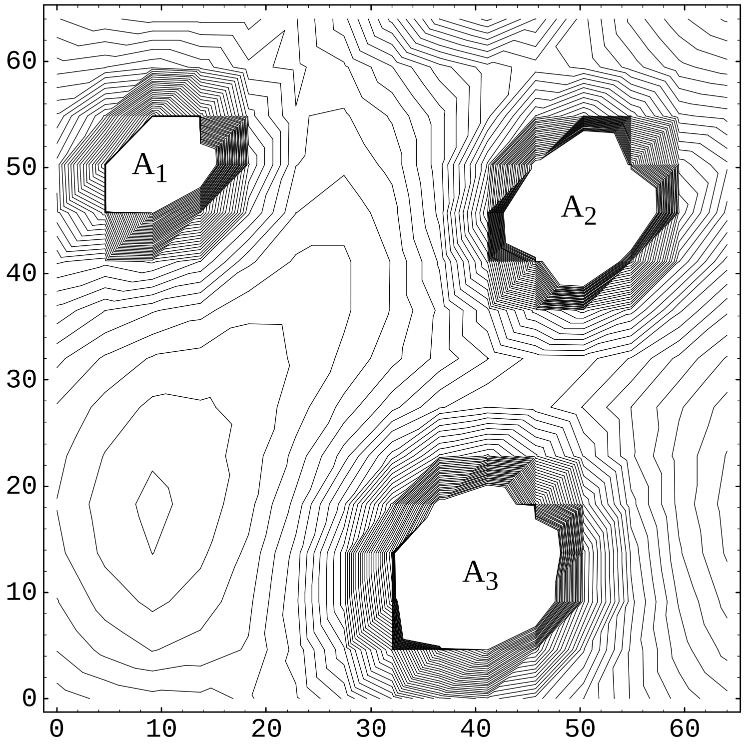

Figure 6: The sound-like model of signal propagation created by Equation 3. Agent 1 outputs a relative signal strength of 1; agent 2 outputs a relative signal strength of 2; and agent 3 outputs a relative signal strength of 3. The topological map indicates how these signals interact, with the highest elevations corresponding to the strongest signals.

3 Results, Part I

Our reason for studying communication as an adaptation was our belief that task-specific communication schemes can indeed be evolved for cooperative multi-agent systems. Unfortunately, our first set of results does not support this belief, for the architecture of Figure 5 offers an easier solution for the modified Tracker task than communication, as we will see below. We present these results in detail for three reasons: first, the negative results reflect our experimental history in exploring the evolution of communication; second, they will motivate the architecture which does support our claim (Section 5); and finally, they illustrate the opportunistic nature of adaptive agent systems.

3.1 Experiment 1: The Curmudgeon Strategy

We begin with a very simple case: 2 agents, each with one hidden unit, capable of passing one real number between each other, with no noise ($u_0 = 0$, see Equation 3). Figure 7 shows a series of snapshots as the agents collect food. The radii of the circles correspond to the strength of communication.[5] These agents do fairly well, scoring 196 on the task.

To investigate whether these agents had learned to communicate the presence of food, we plotted the status of each agent over the course of the run (Figures 8 and 9). These and similar figures describe all there is to know about a particular agent. The lowest graph represents the presence of food: when food is detected, the value spikes to one; otherwise it is zero. The next graph(s) represent the input signal(s) from the other agent(s), one graph per communication channel. The “hidden unit” graph(s) indicates the agent’s internal state. The “behavior” graph shows the agent’s behavior in terms of the output units of the architecture of Figure 5: move is 0; left is 1; right is 2; and noop is 3. Finally, the top graph(s) represent the output signal(s) of the agent, one graph for each communication channel.

At first glance, it appears that the agents have indeed learned to communicate the presence of food. The strength of an agent’s output signal is negatively correlated with the presence of food: when the agent sees food, its output signal is low; otherwise, it is high. Behaviorally, Figure 7b appears to show recruitment, i.e., the black agent attracting the white agent to the food.

In reality, however, these agents are *not* communicating the presence of food. The difference between Figures 7b and c shows that recruitment is not really occurring. In the former, the white agent appears to be moving towards the black agent, but in the latter, the white agent has made no progress towards the food. Instead, it has moved eastward in a path parallel to that of the

---

[5] The circles denote not signal range, but the radius at which signal strength (i.e., the summand in Equation 3) is one.

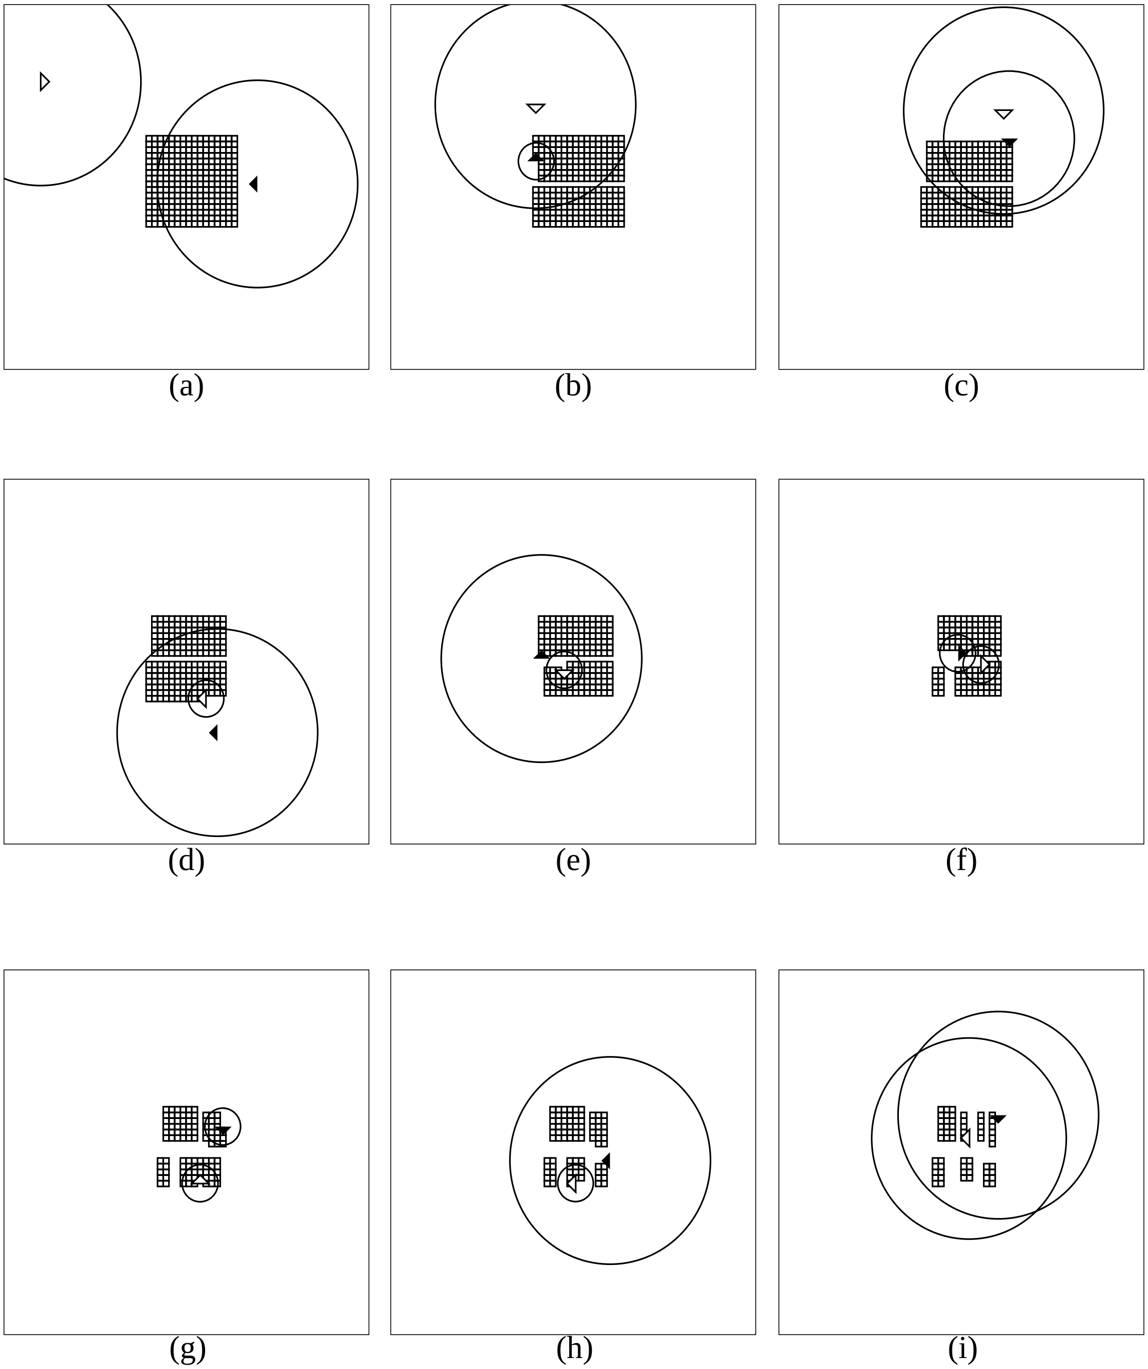

Figure 7: The “Curmudgeon” strategy (2 agents, 1 communication channel, no noise). Primitives are “Move,” “Left,” “Right,” and “Noop.” (a) t=1; (b) t=25; (c) t=50; (d) t=75; (e) t=100; (f) t=125; (g) t=150; (h) t=175; (i) t=200.

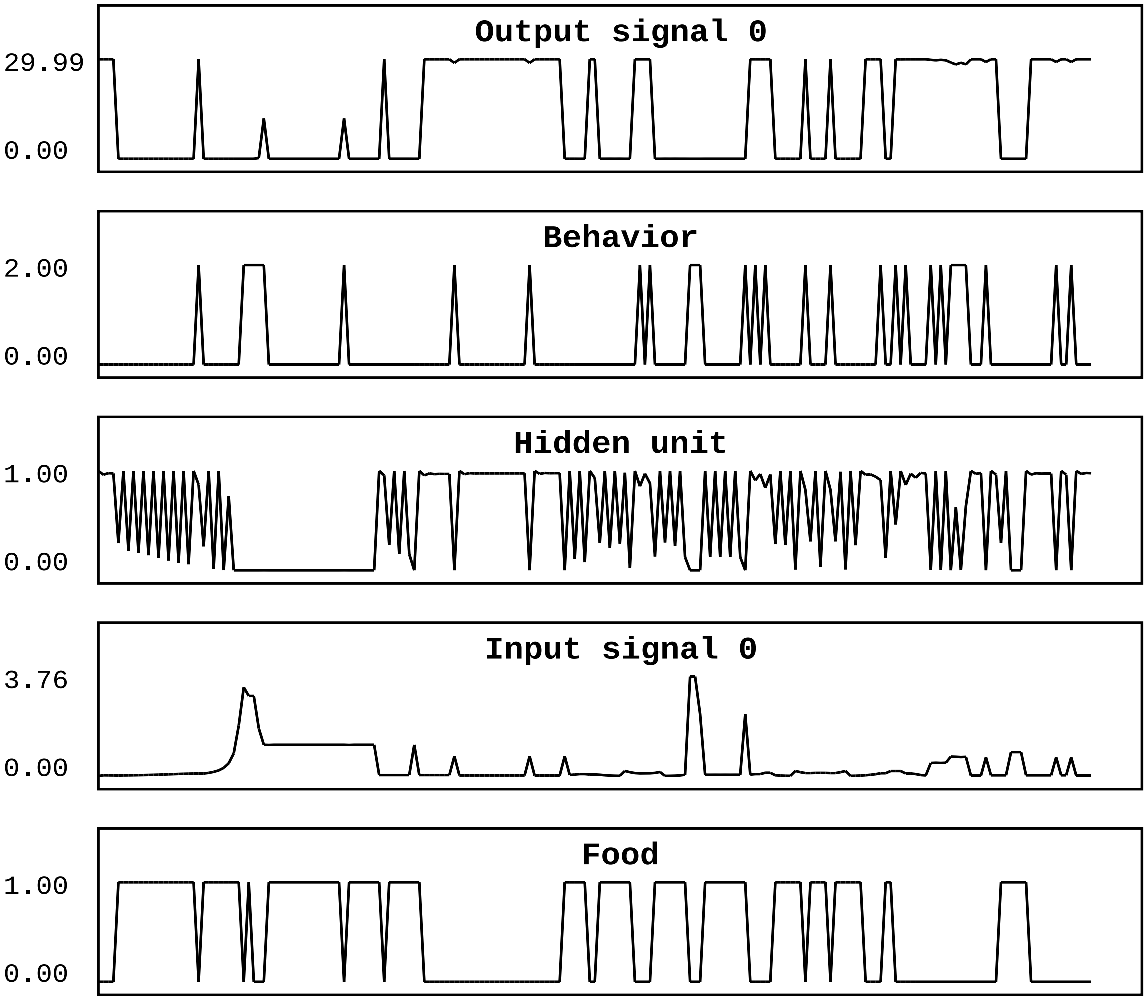

Figure 8: Profile of curmudgeon 0 (200 time steps). The lowest graph is the food input: when food is detected, the value spikes to one; otherwise it is zero. The input signal reflects communication from the other agent. The value of the hidden unit reflects internal state. Behaviors are “Move,” “Turn left,” “Turn right,” and “Noop,” with values 0-3, respectively.

black agent. Figure 7d is similarly surprising: here, the black agent has moved away from the food patch, and is traveling westward in a path parallel to that of the white agent.

In short, these agents have adopted what might be called “the curmudgeon strategy.” Each agent prefers to be alone – rather than attraction, what we see here is repulsion! When run on an empty grid, the agents move about avoiding each other. Food coerces agents to move closer together, but they still often maintain a separation, as shown above. In retrospect, one can see why this

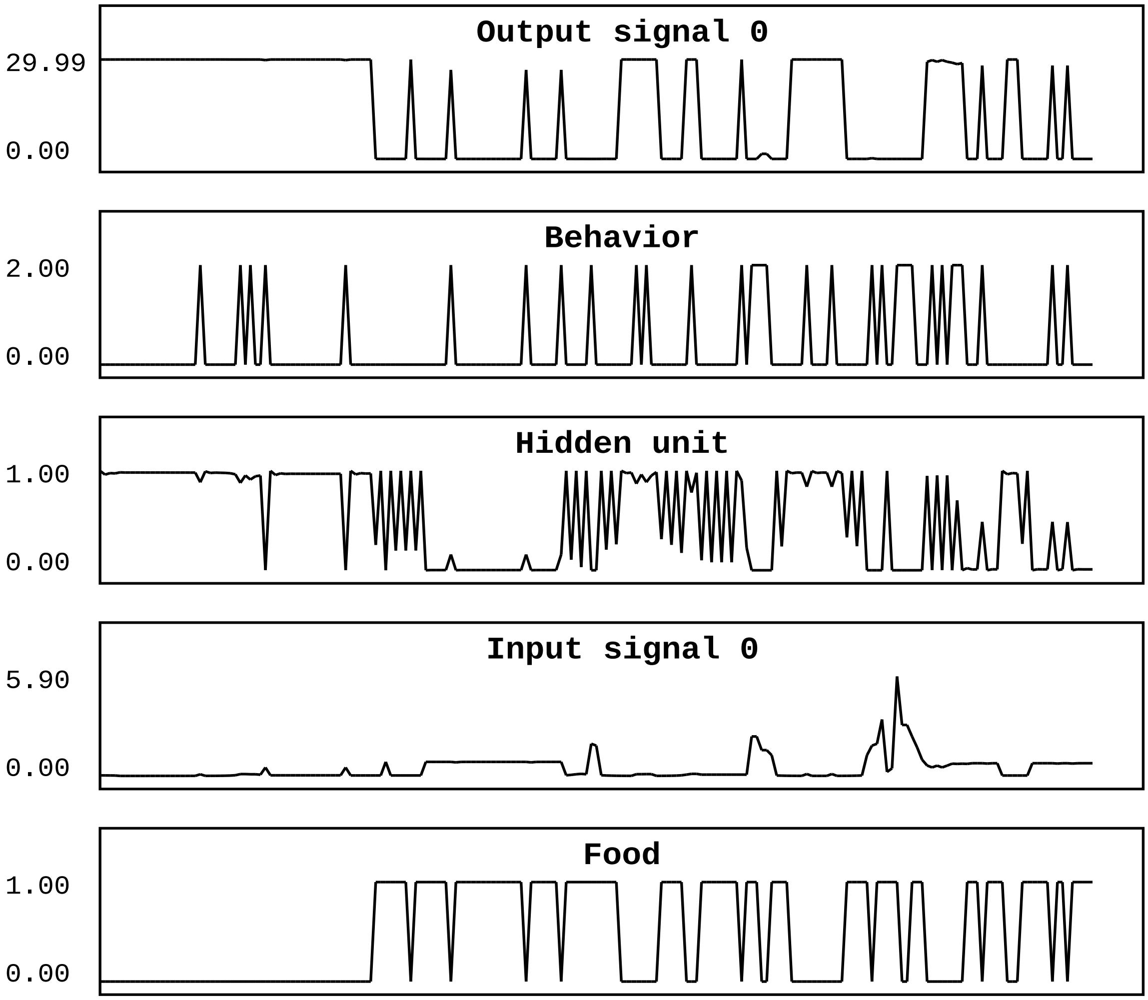

Figure 9: Profile of curmudgeon 1 (200 time steps). This agents behavior is sensitive to the input signal from curmudgeon 0, but not in a way that qualifies as communication. See text for details.

strategy was selected. By pushing each other away, the agents are able to explore a large area of their environment.

3.2 Experiment 2: The Diagonal Strategy

After several similar runs, we developed a hypothesis about why the agents refused to communicate the presence of food: it is far easier to use signals to implement a pseudo-random search strategy than for each agent to develop a good search strategy alone and communicate the results.

To test this hypothesis, we decided to endow each agent with a good initial search strategy, and let them evolve communication from there. We chose as our starting point a strategy discovered experimentally by GNARL for the normal Tracker task: the five-state FSA of Figure 3. We built this FSA into each initial network, and then evolved solutions from there.

To clarify just how much this initial starting point helps, we first present the results of running just the FSA-network on the modified Tracker task (i.e., with no evolution). Figure 10 shows the result for the case of 3 agents. Because the black agent does not communicate the presence of food to the others, the white agent moves continually eastward, following the “search in a circle then move forward” strategy of the FSA. Similarly, the gray agent moves continually southward. These agents receive a score of 128.

Figure 11 shows the performance after evolutionary search for the three agent case. Fitness is now 195. Once again, to investigate whether these agents had learned to communicate the presence of food, we plotted the status of each agent over the course of the run (Figures 12-14); once again we learned the agents are not communicating. In the presence of food, each agent executes the “Move forward” strategy of the FSA. But in the absence of food, each agent ignores its input signal, and executes the strategy “Left, left, left, move, left, move,” creating a diagonal walk through the search space. This is evident in the behavior profiles of agents 1 and 2, and even in the snapshots of the three agents. Note in Figure 11b that the black agent enters the food patch north of its initial position; it is moving northwest. Comparing Figures 11a-e shows that the white agent is moving southeast. Finally, the position of the gray agent throughout all the snapshots shows that it is simply moving southwest.

We repeated this experiment several times varying the number of agents (2, 3, and 5), and always obtained similar results. The “diagonal strategy” was the most pervasive, though a few other perturbations of the FSA’s behavior were found.

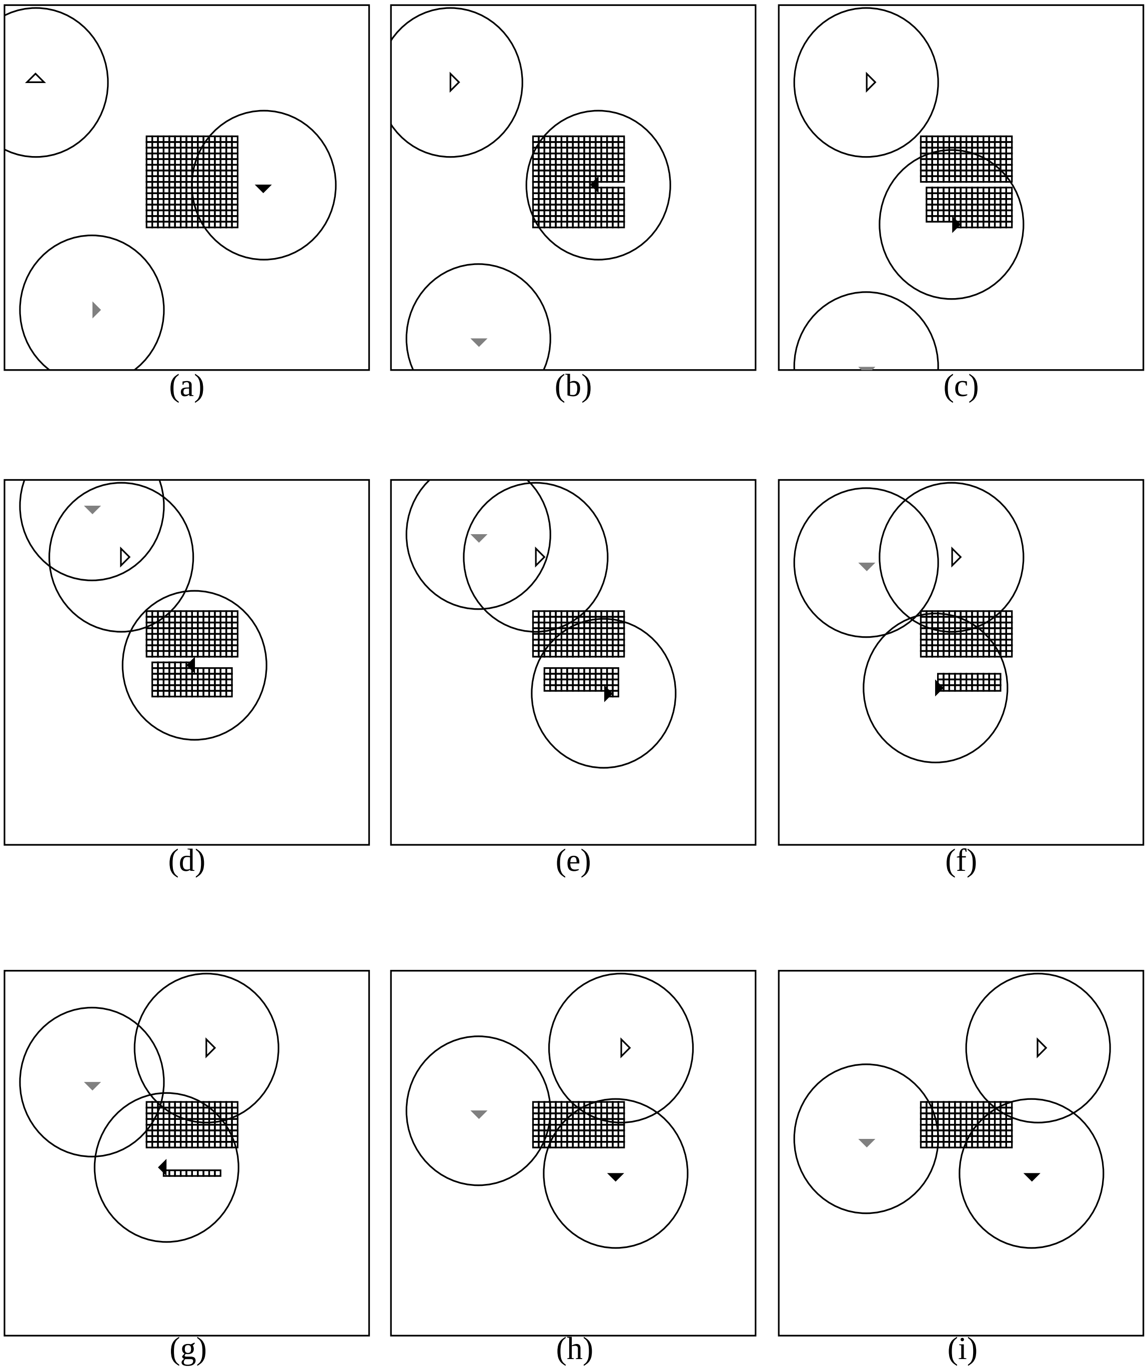

Figure 10: Three "FSA-initialized" agents, before any evolution (one communication channel, no noise). Primitives are "Move," "Left," "Right," and "Noop." (a) t=1; (b) t=25; (c) t=50; (d) t=75; (e) t=100; (f) t=125; (g) t=150; (h) t=175; (i) t=200.

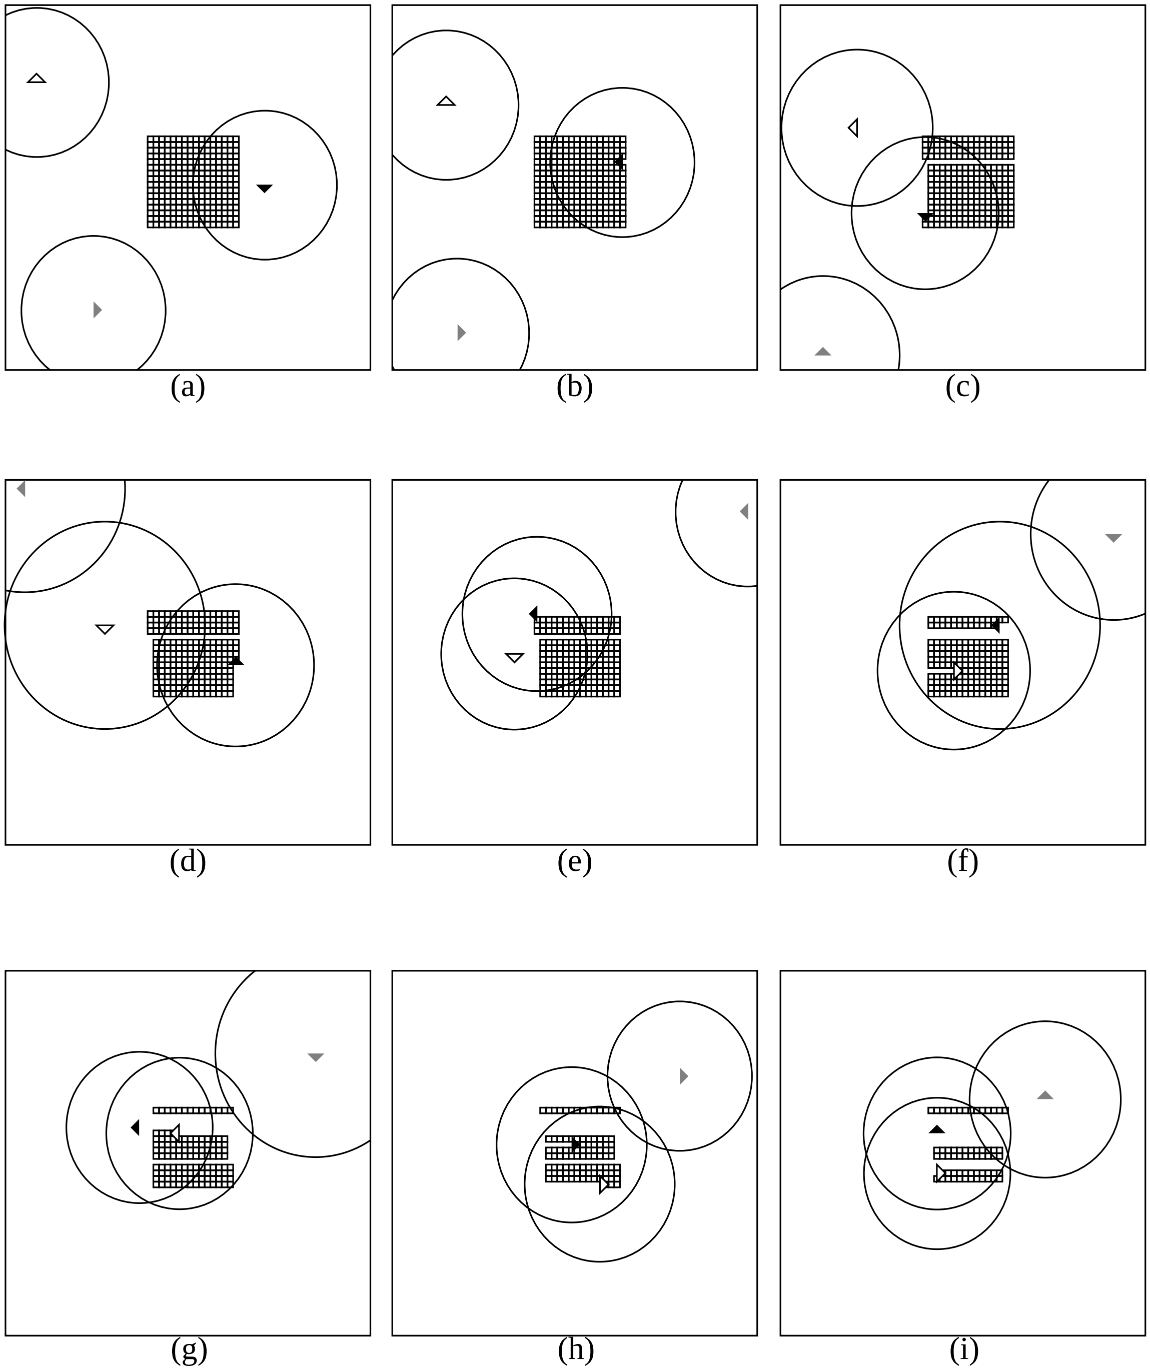

Figure 11: Three "FSA-initialized" agents, after evolution (one communication channel, no noise). Primitives are "Move," "Left," "Right," and "Noop." (a) t=1; (b) t=25; (c) t=50; (d) t=75; (e) t=100; (f) t=125; (g) t=150; (h) t=175; (i) t=200.

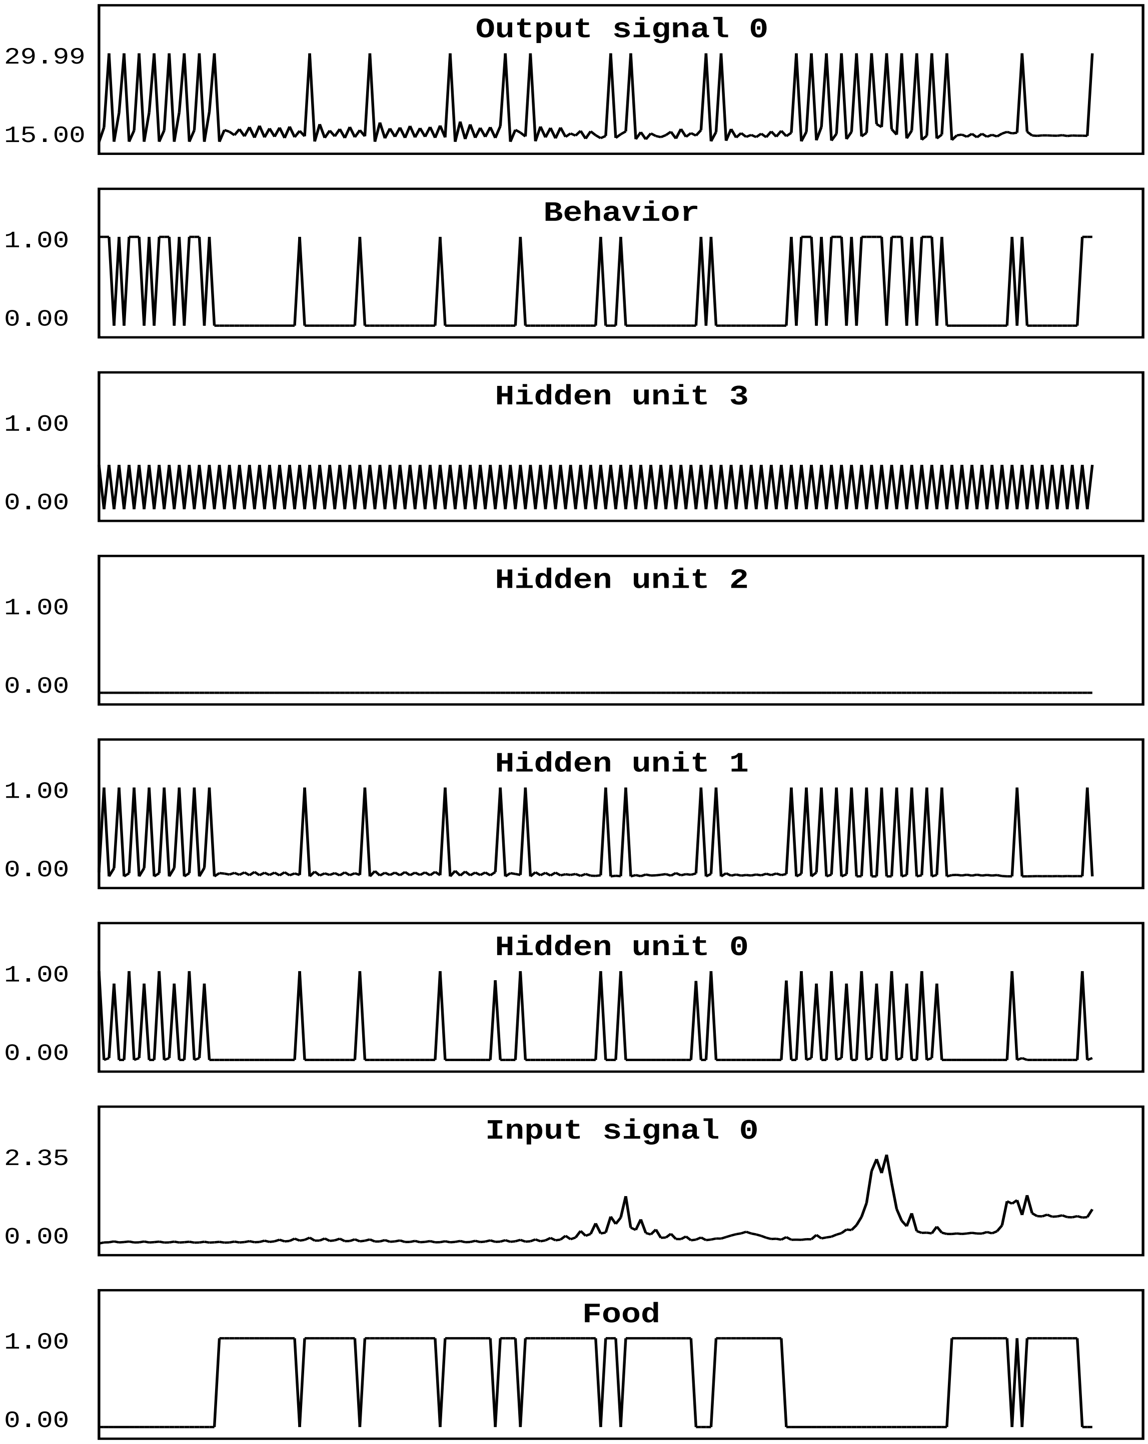

Figure 12: Profile of "FSA-initialized" the black agent (200 time steps). Behaviors are "Move," "Turn left," "Turn right," and "Noop," with values 0-3, respectively. This agent has learned to correlate oscillation of its output signal with the absence of food.

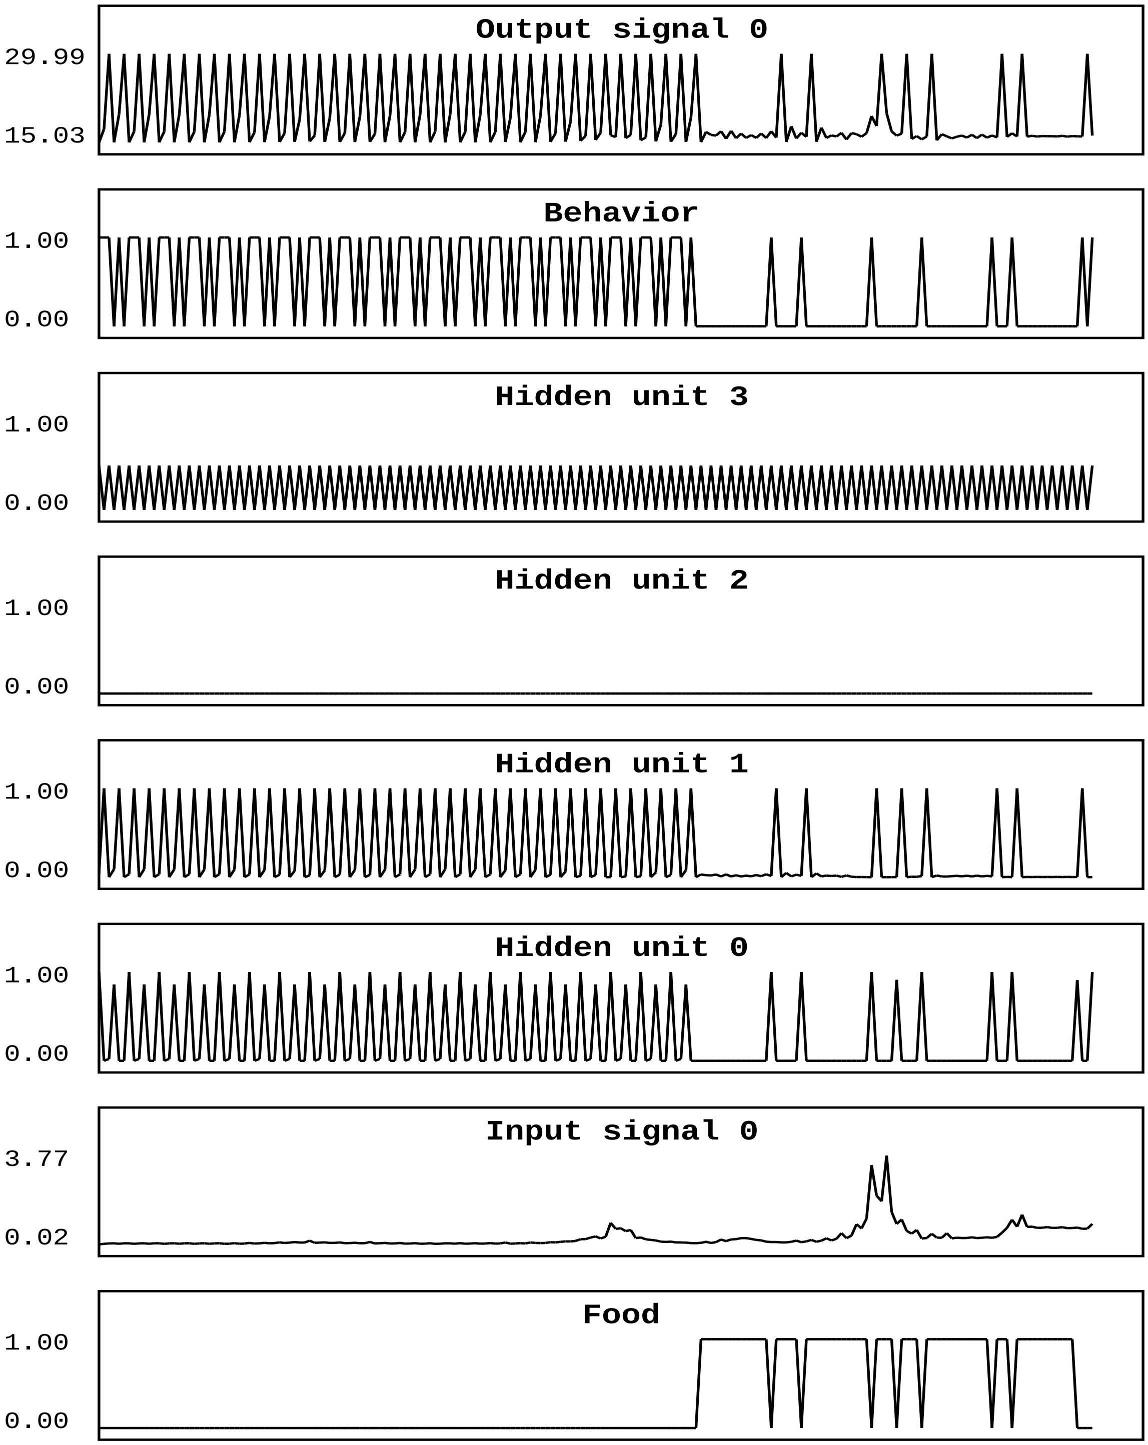

Figure 13: Profile of “FSA-initialized” the white agent (200 time steps). The behavior of this agent depends only upon the presence of food, and not upon the varying input signal.

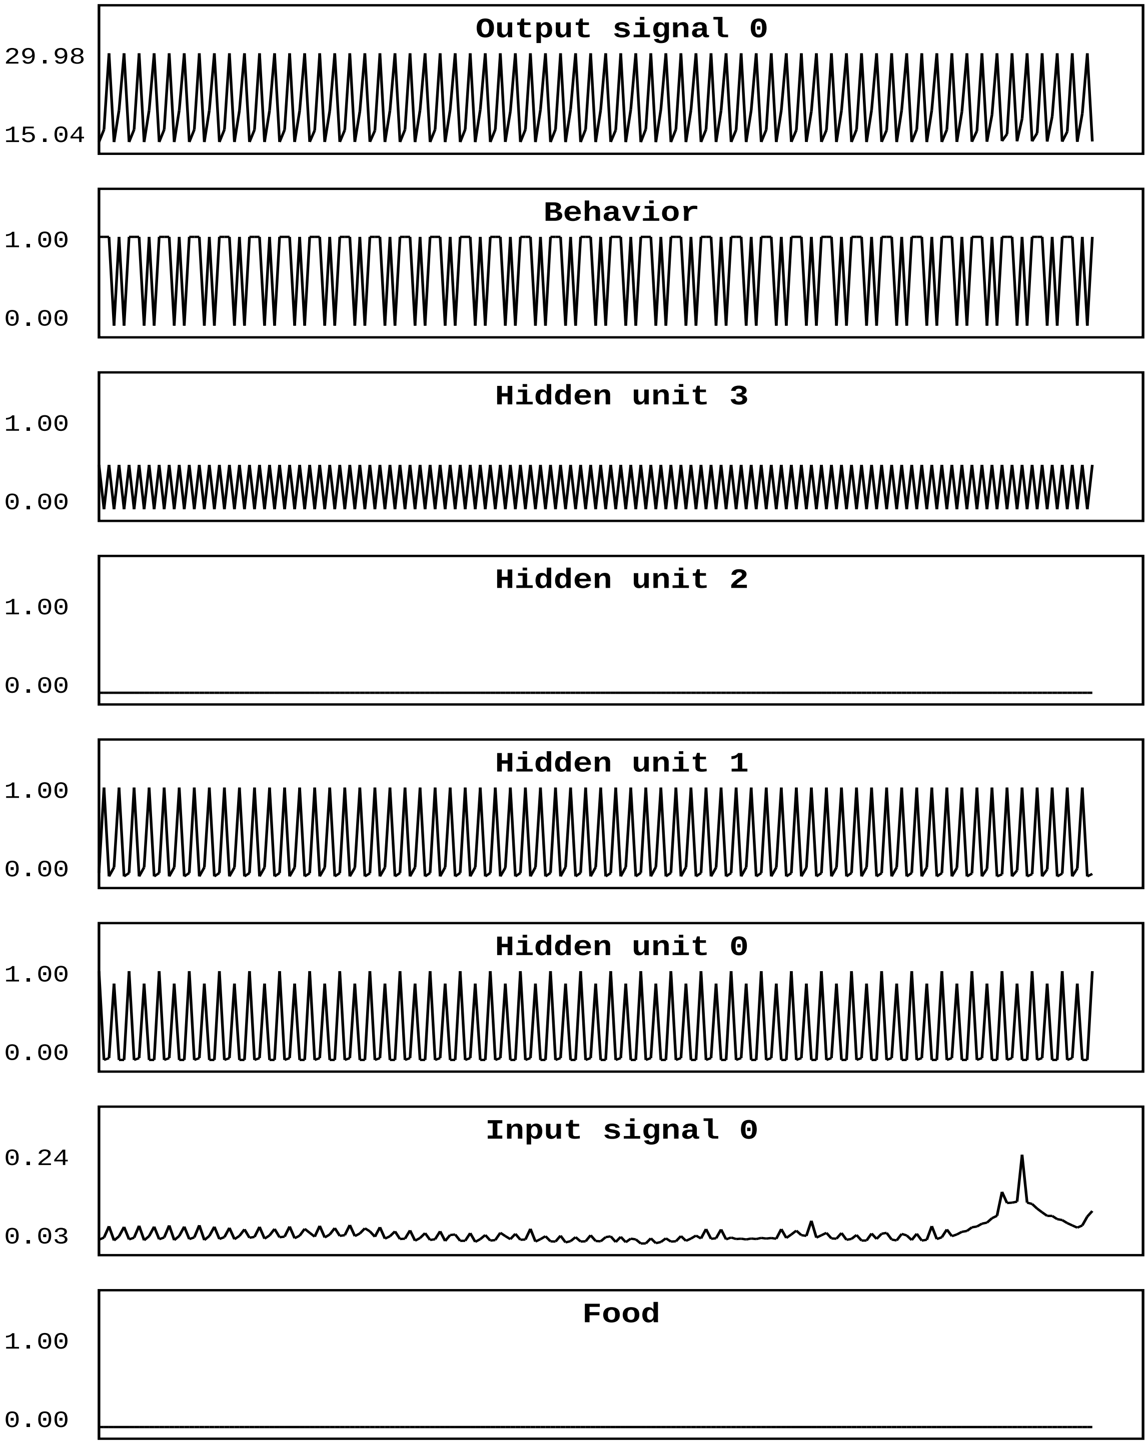

Figure 14: Profile of “FSA-initialized” the gray agent (200 time steps). Despite the large variation in its input signal, this agent has constant behavior. This is due to the fact that it never encounters food.

4 The Tracker Task, Rethought

This section examines why communication was not evolved in the experiments above, and proposes a slightly different direction of study.

4.1 Analysis of First Results

Why should the evolutionary algorithm favor strategies like “curmudgeon” and “diagonal” as opposed to solutions which involve the communication of food? The problem is twofold: the chosen primitives do not easily support communication, but they do easily support pseudo-random search.

First, following a communication signal is extremely difficult with the primitives “Move, Left, Right, Noop.” Because the agent only has access to the value of the signal at one particular grid position, to follow a gradient it would have to first turn four times, sampling and storing the signal value in four positions, then compare the stored results, turn to the appropriate direction, and finally move.

Second, the same “Move, Left, Right, Noop” primitives do easily support pseudo-random search. We have already seen that for a single agent with internal state can evolve a network which supports complex exploration of the search space (Angeline, Saunders, and Pollack, 1994; Saunders, Angeline, and Pollack, 1994). With multiple interacting agents, such solutions are even easier to evolve.

In short, the reason for the lack of communication is that the presence of a simpler solution with the “Move, Left, Right, Noop” primitives: wander around the search space in a pseudo-random manner, then eat food when it is discovered. This type of solution was commonplace starting from either random networks (“curmudgeon” strategy, Figure 7) or from FSA-initialized networks (“diagonal” strategy, Figure 11).

4.2 A Shift in Architectures

The results above were more surprising than disappointing. Adaptive agent algorithms are inherently opportunistic, and they will always find the simplest solution to a task. Our original claim was that task-specific communication schemes can indeed be evolved for cooperative multi-agent systems. The real problem is not that this claim is false; rather, the problem is that the "Move, Left, Right, Noop" primitives make the modified Tracker task too simple: solutions with pseudo-random walks of non-communicating agents abound.

But because the real concern is *communication*, pseudo-random solutions can be inhibited by simply shifting the task slightly; specifically, by changing the base architecture. To enable agents to exploit each other's signals, we give them the ability to follow communication gradients. To prevent the agents from alternating between Move, Left, and Right to effect a pseudo-random search, we force them to follow a particular search strategy; namely, that of the Jefferson FSA of Figure 3 (also evolved experimentally – the simpler of the two results reported in Saunders, Angeline, and Pollack, 1994).

The new agent architecture, shown in Figure 15, allows a clear separation between complexity arising from communication, and complexity arising from clever activation of the output nodes. The n additional output units represent an agent's actions relative to the n communication channels. When the $i^{\underline{th}}$ "Follow gradient" node receives highest activation, the agent follows the gradient of communication channel i.

These modifications, though not essential to the results, greatly facilitate their analysis. The food collection strategy of the FSA is indeed quite simple; if activated repeatedly on a grid containing no food, the agent traverses its environment, turning in circles, but never veering from a straight line. Thus any agents moving non-linearly in the absence of food *must* be following a communication signal. Furthermore, because of the implicit winner-take-all network, it is easy to ob-

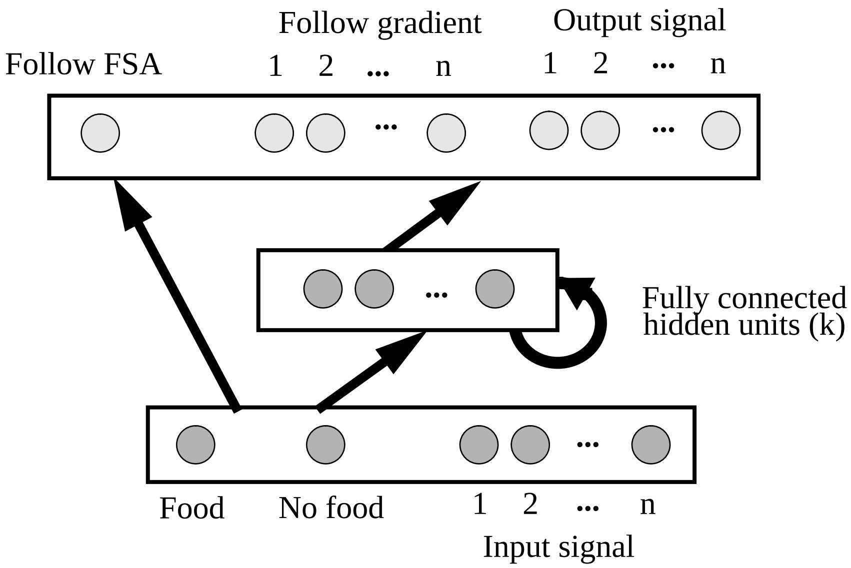

Figure 15: An agent architecture which prevents pseudo-random solutions of Section 3. The “Move, Left, Right, and Noop” actions are condensed into “Follow FSA,” one particular strategy found by GNARL. The “Follow gradient” nodes give the agent an Addam-like ability to respond to environmental signals.

serve *which* communication signal the agent is pursuing by simply comparing activations across the output nodes.

5 Results, Part II

This section describes the results with the new agent architecture. In spite of the fact that these agents possess just two types of possible actions (“follow FSA” or “follow gradient”), the task is more difficult than with the “Move, Left, Right, Noop” primitives. The reason is that pseudo-random solutions are now disallowed, and that to search for food effectively, the agents must learn to communicate. The difficulty here is in finding a way to do this over continuous channels.

5.1 Experiment 3: Constant Communication Schemes

We once again begin with a very simple case: 2 agents, each with one hidden unit, capable of passing one real number between each other, with no noise ($u_0 = 0$, see Equation 3). Figure 16a shows

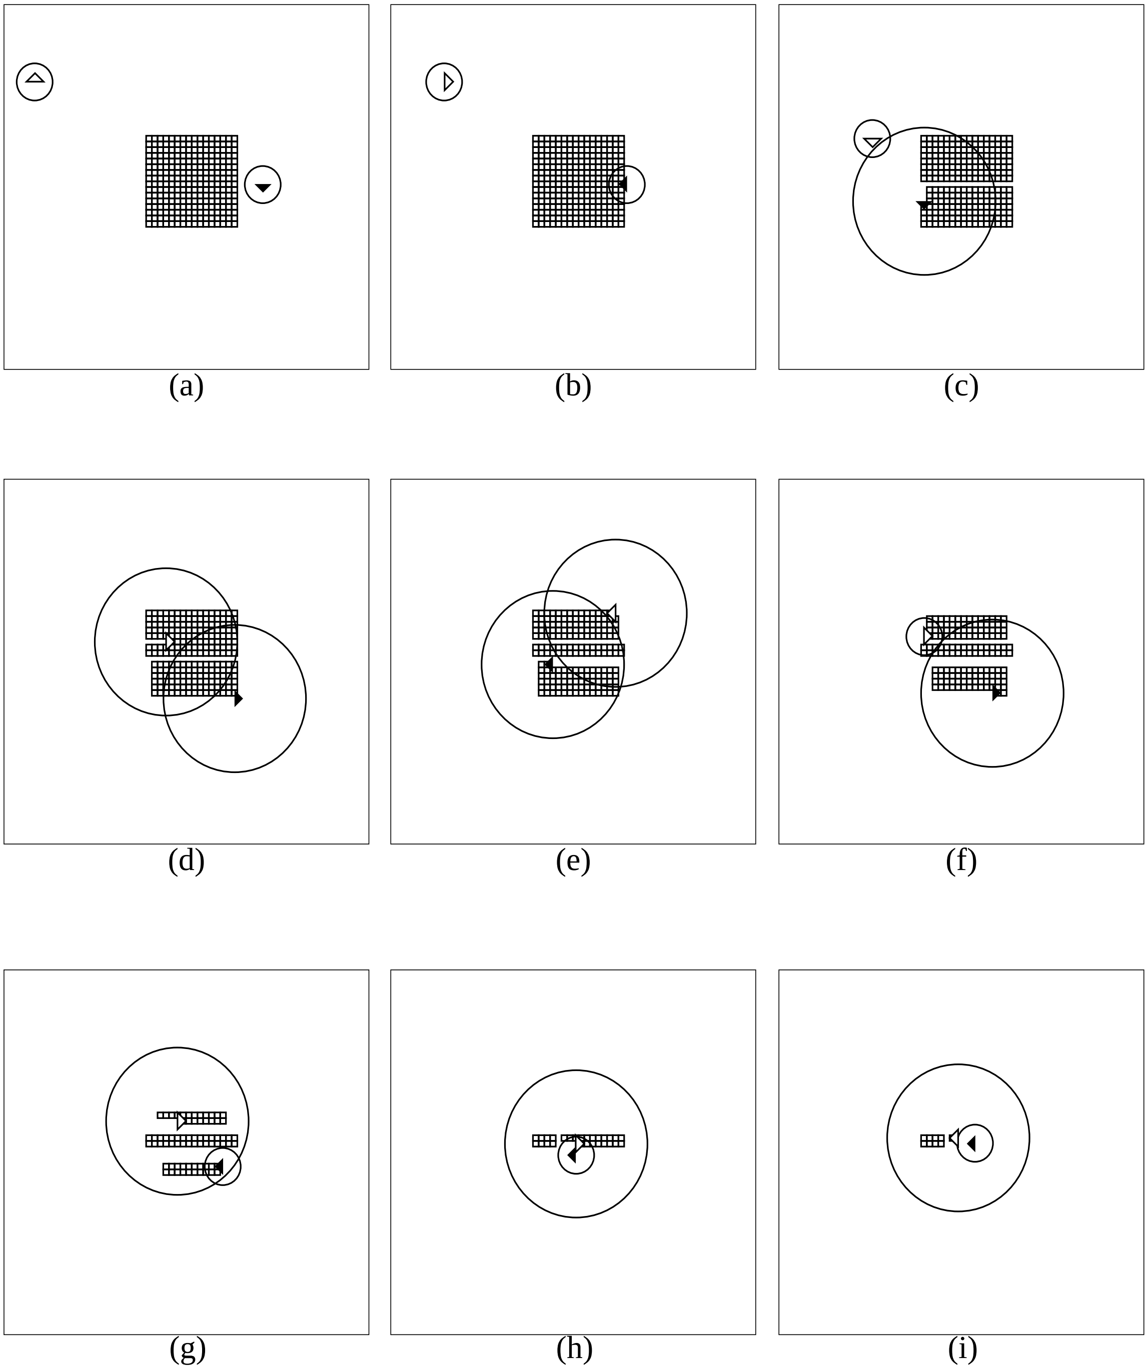

Figure 16: Scenes of evolved communication, 2 agents, 1 communication channel, no noise. Primitives are "Follow FSA" or "Follow signal." (a) Initial positions: neither agent can sense food; (b) The black agent just reaches food, time is t=20; (c) Recruitment – first agent attracting the second, t=40; (d-g) Scenes at t=60, 80, 100, 140, respectively; (h-i) Recruitment again, though now in reverse, t=180 and 200, respectively.

the initial environment. Without communication, each agent would follow the FSA, and the white agent would move in a straight line, finding no food.

With communication, however, the story is quite different. In 700 generations, GNARL discovered a pair of agents (from a population of 50) which had learned to communicate the presence of food. Figure 16b shows the case just as the black agent reaches the food. Figure 16c shows recruitment: the black agent’s strong signal, due to the food, attracts the white agent. Figures 16d and e show both agents are emitting high signals while eating. Between Figures 16f and i, the black agent finishes his “half” of the food, and is recruited to help the white agent finish what is left.

Figures 17 and 18 show the agents’ signals which allow them to produce this behavior. These and similar figures describe all there is to know about a particular agent. The “behavior graph” now represents the agent’s action in terms of the architecture of Figure 15: zero indicates the agent is following the Jefferson FSA (Jefferson, et al., 1991); a value of *i* indicates that the agent is pursuing the gradient of signal *i-1*.

The fact that only one communication channel was involved in the first experiment simplifies the interpretation of Figures 17 and 18. The black agent has learned to correlate the magnitude of its output signal with the presence of food. Of course, this correlation would be meaningless without a suitable response from the white agent: when the white agent “hears” a large input signal, it follows the signal to find food.

We chose this case as a demonstration for several reasons. First, snapshots easily capture the evolved communication scheme: larger circles imply a higher signal. Second, the communication scheme is fairly intuitive: each agent “yells” when it finds food by increasing the strength of its output signal; upon “hearing” such a signal, the second agent follows it to the source of food. We have also observed other implementations of the same behavior, e.g., “Yell constantly when

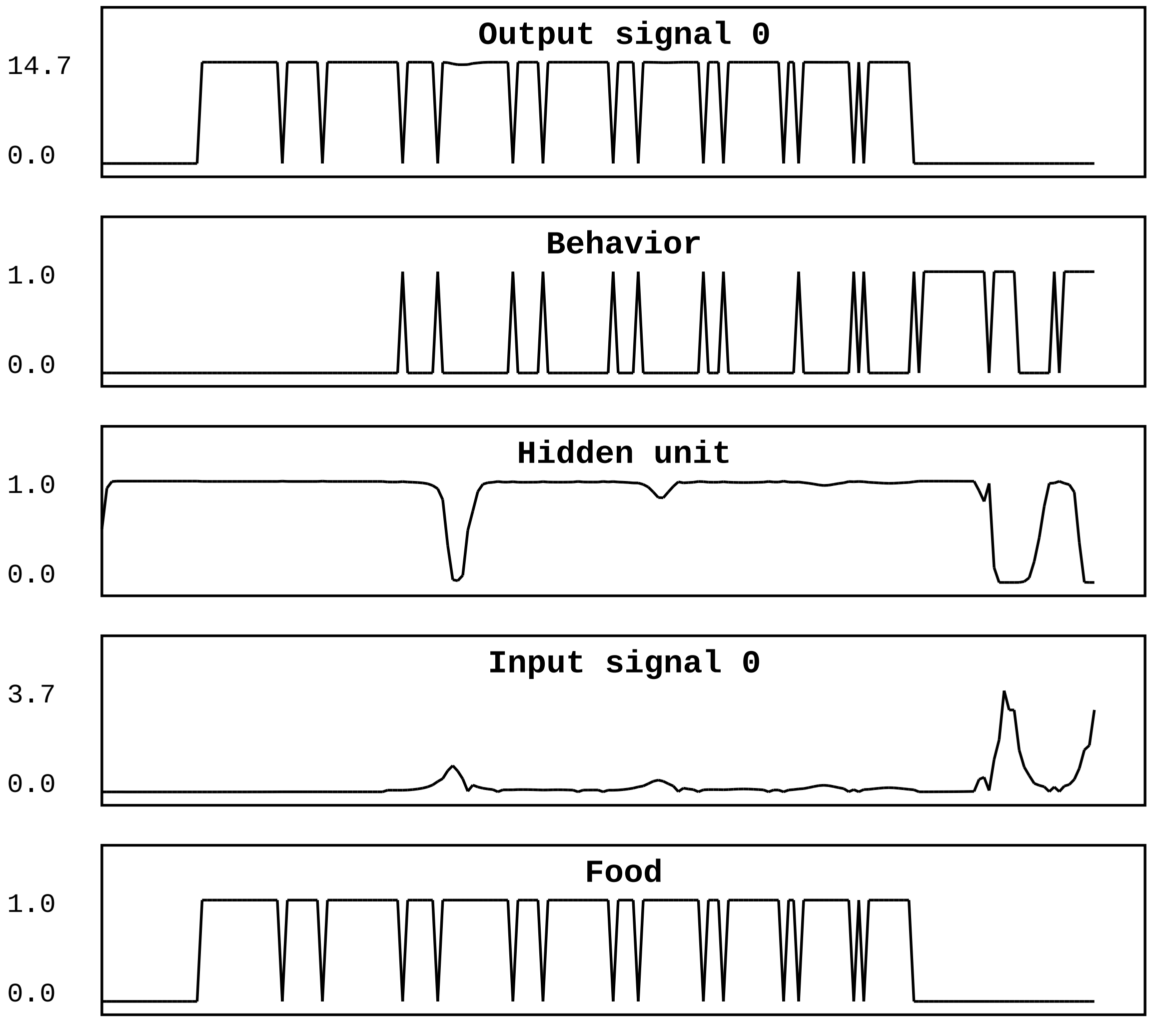

Figure 17: Profile of the black agent (200 time steps). Behavior is either "Follow FSA" (low) or "Follow Signal" (high). This agent has learned to correlate the magnitude of its output signal with the presence of food.

you're searching for food, but then grow quiet when eating." In this case, agents learn to respond to silence. But such constant signalling behavior by no means exhausts the possible means of communication.

5.2 Experiment 4: Oscillatory Communication Schemes

In contrast to the cases described above, the next example shows how oscillatory communication schemes may evolve. For this experiment, we used the same food distribution, increased the num-

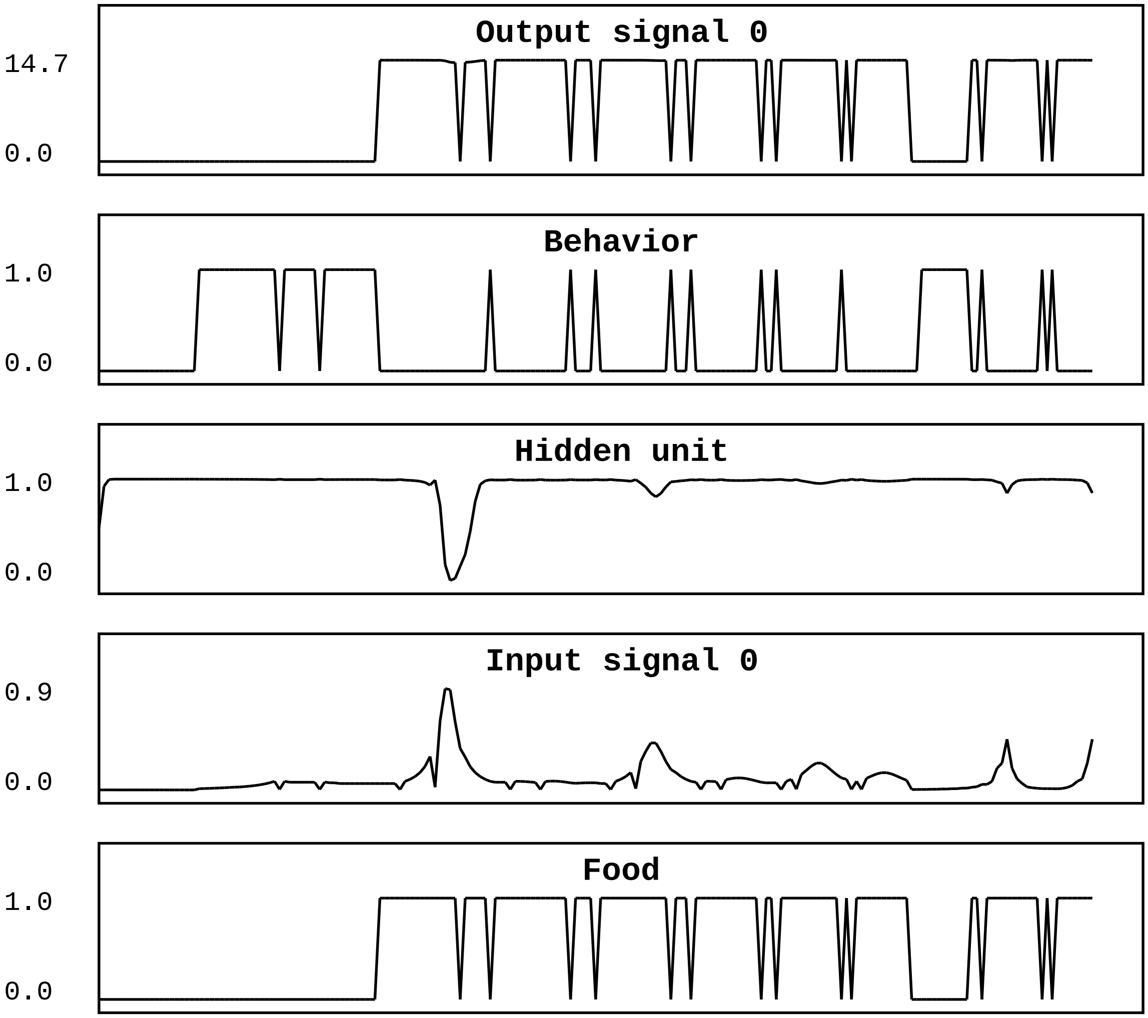

Figure 18: Profile of the white agent (200 time steps). The point where recruitment occurs (Figure 16c) corresponds to the first spike in the behavior profile of this agent. When this agent reaches food (between Figures 16c and d), the behavior reverts to following the FSA.

ber of agents to three, and retained a a single hidden unit for each agent. To investigate how the agents would respond to noise, we gave them two communication channels, the first clear (u_0=0), the second noisy (u_1=10).

Figure 19a shows the initial environment. The circles reflect the strength of signal 0. We omit signal 1, transmitted along a noisy channel, because it is not used by the agents (more on this

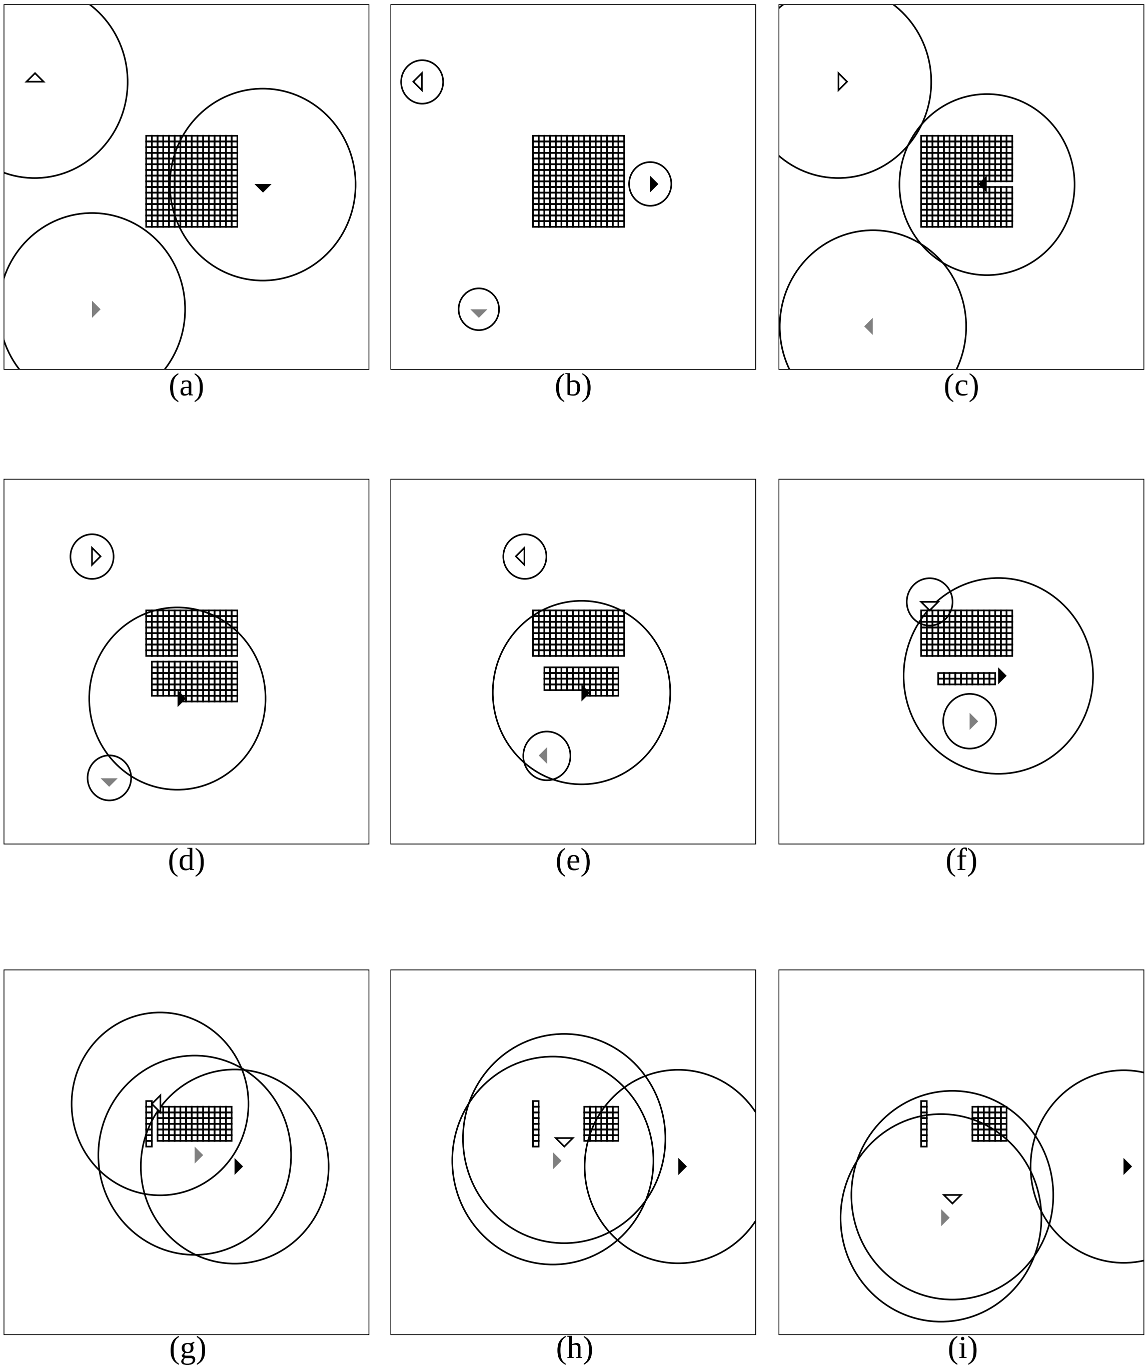

(a) (b) (c)

(d) (e) (f)

(g) (h) (i)

Figure 19: Scenes of evolved communication, 3 agents, 1 communication channel, no noise. Primitives are “Follow FSA” or “Follow signal.” (a) Initial condition, circles denote signal 0; (b) After one time step, signal 0 has shrunk to its minimum value. It oscillates between the two extremes when no food is present; (c-i) t=25, 50, ..., 300.

below). After just *one* time step, the signals along channel 0 have shrunk to their size in Figure 19b. In the absence of food, signals in this channel oscillate between these two extreme values.

Because overall behavior is difficult to discern from just snapshots, Figure 20 abstracts just the agents' paths from Figure 19. Figure 20a shows recruitment by agent 0; Figure 20b shows that recruitment is not permanent: when the food has been consumed, agent 0 strikes out on its own.

Figures 21-23 show the signals by which this behavior is accomplished. Figure 21 gives the profile of the black agent over the run. Note how its output signal 0 oscillates in the absence of food. Figure 22 shows the profile of the white agent throughout the run. The lack of oscillation in the black agent's output is enough to turn the white agent towards the food. (The 5 spikes in the behavioral profile indicate "Follow signal 0" behavior.)

The gray agent, however, is slightly different (Figure 23). Note the oscillation in its behavior, as it alternates between following the gradient of signal 0 and following the FSA. At first glance, this seems incorrect, because the inputs to agents 1 and 2 look identical, but their output behaviors are very different. The problem, however, is simply one of scale. Figure 24 zooms in on the first 50 time steps of the signal 0 input to agents 1 and 2. It is the phase difference between these two signals which is responsible for the difference in the agents' behaviors.

From the agent profiles (Figures 21-23), it appears that the evolved agents are relying solely on channel 0, the clear channel. To test this, we blocked the agents' signals by shunting the channel with various constant values. In all cases, removal of channel 0 drastically reduced fitness, yet the removal of channel 1 failed to hamper the search behavior of the agents, confirming our expectations that the agents had learned to ignore the noisy channel.

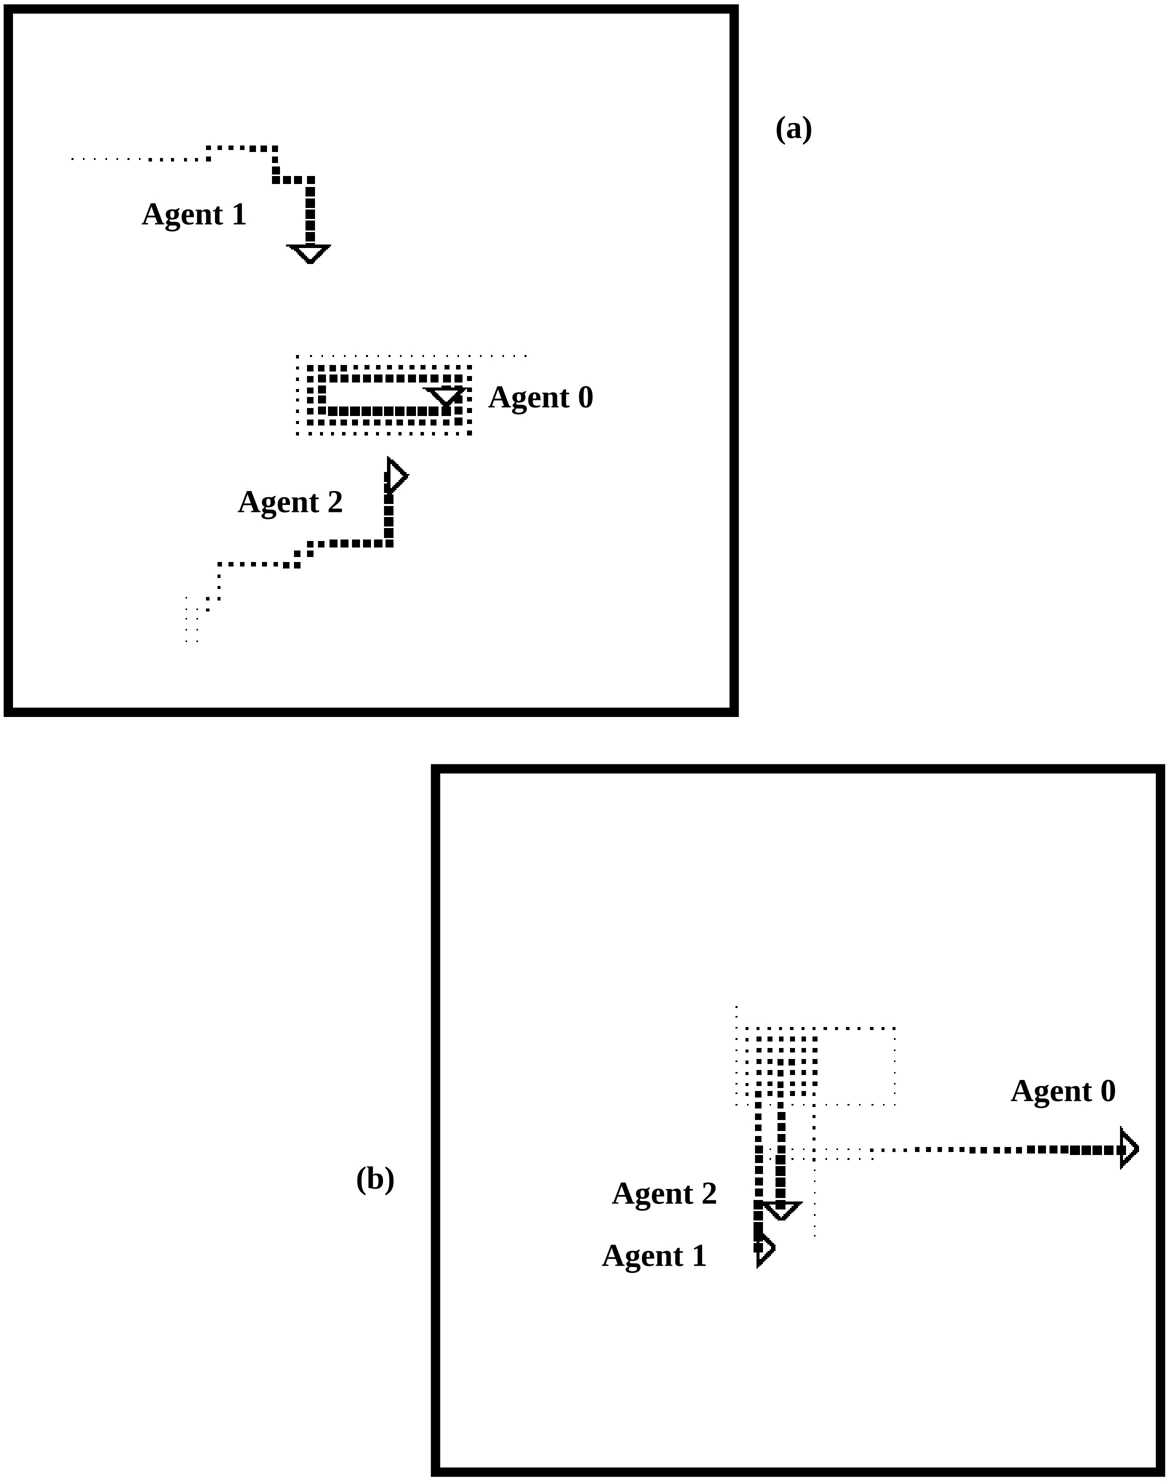

Figure 20: Agent paths for the 3-agent case. Dots indicate path of agents (food and signals have been removed for clarity). (a) Agent 0 recruiting the others. After food has been consumed, agents 1 and 2 stay together, but agent 0 strikes out on a different path.

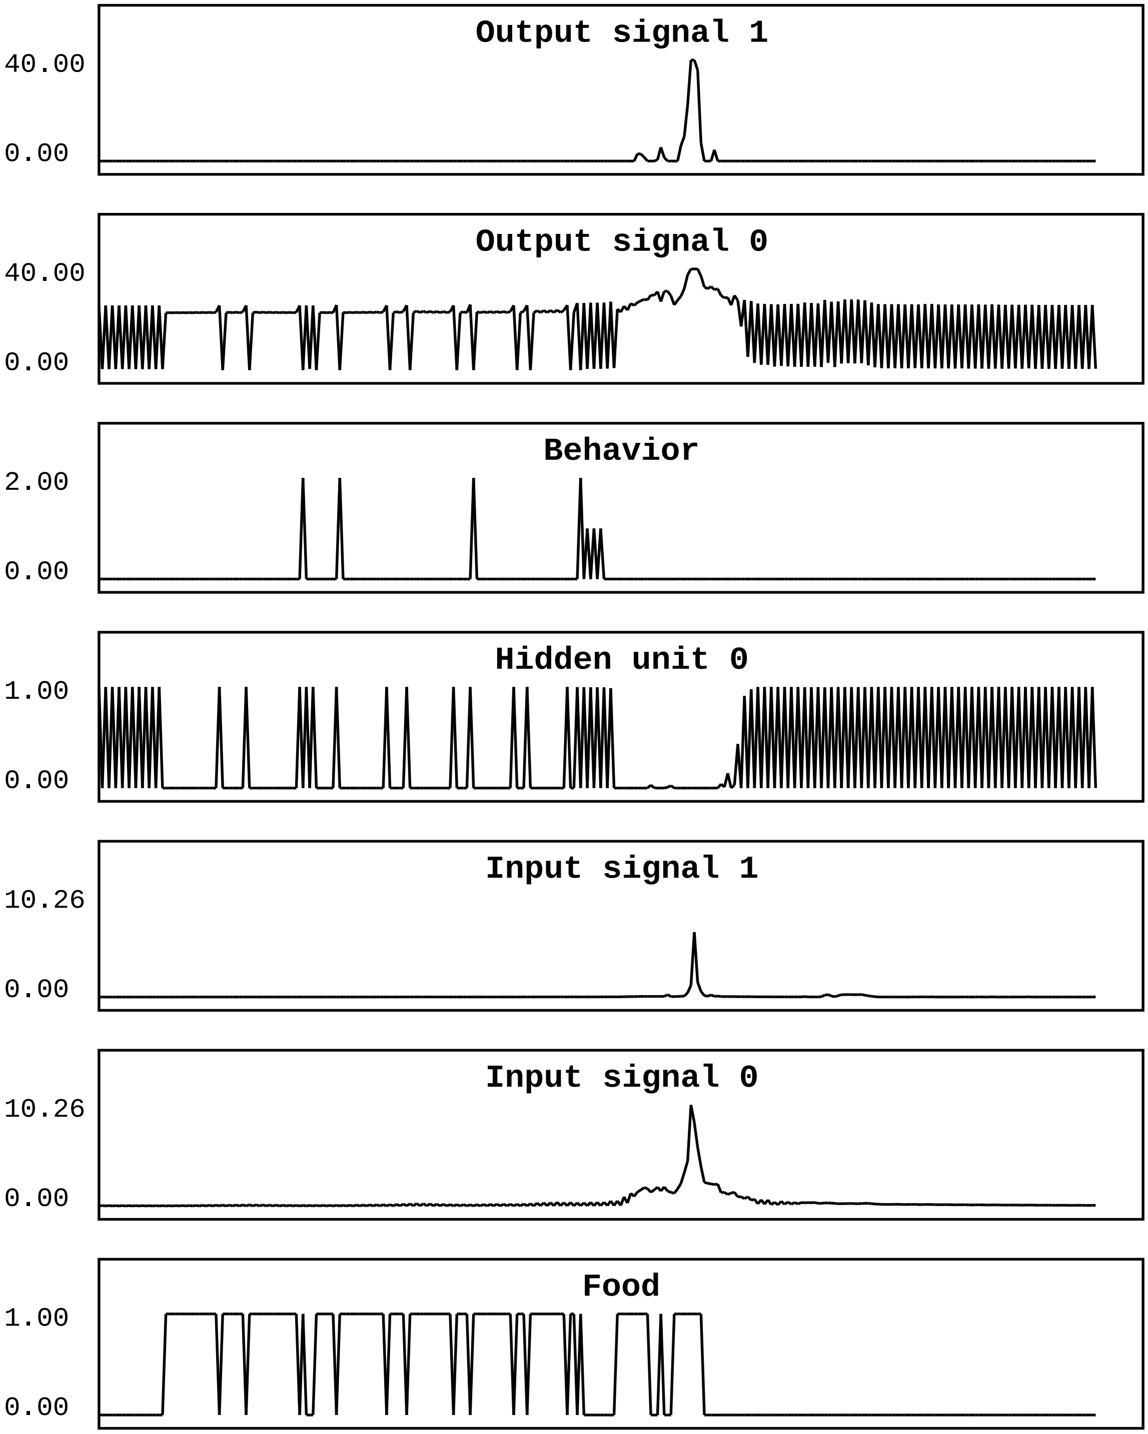

Figure 21: Profile of the black agent (300 time steps). This agent has learned to correlate oscillation of its output signal 0 with the presence of food.

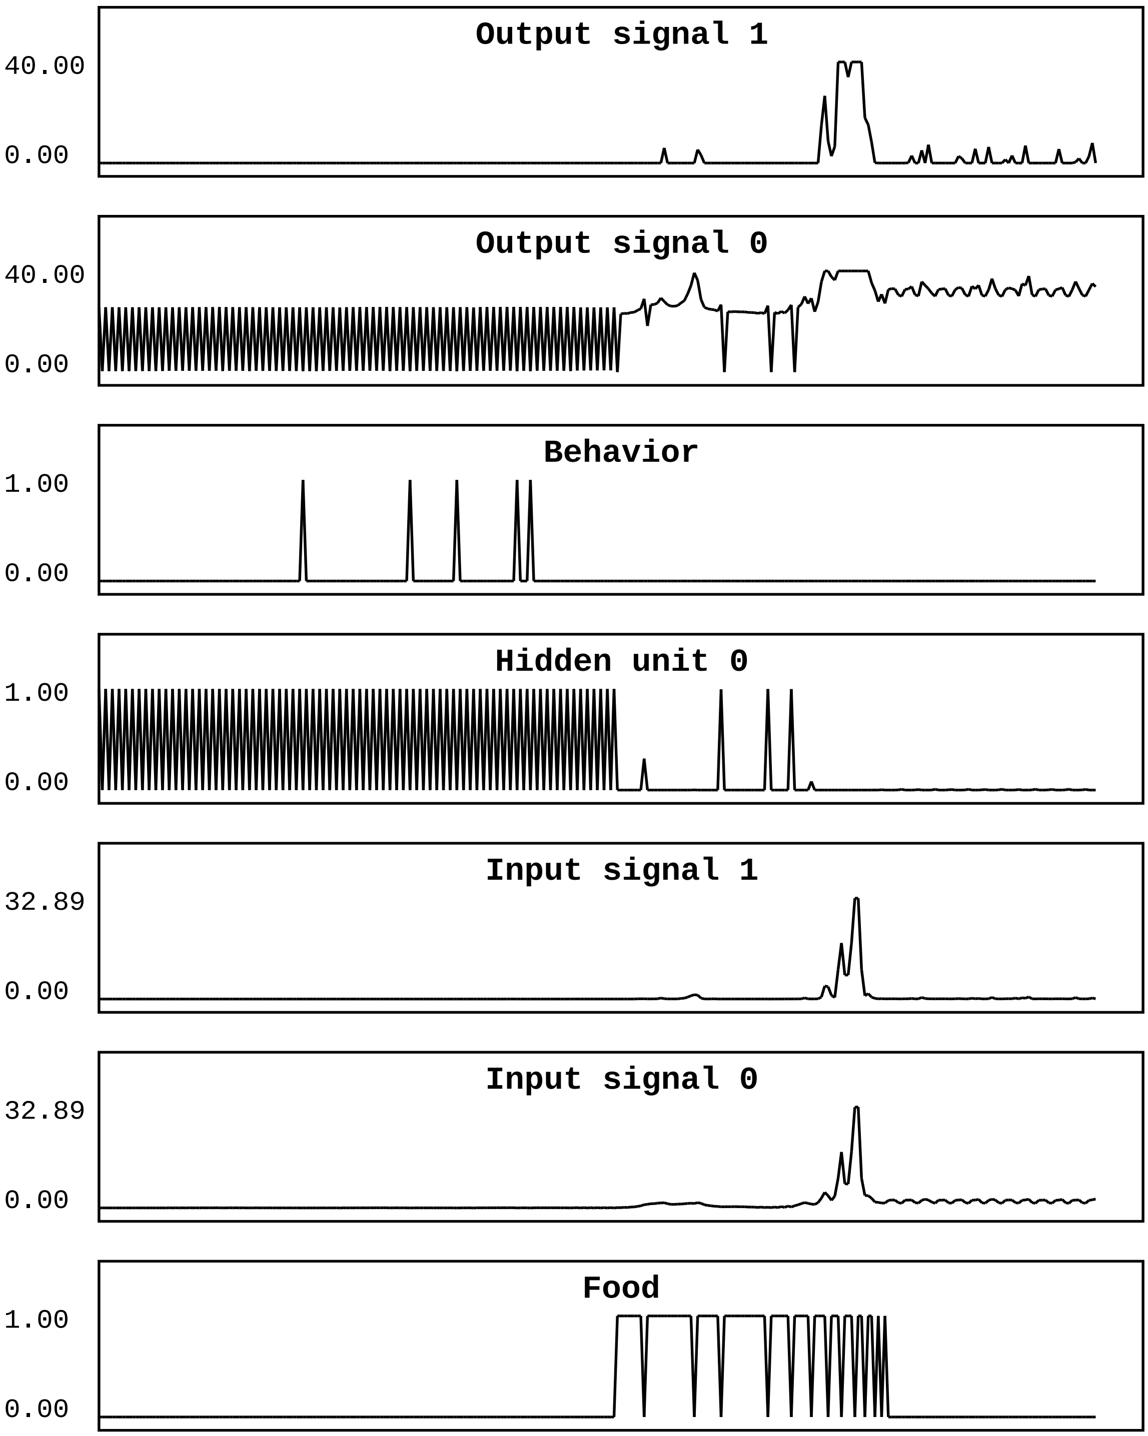

Figure 22: Profile of the white agent (300 time steps). The five spikes in behavior indicate points where the agent follows signal 0, as can be seen in Figure 19. Because the agent perceives no food during this time, the resulting behavior occurs due to the agent's input signals.

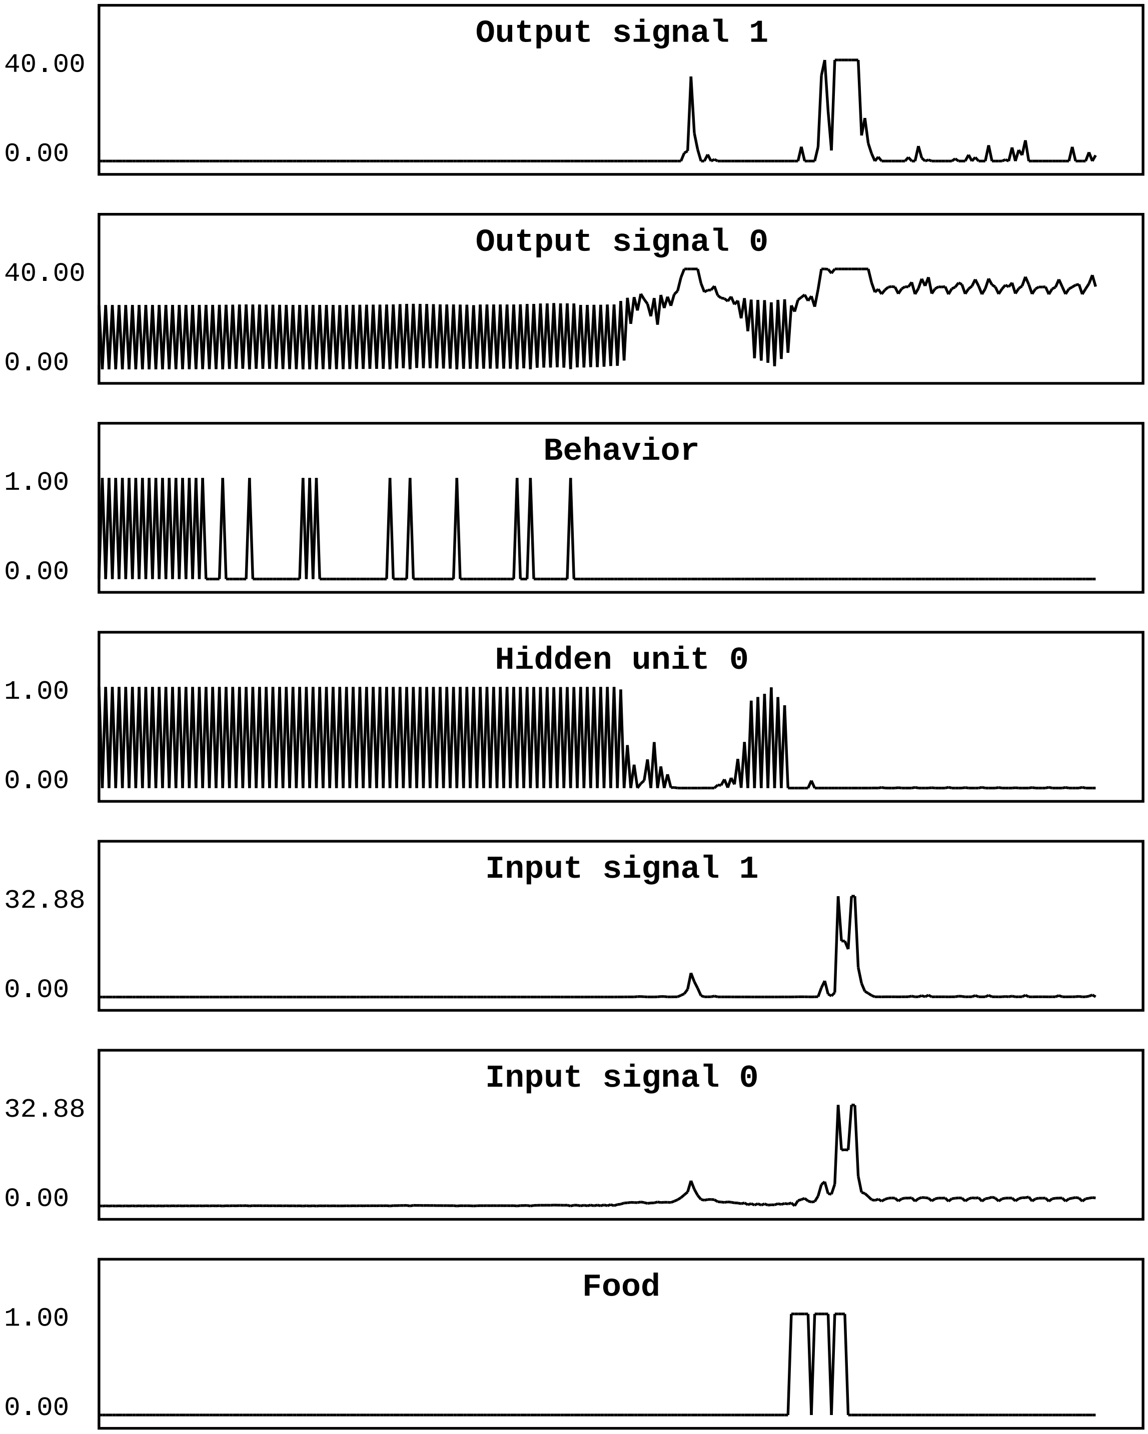

Figure 23: Profile of the gray agent (300 time steps). Although its initial inputs (food & signals) look identical to that of the white agent, this agent's initial behavior oscillates between "Follow food" and "Follow signal." The difference is resolved in Figure 24.

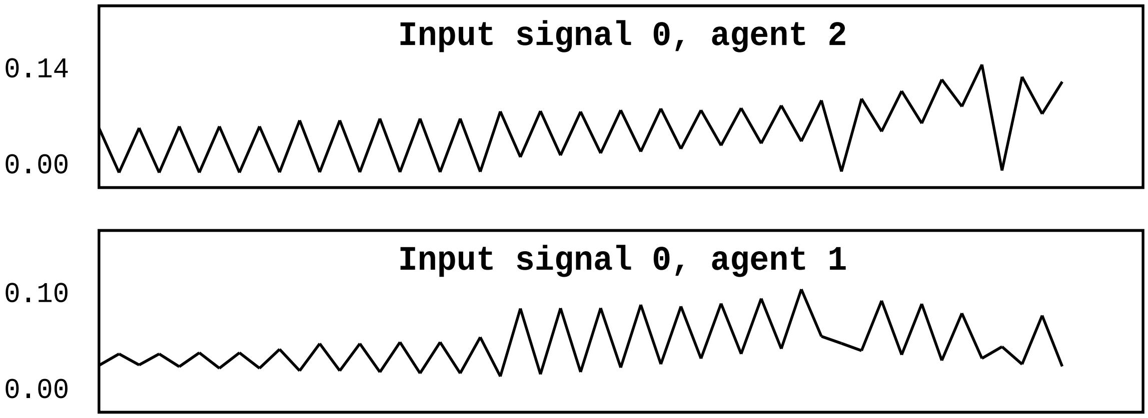

Figure 24: Magnified view of the first input signal of agents 1 and 2 (50 time steps). The white agent's input begins oscillating between .03 and .04. The gray agent's input begins oscillating between .06 and 0. It is not the magnitude, but the difference in phase which is responsible for the agents' different behaviors.

5.3 The Evolution of Communication

The previous subsections focused on how the evolved systems of signals and responses allowed the agents to communicate the presence of food. This section, in contrast, focuses on an orthogonal question: how the system of communication *evolves*.

Section 5.1 described an evolved communication scheme in which a one agent signalled the presence of food by "yelling," and the other agent responded to the louder signal by moving towards the food. As stated earlier, this solution was discovered in 700 generations. To determine how it came about, we compared final agents with their parents in generation 699, and then their grandparents, etc., tracing their ancestry all the way back to randomly-generated agents in the initial generation. Our first step in analyzing this lineage was to determine if evolution proceeded smoothly or by phase transition (Huberman and Hogg, 1987; Pollack, 1991). To do this, we compared gross characteristics of the agents over the course of the run: fitness, number of links in the networks, and total absolute value of the network weights (Figure 25). We discovered a phase tran-

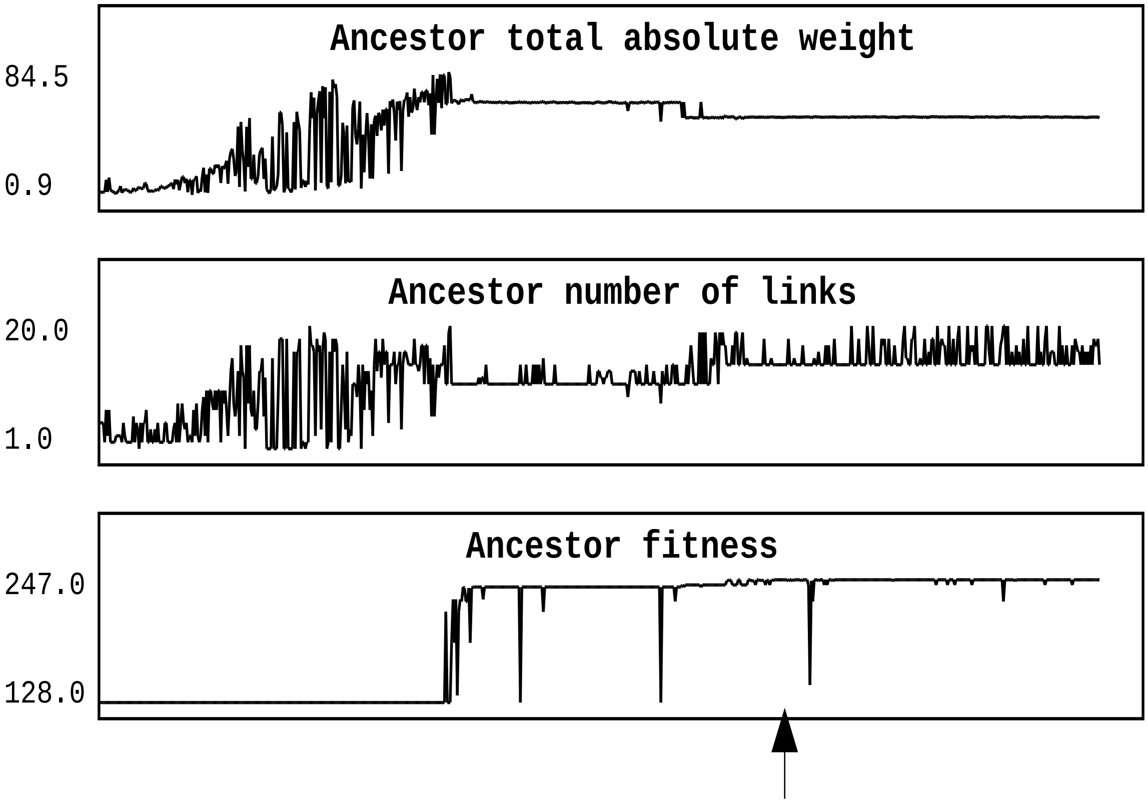

Figure 25: Gross characteristics of the ancestors of the agents which demonstrated a constant communication scheme (Section 5.1). This figure shows generations 0 - 700. (The arrow denotes generation 475, used later.)

sition before which fitness was low and gross characteristics varied widely, and after which fitness was high and gross characteristics were fairly stable.

We next moved from gross characteristics to specific agent features; specifically, we began examining the networks which comprised the agents. Because the agents were fairly stable towards the end of the run, we plotted only through generation 475 (the location of the arrow in Figure 25). Each network consisted of a maximum of 20 links in a 5x4 matrix (w_{00} - w_{43}). This result is split across Figures 26 and 27. Once again, we observed a phase transition in the weight values.

Unfortunately, just seeing the weight changes in the networks over time is of little help in understanding how the varying structure effects an increase in fitness. We proceeded with a *func-*

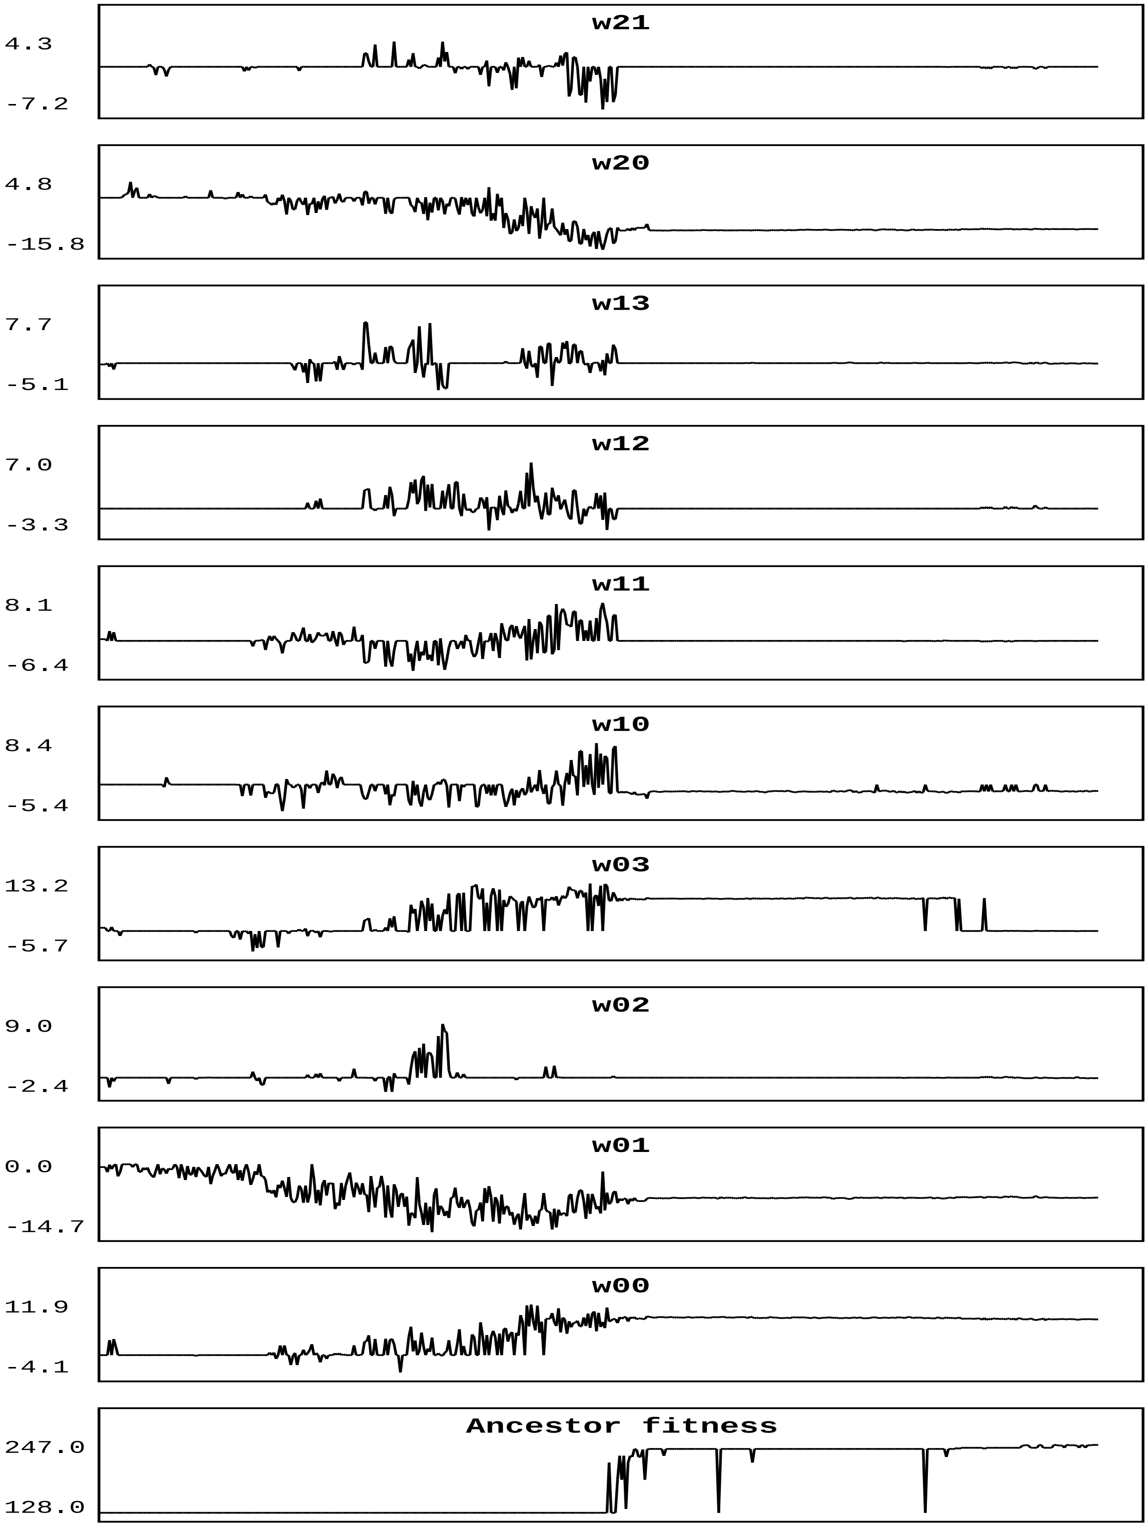

Figure 26: Evolution of network weights w_{00} - w_{21}, generations 0-475. Non-links are shown as a zero weight.

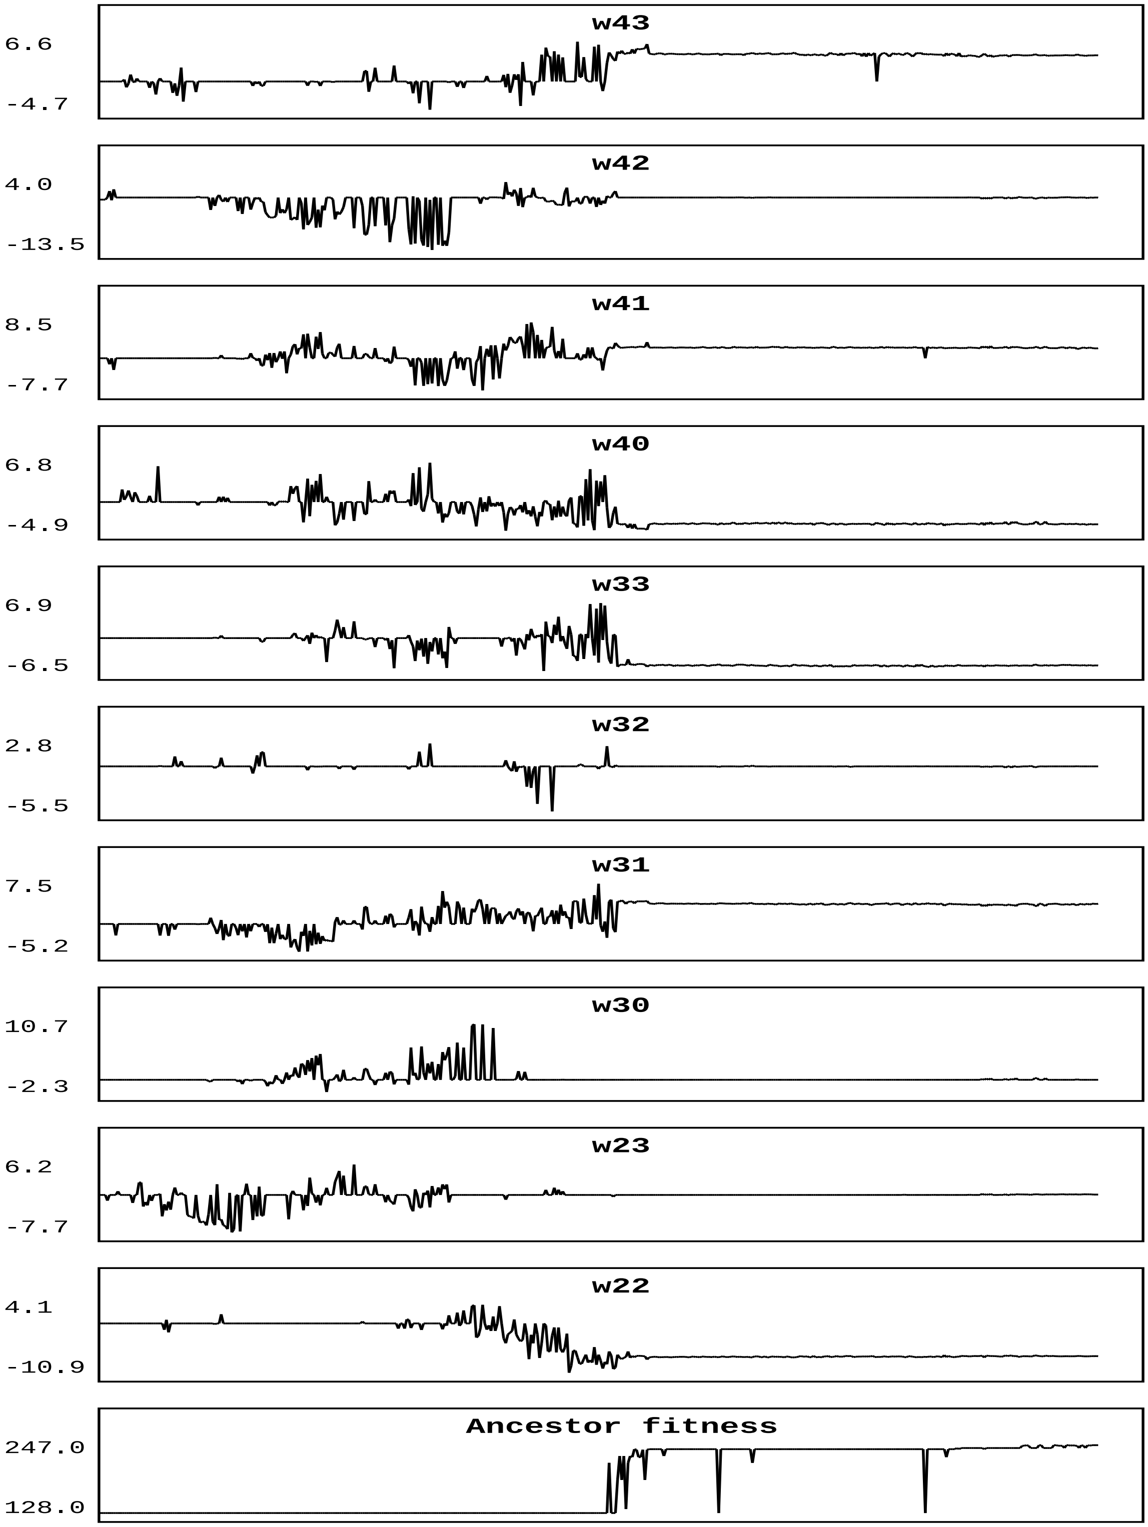

Figure 27: Evolution of network weights w_{22} - w_{43}, generations 0-475. Non-links are shown as a zero weight.

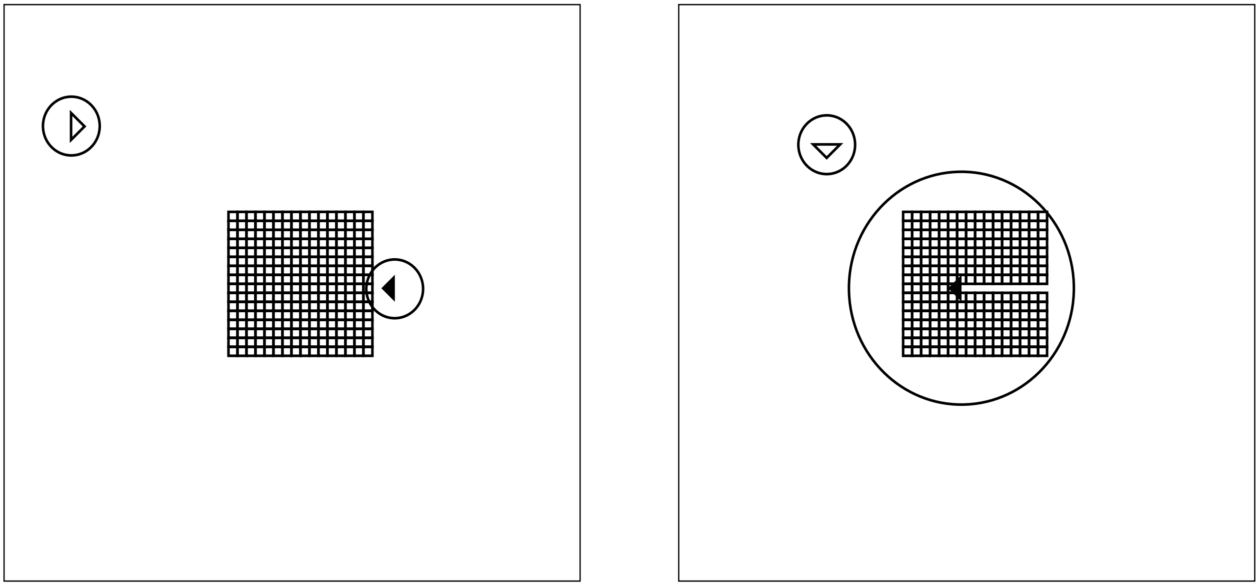
(a) (b)

Figure 28: Snapshots of the agents in their environment at t=10 and t=30 (cf. Figure 16). The agents come from the final solution at generation 700. The particular time steps, chosen arbitrarily, simply capture the situation both before and after the black agent finds the food.

tional analysis as follows. The key to the solution in the agents of generation700 is that *they have learned to communicate the presence of food*. In particular, if neither agent has discovered the food patch, both agents search; if one agent has discovered the patch, it communicates this knowledge to the other. Figure 28, where the agents are taken from the final solution at generation 700, illustrates these two critical points in the task. To determine how the agents evolve functionally, we compared how they behave in these two situations. Specifically, we compared four features: the output signal of the black agent before it encounters food, the response of the white agent to this signal, the output signal of the black agent after it has discovered the food patch, and the response of the white agent to this signal. Furthermore, we graphed the absolute difference in values between the output signals of the black agent in the food/no-food cases.

Figure 29 shows the results. Let the black agent be 0; the white agent be 1; the food signal of agent 0 be S_f, and the no-food signal of agent 0 be S_n. Behaviorally, the critical aspect of the

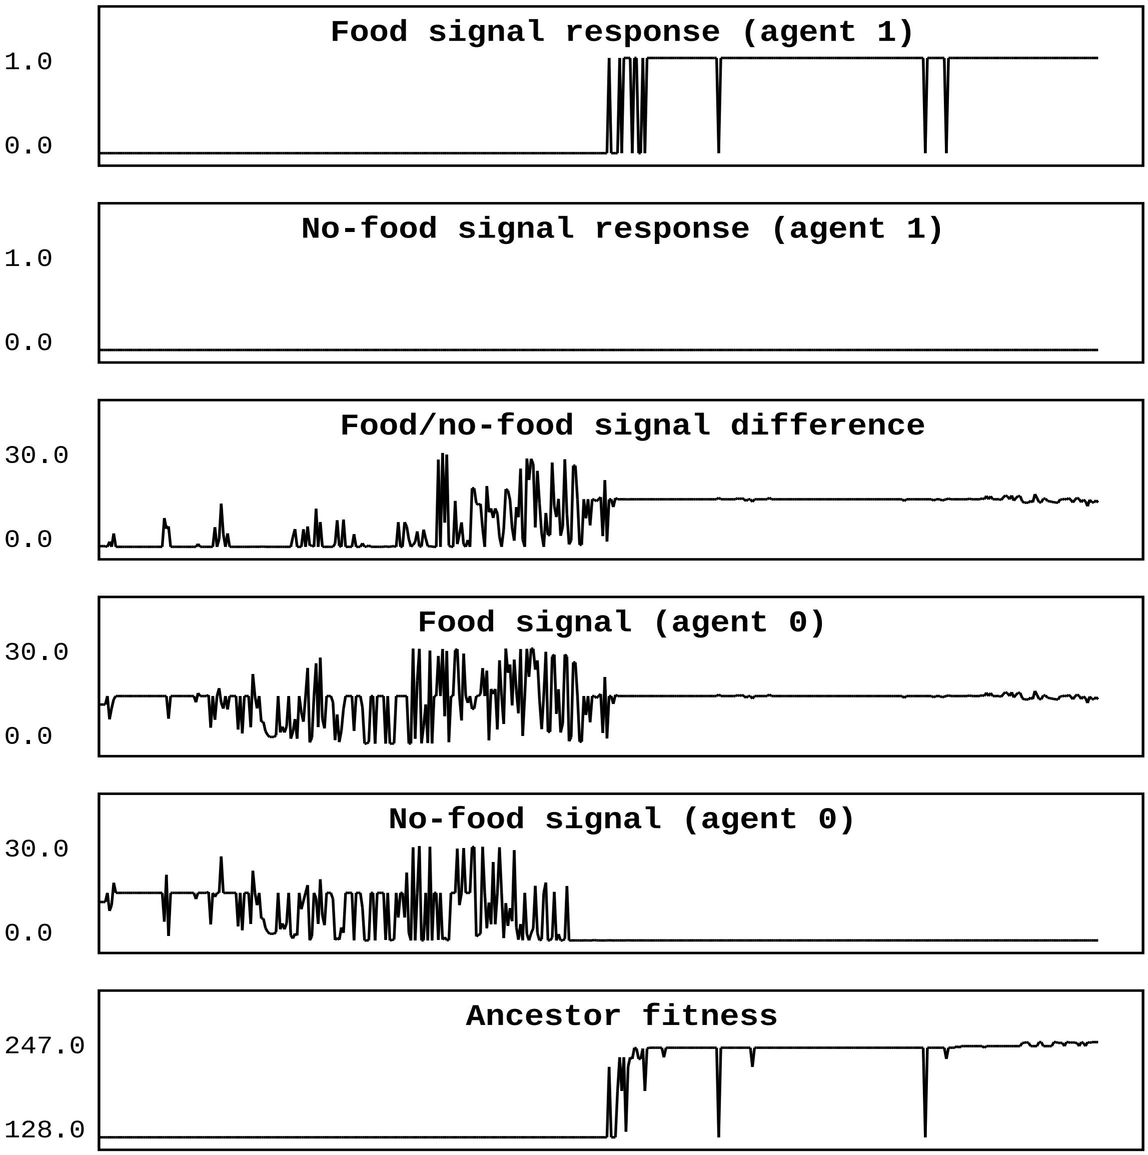

Figure 29: Functional difference of the agents over time (generations 0 - 475). Two events precede the phase transition in fitness: a distinction in the food/no-food signals by agent 0, and a recognition of this distinction by agent 1.

figure is the response agent 1 to S_f. The phase transition in fitness corresponds exactly to a phase transition in this behavior.

Why does it take 475 generations for these phase transitions to occur? From the middle graph, one can observe a non-linear increase in $| S_f - S_n |$, representing a distinction in the food /

no-food cases by agent 0. Interestingly, however, the phase transition in fitness is not forthcoming despite many large spikes in this value, since two conditions must be met for the jump in fitness:

C_0 – agent 0 must effect a distinction between S_n and S_f,
C_1 – agent 1 must respond to this distinction appropriately.

Consequently, search is hindered by two factors. First, C_0 must logically precede C_1, and second, the simple fitness function used for this task (number of pieces of food consumed) offers no reward for satisfying C_0 if C_1 is violated.

6 Remarks

We began with very few assumptions about the nature of communication, essentially stripping away the information-theory veneer that has made previous systems easy to understand. First we replaced the engineer with evolutionary search. Second, we eliminated discrete events and allowed the agents to modify channels with continuous values. These assumptions did not prevent solutions to the modified Tracker problem; in fact some novel approaches were discovered. We were able to evolve agents which demonstrated such task-specific behaviors as recruitment. In this section, we discuss implications of our results.

6.1 The Communicative Stance

Detecting the presence of communication is more difficult than it sounds. Communication can occur across long and short distances of both space and time, as can random noise. As examples of evolved communication, we rejected the results in Section 3, but embraced those in Section 5. What justifies the unequal treatment of the two?

Beer (1992) suggests dynamical systems theory as a unifying metaphor in which to understand the behavior of adaptive agents. Certainly, this language applies to all the results above: in each case, the agents evolved a “structural congruence” (Beer, 1990) between their internal dynamics and the dynamics of their environment. It is congruence that permitted the interlocking patterns of behavior (Winograd and Flores, 1986) which we termed either pseudo-random search (Section 3) or communication (Section 5).

McDermott (1981) warns against *creative naming*, which in our case means thinking that a set of agents is communicating simply because we adopt the label “communication” to describe their actions. If the agents’ behaviors can be explained at the level of dynamical systems, we must carefully justify adopting the language of *communication*.

Dennett (1987) deals with a similar problem regarding the appropriate language in which to discuss *mind*, and designates two very different levels of explanation:

> Suppose, for example, that we adopt the intentional stance toward bees, and note with wonder that they seem to *know* that dead bees are a hygiene problem in a hive; when a bee dies its sisters *recognize* that it has died, and *believing* that dead bees are a health hazard and *wanting*, rationally enough, to avoid health hazards, they *decide* they must remove the dead bee immediately. Thereupon they do just that. Now if that fancy an intentional story were confirmed, the bee-system designer would be faced with an enormously difficult job. Happily for the designer (if sadly for bee romantics), it turns out that a much lower-order explanation suffices: dead bees secrete oleic acid; the smell of oleic acid turns on the “remove it” subroutine in the other bees; put a dab of oleic acid on a live, healthy bee, and it will be dragged, kicking and screaming, out of the hive (Dennett, 1987, p. 256).

He goes on to justify the “bee romantic” level of explanation by the fact that it may lead to a “better” description of a system, where “better” is realized in terms of predictive power or generality (Dennett, 1987, p. 139). He labels this description in terms of beliefs, desires, and the like as the *intentional stance.*

Similarly, saying that the agents of Section 5 are *communicating* is a stance we adopt towards the agents, one which leads to certain predictions and generalizations about the agents’ behaviors. In repeating the experiments, we showed how each time a set of agents evolved a communication scheme which somehow helped perform the task. Sometimes they evolved the “scream at food” communication scheme; other times they evolved “be quiet when eating” scheme; still other times they evolved “oscillate until finding food” scheme. Although these schemes appear quite different, an underlying commonality binds them together: namely, in the types of information the agents transmit, receive, and use. In all cases, the agents have learned to communicate to each other the presence or absence of food, the most critical object in their fitness environment. Interpreting the signals as communicating information allows not only this generalization, but also the prediction that any evolved communication scheme will convey similar information, regardless of its surface appearance.

We avoid calling the pseudo-random search of the agents in Section 3 “communication” because the agents appear quite different at the information processing level. Their signals and responses remain invariant when food was removed from their environment – thus failing the test which we used to operationalize our notion of communication.

6.2 Relativized Semantics

Dynamical systems theory and the communicative stance provide two levels at which to describe the agents’ behavior, but of course this does not exhaust the possible descriptions. So far, we have avoided calling the evolved communication schemes *languages*. Now we adopt that perspective.

As languages, the *constant* and *oscillatory* systems of communication in Section 5 appear quite impoverished. They possess no syntax by which different signals can combine into higher-level signals, and thus they lack both compositionality (Chandrasekaran, 1990) and generative capacity (Fodor and Pylyshyn, 1988; Pollack, 1990). But, interestingly, these systems of communication do possess a semantics. Consider the agents in experiment 3 (Section 5.1). They learned to generate signals of varying magnitude to indicate the presence of food. From the perspective of one of these agents, the signal “High, high, high...” *means* that food is present, while the signal “Low,

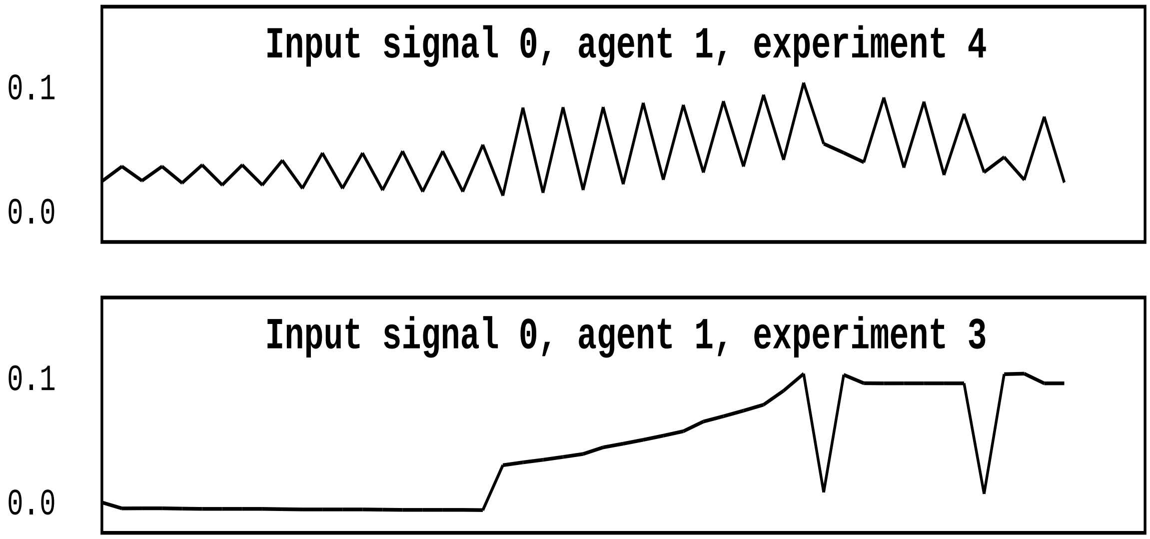

Figure 30: Comparison of input streams across experiment 3 (Section 5.5.1) and experiment 4 (Section 5.5.2) for the evolved white agent, t=0-50.

low, low" *means* no food has been found. In other words, the agents have developed a language with meaning relative to the task domain; consequently, we propose that these signals possess a semantics, which we call *relativized semantics.*

These semantics emerge independently of the programmer's notion of how the agents should communicate. Consider the inputs to the white agents in experiment 3 (Section 5.1) and experiment 4 (Section 5.2), shown together in Figure 30. Although both were evolved in an identical fashion, the languages are radically different: the first is based upon the input signal's magnitude; the second is based upon its phase.

6.3 Conclusions

The most difficult aspect of speech recognition is determining how to carve a continuous sound wave into meaningful pieces – if the speaker pauses between words, the problem is trivial. An infant just learning language or an adult learning a foreign tongue faces the same problem: just where does one word stop and the next begin (if indeed there even is a clear boundary)? The discrete communication systems in Section 1 never address this issue. Indeed, they never address how an agent detects communication as opposed to noise, how the discrete symbols arise, or how the discrete symbols may be integrated with the continuous substrate of sound.

Our agents perform differently – at no time is the input stream partitioned, normalized, or recognized; it simply modulates the behavior of the agent (as in Kolen, 1994b). Yet it does so in a way which admits a meaningful explanation at the information processing level, one which points to a integration of the continuous substrate of sound with the discrete system of words. Furthermore, it allows not just an integration of syntax and semantics, but a grounding of semantics in both overt behavior and changes to internal state.

While a set of agents evolving a way to signal the presence of food is a far cry from natural language, the evolved communication – and we really mean *communication* in the sense of Dennett – is also a far cry from the way communication has been studied as matching of isolated symbols. And we feel it is a step in the right direction towards a true understanding of the evolution of communication.

7 Acknowledgments

The authors thank John Kolen for help with an earlier version of this work. His insights in communication helped shaped many of the ideas above. We thank Pete Angeline, without whom GNARL would not exist. Finally, we thank Barbara Becker for feedback on an earlier draft of this paper.

8 References

Angeline, P. J., Saunders, G. M., and Pollack, J. B. (1994). An evolutionary algorithm that constructs recurrent neural networks. *IEEE Transactions on Neural Networks*, 5(1):54-65.

Arbib, M. A., Schweighofer, N., and Thach, W. T. (1994). Modelling the role of the cerebellum in prism adaptation. In Cliff, D., Husbands, P., Meyer, J. A., and Wilson, S. W., editors, *From*

Animals to Animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior, pages 36-44, Cambridge, Massachusetts. MIT Press.

Arkin, R. C. (1992). Cooperation without communication: Multiagent schema-based robot navigation. *Journal of Robotic Systems*, 9(3):331-364.

Arkin, R. C. and Hobbs, J. D. (1993). Dimensions of communication and social organization in multi-agent robotic systems. In Meyer, J. A., Roitblat, H. L., and Wilson, S. W., editors, *From Animals to Animats 2: Proceedings of the Second International Conference on Simulation of Adaptive Behavior*, pages 486-493, Cambridge, MA. MIT Press.

Barto, A. G. (1990). Connectionist learning for control. In Miller-III, W. T., Sutton, R. S., and Werbos, P. J., editors, *Neural Networks for Control*, chapter 1, pages 5-58. MIT Press, Cambridge.

Beer, R. D. (1990). *Intelligence as Adaptive Behavior: An Experiment in Computational Neuroethology,* volume 6 of *Perspectives in Artificial Intelligence*. Academic Press, Inc.

Beer, R. D. (1992). A dynamical systems perspective on autonomous agents. Technical Report CES-92-11, Case Western Reserve University, Cleveland, Ohio.

Beer, R. D. and Gallagher, J. C. (1992). Evolving dynamical neural networks for adaptive behavior. *Adaptive Behavior*, 1(1):91-122.

Brooks, R. A. (1986). A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, 2(1):14-23.

Brooks, R. A. (1991a). Intelligence without representations. *Artificial Intelligence*, 47:139-159.

Brooks, R. A. (1991b). Challenges for complete creature architectures. In Meyer, J. A. and Wilson, S. W., editors, *From Animals to Animats: Proceedings of the First International Conference on Simulation of Adaptive Behavior*, pages 434-443. MIT Press.

Chandrasekaran, B. (1990). What kind of information processing is intelligence: A perspective on ai paradigms and a proposal. In Partridge, D. and Wilks, Y., editors, *The foundations of artificial intelligence: A sourcebook*, pages 14-46. Cambridge University Press.

Collins, R. J. and Jefferson, D. R. (1991). Representations for artificial organisms. In Meyer, J. A. and Wilson, S. W., editors, *From Animals to Animats: Proceedings of the First International Conference on Simulation of Adaptive Behavior*, pages 382-390, Cambridge, MA. MIT Press.

Collins, R. J. and Jefferson, D. R. (1992). Antfarm: Towards simulated evolution. In Langton, C. G., Taylor, C., Farmer, J. D., and Rasmussen, S., editors, *Artificial Life II*, pages 579-601. Addison-Wesley.

de Bourcier, P. and Wheeler, M. (1994). Signalling and territorial aggression: An investigation by means of synthetic behavioral ecology. In Cliff, D., Husbands, P., Meyer, J. A., and Wilson, S.

W., editors, *From Animals to Animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior*, pages 463-472, Cambridge, Massachusetts. MIT Press.

Deneubourg, J. L., Goss, S., Franks, N., Sendova-Franks, A., Detrain, C., and Chrétien, L. (1991). The dynamics of collective sorting: Robot-like ants and ant-like robots. In Meyer, J. A. and Wilson, S. W., editors, *From Animals to Animats: Proceedings of the First International Conference on Simulation of Adaptive Behavior*, pages 356-363. MIT Press.

Dennett, D. C. (1987). Intentional systems in cognitive ethology: The Panglossian paradigm defended. In *The Intentional Stance*, pages 237-286. MIT Press, Cambridge, MA.

Donnart, J. Y. and Meyer, J. A. (1994). An hierarchical classifier system implementing a motivationally autonomous animat. In Cliff, D., Husbands, P., Meyer, J. A., and Wilson, S. W., editors, *From Animals to Animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior*, pages 144-153, Cambridge, Massachusetts. MIT Press.

Fodor, J. A. and Pylyshyn, Z. W. (1988). Connectionism and cognitive architecture: A critical analysis. *Cognition*, 28:3-71.

Fogel, D. B. (1992). *Evolving Artificial Intelligence*. Ph.D. thesis, University of California, San Diego.

Grefenstette, J. J. (1992). The evolution of strategies for multiagent environments. *Adaptive Behavior*, 1(1):65-90.

Harvey, I., Husbands, P., and Cliff, D. (1994). Seeing the light: Artificial evolution, real vision. In Cliff, D., Husbands, P., Meyer, J. A., and Wilson, S. W., editors, *From Animals to Animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior*, pages 392-401, Cambridge, Massachusetts. MIT Press.

Huberman, B. A. and Hogg, T. (1987). Phase transitions in artificial intelligence systems. *Artificial Intelligence*, 33:155-171. North-Holland Publishers.

Jefferson, D., Collins, R., Cooper, C., Dyer, M., Flowers, M., Korf, R., Taylor, C., and Wang, A. (1992). Evolution as a theme in artificial life: The Genesys/Tracker system. In Langton, C. G., Taylor, C., Farmer, J. D., and Rasmussen, S., editors, *Artificial Life II: Proceedings of the Workshop on Artificial Life*, pages 549-577. Addison-Wesley.

Kirkpatrick, S., Gelatt, C. D., and Vecchi, M. P. (1983). Optimization by simulated annealing. *Science*, 220:671-680.

Kolen, J. F. (1994a). *Explorations of Back-Propagation and the Computational Capabilities of Recurrent Neural Networks*. Ph.D. thesis, Ohio State University.

Kolen, J. F. (1994b). Fool’s gold: Extracting finite state machines from recurrent neural networks. In Cowan, J. D., Tesauro, G., and Alspector, J., editors, *Advances in Neural Information Processing 6*. Morgan Kaufmann.

MacLennan, B. (1992). Synthetic ecology: An approach to the study of communication. In Langton, C. G., Taylor, C., Farmer, J. D., and Rasmussen, S., editors, *Artificial Life II*, pages 631-658. Addison-Wesley.

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

Maes, P. (1991). Adaptive action selection. In *Proceedings of the Thirteenth Annual Conference of the Cognitive Science Society.*

Maes, P. (1992). Learning behavior networks from experience. In Varela, F. J. and Bourgine, P., editors, *Toward a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life*, pages 48-57, Cambridge, MA. MIT Press.

Mataric, M. J. (1993). Designing emergent behaviors: From local interactions to collective intelligence. In Meyer, J. A., Roitblat, H. L., and Wilson, S. W., editors, *From Animals to Animats 2: Proceedings of the Second International Conference on Simulation of Adaptive Behavior*, pages 432-441, Cambridge, MA. MIT Press.

Mataric, M. J. (1994). Learning to behave socially. In Cliff, D., Husbands, P., Meyer, J. A., and Wilson, S. W., editors, *From Animals to Animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior*, pages 453-462, Cambridge, Massachusetts. MIT Press.

McDermott, D. (1981). Artificial intelligence meets natural stupidity. In Haugeland, J., editor, *Mind Design: Philosophy, Psychology, Artificial Intelligence*, chapter 5, pages 143-160. MIT Press, Cambridge, MA, reproduced from *SIGART Newsletter*, No. 57, April 1976.

McFarland, D. (1994). Towards robot cooperation. In Cliff, D., Husbands, P., Meyer, J. A., and Wilson, S. W., editors, *From Animals to Animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior*, pages 440-444, Cambridge, Massachusetts. MIT Press.

Meyer, J. A. and Guillot, A. (1991). Simulation of adaptive behavior in animats: Review and prospect. In Meyer, J. A. and Wilson, S. W., editors, *From Animals to Animats: Proceedings of the First International Conference on Simulation of Adaptive Behavior*, pages 2-14, Cambridge, MA. MIT Press.

Meyer, J. A. and Guillot, A. (1994). From SAB90 to SAB94: Four years of Animat research. In Cliff, D., Husbands, P., Meyer, J. A., and Wilson, S. W., editors, *From Animals to Animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior*, pages 2-11, Cambridge, Massachusetts. MIT Press.

Pollack, J. B. (1990). Recursive distributed representations. *Artificial Intelligence*, 46:77-105.

Pollack, J. B. (1991). The induction of dynamical recognizers. *Machine Learning*, 7:227-252.

Robbins, P. (1994). The effect of parasitism on the evolution of a communication protocol in an artificial life simulation. In Cliff, D., Husbands, P., Meyer, J. A., and Wilson, S. W., editors, *From Animals to Animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior*, pages 431-439, Cambridge, Massachusetts. MIT Press.

Saunders, G. M., Kolen, J. F., and Pollack, J. B. (1994). The importance of leaky levels to behavior-based AI. In Cliff, D., Husbands, P., Meyer, J. A., and Wilson, S. W., editors, *From Animals to Animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior*, pages 275-281, Cambridge, Massachusetts. MIT Press.

Saunders, G. M., Angeline, P. J., and Pollack, J. B. (1994). Structural and behavioral evolution of recurrent networks. In Cowan, J. D., Tesauro, G., and Alspector, J., editors, *Advances in Neural Information Processing 7*, pages 88-95. Morgan Kaufmann.

Schmajuk, N. A. (1994). Behavioral dynamics of escape and avoidance: A neural network approach. In Cliff, D., Husbands, P., Meyer, J. A., and Wilson, S. W., editors, *From Animals to Animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior*, pages 118-127, Cambridge, Massachusetts. MIT Press.

Shannon, C. E. and Weaver, W. (1949). *The Mathematical Theory of Communication*. The University of Illinois Press, Urbana, Illinois.

Steels, L. (1994). A case study in the behavior-oriented design of autonomous agents. In Cliff, D., Husbands, P., Meyer, J. A., and Wilson, S. W., editors, *From Animals to Animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior*, pages 445-452, Cambridge, Massachusetts. MIT Press.

Werner, G. M. and Dyer, M. G. (1992). Evolution of communication in artificial organisms. In Langton, C. G., Taylor, C., Farmer, J. D., and Rasmussen, S., editors, *Artificial Life II*, pages 659-687. Addison-Wesley.

Werner, G. M. and Dyer, M. G. (1993). Evolution of herding behavior in artificial animals. In Meyer, J. A., Roitblat, H. L., and Wilson, S. W., editors, *From Animals to Animats 2: Proceedings of the Second International Conference on Simulation of Adaptive Behavior*, pages 393-399, Cambridge, MA. MIT Press.

Winograd, T. and Flores, F. (1986). *Understanding Computers and Cognition*. Addison-Wesley, Reading, MA.

Yanco, H. and Stein, L. A. (1993). An adaptive communication protocol for cooperating mobile robots. In Meyer, J. A., Roitblat, H. L., and Wilson, S. W., editors, *From Animals to Animats 2: Proceedings of the Second International Conference on Simulation of Adaptive Behavior*, pages 478-485, Cambridge, MA. MIT Press.