

Making Meaning Happen

PATRICK GRIM, TRINA KOKALIS, ALI ALAI-TAFTI,
NICHOLAS KILB and PAUL ST DENIS

*Group for Logic & Formal Semantics, Dept. of Philosophy, SUNY
at Stony Brook, Stony Brook, NY 11794, USA
e-mail: pgrim@notes.cc.sunysb.edu*

Abstract. What is it for a sound or gesture to have a meaning, and how does it come to have one? In this paper, a range of simulations are used to extend the tradition of theories of meaning as use. The authors work throughout with large spatialized arrays of sessile individuals in an environment of wandering food sources and predators. Individuals gain points by feeding and lose points when they are hit by a predator and are not hiding. They can also make sounds heard by immediate neighbours in the array, and can respond to sounds from immediate neighbours. No inherent meaning for these sounds is built into the simulation; under what circumstances they are sent, if any, and what the response to them is, if any, vary initially with the strategies randomized across the array. These sounds do take on a specific function for communities of individuals, however, with any of three forms of strategy change: direct imitation of strategies of successful neighbours, a localized genetic algorithm in which strategies are ‘crossed’ with those of successful neighbours, and neural net training on the behaviour of successful neighbours. Starting from an array randomized across a large number of strategies, and using any of these modes of strategy change, communities of ‘communicators’ emerge. Within these evolving communities the sounds heard from immediate neighbours, initially arbitrary across the array, come to be used for very specific communicative functions. ‘Communicators’ make a particular sound on feeding and respond to that same sound from neighbours by opening their mouths; they make a different sound when hit with a predator and respond to that sound by hiding. Robustly and persistently, even in simple computer models of communities of self-interested agents, something suggestively like signalling emerges and spreads.

Keywords: meaning, communication, genetic algorithms, neural nets, multi-agent modeling

Received October 2003; revision requested April 2004; final version accepted 1 May 2004

1. Introduction

Colonies of sessile *compuzoans* are commonly found in simple environments of drifting food sources and small wandering predators. In order to feed, the *compuzoan*

must simply open its mouth when a food source comes its way. In order to avoid harm, it must hide when a predator is in the area. Like many cave creatures, however, *compuzoans* are blind; they cannot see food or predators coming. They cannot even perceive that a neighbour has successfully fed, or that it has been hit by a predator. But *compuzoans* are enabled with an ability to make simple sounds that carry as far as their immediate neighbours in the colony.

Could *compuzoans* come to use these sounds as *signals*? This is a philosophical thought experiment that we have explored using a range of computer models. In our simplest models, the strategies of our *compuzoans* change by simple imitation in the direction of more successful neighbours. Less successful strategies are replaced by strategies that have proven more successful in their immediate neighbourhood.¹ In more complicated variations, behavioural strategies are replaced with variations formed by genetic algorithm recombination with the strategies of more successful neighbours. In the most sophisticated models we offer here, the nervous systems of our *compuzoans* consist of neural nets, changed by backpropagation training on the behaviour of successful neighbours.

In all of these models, communities of ‘communicators’ emerge that make a particular sound on feeding and react to that same sound from a neighbour by opening their mouths. Communicators make a different sound when hurt by a predator, and respond to that sound from a neighbour by hiding. It should be emphasized that *compuzoans* benefit only individually by feeding and avoiding predation; there is no reward for communication *per se*. Consistently and robustly, across mechanisms of strategy change by imitation, genetic algorithm and neural net learning, the result is a co-evolution of strategies coordinated in terms of signalling and response.

2. The philosophical motivation

Many are the philosophical theories of meaning, and many are their aims. It is probably wrong to think of all such theories—Fregean, Gricean, Davidsonian—as theories of meaning in the same sense. It is perhaps wrong to think of them as trying to answer the same questions about meaning.

We consider the question we are pursuing to be the simplest and most basic in the philosophy of meaning: What is it for a sound or gesture to have meaning at all? For the present our interests in the theory of meaning are restricted to this basic question. We want to know what meaning amounts to and how something can take on a meaning: what meaning is and how it happens.

The models we have to offer fall in the general category of theories of meaning as use. Essential to meaning in these models is behavioural co-ordination across a community; meaning is not something psychological in an individual but something developed socially across a community. Wittgensteinian theories of meaning as use, however, are notorious for teasing vagueness and obscurity. One way of pursuing the basic insight of such theories without incurring these vices is to follow the game-theoretic approach first suggested in Lewis’s *Convention* (1969) and further developed in Skyrms’s *Evolution of the Social Contract* (1996). It is these that we take as our philosophical precursors here.²

With the notable exception of Bennett’s *Linguistic Behaviour* (1976), philosophers have generally been wary of speculations as to how language might begin or how systems of signalling or communication might emerge. We on the contrary take the

question of how communication might arise to be crucial to an understanding of what it is. In our less self-critical moments, we think of what we are doing as trying to understand the basic phenomenon of meaning by making meaning *happen*.

Our computer models track spatialized arrays of self-interested individuals in an environment of wandering food sources and predators. Our individuals gain points by feeding and avoid harm by hiding. Their range of perception is limited: They know only when they themselves are fed, when they are hurt and when an immediate neighbour has made a particular sound. Their responses to particular stimuli—resting in neutral, opening their mouths, hiding or making a particular sound—are dictated by a simple behavioural strategy or an elementary neural net. We begin with a large randomized array of such individuals, and have strategies change in the direction of that neighbouring strategy that has proven most successful. In various forms of the model, we have explored forms of strategy change in which the strategy of successful neighbours is imitated in full, in which strategies are cross-bred locally by genetic algorithm, and in which strategies are changed by partial training in neural nets.

The result in all of these models is that flourishing communities appear that are behaviourally co-ordinated so as to uniformly generate a particular sound at feeding and to react to that sound by opening their mouths, to generate a particular sound when hurt and to react to that sound by hiding. In actually working with these models, it is almost impossible not to think of what emerges as a simple signalling system: a simple pattern of communication in which particular sounds carry a particular *meaning*. From that perspective, ours are simple models in which we can see meaning as an aspect of behavioural co-ordination across a community, and in which we watch meaning arise.

It should be noted that on each round our individuals can perceive only that they themselves successfully feed or are hit by a predator; they cannot in general tell even when an immediate neighbour has fed or been hurt. Within evolving communities of ‘communicators’, in contrast, the sounds that agents can make and can hear from immediate neighbours are co-ordinated so as to serve a communicative function. Two clear indicators that this use of sound does not qualify as a form of perception are (1) that it is mediated by sounds which may be sent by one agent (or not) and received by another, and (2) because information transfer can still go wrong through that mediation. A ‘communicator’ may act in response to a signal from a different strategy, or from a ‘communicator’ with a different behavioural co-ordination of sounds, in a radically inappropriate way—opening its mouth just in time to be hit by a predator, for example. All of this marks the emerging use of sounds in a community of ‘communicators’ as signal-like rather than perception-like, giving us a simple simulational instantiation of a theory of meaning as use.

In our more sober moments, we have to admit that our ‘individuals’ are mere blips on a computer screen, controlled *in silico* by deterministic strategies or neural nets. Despite our working tendency to think of these as forming communities of communicators in which certain signals take on a particular meaning, we have to admit that there is no real entity before us that literally *means* anything by anything. On this more dignified characterization of what we are up to, meaning is not something we *make* but merely something we *model*. From even this deflationary perspective, however, the model may afford a better understanding of meaning. A model that tracks the dynamics of meaning may give us a deeper understanding

of the phenomenon even if real meaning does not literally appear within the model at all.

It must also be admitted that those elements the meaning of which we attempt to study here are at best simple signals, so simple as to lie beneath syntax and thus well beneath any theory of compositionality. Compositional theories, however, do assume meaning at the terminal nodes; what we want to understand is what it is for even terminal nodes to have a meaning, and how such meaning can arise. Contemporary work by Martin Nowak and others, moreover, suggests that models like those offered here might eventually be extended to explain emergence of elementary syntax as well (Nowak and Krakauer 1999, Nowak *et al.* 1999, 2000).

3. Modelling background

What is it for a sound or gesture to have a meaning? Theories of meaning as use are standardly introduced in opposition to classical approaches in which meaning is taken to be a relation. A sound or gesture is meaningful, on classical accounts, because it stands in a particular relation to something, and the thing to which it stands in the proper relation is taken to be its meaning. The question for any relational theory of meaning, then, is precisely what the crucial relation is and what it is a relation to. *Referential* theories take the essential relation to be reference to things in the world. Words have meanings because they have referents, and the meaning of a word is the thing to which it refers (Augustine c. 400 AD, Mill 1884, Russell 1921, 1940). *Ideational* theories take meaning to be a relation between a sound or gesture and the images, ideas, or internal representations it is used to express. The meaning of the word is the thing in the head it is used to convey, and communication becomes an attempt to transfer the contents of my head into yours, or to make the contents of your head match mine (Aristotle c. 330 BC, Hobbes 1651, Locke 1689). *Abstractual* theories portray meaning as a relation neither to things in the world nor to the contents of heads but to some third form of object, removed from the world and yet non-psychological (Frege 1918).

Outside of sophisticated forms of representationalism, perhaps, philosophers tend to think of at least simple relational theories as easy targets in the first few steps of philosophy of language.³ Within other disciplines now actively engaged in exploring models for communication, however—theoretical biology, evolutionary linguistics, computer science and artificial life—relational theories of meaning are clearly alive and well. We take it as a sure sign that the theory in play is of this sort when the measure of ‘identity of meaning’ or ‘successful communication’ between two individuals is a match between their representation maps or signal matrices. A referential theory, in which the meaning of a term is taken to be the object or situation it applies to, is more or less explicit in the work of Bruce MacLennan and Gordon Burghardt (MacLennan 1991, MacLennan and Burghardt 1994), John Batali and Michael Oliphant (Batali 1995, Oliphant and Batali 1997) and Kyle Wagner (2000). An ideational theory, in which communication involves a match of internal representations, is a clear theme in the work of Michael Levin (1995), Domenico Parisi (1997), Edward Hutchins and Brian Hazlehurst (1991, 1995), Daniel Livingstone and Colin Fyfe (Livingstone 2000, Livingstone and Fyfe 1999) and Martin Nowak and his collaborators (Nowak *et al.* 1999, 2000).

Perhaps those relational theories of meaning that have the most currency among contemporary philosophers are forms of representationalism. On such a view,

meanings are in the head, correspond to the outside world by representing it, may themselves have a linguistic structure and are communicated by way of a complex pattern of co-ordinated intentions (Fodor 1975, 1981). Our goal here is less to refute the representational picture than to develop an alternative using the formal tools of simulation: an alternative in which meaning is crucially pragmatic and fundamentally social.

The central tenet of any theory of *meaning as use* is that a grasp of meaning will come not by looking for the right relation to the right kind of object but by attention to the co-ordinated interaction of agents in a community. In practical terms, the measure of communication will be functional co-ordination alone, rather than any attempt to find matches between internal representations or referential matrices. The understanding of meaning that we seek may thus come with an understanding of the development of patterns of functional communication, but without our being able at any stage to identify a particular relation as the ‘meaning’ relation or a particular object—concrete, ideational or abstract—as the ‘meaning’ of a particular term. Although the cross-disciplinary modelling literature is clearly dominated by relational views of meaning, this more dynamical approach also has its representatives: we note with satisfaction comments in that direction in the theoretical and robotics work of Luc Steels (1996, 1998).

There are also other ways in which the models we explore here differ from their predecessors in the technical literature. An essential aspect of our models is spatialization, carried over from previous work in co-operation (Grim 1995, 1996, Grim *et al.* 1998). Our community is modelled as a two-dimensional cellular automata array, with all essential interactions functioning purely locally; each individual interacts with its immediate neighbors, but no individual interacts with all members of the community as a whole. The sounds made by individuals in the array are heard only by their immediate neighbours. Fitness is measured purely locally, and strategy change proceeds locally as well: it is always the locally successful neighbour that is used in strategy change by full imitation, genetic recombination or backpropagation learning. Spatialization of this thorough-going sort has not been exploited in earlier models for communication. Some studies employ a spatial task of some sort, but both communication and reproduction proceed globally across random selections from the population as a whole (MacLennan and Burghardt 1994, Cangelosi and Parisi 1998, Wagner 2000). In some cases both co-operation and communication are conceived spatially, but new strategies arise by mutation using a fitness algorithm applied globally across the population as a whole.⁴ Aside from our models, that by Ackley and Littman (1994) is perhaps the most consistently spatialized to date.⁵

In almost all previous models in the literature, communication is rewarded as an end in itself. Both ‘senders’ and ‘receivers’ are simultaneously rewarded for each case of ‘successful communication’, rather than a far more realistic economy in which immediate benefits can be expected to accrue to the receiver alone. An assumption of mutual benefit from communicative exchanges is explicitly made in Lewis (1969). The models in MacLennan (1991) reward both ‘senders’ and ‘receivers’ for successful communication, with strategies then perfected through genetic algorithm. As Ackley and Littman note, the resulting models reflect an artificial environment ‘where “truthful speech” by a speaker and “right action” by a listener cause food to rain down on both’ (Ackley and Litman 1994: 40). In later work, MacLennan and Burghardt (1994) and Wagner (2000) continue to use symmetrical reward models,

though these are underwritten with the explicit proviso that their investigations are limited to communication regarding co-operative tasks, motivated by a story of communication for co-operation in bringing down a large animal. That proviso also limits generalizability of their models to communication in general, however. Through explicit assumption or shared tasks, symmetrical rewards for communication also characterize the work of Werner and Dyer (1994), Saunders and Pollack (1996) and Hutchins and Hazlehurst (1991, 1995).

The need for a model of how communication regarding *non*-shared tasks might originate is noted explicitly by a number of authors (Ackley and Littman 1994, Batali 1995, Dyer 1995, Noble and Cliff 1996, Cangelosi and Parisi 1998). Batali writes,

While it is of clear benefit for the members of a population to be able to make use of information made available to others, it is not as obvious that any benefit accrues to the sender of informative signals. A good strategy, in fact, might be for an individual to exploit signals sent by others, but to send no informative signals itself. Thus there is a puzzle as to how coordinated systems of signal production and response could have evolved. (Batali 1995: 2)

In an overview of approaches, Domenico Parisi is still more explicit:

In the food and danger simulations the organism acts only as a receiver of signals and it evolves an ability to respond appropriately to these signals. It is interesting to ask, however, where these signals come from. . . Why should the second individual bother to generate signals in the presence of the first individual? The evolutionary 'goal' of the first individual is quite clear. Individuals who respond to the signal 'food' ('danger') by approaching (avoiding) the object they currently perceive are more likely to reproduce than individuals who do not do so. Hence, the evolutionary emergence of an ability to understand these signals. . . But why should individuals who perceive food or danger objects in the presence of another individual develop a tendency to respond by emitting the signal 'food' or 'danger'? (Parisi 1997: 129)

The models we offer here are intended to close this explanatory gap, showing how communication can emerge given precisely the reward structure Batali and Parisi call for. Here individuals develop patterns of communication in an environment in which they benefit only individually from capture of food and avoidance of predators, and indeed in which there is an assigned cost for generating signals. Without any restriction to shared tasks, the model offers steps toward understanding how signalling in general can emerge.

4. The environment of the common compuzoan

We use a randomized 64×64 two-dimensional cellular automata array of individuals carrying different strategies. Technically, the array forms a torus, 'wrapping around' so that individuals on the bottom edge have neighbors at the top edge and those at the left edge have neighbours on the right (figure 1). Ours are sessile *compuzoans*: like coral in a reef, the individuals themselves do not move. What does move are food sources and predators, each of which migrate in a random walk across the array.

If a predator lands on an individual that is not hiding, that individual is 'harmed' and loses a point. If a food source lands on an individual with its mouth open, that individual 'feeds' and gains a point. We should emphasize that the latter are food




Figure 1. A typical array of *compuzoans*, with different strategies shown as different colours and open mouths as central black squares.

sources: individuals feed from them, but the food sources themselves are not consumed and do not disappear. Like a cloud of plankton or a school of fish, perhaps, they continue their random walk offering nourishment for the next guy down the line.

On any given round, an individual's strategy may dictate that it opens its mouth or that it hides. Individuals are also capable of making one of two sounds on a given round, which we term sound *s1* and sound *s2*. But *compuzoans*' sounds are weak, detected only by themselves and their eight immediate neighbours.⁶

The perceptual world of the *compuzoan* is severely limited: they know only when they are fed, when they are hurt and when someone in their immediate neighbourhood has made a particular sound. Their behavioural repertoire is simple as well: a basic strategy dictates when they will open their mouths, hide or make sound *s1* or *s2*.

For even such simple creatures in such a simple environment, it seems, there *is* a behavioural strategy that would qualify as an elementary form of signalling or communication. Imagine a community of *compuzoans* that share the following behavioural strategy:

They make sound *s1* when they successfully feed.

They react to hearing sound *s1* by opening their mouths.




Figure 2. Migration of a single food source across a hypothetical array of communicators. In the left frame, a food source dot lands on an open mouth, indicated by grey shading. That central individual makes a sound * heard by its immediate neighbours, which in the second frame open their mouths in response. One of them feeds successfully, making a sound heard by its immediate neighbours, which are shown opening their mouths in the third frame. The result in a community of communicators is a chain reaction of efficient feeding.

They make sound s2 when they are hurt.

They react to sound s2 by hiding.

When a particular individual of such a community feeds, it makes sound s1. Its immediate neighbours open their mouths in response. Since the food source follows a random walk, it will then fall on an open mouth on the next round. That individual, having successfully fed, will make sound s1 and its neighbours will open their mouths. The result, in a community sharing such a strategy, is a chain reaction in which the food source is successfully exploited on each round (figure 2).

The dynamics of predation in such a community is interestingly different. When an individual is hurt, it makes sound s2. Its immediate neighbours then hide, avoiding predation on the next round. Because none of them is hurt, however, they make no sound. The predator thus finds a target on the following round. Here again an individual is hurt, resulting in sound s2 and predator avoidance on the next round (figure 3). In the environment we have created, in other words, communication regarding food can result in a chain reaction of food source exploitation on each round. Communication regarding predation, on the other hand, results in avoidance of predators only on every second round.⁷

We term *compuzoans* with this behavioural strategy ‘perfect communicators’. There are, of course, two forms of ‘perfect communicators’: those that use sound s1 for food and sound s2 for hiding, as above, and the symmetrical variation that uses sound s2 for food and sound s1 for hiding. The dynamics outlined make clear the potential benefit to any individual lucky enough to find itself in an established community of perfect communicators. The question we are interested in here, however, is whether communities of communicators will emerge if we start with initial arrays randomized across a range of different strategies. We take meaning to be a phenomenon of co-operative communicative behaviour. How might that co-operative communicative behaviour develop?

In the models that follow, we explore variations on this basic environment for the common *compuzoan*. In particular, we explore variations in the mode of strategy change. Over the course of 100 rounds, individuals total their gains and losses from feeding, predation and energy expenditure in mouth-opening, hiding and making sounds. At that point the crucial question is whether any of a cell’s




Figure 3. Migration of a single predator across a hypothetical array of communicators. In the left frame, a predator is shown as a hollow dot landing on a cell that is not hiding. The victim makes a sound shown as a cross, heard by itself and its immediate neighbours. In the second frame, all who hear the warning hide, shown in black, as the predator moves on. The cell on which it lands in the second frame has avoided victimization by hiding, however, and so does not make a sound. The predator thus lands on an unwarned victim again in the third frame, and that victim sends out a warning call. The result in a community of communicators is avoidance of predation only on every other round.

immediate neighbours has accumulated a higher score. In our simplest models, if any neighbour has a higher score, the strategy of our *compuzoan* is replaced wholesale with that of its most successful neighbour. In a second set of models, strategy change is by localized genetic algorithm: the strategies of our *compuzoans* are cross-bred with the strategies of their most successful neighbour. In the most complex models offered here, strategies are instantiated in simple neural nets and strategy change is by a partial backpropagation training on the behaviour of the most successful neighbour. Within the environment outlined for the common *compuzoan*, our core result is that simple signalling can emerge using any of these mechanisms of strategy change.

5. Emergence of communication by imitation

In our earliest studies, using strategy change by imitation, we kept our behavioural repertoires simple and thus our sample space of strategies small. Here each individual is envisaged as either opening its mouth or hiding on each round, with no possibility of doing both and no provision yet for coasting in neutral. Our individuals were allowed to make sound *s1* or sound *s2*, but could not make both. Given those limitations, our strategies can be represented as four-tuples $\langle f, h, s1, s2 \rangle$, with three possibilities for each variable. A strategy description $\langle f, h, s1, s2 \rangle$ specifies what the individual does when fed *f* (make no sound, make sound *s1* or make sound *s2*), what it does when hurt *h* (the same three options), what it does when it hears a sound *s1* (open its mouth, hide, or a random selection between the two), and what it does when it hears a sound *s2* (the same options). When two different sounds are heard from immediate neighbours, a random flip of a coin determines which sound the individual will respond to. When no sound is heard, in this simple model, a random flip of a coin dictates whether a cell will open its mouth or hide.

This first model is therefore structured in terms of 81 strategies. These can be envisaged in base 3 as follow:

- (0,0,0,0) never sounds and continues to behave randomly (hiding or opening its mouth) on hearing any sound.
- (0,0,0,1) never sounds, behaves randomly on hearing s1, but opens its mouth on hearing s2
- (0,0,0,2) never sounds, behaves randomly on hearing s1, but hides on hearing s2
- (0,0,1,0) never sounds, opens its mouth on hearing s1, behaves randomly on hearing s2
- (0,0,1,1) never sounds, opens its mouth on hearing either s1 or s2
- (0,0,1,2) never sounds, opens its mouth on hearing s1 and hides on hearing s2
- (0,0,2,0) never sounds, hides on hearing s1 and behaves randomly on hearing s2
- (0,0,2,1) never sounds, hides on hearing s1 but opens its mouth on hearing s2
- (0,0,2,2) never sounds, hides on hearing any sound
- (0,1,0,0) sounds s1 when hurt, behaves randomly on hearing any sound.
- (0,1,0,1) sounds s1 when hurt, behaves randomly on hearing s1, but opens its mouth on hearing s2
- (0,1,0,2) sounds s1 when hurt, behaves randomly on hearing s1, but hides on hearing s2
- (0,1,1,0) sounds s1 when hurt, opens its mouth on hearing s1, behaves randomly on hearing s2
- (0,1,1,1) sounds s1 when hurt, opens its mouth on hearing either s1 or s2
- (0,1,1,2) sounds s1 when hurt, opens its mouth on hearing s1 and hides on hearing s2
- (0,1,2,0) sounds s1 when hurt, hides on hearing s1 and behaves randomly on hearing s2
- (0,1,2,1) sounds s1 when hurt, hides on hearing s1 but eats on hearing s2
- (0,1,2,2) sounds s1 when hurt, hides on hearing any sound
- (0,2,0,0) sounds s2 when hurt, behaves randomly on hearing any sound.
- (0,2,0,1) sounds s2 when hurt, behaves randomly on hearing s1, but opens its mouth on hearing s2
- (0,2,0,2) sounds s2 when hurt, behaves randomly on hearing s1, but hides on hearing s2
- (0,2,1,0) sounds s2 when hurt, opens its mouth on hearing s1, behaves randomly on hearing s2
- (0,2,1,1) sounds s2 when hurt, opens its mouth on hearing either s1 or s2
- (0,2,1,2) sounds s2 when hurt, opens its mouth on hearing s1 and hides on hearing s2
- (0,2,2,0) sounds s2 when hurt, hides on hearing s1 and behaves randomly on hearing s2
- (0,2,2,1) sounds s2 when hurt, hides on hearing s1 but eats on hearing s2
- (0,2,2,2) sounds s2 when hurt, hides on hearing any sound
- (1,0,0,0) sounds s1 when fed, behaves randomly on hearing any sound
- (1,0,0,1) sounds s1 when fed, behaves randomly on hearing s1, opens its mouth on hearing s2
- (1,0,0,2) sounds s1 when fed, behaves randomly on hearing s1, hides on hearing s2
- (1,0,1,0) sounds s1 when fed, opens its mouth on hearing s1, behaves randomly on hearing s2
- (1,0,1,1) sounds s1 when fed, opens its mouth on hearing either s1 or s2
- (1,0,1,2) sounds s1 when fed, opens its mouth on hearing s1 and hides on hearing s2

- (1,0,2,0) sounds s1 when fed, hides on hearing s1 and behaves randomly on hearing s2
- (1,0,2,1) sounds s1 when fed, hides on hearing s1 but eats on hearing s2
- (1,0,2,2) sounds s1 when fed, hides on hearing any sound
- (1,1,0,0) sounds s1 when fed or hurt, behaves randomly on hearing any sound.
- (1,1,0,1) sounds s1 when fed or hurt, behaves randomly on hearing s1, opens its mouth on hearing s2
- (1,1,0,2) sounds s1 when fed or hurt, behaves randomly on hearing s1, hides on hearing s2
- (1,1,1,0) sounds s1 when fed or hurt, opens its mouth on hearing s1, behaves randomly on hearing s2
- (1,1,1,1) sounds s1 when fed or hurt, opens its mouth on hearing either s1 or s2
- (1,1,1,2) sounds s1 when fed or hurt, opens its mouth on hearing s1 and hides on hearing s2
- (1,1,2,0) sounds s1 when fed or hurt, hides on hearing s1 and behaves randomly on hearing s2
- (1,1,2,1) sounds s1 when fed or hurt, hides on hearing s1 but eats on hearing s2
- (1,1,2,2) sounds s1 when fed or hurt, hides on hearing any sound
- (1,2,0,0) sounds s1 when fed and s2 when hurt, behaves randomly on hearing any sound.
- (1,2,0,1) sounds s1 when fed and s2 when hurt, behaves randomly on hearing s1, opens its mouth on hearing s2
- (1,2,0,2) sounds s1 when fed and s2 when hurt, behaves randomly on hearing s1, hides on hearing s2
- (1,2,1,0) sounds s1 when fed and s2 when hurt, opens its mouth on hearing s1, behaves randomly on hearing s2
- (1,2,1,1) sounds s1 when fed and s2 when hurt, opens its mouth on hearing either s1 or s2
- (1,2,1,2) sounds s1 when fed and s2 when hurt, opens its mouth on hearing s1 and hides on hearing s2
- (1,2,2,0) sounds s1 when fed and s2 when hurt, hides on hearing s1 and behaves randomly on hearing s2
- (1,2,2,1) sounds s1 when fed and s2 when hurt, hides on hearing s1 but eats on hearing s2
- (1,2,2,2) sounds s1 when fed and s2 when hurt, hides on hearing any sound
- (2,0,0,0) sounds s2 when fed, behaves randomly on hearing any sound
- (2,0,0,1) sounds s2 when fed, behaves randomly on hearing s1, opens its mouth on hearing s2
- (2,0,0,2) sounds s2 when fed, behaves randomly on hearing s1, hides on hearing s2
- (2,0,1,0) sounds s2 when fed, opens its mouth on hearing s1, behaves randomly on hearing s2
- (2,0,1,1) sounds s2 when fed, opens its mouth on hearing either s1 or s2
- (2,0,1,2) sounds s2 when fed, opens its mouth on hearing s1 and hides on hearing s2
- (2,0,2,0) sounds s2 when fed, hides on hearing s1 and behaves randomly on hearing s2
- (2,0,2,1) sounds s2 when fed, hides on hearing s1 but eats on hearing s2
- (2,0,2,2) sounds s2 when fed, hides on hearing any sound

- (2,1,0,0) sounds s2 when fed, s1 when hurt, behaves randomly on hearing any sound.
- (2,1,0,1) sounds s2 when fed, s1 when hurt, behaves randomly on hearing s1, opens its mouth on hearing s2
- (2,1,0,2) sounds s2 when fed, s1 when hurt, behaves randomly on hearing s1, hides on hearing s2
- (2,1,1,0) sounds s2 when fed, s1 when hurt, opens its mouth on hearing s1, behaves randomly on hearing s2
- (2,1,1,1) sounds s2 when fed, s1 when hurt, opens its mouth on hearing either s1 or s2
- (2,1,1,2) sounds s2 when fed, s1 when hurt, opens its mouth on hearing s1 and hides on hearing s2
- (2,1,2,0) sounds s2 when fed, s1 when hurt, hides on hearing s1 and behaves randomly on hearing s2
- (2,1,2,1) sounds s2 when fed, s1 when hurt, hides on hearing s1 but eats on hearing s2
- (2,1,2,2) sounds s2 when fed, s1 when hurt, hides on hearing any sound
- (2,2,0,0) sounds s2 when fed or hurt, behaves randomly on hearing any sound.
- (2,2,0,1) sounds s2 when fed or hurt, behaves randomly on hearing s1, opens its mouth on hearing s2
- (2,2,0,2) sounds s2 when fed or hurt, behaves randomly on hearing s1, hides on hearing s2
- (2,2,1,0) sounds s2 when fed or hurt, opens its mouth on hearing s1, behaves randomly on hearing s2
- (2,2,1,1) sounds s2 when fed or hurt, opens its mouth on hearing either s1 or s2
- (2,2,1,2) sounds s2 when fed or hurt, opens its mouth on hearing s1 and hides on hearing s2
- (2,2,2,0) sounds s2 when fed or hurt, hides on hearing s1 and behaves randomly on hearing s2
- (2,2,2,1) sounds s2 when fed or hurt, hides on hearing s1 but eats on hearing s2
- (2,2,2,2) sounds s2 when fed or hurt, hides on hearing any sound

Among these 81 strategies there are only two that qualify as ‘perfect communicators’: strategies $\langle 1,2,1,2 \rangle$ and $\langle 2,1,2,1 \rangle$. Strategy $\langle 1,2,1,2 \rangle$ makes sound s1 when fed and sound s2 when hurt, responding to s1 symmetrically by opening its mouth and to s2 by hiding. Strategy $\langle 2,1,2,1 \rangle$ follows the same pattern with the role of sounds s1 and s2 interchanged. Among the 79 strategies that do not qualify as perfect communicators are ‘free riders’, which exploit communication received without responding with symmetrical communication in return. A ‘free rider’ may react to a particular sound by opening its mouth, for example—thereby benefiting from a neighbour’s signal—but when fed will not send a corresponding signal from which its neighbour might benefit in return.

In these runs we set the gain for ‘feeding’ at 1 point and the loss for being ‘hurt’ at a negative 1 point. Cells accumulate a total (perhaps negative) over a run of 100 rounds—a ‘generation’. At the end of each generation, the crucial question is whether a cell has a more successful immediate neighbour. If any neighbour has a higher score, the cell’s strategy is replaced wholesale with the strategy of that neighbour with the highest score—a mechanism of strategy change by straight

imitation.⁹ In the case of two neighbours with tied higher scores, one is chosen at random.

We start with a randomization of all 81 strategies across our 4096 individuals. With an eye to the different dynamics of feeding and predation noted above, we use 50 food items and 100 predators.¹⁰ Food items and predators wander in a random walk as outlined, and our individuals periodically adopt the strategy of their most successful neighbours.




Figure 4. Development of a typical array of imitators, with perfect communicators shown in pure black and pure white. Generations 1, 10, 25, 50, 100 and 200 shown. For an animation of the entire development, please go to <http://129.49.17.140/mmh/mmh.htm>

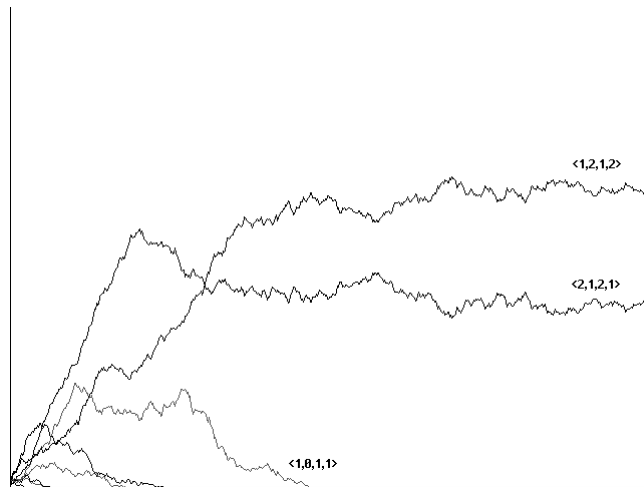


Figure 5. Conquest by ‘perfect communicators’ with 100 predators and 50 food items; 600 generations shown.

Can simple patterns of communication emerge in such an environment? As it turns out, emergence of communication under such circumstances is quite nearly inevitable. The development of a typical array is shown in figure 4. In such an environment, communities of perfect communicators slowly develop and spread. Ultimately, the array will be occupied by our two species of perfect communicators alone, divided by a fluctuating border at which individuals occasionally switch from one strategy to the other. A running animation of the full development can be seen at <http://129.49.17.140/mmh/mmh.htm>

In figure 5, we graph the development of a typical array in terms of percentages of the population over time. Quick dominance by our two perfect communicators is clear.

6. Breeding communicators with localized genetic algorithms

In the second set of models, we replace strategy change by simple imitation with strategy change by localized genetic algorithm. Rather than starting with a randomization across our full range of strategies, we begin with a small number of ‘Adam and Eve’ strategies, chosen at random except that we quite deliberately *weed out* any strategies toward which we anticipate convergence. We deliberately avoid any strategies that are in any way close to perfect communicators, for example. Sets of initial Adams and Eves are also chosen so that each possible variable (0, 1 or 2) will be represented in each variable position somewhere within the set. Although strategies can emerge or re-emerge through genetic algorithm despite not being represented in earlier generations, it is only possible to end up with a strategy using a 2 in the final position, for example, if at least one of the strategies we begin with has a 2 in that position. Without all values at least initially possible in all positions we would be unable to explore major portions of the full strategy space.

From an initial randomization of Adams and Eves, strategy change proceeds as follows. Here again we work with 50 food items and 100 predators. Each cell has




Figure 6. General strategy of genetic recombination of strategies.

accumulated a score from feeding and predation over the course of 100 rounds, and the crucial question for strategy updating is whether a cell has any neighbour with a higher score. If it does, it ‘breeds’: its strategy is changed to one that is a genetic cross between its previous strategy and that of its most successful neighbour.

Technically, this is done by choosing two points at random in the series of values that characterize the two strategies at issue: between the first and second in our series of four values, for example, and between the third and fourth. This gives us an ‘inside’ set of values and an ‘outside’ set. We then use the programming equivalent of flipping a coin. If the coin comes up heads, it is the ‘inside’ values of the original strategy that remain and the ‘outside’ values that are changed to those of its most successful neighbour. If the coin comes up tails, it is the ‘outside’ values that remain the same and the ‘inside’ values that are changed (figure 6).

Crossover is often implemented in genetic algorithms using only one random point rather than two (see, for example, Holland 1992). One chooses a random spot at which to cut both strategies A and B, creating a new strategy that combines that section of A to the left of the random cut with that section of B to the right, or vice versa. One consequence of using one-point crossover, however, is that the ends of a strategy are treated differently from the middle. Consider, for example, the digits at the very ends of strategy A. With a single random cut, either of these digits may end up being the only element from A dropped into an otherwise different strategy. This does not hold for digits between the ends, which on a single cut will always reappear with at least one original neighbour. In choosing two points for crossover instead, including the possibility of a cut ‘after the last digit’, we address both this positional bias or ‘end effect’ and several other difficulties noted in the literature (Eshelman *et al.* 1989, Mitchell 1996).

Often genetic algorithms are applied globally, cross-breeding a select few of the ‘most successful’ strategies across an entire population. It should be noted that our genetic algorithm is here applied purely locally; it is the most successful immediate neighbour with which a cell’s strategy is crossed. Here all reproduction, like all interaction and all fitness measurement, proceeds purely locally.

In this model, a successful strategy can be expected to contribute only part of its genetic code to those strategies that replace its less successful neighbours. Only by progressive approximation can a successful strategy be literally duplicated or fully imitated, then, and only by being as successful as any strategy in its immediate neighbourhood, or by cross-breeding with a genetically identical neighbour, can the strategy of any given cell in the array remain unchanged. In the course of genetic

recombination, we can expect a number of different strategies to be reproduced, and these might or might not prove more successful than their neighbours. The result is a wide sampling across possible strategies through local action, despite the fact that we start with only a manageable handful of Adams and Eves.¹¹

Can ‘perfect communicators’ emerge and spread through the mechanisms of localized genetic algorithm? Yes. The emergence of communities of perfect communicators in a typical run is shown in figure 7, using the same environment and the same 81 strategies as in the previous section. We begin with a randomization of a small number of Adams and Eves across the array—in this case, only six. Territory is progressively occupied by the two perfect communicators that develop, though their border continues to be marked by a range of experimental hybrids.

Figure 8 shows proportions of the population for three different runs, with three different sets of six Adams and Eves over a varying number of generations. All show the progressive emergence of our two ‘perfect communicators’ and the eventual triumph of one. Here we have applied a new form of strategy change to the same environment and the same range of strategies used in our imitative model in section 4. One of the benefits of using localized genetic algorithms, however, is that they allow exploration beyond the limitations of that earlier model, expanding our sample space to a larger range of more complex strategies. With only 81 strategies, we still have a number small enough for significant initial representation in a randomized 64×64 array. If we introduce more complex strategies, however, we will have to deal with a much larger sample space of variations. The larger the number of strategies, the less chance that each will be significantly represented in an initial array, particularly where we are dealing with strategies dependent on coordination with neighbours. In a model using strategy change by imitation, moreover, no strategy can appear *de novo*. Any strategy not initially represented cannot later appear, and it is impossible for any strategy eliminated from the population at any point to reappear later in a changed environment.

Both of these obstacles are addressed by starting from a small number of more complex Adams and Eves and allowing successful strategies to emerge through localized genetic recombination. The promise of localized genetic algorithms is thus a model that allows us to explore a larger sample space of more complicated strategies without swamping an initial array with a randomization of all strategies at once and while still modeling strategy change in terms of purely local action. One might think of the model as shifting from one in which samples are introduced by an initial spatial randomization and functionally sorted over time to one in which both introduction and sorting occur temporally as well as spatially.

Let us thus consider a range of more complex behavioural strategies. One important addition is to move beyond the ‘open mouth, hide or random selection’ restraints of the simpler model by adding a genuinely ‘neutral’ state in which an individual neither opens its mouth nor hides. When coasting in ‘neutral’, an individual cannot gain points by feeding but is still vulnerable to losing points from predation. We can also expand the simpler model by including a behavioural specification for the case in which an individual is neither fed nor hurt, and another for the case in which it hears no sound from immediate neighbours.

These more complex strategies can be coded as six-tuples $\langle \phi, f, h, s1, s2, s\phi \rangle$, with variables indicating what an individual does when it is neither fed nor hurt ϕ (make sound s1, make sound s2 or make no sound), what it does when fed f or hurt h (the same three options), and what it does when it hears sound s1, sound s2, or no sound




Figure 7. Typical evolution of a spatialized array by genetic algorithm, starting with a random distribution of six Adams and Eves. Initial strategies quickly proliferate, eventually converging to a small number of perfect communicators with continued genetic recombination at their borders.




Figure 8a. Genetic algorithm evolution to a perfect communicator $\langle 1,2,1,2 \rangle$ over 1724 generations.




Figure 8b. Genetic algorithm evolution to a perfect communicator $\langle 2,1,2,1 \rangle$ over 1998 generations.




Figure 8c. Genetic algorithm evolution to two perfect communicators, with the eventual triumph of one; 5500 generations shown.

ϕ (open its mouth, hide or ‘coast’ in neutral). Those variations expand our sample space significantly, from 81 possible strategies to 729.

One final wrinkle: following earlier work on co-operation, itself based on that of Nowak and Sigmund, the world we model here is a stochastically imperfect one (Nowak and Sigmund 1990, 1992, Grim 1995, 1996, Grim *et al.* 2000). In only 90% of cases do individuals respond to sounds precisely as their strategy specifies. In a random 5% of cases, an individual with a strategy that specifies an open mouth will hide instead; in 5% of cases, it will coast in neutral. Other responses to sound reflect similar stochastic noise.

Here we use an initial randomization of 18 Adam and Eves, chosen to include each of our three variables in each of our six positions somewhere in the set and sorted so as to avoid initial communicators. We maintain the same ratio of food and predators, but for speed of computation raise the numbers to 75 foods and 150 predators. Here for the first time we also assign an energy tax of 0.02 points for opening one’s mouth or for hiding. We exact a tax of 0.01 points for making any sound, on the theory that sound-making requires energy as well. Each cell has accumulated a score from feeding, predation and energy taxes over the course of 100 rounds. If there is an immediate neighbour with a higher score, the cell’s strategy is replaced with a genetic cross between its previous strategy and that of its most successful neighbour. The technical aspects of recombination are as before; applied to six variables, our genetic algorithm is illustrated in figure 9.

Of our 729 strategies, there are still only two that qualify as perfect communicators: strategies $\langle 0,2,1,2,1,0 \rangle$ and $\langle 0,1,2,1,2,0 \rangle$. The first of these makes sound s2 on feeding and sound s1 on being hurt, responding to sound s1 by hiding and sound s2 by feeding. The second makes sound s1 on feeding and sound s2 on being hurt, responding to sound s2 by hiding and to sound s1 by feeding. Both ‘coast’ in neutral on hearing no sound, and neither makes any sound when not being either hurt or fed.

Does communication emerge from this larger sample space by the mechanism of localized genetic algorithm as well? Yes. Figure 10a shows evolution to a single perfect communicator in 2677 generations, starting from a randomization of 18 Adams and Eves. Figure 10b shows evolution to a perfect communicator in 995 generations, starting from a sample of 18 different Adam and Eves.




Figure 9. Genetic recombination applied to strategies of six variables.

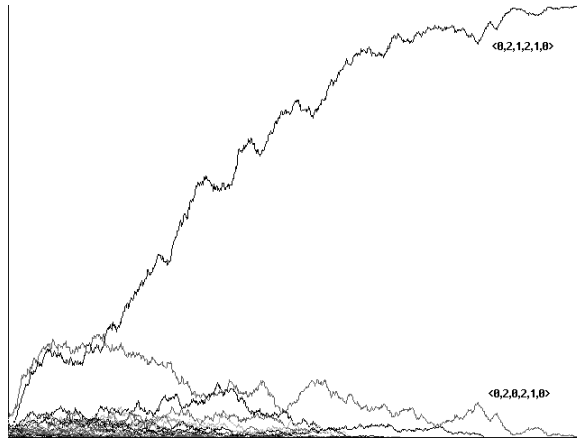


Figure 10a. Evolution to a perfect communicator in 2677 generations, with 75 foods, 150 predators, 0.02 tax for action, 0.01 tax for sounding.

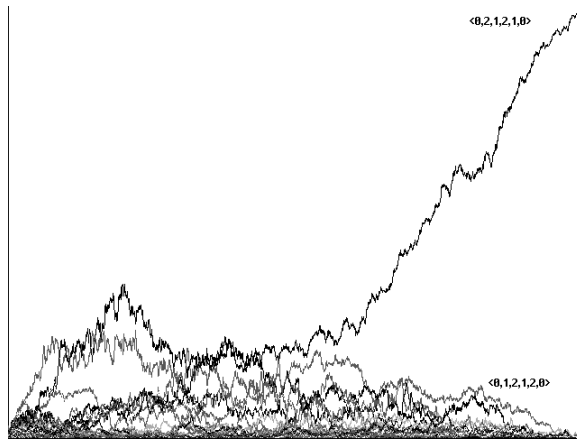


Figure 10b. Evolution to a perfect communicator in 995 generations with different initial Adams and Eves, 75 foods, 150 predators, 0.02 tax for action, 0.01 tax for sounding.

7. Learning to communicate in arrays of perceptrons

We have also investigated a third species of *compuzoans* that have simple nervous systems composed of feed-forward neural nets. In the simpler of two subspecies, these neural nets have no hidden layers: the behaviour of each individual is generated by a two-layer perceptron of the form shown in figure 11.

In this form, each of the 4096 individuals in our 64×64 array is coded using 12 weights and biases. In order to keep computation manageable in an array this large, we use discrete weights each of which carries one of the following values: -3.5 , -2.5 , -1.5 , -0.5 , $+0.5$, $+1.5$, $+2.5$ and $+3.5$. Discrete weights do just fine with the simple delta rule we will use here for training.¹²

The basic neural component of our perceptrons is shown in figure 12. The structure of this 'quadrant' is repeated four times in the complete structure of




Figure 11. The perceptron architecture.




Figure 12. A single quadrant of the perceptron.

figure 11, with two quadrants sharing inputs in each of two ‘lobes’. This ‘two lobe’ configuration for communication has been re-invented or re-discovered repeatedly in the history of the literature. Since de Saussure (1916), many have noted the intrinsic distinction between the kinds of action represented here by (1) making sounds and (2) mouth-opening or hiding in response to sounds heard, and it seems natural to embody that distinction in the neural architecture of the individuals modelled.¹³

We use a bipolar coding for inputs, so that ‘hear sound s1’, for example, takes a value of +1 if the individual hears sound s1 from an immediate neighbour on the previous round, and takes a value of -1 if it does not. Each input is multiplied by the weight shown on arrows from it, and the weighted inputs are then summed at the output node. To that is added the value (positive or negative) of the bias, which might alternatively be thought of as a third weight with a constant input of 1. If the total at the output node is greater than 0, we take our output to be +1 and the individual opens its mouth, for example; if the weighted total is less than or equal to 0, we take our output to be -1 and the individual keeps its mouth closed. Here, as before, an element of noise is built in: in a random 5% of cases each individual will open its mouth regardless of weights and inputs. On the other side of the lobe, individuals also hide in a random 5% of cases.

There are four possible sets of inputs for each quadrant: (-1,-1), (-1,+1), (+1,-1) and (+1,+1). In principle, the output in each case might be either -1 or +1, giving us the standard 16 Boolean functions. But not all net architectures can represent all 16 Booleans, and it is well known that perceptrons are limited in this regard (Minsky and Papert 1969). Each quadrant of our perceptrons can in fact handle only 14 of the 16 Booleans. A quadrant is capable of giving an output for ‘open mouth’, for example:

- never, regardless of sound input
- only when both sounds are heard
- when only sound s2 is heard
- when sound s2 or both sounds are heard
- when only sound s1 is heard

when sound s1 or both sounds are heard
 when either sound or both are heard
 only when neither sound is heard
 precisely when sound s1 is not heard
 when sound s1 is not heard or both sounds are heard
 precisely when sound s2 is not heard
 when sound s2 is not heard or both sounds are heard
 whenever it is not the case that both sounds are heard
 always, regardless of input.

The two Booleans that cannot be captured within such a structure are exclusive 'or' (xor) and the biconditional. Such a net has no way of giving an output just in case:

Either sound s1 is heard or sound s2 is heard, but not both
 Either both are heard or neither is heard.

Even with those limitations, however, our perceptrons offer a sample space of strategies far larger than the 729 considered in terms of genetic algorithms above. With a total of 12 discrete weights we have a sample space of 68 719 476 736 numerically distinct nets. Not every numerical difference makes a behavioural difference, of course: the same behaviour may be produced by webs with different weight and bias assignments and indeed by nets with major differences in weight ratio and balance at different points. With 14 of 16 Booleans represented in each quadrant of our nets, we are nonetheless dealing with a range of 38 416 possible behavioural strategies.

We can code these behavioural strategies in terms of output for different pairs of inputs. The possible inputs at 'hear sound s1' and 'hear sound s2' for the left lobe of our structure are $(-1, -1)$, $(-1, +1)$, $(+1, -1)$, and $(+1, +1)$. Outputs for a given strategy will be pairs representing the output values for 'open mouth' and 'hide' for each of these pairs of inputs. We might thus encode the left-lobe behaviour of a given strategy as a series of 8 binary digits. The string 00 00 00 11, for example, represents a behaviour that outputs an open mouth or a hide only if both sounds are heard, and then outputs both. The string 00 01 01 01 characterizes a cell that never opens its mouth, but hides if it hears either sound or both. We can use a similar pattern of behavioural coding for the right lobe, and thus encode the entire behaviour of a net in terms of 16 binary digits. We will represent the behaviour for a complete net using a single separation between representations for the two lobes, like this: 00110011 11001100.

Of the 38 416 behavioural strategies in our sample space, there are still only two that qualify as 'perfect communicators'. Pattern 00011011 00011011 represents an individual that hides whenever it hears sound s2, opens its mouth whenever it hears sound s1, makes sound s2 whenever it is hurt and makes sound s1 whenever it is fed. The 'whenever' indicates that it will both hide and open its mouth when it hears both sounds and will make both sounds when both hurt and fed. The pattern 00100111 00100111 represents an individual with a symmetrical behaviour in which only the sound correlations are changed: it reacts to sound s2 by eating and responds to being fed by making sound s2, reacts to sound s1 by hiding and responds to being hurt by making sound s1.

We initially populate our array with neural nets carrying 12 random weights, randomizing over our 68 billion numerical strategies. A total of 100 food sources and

200 predators drift in a random walk across the array, without being consumed or satiated at any point.¹⁴ Whenever a cell has its mouth open and a food source lands on it, it feeds and gains one point. Whenever a predator lands on a cell that is not hiding, that cell is ‘hurt’ and loses one point. In our neural net runs, we raised the energy cost for opening one’s mouth or hiding to 0.05 points, with energy expenditure for making any sound at 0.05 points as well.

Over the course of 100 rounds, our individuals total their points as before. Here again the question is whether a cell has a higher-scoring neighbour. If so, its strategy is changed through partial training on the behaviour of its highest-scoring neighbour.

For perceptrons we use the standard delta rule as our learning algorithm. For a set of four random inputs, the cell compares its outputs with those of its target neighbour. At any point at which the behaviour of the training cell differs from its target, we nudge each of the responsible weights and biases one unit positively or negatively. Within the limits of our value scale, use of bipolar values for target and input allow us to calculate this simply as $w_{\text{new}} = w_{\text{old}} + (\text{target} \times \text{input})$ and $\text{bias}_{\text{new}} = \text{bias}_{\text{old}} + \text{target}$.

Our training run consists of only four random sets of inputs, with no provision against duplication. Training will thus clearly be partial: only four sets of inputs are sampled, rather than the full 16 possible, and indeed the same set may be sampled repeatedly. The learning algorithm is applied using each set of inputs only once, moreover, leaving no guarantee that each weight will be shifted enough to make the behavioural difference that would be observable in a complete training. Partial training is quite deliberately built into the model in order to allow numerical combinations and behavioural strategies to emerge from training that might not previously have existed in either teacher or learner, thereby allowing a wider exploration of the sample space of possible strategies.

Suppose we start with an array of perceptrons with randomized weights and biases. Will communities of cells learn to communicate? For these simple neural nets, as for strategy change by localized genetic algorithm and by simple imitation, the answer is ‘yes’. Figure 13 shows the emergence of communication in a typical array over 300 generations. Figure 14 shows the same emergence in terms of proportions of the population. A full animation of the development can be seen at <http://129.49.17.140/mmh/mmh.htm>

8. Backpropagation in more complex neural nets

It has long been known that neural nets of just two layers are incapable of representing all Boolean functions: we’ve noted the exclusive ‘or’ and biconditional as exceptions. This crucial limitation dulls the impact of the otherwise remarkable perceptron learning convergence theorem: that the simple delta rule is adequate to train any perceptron, in a finite number of steps, to any function it can represent (Rosenblatt 1959, 1962, Minsky and Papert 1969, Fausett 1994). In the 1970s, that limitation posed a significant stumbling block to the further development of neural nets. It was known even then that the addition of intermediate layers to perceptrons would result in multiple layer neural nets that could model the full spectrum of Boolean functions, but the simple delta rule was known to be inadequate for training multiple-layer nets.




Figure 13. Emergence of two dialects of perfect communicators, shown in solid black and white, in a randomized array of perceptrons with partial training on successful neighbours. All other behavioural strategies coded using shades of grey for backgrounds and central dots. Centuries 1, 10, 50, 100, 200 and 300 shown. A full animation of the development can be seen at <http://129.49.17.140/mmh.mmh.htm>




Figure 14. Learning to communicate in a randomized array of perceptrons with partial training on successful neighbours in a sample space of 38 416 behavioural strategies. Percentages of population for different strategies graphed over 300 generations.




Figure 15. The quadrant structure of our backpropagation nets.

With the use of continuous and differentiable activation functions, however, multiple-layer neural nets can be trained by backpropagation of errors using a generalized delta function. This discovery signalled the re-emergence of active research on neural nets in the 1980s (McClelland and Rumelhart 1988). Here again there is a convergence theorem: it can be shown that any continuous mapping can be approximated to any arbitrary accuracy by using backpropagation on a net with some number of neurons in a single hidden layer (White 1990, Fausett 1994).

The most complicated neural nets we have to offer here use backpropagation in order to train to the full range of Boolean functions of their inputs. Each of our nets is again divided into two lobes, with inputs of two different sounds on the left side and outputs of mouth-opening or hiding, inputs of ‘fed’ and ‘hurt’ on the right side with outputs of two different sounds made. Each lobe is again divided into two quadrants, but our quadrants are now structured as neural nets with a single hidden node (figure 15). In order to apply backpropagation training, weights and biases are no longer discrete, taking instead any real values in the $[-3.5, +3.5]$ interval.




Figure 16. Activation function.

The feedforward neural nets most commonly illustrated in the literature have hierarchically uniform levels—all inputs feed to a hidden layer, for example, and only the hidden layer feeds to output. For reasons of economy in the number of nodes and weights to be carried in memory over a large array, the design of our nets is not hierarchically uniform. As is clear from figure 15, inputs feed through weights w_1 and w_4 directly to the output node as well as through weights w_2 and w_3 to a hidden node. The output node receives signals both from inputs directly and through weight w_5 from the hidden node.

At both the hidden node and the output node we use a sigmoid activation function

$$f(x) = \frac{2}{1 + \exp(-x)} - 1 \quad \text{equivalent to} \quad \frac{1 - \exp(-x)}{1 + \exp(-x)}$$

graphed in figure 16. In our sample quadrant, bipolar inputs -1 or $+1$ from ‘hear sound s1’ and ‘hear sound s2’ are first multiplied by weights w_2 and w_3 , initially set between -3.5 and $+3.5$. At the hidden node, those products are added to a constant bias 2 set initially in the same range. The total is then treated as input to the activation function, generating an output somewhere between -1 and $+1$ that is sent down the line to the output node.

The signal from the hidden node is multiplied by weight w_5 , which is added at the output node to the product of the initial inputs multiplied by weights w_1 and w_4 . Bias 1 is also added to the sum. Here again all initial weights and biases are set between -3.5 and $+3.5$. This sum is finally passed through our activation function once again, with an output >0 treated as a signal for an open mouth, for example, and an output ≤ 0 as a signal for a closed mouth. With different weight settings, this simple multi-layered structure is adequate to represent all 16 Boolean functions.

We employ a form of backpropagation appropriate to nets with this structure, using the derivative of our activation function $f'(x) = [1 + f(x)][1 - f(x)]/2$.¹⁵ Training can be illustrated in terms of a single quadrant.

For a particular pair of inputs, we will at training have a particular target t : the output (-1 or $+1$) toward which our net is to be trained for those inputs. We operate our net feedforward, as outlined above, to obtain a final output o of -1 or $+1$. We calculate an output error information term $\delta_o = (t - o)$.

δ_o is applied directly to calculate changes in weights w_1 and w_4 on lines feeding straight from inputs. In each case, the weight change Δ is a learning rate lr times δ_o times the input signal that was fed down that line. Our learning rate is set at a constant 0.02 throughout.

$$\Delta w1 = lr \times \delta_o \times \text{input}(\text{sound s1})$$

$$\Delta w4 = lr \times \delta_o \times \text{input}(\text{sound s2}).$$

A bias can be thought of as a weight with a constant input of 1, and the change for bias 1 is calculated in the same way:

$$\Delta b1 = lr \times \delta_o.$$

The weight change for w5, from hidden node to output, follows a similar pattern, though $\Delta w5$ is calculated in terms of the signal which was sent down the line from hidden to output node in the feedforward operation of the net:

$$\Delta w5 = lr \times \delta_o \times \text{output}(h).$$

Weight changes for w2 and w3 are calculated by backpropagation. Here we first calculate a new error information term $\delta_h = w5 \times \delta_o \times f'(\text{inp}_h)$, where $f'(\text{inp}_h)$ is the derivative of our activation function applied to the sum of weighted inputs at our hidden node. Changes in weights w2 and w3 are then calculated in terms of δ_h and our initial inputs:

$$\Delta w2 = lr \times \delta_h \times \text{input}(\text{sound s1})$$

$$\Delta w3 = lr \times \delta_h \times \text{input}(\text{sound s2}).$$

The change in bias 2 will be simply $lr \times \delta_h$.

Once all weight and bias changes are calculated, they are simultaneously put into play: $w_{\text{new}} = w_{\text{old}} + \Delta w$ and $\text{bias}_{\text{new}} = \text{bias}_{\text{old}} + \Delta b1$ for each of our weights and biases.

We wanted to assure ourselves that our net structure was satisfactorily trainable to the full range of Booleans. The convergence theorem for standard backpropagation on multiple-layered and hierarchically uniform neural nets shows that a neural net with a sufficient number of nodes in a hidden layer can be trained to approximate any continuous function to any arbitrary accuracy (White 1990, Fausett 1994). Our nets are not hierarchically uniform, however, they employ only a single hidden node, and our training is to the Booleans rather than a continuous function. Is the training algorithm outlined here adequate to the task?

With minor qualification, the answer is 'yes'. We ran sets of 4000 initial random sets of weights in the interval between -3.5 and $+3.5$ for a quadrant of our net. Training for each set of weights was to each of the 16 Boolean functions, giving 64000 training tests. Trainings were measured in terms of 'epochs', sets of all possible input configurations in a randomized order. Our results showed successful training to require an average of 16 epochs, though in a set of 64000 training tests there were on average approximately six tests, or 0.01%, in which a particular weight set would not train to a particular Boolean in less than 3000 epochs.¹⁶ As those familiar with practical application of neural nets are aware, some weight sets simply 'do not train well'. The algorithm outlined did prove adequate for training in 99.99% of cases involving random initial weight sets and arbitrary Booleans.

For the sake of simplicity we have outlined the basic structure of our nets and our training algorithm above in terms of an isolated quadrant. Our nets as a whole are four times as complicated, of course, with two lobes of two quadrants each (figure 17). Each of these more complex nets employs a total of 20 weights, plus eight biases, requiring a total of 28 variable specifications for each net at a given time. As noted, backpropagation requires a continuous and differentiable activation function,




Figure 17. The full architecture of our neural nets.

and will not work properly with the discrete approximations used for perceptrons above. Here our individual nets are specified at any time in terms of 28 real values in the range between -3.5 and $+3.5$.

Had we used discrete weights at unit intervals, we would have been dealing with 8^{28} different numerical strategies. With real values for weights, the ceiling is lifted and we are dealing at least theoretically with a non-denumerably infinite number of numerical strategies.¹⁷ But of course not every difference in weights makes a difference in output behaviour. Each quadrant is capable of 16 different output patterns for a complete cycle of possible inputs. Our sample space is therefore one of 65 536 distinct behavioural strategies.

Here as before we can code our behavioural strategies in terms of binary strings. Pairs of digits such as 01 represent a lobe's output for a single pair of inputs. A coding 00 01 01 11 can thus be used to represent output over all possible pairs of inputs to a lobe: $(-1, -1)$, $(-1, +1)$, $(+1, -1)$, and $(+1, +1)$. A double set 01111000 00100011 serves to represent the behaviour of both lobes in a network as a whole.

Of the 65 536 distinct behavioural strategies that can thus be encoded, there are still precisely two that qualify as perfect communicators. The pattern 00011011 00011011 represents an individual that makes sound s1 whenever it is fed and reacts to sound s1 by opening its mouth, makes sound s2 whenever it is hurt and reacts to sound s2 by hiding. It will both hide and open its mouth when it hears both sounds and will make both sounds when both hurt and fed. Pattern 00100111 00100111 represents an individual with a symmetrical behaviour in which only the sound correlations are changed. This second individual makes sound s2 when it is fed and reacts to sound s2 by opening its mouth, makes sound s1 when hurt and reacts to sound s1 by hiding.

There are also variants on the pattern of perfect communicators that differ by a single digit in their encoding. The most significant, it turns out, are 'right-hand variants', which differ from one or the other of our perfect communicators in just one of the last two digits, applicable only on those rare occasions when an individual is both fed and hurt at the same time. Patterns 00011011 00011010 and 00011011 00011001 differ from a perfect communicator in that they each make just one sound rather than two in the case that they are simultaneously fed and hurt. Patterns 00100111 00100110 and 00100111 00100101 vary from our other perfect communicator in the same way. For our two 'perfect communicators' there are thus also four minimally distinct 'right-hand variants' out of our 65 000 behavioural strategies.

We initially randomize all 28 weights as real values between -3.5 and $+3.5$ for each of the 4096 neural nets embedded in our 64×64 array. A total of 100 food sources and 200 predators wander across the array as before. When a cell has a

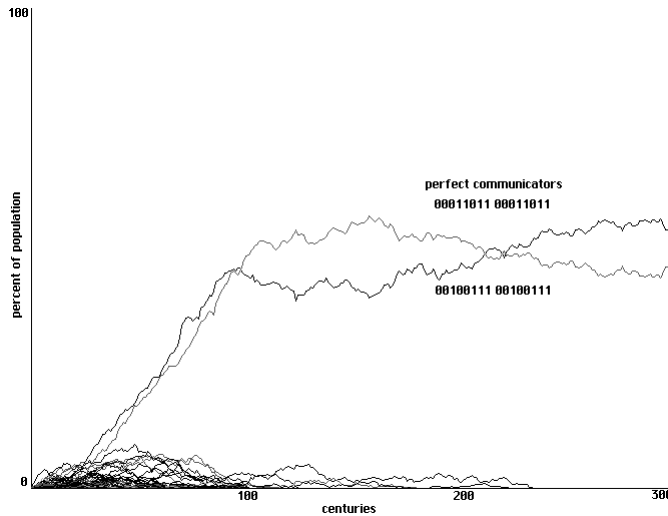


Figure 18. Emergence of perfect communication using backpropagation in an array of randomized neural nets. One training epoch used, 300 generations shown.

predator on it and is not hiding, it is ‘hurt’ and loses 1 point; when it has food on it and has its mouth open, it is ‘fed’ and gains 1 point. Each of these activities carries an energy expenditure of 0.05 points. In ‘neutral’ an individual is neither hiding nor has its mouth open, which saves it from energy expenditure but makes it incapable of capturing food and leaves it vulnerable to victimization by a predator. When hurt or fed, an individual makes one sound, both, or neither, where making a sound also carries an energy expenditure of 0.05 points. Here as before, it should be noted, ours is a stochastically imperfect world: individuals open their mouths and hide in a random 5% of all cases regardless of inputs and internal structure.

Over the course of 100 rounds (a generation), our individuals collect points from successful feeding and lose points as victims of predation. At the end of each generation the crucial question is whether any immediate neighbour has garnered more points. If so, our cells do a partial training on the behaviour of the highest-scoring neighbor. All this is as it was before; what differs is the structure of the nets themselves, the full sample space of behavioural strategies, and training by the backpropagation algorithm outlined above.

Full training by backpropagation standardly requires a large number of epochs, each consisting of the complete training set in a randomized order. Here, however, we use only a single training epoch. Training uses a complete but randomized set of possible inputs for each quadrant and takes the more successful neighbour’s behavioural output for each pair of inputs as target. This cannot be expected to be a full training in the sense that would make behaviours match; training using a single epoch will typically shift weights only to some degree in a direction that accords with the successful neighbour’s behaviour. Often the resulting behaviour will match neither the initial behaviour of the ‘trainee’ nor the full behaviour of its more successful neighbour.

Can communication emerge by backpropagation training in an environment this simple and from a strategy sample this large? Figure 18 shows a typical result with

one epoch of training over the course of 300 generations. Rather than plotting all 65 000 behavioural strategies, we have simplified the graph by showing only those strategies that at one point or another appeared among the top 20 in the array.

Here the two strategies that emerge from a sample space of 65 000 are our two perfect communicators. In some cases, starting from a randomized configuration, one or another ‘right-hand variant’ can play a significant role as well. The basic pattern we have tracked in earlier sections with wholesale imitation, genetic algorithms and perceptrons clearly appears here as well: random arrays of neural nets can learn to communicate (Grim *et al.* 2002).

9. Conclusion

We imagine sessile *compuzoans* as blind in an environment of wandering food sources and predators, but able to make sounds heard by their immediate neighbours. Our initial question was whether colonies of *compuzoans* could come to use those sounds as *signals*.

Across a range of modes of strategy change geared toward the behaviour of most successful neighbours—wholesale imitation, hybridization by genetic algorithm, perceptron training and backpropagation learning in neural nets—the answer appears to be a resounding ‘yes’. Even in an environment this simple one can see a pattern of basic signalling emerge.

What do models like these have to tell us about communication and ultimately about meaning? First and foremost, they offer an existence proof. Basic signalling *can* emerge in a simple spatialized environment of food gain and predator loss, and can emerge among egoistic agents in an economy of purely individualistic gains and losses. The result requires no assumption of symmetrical gains from each act of communication nor any hypothesis of co-operative tasks.

The results above also suggest that the phenomenon at issue is not picky about its methods: signalling behaviour can appear in simple spatialized environments with strategy change by wholesale imitation, by genetic hybridization, or by neural net learning. In *Jurassic Park*, Ian Malcolm insists that ‘Life will find a way.’ What these models suggest is that the same may be true of at least simple patterns of communication.

Genetic algorithms are often thought of as analogues for physical genetics, while the delta rule and backpropagation are thought of as models for learning. If thought of in these terms, the lesson seems to be that simple patterns of communication can emerge either by physical genetics or cultural learning. We are not convinced, however, that the formal mechanism of genetic algorithms need be thought of as applicable solely to physical genetics. Codes in recombination might be taken instead to represent cultural strategies (‘memes’) that are partially transmitted and combined. Nor are we convinced that the learning algorithms typical of neural nets must always be thought of as analogues to cultural learning. At this point it might be better to view application of the delta rule and backpropagation simply as alternative techniques for the exploration of a sample space of available strategies. What appears crucial to emergence of communication across our models is not the particular mode of strategy change but its direction, toward the strategies of locally successful neighbours—the general dynamics of spatialized strategy change in a spatialized environment.

... you think of the meaning as a thing of the same kind as the word, though also different from the word. Here the word, there the meaning. The money, and the cow that you can buy with it. (But contrast: money, and its use.) (Wittgenstein, *Philosophical Investigations*: 120).

As indicated in the introduction, we consider these models a contribution to the general tradition of theories of meaning as use. Evident across much of contemporary modeling for communication, in a range of disciplines, is a Lockean portrayal of meaning as a relation between a sound and some corresponding internal representation.¹⁸ That picture of meaning is much less plausible here, particularly in the case of the neural net models. In the neural net training employed above, strategy change proceeds by weight-shifting toward the behavioural strategy of a successful neighbour. When a community of communicators emerges from an array of randomized neural nets, it is convergence to a behavioural strategy that is crucial.

In this model there is no guarantee that the internal workings of behaviourally identical strategies in two individuals are themselves identical. There are in principle non-denumerably many neural configurations that may show the same behavioural strategy. In training to match a neighbouring 'perfect communicator', a neural net may not only fail to match the absolute values of its neighbour's weights, but may also differ significantly in its over-all structure of relative weight balances. What arises in a community is a pattern of coordinated behaviour, but in developing from an initially randomized array that co-ordinated behaviour need not be built on any uniform understructure in the nets themselves. If you open up the neural nets in an array that has converged on communicators, you will not be able to identify the same meanings inside. Moving from individual to individual, you will not find the same numbers or even the same ratios of numbers inside. Ours is a model in which meaning is evident in behavioural coordination but simply isn't in the head.

But what about this: is the call 'Slab!' ... a sentence or a word? (Wittgenstein, *Philosophical Investigations*: 19).

Ours is also a model in which meaning remains indeterminate in an important sense. If you ask what sound s1 means in a developed community of communicators, there is no reason to favour 'I'm eating' over 'Yum' or 'There's food here' over 'Open your mouth'. It would be wrong to think that all we need to do is look more closely in order to figure out precisely which of these is what is meant. Perhaps indeterminacy arises because the language in which we distinguish these alternative readings is so much more complex than the signals we wish to represent. Perhaps indeterminacy of some sort is an inherent feature of language in general; perhaps we can always, for any utterance, offer alternative readings of this sort.

A conclusion is an appropriate place for concessions. What we have attempted to offer here are models for simple forms of communication, suggesting simple forms of meaning. We have to admit, of course, that what we have to offer are *merely* models. In our uncritical moments, we think of what we are doing as making meaning happen. But what we have written are computer programs that generate blips on the screen. Our 'agents' are virtual or artificial, and for that reason alone cannot literally mean anything by anything. If these models prove useful, it will not be because they contain real communication or real meaning but because they successfully *model* important aspects of the real thing. Soberly speaking, what these models suggest is that a very simple dynamics can explain the emergence of basic forms of communication.

Even if successful, of course, the forms of signalling, communication and meaning that these models succeed in modelling are excruciatingly simple forms of signalling, communication and meaning. Although we have helped ourselves freely to intentional terminology, we have shied away from calling what is at issue a 'language'. Even if our models do capture something about simple meaning, it must be admitted that meaning is not all simple. It may not be a single phenomenon at all. Different kinds of theories of meaning may be required to address different aspects of meaning or different questions about it.

To what extent might models like these be extended, for example, to simple syntax and an infinite set of possible messages?¹⁹ That remains an open question for further work.

Notes

1. Strategy replacement by 'imitation' need not be thought of as in any way intentional, nor even as involving any perception of successful neighbours. If a cell has a more successful neighbour, its strategy is replaced with the full form of the strategy of its most successful neighbour (hence full 'imitation'). In the case of two 'most successful' higher-performing neighbours, one is chosen at random. This process of strategy change might alternatively be thought of as a form of reproduction by more successful strategies into the niches of their less successful neighbours.
2. Recent non-game-theoretic attempts to develop more adequate theories of meaning as use include Peacocke (1992) and Horwich (1998).
3. See, for example, Ludlow (1997).
4. Saunders and Pollock (1996). Werner and Dyer (1991) use a model in which blind 'males' and signalling 'females' are thought to find each other spatially, but random relocation of offspring results in an algorithm identical to a global breeding of those above a success threshold on a given task.
5. Ackley and Littman (1994) do use local communication and limit reproduction to those individuals in a 'quad' with the highest fitness rating. There is also a model complicated with a blizzard of further interacting factors, however, including reproductive 'festivals' and a peculiar wind-driven strategy diffusion.
6. Although we speak of 'sounds' throughout, what individuals send and receive might also be thought of in other terms: as chemical signals, for example.
7. In our initial studies we used the same number of food items and predators, and until the differences in dynamics became clear were puzzled by the resulting tilt toward communication regarding food in particular. In later studies, we explored (1) behaviours that signal an alarm when a predator is present, whether or not the individual is harmed, and (2) models in which twice as many predators as food sources are used. The latter is the environment emphasized here—see also Grim *et al.* (2001).
8. 'Free riders' of this sort play a major role in the puzzles quoted from Batali (1995) and Parisi (1997) in section 2.
9. As indicated in note 1, strategy replacement by 'imitation' need not be thought of as in any way intentional, nor even as involving any perception of successful neighbours. The process might alternatively be thought of as a form of reproduction by more successful strategies into the niches of their less successful neighbours.
10. At any point there might be fewer than 50 spaces occupied by food items or 100 occupied by predators, however, since it is possible for multiple food items or predators to occupy the same space simultaneously. In the unusual circumstance that both a food source and a predator happen to land on an individual simultaneously, points total: if its mouth is open, it will both gain points for feeding and lose points for being hurt, for example. It is only in that remote circumstance that an individual is treated as making both sounds at its disposal.

11. This is not to say that the sampling process of such an array is either equivalent to or as broad as some fully randomized sampling with replacement would be. This process of genetic change carves exploratory tracks through the sample space, and given particular histories there are areas that these tracks will not explore. Genetic algorithms, like many sampling methods—and like evolution itself—are subject to difficulties of ‘local maxima’, in which convergence is to that strategy most successful within a certain range of genetic variation, though strategies outside that range would outperform it if they could be introduced.
12. See Fausett (1994). Discrete values of this type also appear in Werner and Dyer (1991) and in Plagianakos and Vrahatis (1999).
13. de Saussure (1916). This maps precisely onto the distinction between ‘emissions’ and ‘actions’ in MacLennan (1991) and between ‘transmission behaviour’ and ‘reception behaviour’ in Oliphant and Batali (1997). These two functions are separated between two different sexes in Werner and Dyer (1991) and between two separate sets of connection weights in the neural nets of Cangelosi and Parisi (1998). Martin Nowak notes that an active matrix for signal-sending and a passive matrix for signal-reading can be treated as completely independent in Nowak *et al.* (1999) and in Nowak *et al.* (2000). One gain from distinguishing the two abilities in this way is that there is then no built-in presumption that individuals will treat signals as bi-directional in the sense of de Saussure (1916): there is no presupposition that a signal will be read in the same way that it is sent. If bi-directionality nonetheless emerges, as indeed it does in our communities of ‘communicators’, it will be as a consequence of learning in an environment rather than as a structural constraint assumed from the start (see also Oliphant and Batali 1997).
14. The observant reader will note that the absolute number of foods and predators has progressively increased in our models, though the proportion remains the same. This is motivated purely by computational time constraints.
15. We are indebted to Laurene Fausett for helpful correspondence regarding training algorithms for nets with the structure used here. Our simple net combines perceptron-like connections (along weights w_1 and w_4) with crucial use of a single hidden node; it will be noted that the training algorithm also combines a perceptron-like training for w_1 , w_4 and w_5 with full back-propagation to update w_2 and w_3 .
16. Those Booleans to which training was not possible were in all cases exclusive ‘or’ or the biconditional. We also explored non-standard forms of backpropagation that did prove adequate for training 100% of our initial weight sets to each of the 16 Booleans. Final results were very similar to those outlined below.
17. ‘At least theoretically’ because of course computer instantiation does not deal with true reals.
18. See, for example: MacLennan (1991); MacLennan and Burghardt (1994); Oliphant and Batali (1997); Wagner (2000); Levin (1995); Parisi (1997); Hutchins and Hazlehurst (1991, 1995); Livingstone and Fyfe (1999); Livingstone (2000); and Nowak *et al.* (1999, 2000).
19. As indicated in introduction, there are suggestive leads towards the question of syntax in the current work of Martin Nowak and his collaborators. See Nowak *et al.* (1999), Nowak and Krakauer (1999) and Nowak *et al.* (2000).

References

- Ackley, D., and Littman, M., 1994, Altruism in the evolution of communication. In R. A. Brooks and P. Maes (eds) *Artificial Life IV: Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems* (Cambridge, MA: MIT Press), pp. 40–48.
- Aristotle, c. 330 BC, *De Interpretatione*. trans. J. L. Ackrill, 1963, (Oxford: Clarendon Press).
- Augustine, c. 400 AD, *Confessions*. (1992) (Oxford: Clarendon Press; New York: Oxford University Press).
- Batali, John, 1995, Small signaling systems can evolve in the absence of benefit to the information sender. Unpublished paper available online: <http://cogsci.ucsd.edu/~batali/papers/small.html>

- Bennett, J., 1976, *Linguistic Behaviour* (Cambridge, MA: Cambridge University Press).
- Brandom, R., 1994, *Making It Explicit* (Cambridge, MA: Harvard University Press).
- Cangelosi, A., and Parisi, D., 1998, The emergence of a 'language' in an evolving population of neural networks. *Connection Science*, **10**: 83–97.
- de Saussure, F., 1916, *Cours de Linguistique Generale*, trans. R. Harris, 1983, as *Course in General Linguistics* (Lubbock, TX: Duckworth).
- Dyer, M., 1995, Toward synthesizing artificial networks that exhibit cooperative intelligent behaviour: some open issues in artificial life. In C. G. Langton (ed.) *Artificial Life: An Overview* (Cambridge, MA: MIT Press), pp. 111–134.
- Eshelman, L., Caruana, R., and Shaffer, J., 1989, Biases in the crossover landscape. In J. D. Shaffer (ed.) *Proceedings of the Third International Conference on Genetic Algorithms* (Los Altos, CA: Morgan Kaufman).
- Fausett, L., 1994, *Fundamentals of Neural Networks* (Upper Saddle River, NJ: Prentice Hall).
- Fodor, J. A., 1975, *The Language of Thought* (Cambridge, MA: Harvard University Press).
- Fodor, J. A., 1981, *Representations* (Cambridge, MA: MIT Press).
- Frege, G., 1918, The thought: a logical inquiry, trans. A. M. Quinton and M. Quinton, 1956, *Mind*, **65**: 289–311.
- Grim, P., 1995, Greater generosity in the spatialized prisoner's dilemma. *Journal of Theoretical Biology* **173**: 353–359.
- Grim, P., 1996, Spatialization and greater generosity in the stochastic prisoner's dilemma. *Biosystems* **37**: 3–17.
- Grim, P., Kokalis, T., and Kilb, N., 2001, Evolution of communication with a spatialized genetic algorithm. *Evolution of Communication*, **3** (2): 105–134.
- Grim, P., Kokalis, T., Tafti, A., and Kilb, N., 2000, Evolution of communication in perfect and imperfect worlds. *World Futures* **56**: 179–197.
- Grim, P., Mar, G., and St. Denis, P., 1998, *The Philosophical Computer: Exploratory Essays in Philosophical Computer Modeling* (Cambridge, MA: MIT Press).
- Grim, P., St. Denis, P., and Kokalis, T., 2002, Learning to communicate: the emergence of signaling in spatialized arrays of neural nets. *Adaptive Behavior* **10**: 45–70.
- Hobbes, T., 1651, *Leviathan*, Norton Critical Edition, 1997 (New York: W. W. Norton & Co.).
- Holland, J. H., 1992, *Adaptation in Natural and Artificial Systems* (Cambridge, MA: MIT Press).
- Horwich, P., 1998, *Meaning* (Oxford: Clarendon Press).
- Hutchins, E., and Hazlehurst, B., 1991, Learning in the cultural process. In C. G. Langton, C. Taylor, J. D. Farmer and S. Rasmussen (eds) *Artificial Life II, SFI Studies in the Sciences of Complexity, vol. X* (Redwood City, CA: Addison-Wesley), pp. 689–708.
- Hutchins, E., and Hazlehurst, B., 1995, How to invent a lexicon: the development of shared symbols in interaction. In N. Gilbert and R. Conte (eds) *Artificial Societies: The Computer Simulation of Social Life* (London: UCL Press), pp. 157–189.
- Levin, M., 1995, The evolution of understanding: a genetic algorithm model of the evolution of communication. *BioSystems*, **36**: 167–178.
- Lewis, D., 1969, *Convention: A Philosophical Study* (Cambridge, MA: Harvard University Press).
- Livingstone, D., 2000, Neutral evolution and linguistic diversity. *Computing and Information Systems Technical Reports*, 9, University of Paisley.
- Livingstone, D., and Fyfe, C., 1999, Diversity in learned communication. In K. Dautenhahn and C. Nehaniv (eds) *AISB '99 Symposium on Imitation in Animals and Artifacts* (London: Society for the Study of Artificial Intelligence and Simulation of Behavior, AISB Press), pp. 139–146.
- Locke, J., 1689, *Essay Concerning Human Understanding* (Oxford: Clarendon Press; New York: Oxford University Press, 1979).
- Ludlow, P. (ed.), 1997, *Readings in the Philosophy of Language* (Cambridge, MA: MIT Press).
- MacLennan, B., 1991, Synthetic ethology: an approach to the study of communication. In C. G. Langton, C. Taylor, J. D. Farmer, and S. Rasmussen, (eds) *Artificial Life II, SFI Studies in the Sciences of Complexity, vol. X* (Redwood City, CA: Addison-Wesley), pp. 631–655.
- MacLennan, B., and Burghardt, G., 1994, Synthetic ethology and the evolution of cooperative communication. *Adaptive Behavior*, **2**: 161–188.
- McClelland, J. L. and Rumelhart, D. E., 1988, *Explorations in Parallel Distributed Processing* (Cambridge, MA: MIT Press).
- Mill, J. S., 1884, *A System of Logic* (London: Longmans, Green and Co).
- Minsky, M., and Papert, S., 1969, *Perceptrons: An Introduction to Computational Geometry* (Cambridge, MA: MIT Press, expanded edition 1990).
- Mitchell, M., 1996, *An Introduction to Genetic Algorithms* (Cambridge, MA: MIT Press).
- Noble, J., and Cliff, D., 1996, On simulating the evolution of communication. In P. Maes, M. Mataric, J. Meyer, J. Pollack and S. Wilson (eds) *From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior* (Cambridge, MA: MIT Press), pp. 608–617.

- Nowak, M., and Krakauer, D., 1999, The evolution of language. *Proceedings of the National Academy of Sciences USA*, **96**: 8028–8033.
- Nowak, M., and Sigmund, K., 1990, The evolution of stochastic strategies in the prisoner's dilemma. *Acta Applicandae Mathematicae*, **20**: 247–265.
- Nowak, M., and Sigmund, K., 1992, Tit for tat in heterogeneous populations. *Nature*, **355**: 250–252.
- Nowak, M., Plotkin, J., and Krakauer, D., 1999, The evolutionary language game. *Journal of Theoretical Biology*, **200**: 147–162.
- Nowak, M., Plotkin, J., and Jansen, V., 2000, The evolution of syntactic communication. *Nature*, **404**: 495–498.
- Oliphant, M., and Batali, J., 1997, Learning and the emergence of coordinated communication. *Center for Research on Language Newsletter*, **11** (1), University of San Diego.
- Parisi, D., 1997, An artificial life approach to language. *Brain and Language*, **59**: 121–146.
- Peacocke, C., 1992, *A Study of Concepts* (Cambridge, MA.: MIT Press).
- Plagianakos, V. P., and Vrahatis, M. N., 1999, Training neural networks with 3-bit integer weights, in *Proceedings of the Genetic and Evolutionary Computation Conference (Gecco '99)*, vol. 1, (Los Altos, CA.; Morgan Kaufman), pp. 910–915.
- Rosenblatt, F., 1959, Two theorems of statistical separability in the perceptron, in *Mechanization of Thought Processes: Proceedings of a Symposium Held at the National Physical Laboratory, November 1958* (London: HM Stationery Office), pp. 421–456.
- Rosenblatt, F., 1962, *Principles of Neurodynamics* (New York: Spartan Press).
- Russell, B., 1921, *The Analysis of Mind* (London: G. Allen and Unwin Ltd).
- Russell, B., 1940, *An Inquiry into Meaning and Truth* (London: G. Allen and Unwin Ltd; New York: Macmillan).
- Saunders, G. M., and Pollack, J. B., 1996, The evolution of communication schemes over continuous channels. In P. Maes, M. Mataric, J. Meyer, J. Pollack and S. Wilson (eds) *From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior* (Cambridge, MA: MIT Press), pp. 580–589.
- Skyrms, B., 1996, *Evolution of the Social Contract* (Cambridge, MA: Cambridge University Press).
- Steels, L., 1996, Emergent adaptive lexicons. In P. Maes, M. Mataric, J. Meyer, J. Pollack and S. Wilson (eds) *From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior* (Cambridge, MA: MIT Press), pp. 562–567.
- Steels, L., 1998, Synthesizing the origins of language and meaning using co-evolution, self-organization and level formation. In J. R. Hurford, M. Studdert-Kennedy and C. Knight (eds) *Approaches to the Evolution of Language: Social and Cognitive Bases* (Cambridge: Cambridge University Press) pp. 384–404.
- Wagner, K., 2000, Cooperative strategies and the evolution of communication. *Artificial Life*, **6**: 149–179.
- Werner, G., and Dyer, M., 1991, Evolution of communication in artificial organisms. In C. G. Langton, C. Taylor, J. D. Farmer and S. Rasmussen (eds) *Artificial Life II, SFI Studies in the Sciences of Complexity, vol. X* (Redwood City, CA: Addison-Wesley), pp. 659–687.
- White, H., 1990, Connectionist nonparametric regression: multilayer feedforward networks can learn arbitrary mappings. *Neural Networks*, **5**: 535–549.
- Wittgenstein, L., 1953, *Philosophical Investigations* (New York: Macmillan).