Simulating Grice: Emergent Pragmatics in Spatialized Game Theory
Patrick Grim

In Anton Benz, Christian Ebert, Gerhard Jäger, and Robert van Rooij, eds., *Language, Games and Evolution: Trends in Current Research on Language and Game Theory*, Springer (2011)

I. Introduction

How do conventions of communication emerge? How do sounds or gestures take on a semantic meaning, and how do pragmatic conventions emerge regarding the passing of adequate, reliable, and relevant information?

My colleagues and I have attempted in earlier work to extend spatialized game theory to questions of semantics. Agent-based simulations indicate that simple signaling systems emerge fairly naturally on the basis of individual information maximization in environments of wandering food sources and predators. Simple signaling emerges by means of any of various forms of updating on the behavior of immediate neighbors: imitation, localized genetic algorithms, and partial training in neural nets.

Here the goal is to apply similar techniques to questions of pragmatics. The motivating idea is the same: the idea that important aspects of pragmatics, like important aspects of semantics, may fall out as a natural results of information maximization in informational networks. The attempt below is to simulate fundamental elements of the Gricean picture: in particular, to show within networks of very simple agents the emergence of behavior in accord with the Gricean maxims. What these simulations suggest is that important features of pragmatics, like important aspects of semantics, don't have to be added in a theory of informational networks. They come for free.[1]

Sections II and III outline some of the background of the current work: emergence of cooperation in spatialized game theory and the emergence of a simple semantics among agents in a simulated environment of wandering food sources and predators. Section IV applies the techniques developed there to the case of pragmatics.

The focus of the current work is the emergence of the Gricean maxims themselves. Communicative exploitation of those maxims in conversational implicature and inference is a further matter, and the simulations used here do not generally take us that far. In the case of scalars, however, Section V shows that coordinated behavior in the general region of Gricean implicature does appear even for agents and environments as simple as those modeled here.

Color illustrations and animations for this chapter can be found at www.pgrim.org/pragmatics.

II. Cooperation in Spatialized Game Theory

In classical game theory, we deal with the interactions of small numbers of idealized agents. In evolutionary game theory we are dealing with whole populations, applying principles of replicator dynamics pioneered by R. A. Fisher in theoretical biology. The result is a model beautiful for its simplicity and mathematical ease. In both theoretical biology and game theory, however, replicator dynamics brings with it what might be termed *global* assumptions: assumptions of randomly mixed interactions within a large population, and assumptions of randomly distributed reproduction.

It is spatialized game theory that is the necessary background for the results below. The focus of spatialized game theory is the dynamics of *local* action within a network structure rather than global interaction across a population. Here agents do not interact randomly with all other agents in the population. They interact (or interact with increased probabilities) with specific agents—their neighbors in the network structure. Local rather than global interaction of this type is addressed in theoretical biology under the term 'viscosity'.

Figure 1 offers an example of spatialized game theory using the simple reactive strategies in the iterated Prisoner's Dilemma: the eight strategies that turn on merely an initial move (cooperate or defect) and what the other player did on the previous round. The eight Boolean combinations of play for initial move, reaction to cooperation, and reaction to defection include All-Defect, All-Cooperate, Cooperate-then-All-Defect, Tit for Tat, and Suspicious Tit for Tat (like Tit for Tat except starting with an initial defection).

In this illustration a randomization of strategies is embedded in the spatialized lattice familiar from cellular automata. All play is local—individual cells play only with immediate neighbors touching them on the sides or the diagonal. After a series of 200 games, players look to see if any neighbor has achieved a higher score. If so, they update to the strategy of their most successful neighbor.

The agents in such a model are of course as egoistic as any in classical game theory. The general question of interest is how conditional cooperation, in the form of a dominance by Tit-for-Tat, might develop.

Updating by imitation in this kind of network turns out to be a powerful mechanism for the emergence of cooperation. Figure 1 shows the evolution of a randomized array, with early success by All Defect (animation at www.pgrim.org/pragmatics). Because of local interaction and local imitation, however—because of 'viscosity'—communities of mutually reinforcing Tit for Tat then expand, eventually occupying the entire array (Grim 1995, 1996; Grim, Mar & St. Denis 1998).[2]

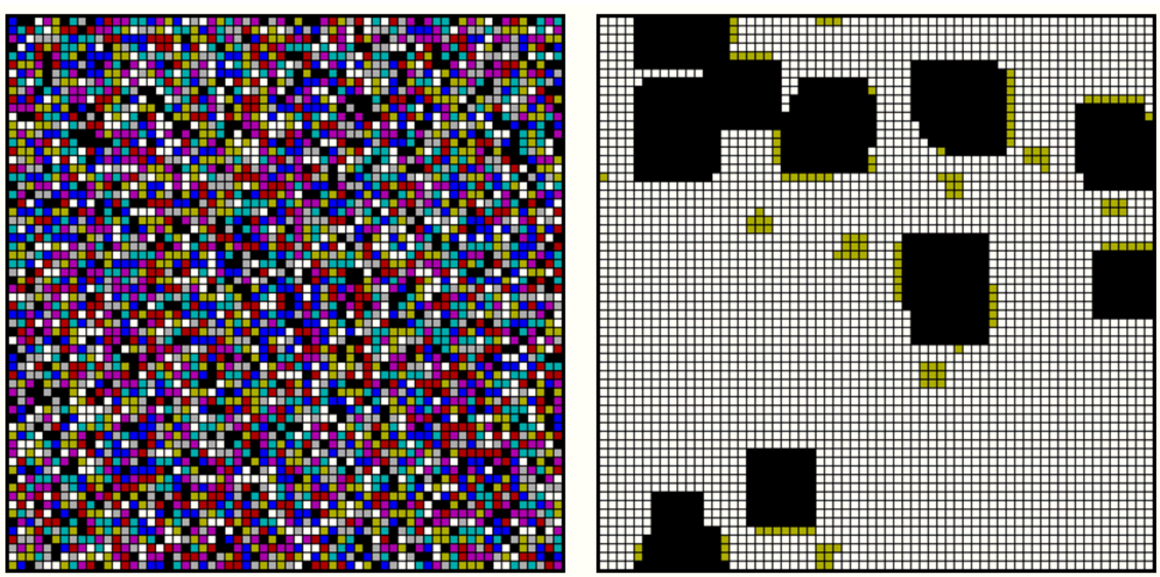

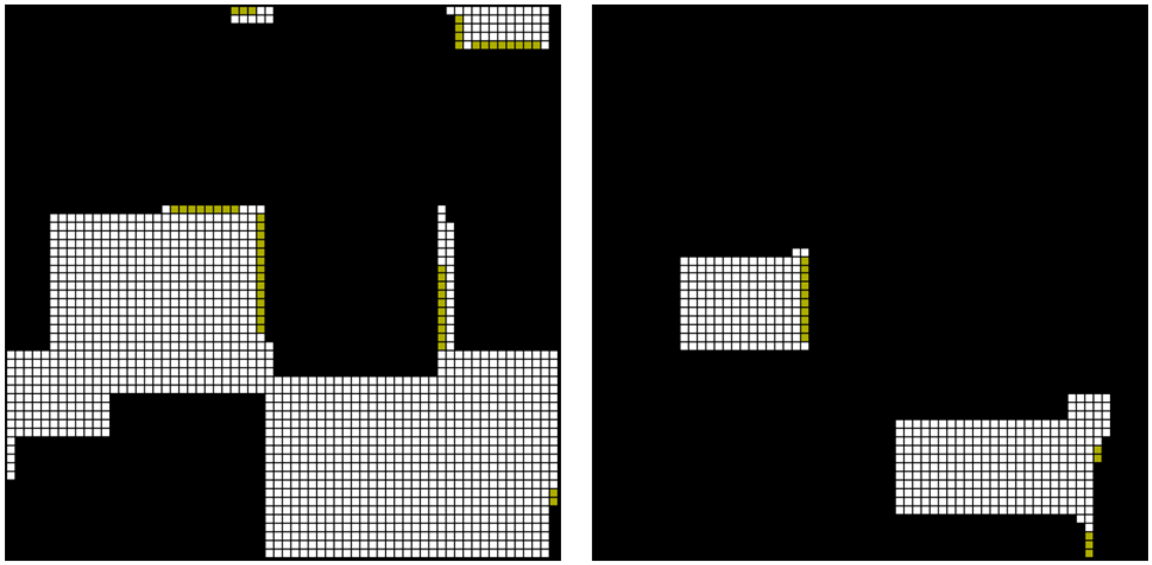

Figure 1. Conquest by TFT in a spatialized environment. Typical evolution of a randomized array of the 8 reactive Prisoner's Dilemma strategies, where cells copy the strategy of their most Successful neighbor. TFT in black, All-D in white. Generations 1, 5, 10 and 15 shown. A full animation of the evolution of the array can be seen at www.prim.org/pragmatics.

It is worthy of emphasis that spatialization is crucial to these results. If one of the mechanisms at issue is made global rather than local, the effect is broken: cooperation no longer proves dominant. Suppose, for example, that each agent updates its strategy not by imitating the strategy of its most successful neighbor but by imitating the strategy of the most successful cell in the entire array. With such an alteration, competitive play remains local but updating becomes global. In that case Tit for Tat is driven to distinction and we have a clear conquest by All-Defect (Figure 2).[3]

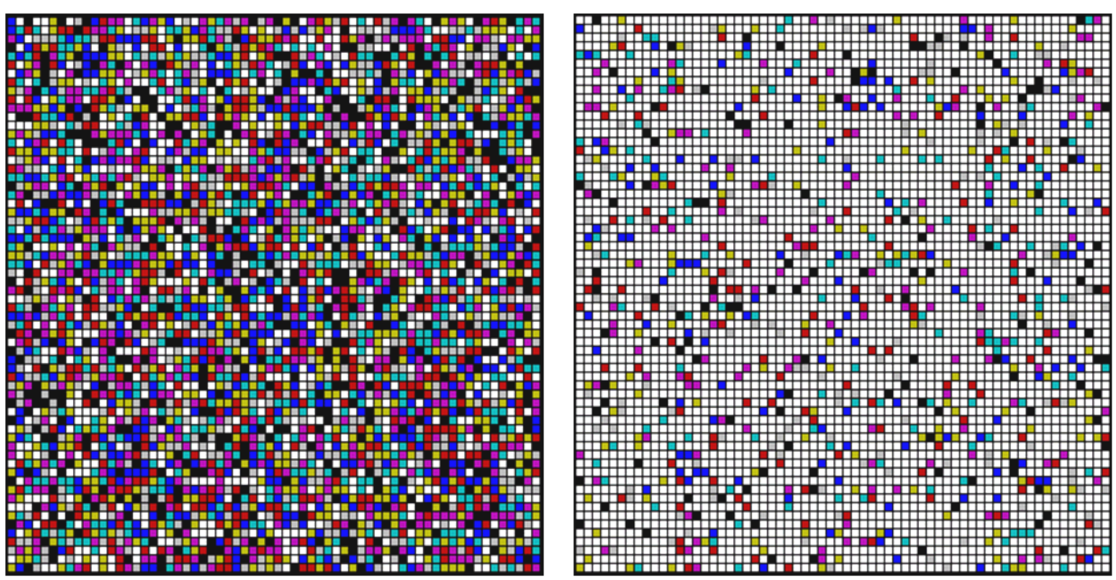

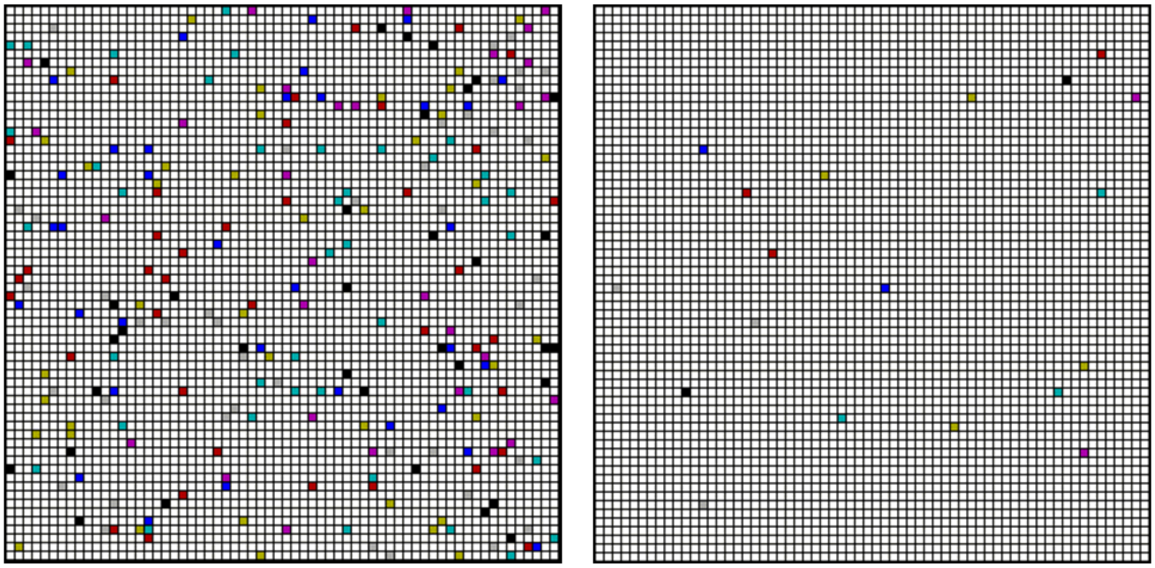

Figure 2. Conquest by All-Defect with a global assumption. Typical evolution of a randomized array of the 8 reactive Prisoner's Dilemma strategies, using global replacement: in each generation 5% of the array is replaced with the strategy of the most successful cell in the array.

III. Studies in Communication

It is clear that game-theoretic cooperation can emerge in distributed networks. Can patterns of simple signaling emerge in a similar way? Can a dynamics this simple produce simple forms of communication? Though the modeling techniques are very similar, the games at issue in this exploration go well beyond the simple Prisoner's Dilemma.

The basic model again uses an array of individuals in a two-dimensional lattice.[4] In order to keep things simple, individuals in this model don't move: they are embedded in the array something like coral in a reef. What does move are food sources and predators, each of which travels in a random walk, cell by cell, across the array. It should be emphasized that the food sources are food *sources*—individuals feed when a food source lands on them, but the food sources themselves are not consumed and do not disappear. Food sources are like a school of fish, perhaps, continuing their random walk and offering nourishment for the next individual down the line. In much the same way, predators in the model are never sated: they continue to pose a danger in their random walk.

Each of the individuals embedded in the array has a small behavioral repertoire. It can open its mouth, hide, stay in neutral, make a sound 1 (heard by itself and immediate neighbors), or make a sound 2. Each individual also has a limited range of perception. It knows when it is fed—when its mouth is open and a food source lands on it—and when it is hurt—when it is not hiding and a predator lands on it. Each agent can hear and distinguish sounds made by itself or immediate neighbors; it knows when someone just made sound 1, for example, or someone just made sound 2.

The behavior of each individual is dictated by a simple strategy code. An agent's code may dictate that it opens its mouth only when it hears sound 1, for example, or hides only when it hears sound 2. An agent's code might also dictate that it *never* opens its mouth. When an individual opens its mouth and food is on it, it gains a point for 'feeding'. When a predator hits an individual that isn't hiding, that individual is 'hurt' and loses a point. But opening one's mouth, hiding, and making a sound all exact an energy cost. There are no costless actions in this model, and there are no costless signals.

This model differs from others in the literature in a number of ways. One difference is a thorough spatialization; in this model, all action is local action. Some of Cecilia and Paolo Di Chio's work puts a similar emphasis on spatialization (Di Chio & Di Chio 2007). A second difference is the fact that the emergence of signaling is ecologically situated. The question is how arbitrary sounds take on a meaning, but the context is an environment with some built-in significance of its own: the importance for individual agents of capturing food and avoiding predators.

Many of the models in the literature employ something like a Lockean theory of meaning, in which meanings correspond to something in individual heads. In contemporary guise, those Lockean meanings often take the form of representation matrices, with convergence on the same matrix in different individuals taken as the measure of 'successful communication.' The model offered here is much closer to a theory of meaning as use. Here there is no matching of matrices in different heads; the measure of communication is simply and solely successful behavioral coordination across a community.

A final distinctive feature of the model is its emphasis on individual gains, tied directly to individual success in capturing food and avoiding predators. Unlike many models—unlike even Lewis's signaling games—there is no 'mutual reward' for successful communication or mental matching, for example. Prashant Parikh's models share this emphasis on individual gains (Parikh 2001, 2006, Parikh & Clark 2007).

What my collaborators and I were interested in was the possibility that patterns of simple signaling might emerge as a result of a situated economics of information: that communication might emerge in response to environmental pressures on the basis of individual gains. We investigated three forms of strategy updating in three stages of the work. In initial studies, strategy change was by simple imitation, just as in the cooperation studies above (Grim, Kokalis, Tafti & Kilb 2000). In a second series of investigations, we cross-bred our strategies using genetic algorithms, but restricted our genetic algorithms—like everything else—to purely local instantiation. Genetic algorithms are generally applied globally; individuals ranking highest on some fitness function are selectively cross-bred and their offspring are re-introduced randomly into the population as a whole. We replaced that global algorithm with a local one in which strategies are 'bred' as hybrids between the strategy of a cell and that of its most successful neighbor; local success results in local hybrid reproduction (Grim, Kokalis, Tafti & Kilb, 2001). In more recent models, our individuals instantiate simple neural nets, doing a partial training on the behavior of successful neighbors (Grim, St. Denis & Kokalis 2003; Grim, Kokalis, Alai-Tafti, Kilb & St. Denis 2004). On the basis of lessons learned regarding the pitfalls of perfect worlds (Grim, P., Kokalis, T., Alai-Tafti, A., & Kilb, N. 2000) and following hints in the work of Martin Nowak (Nowak & Sigmund 1990, 1992), we worked throughout with a stochastically imperfect world. In 5% of cases individuals open their mouths and in 5% of cases they hide regardless of strategy or input.

Results using all three forms of strategy updating showed an emergence of simple signaling in a spatialized environment of wandering food sources and predators. The following sections summarize main points using results from the more recent work with neural nets.

A. The Dynamics of Feeding, Predation, and Simple Signaling

For even creatures as simple as those outlined, there is a behavioral strategy that would seem to qualify as an elementary form of signaling. Consider a spatialized community of individuals which share the following behavioral strategy:

They make sound 1 when they successfully feed
They open their mouths when they hear a neighbor make sound 1

They make sound 2 when they are hurt
They hide when they hear a neighbor make sound 2

We have termed these 'perfect communicators.' Our food sources migrate in a random walk from cell to cell, never being entirely consumed. So suppose a community of ‘perfect communicators’, and suppose that one individual successfully feeds. Because it's a 'perfect communicator', it makes sound 1. All its neighbors are perfect communicators as well; when they hear sound 1, they open their mouths. The food source will migrate, and one of those neighbors will successfully feed. *That* cell will then make sound 1, alerting its neighbors. The result in a community of ‘perfect communicators’ is a chain reaction in which the food source is successfully exploited on each round.

This advantage, of course, demands that we *have* a full community of ‘perfect communicators’. For a single ‘perfect communicator’ there is no advantage to such a strategy at all: given the energy cost for sounding and opening its mouth, in fact, there is a significant *dis*advantage.

The dynamics are slightly different for predators. If an individual is hit with a predator and gives an alarm call, his neighbors hide and therefore are *not* hit by the wandering predator. Those neighbors therefore do not pass on the alarm call. Because there is no alarm call, someone does get hit by the predator on the next round. Even ‘perfect communication’ therefore affords protection against predators only on alternate rounds. Because of that difference in dynamics, and in order to avoid a bias toward one environmental pressure rather than another, we worked throughout with twice as many predators as food sources across the array.

Since our individuals have two sounds at their disposal, of course, there are two versions of ‘perfect communicators’. One uses sound 1 for food and sound 2 for predators; the other uses sound 2 for food and sound 1 for predators.

B. Emergence of Communication in Spatialized Arrays of Neural Nets

In our simplest work with neural nets, our individuals are instantiated as perceptrons: two-layer feed-forward nets of the form shown in Figure 3.[5] This structure has two distinct lobes. The left takes sound as input and outputs whether an individual opens its mouth, hides, does neither or both. The right lobe takes as input whether it has been fed or hurt on a particular round, outputting any sound the agent makes.[6]

We use a bipolar coding for inputs, so that ‘hear sound 1’, for example, takes a value of +1 if the individual hears sound 1 from an immediate neighbor 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 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.

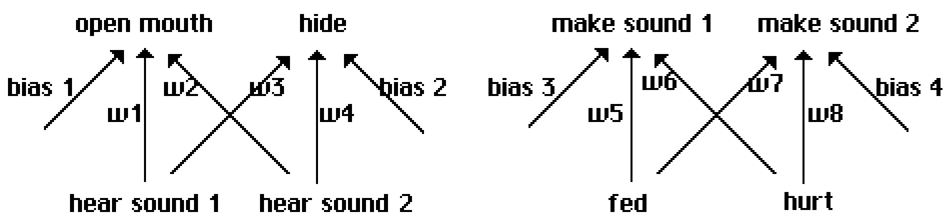

Figure 3 Perceptron structure of each agent embedded in the spatialized environment of wandering food sources and predators with partial training on the behavior of successful neighbors.

We code our behavioral strategies in terms of the outputs they give for possible pairs of inputs. Figure 4 shows the possible inputs at 'hear sound 1' and 'hear sound 2' for the left lobe of the net, with samples of output pairs for 'open mouth' and 'hide' for these pairs of inputs. The left-lobe behavior of a given strategy can thus be coded as a series of 8 binary digits. With a similar pattern of behavioral coding for the right lobe, we can encode the entire behavior of a net with 16 binary digits. The sample space consists of 38,416 strategies in this sample space, but there are only 2 that qualify as 'perfect communicators'.

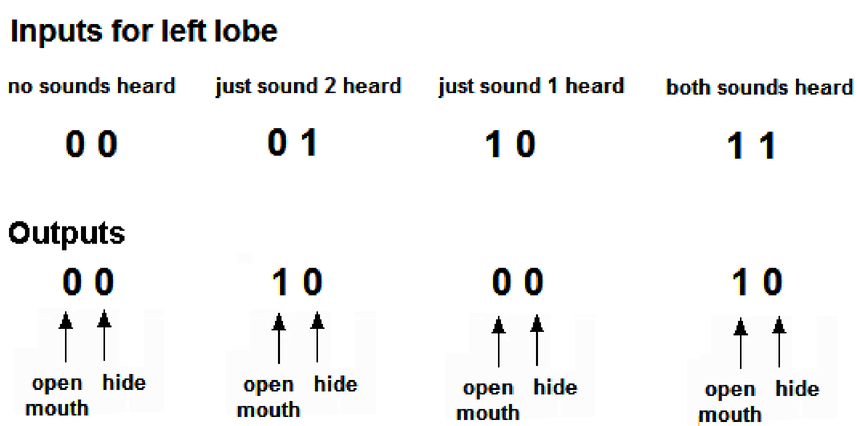

Figure 4 Possible pairs of inputs at 'hear sound 1' and 'hear sound 2' for the left lobe of each agent, with samples of output pairs for 'open mouth' and 'hide.' This strategy opens its mouth whenever it hears sound 2, and never hides. With a similar coding for the right lobe, the behavior of each net can be coded as 16 binary digits.

We populate our array with neural nets carrying twelve random weights. After each 100 rounds, our individuals look around to see if there are more successful neighbors. If so, they do a partial training on the behavior of their most successful neighbor. We use the standard delta rule as the learning algorithm for our perceptrons. For a set of four random inputs, the cell compares its outputs with those of its target neighbor. 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 $w_{new} = w_{old} + (\text{target} \times \text{input})$ and $\text{bias}_{new} = \text{bias}_{old} + \text{target}$.

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 which might not previously have existed in either teacher or learner, thereby allowing a wider exploration of the sample space of possible strategies (for details see Grim, St. Denis, Kokalis, 2002 and Grim, Kokalis, Alai-Tafti, Kilb, & St. Denis, 2004).

We start, then, with an array of perceptrons with randomized weights and biases. Figure 5 shows the evolution of such a randomized array over the course of 300 generations (animation at www.pgrim.org/pragmatics). Here 'perfect communicators' are coded in pure black and white, though it should be noted that they appear here entirely by partial training. The initial display contains no 'perfect communicators' at all: they emerge by the mechanism of partial training on successful neighbors, spreading by the same means. Figure 6 graphs the result. [7]

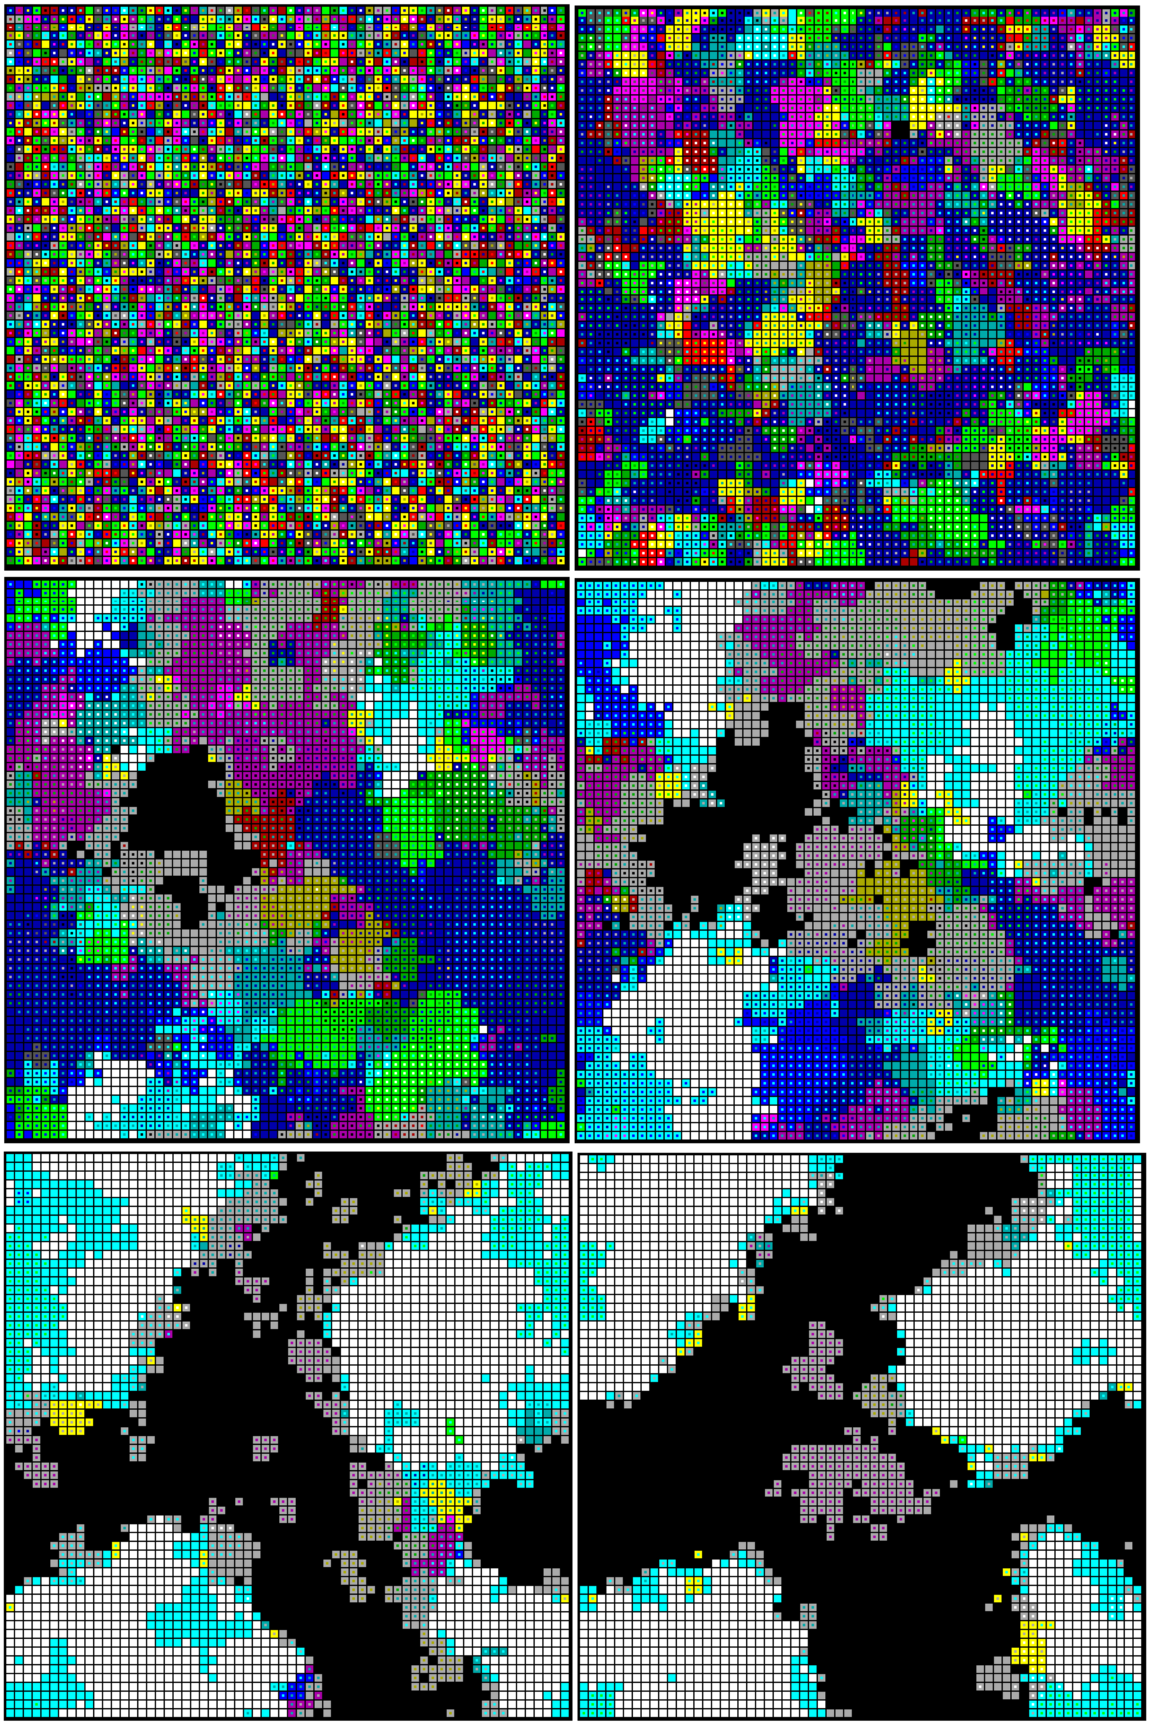

Figure 5 Emergence of two dialects of perfect communicators, shown in solid black and white, in a randomized array of perceptrons with partial training on successful neighbors. All other behavioral strategies are coded using shades of gray for backgrounds and central dots. Centuries 1, 10, 50, 100, 200 and 300 shown. A full animation of the development can be seen at www.prim.org/pragmatics.

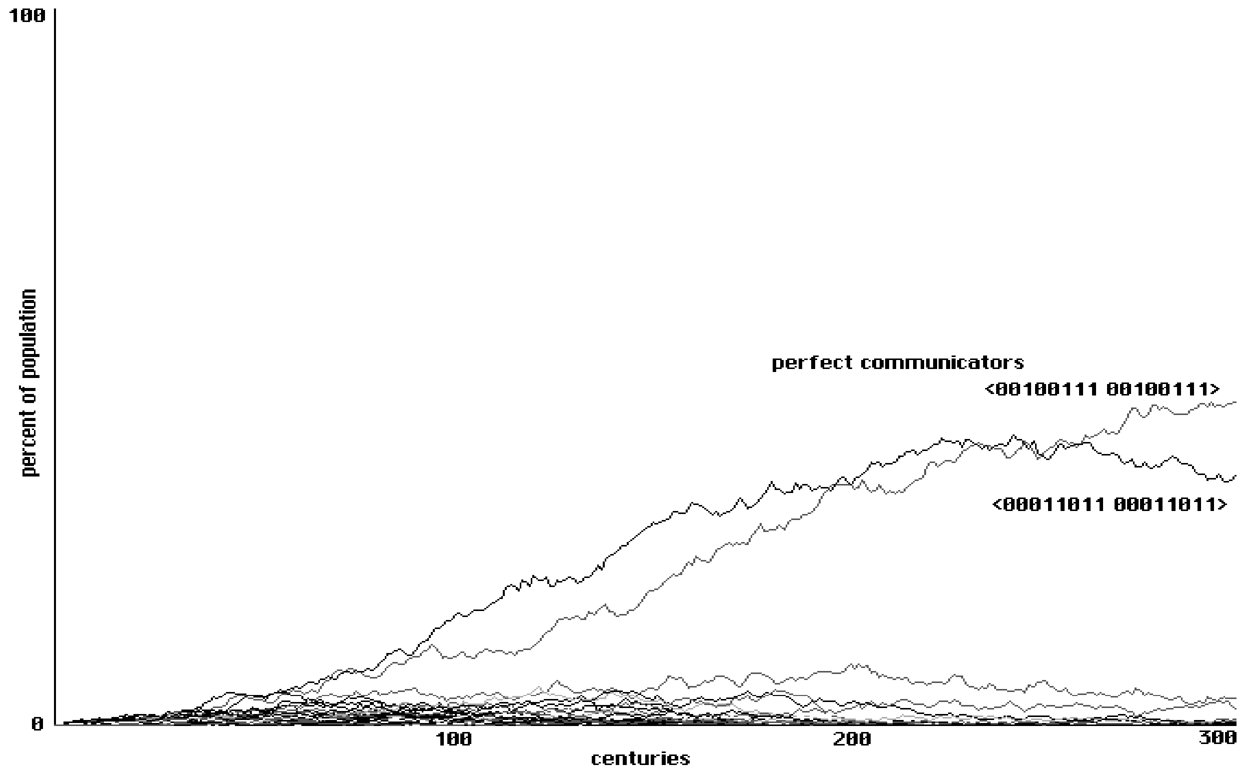

Figure 6 Emergence of two dialects of perfect communicators in an array of perceptrons through partial training on successful neighbors in a spatialized environment of wandering food sources and predators.

The primary lesson of the work in communication is that simple forms of signaling can emerge on the basis of entirely individual gains, geared solely to environmental success, in distributed networks of local interaction. If these simple forms of communication count as 'semantics', what this shows is that quite minimal conditions are sufficient for the emergence of simple semantics.

IV. The Emergence of Pragmatics

Can the reach of this kind of modeling be extended beyond semantics? In what follows the same tools are applied to aspects of pragmatics in an attempt to understand the emergence of the Gricean maxims.

Here as in the semantic case we wanted the information our agents deal with to be real information, rather than simply some arbitrary fitness function. In line with the studies in semantics, we wanted the pragmatics of communication to be contextually situated—realistically tied to individual success in the environment.   Here we return to a two-dimensional lattice, but with an array of agents who have already fixed on a signaling system—an array of agents for whom a system of communication is already in place.

The environment for these studies is not one of wandering food sources and predators, but it does have an essentially spatialized character. This environment, like that of the semantic studies, is an environment of random events of local significance. Here events come in three colors, playing across the array—these are rains of manna, perhaps, or threatening storms, or clouds of locusts. The important thing for each individual is that it act appropriately with regard to the environmental events around it. If an agent finds itself in a magenta cloud (portrayed as a shade of gray in Figure 7, but visible in color at the url for this chapter), it will gain points by acting in an appropriately magenta way. It will lose points if it does not. If an agent finds itself in a bright red cloud, it is red behavior that is advantageous and failing to act in an appropriately red way that will lose points. Color acting is by no means free, however—one cannot hedge one's bets by always acting magenta, red, and yellow simultaneously. If an agent acts red when there is no red cloud it loses points as well.

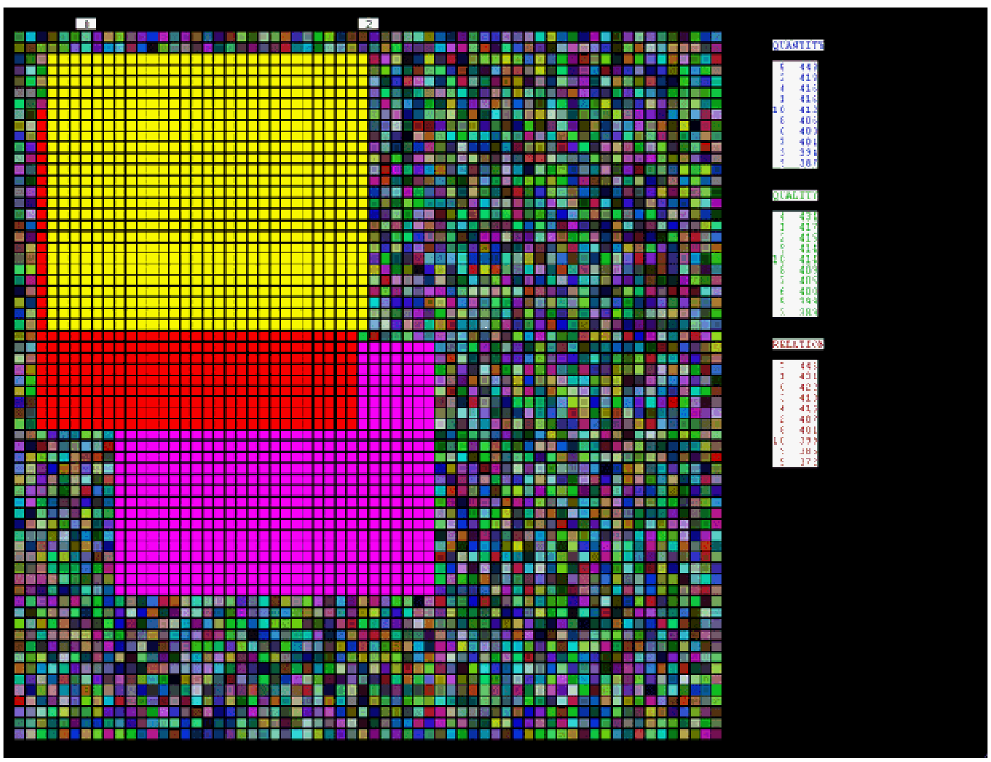

Figure 7   The environment of random events of local significance used in the pragmatics studies.

If all agents could detect their own environmental contingencies at all times, the model would be much less interesting—everyone would do well. In our model, however, agents *cannot* always tell what is happening; only a random 10% of the agents on each round are able to detect their current environment. The others, for that round, are blind.

Our assumption, however, is that a system of communication is already in place. Our agents already speak 'red', 'yellow', and 'magenta', in the following sense: They can give 'red', 'yellow', and 'magenta' signals heard by their immediate neighbors. If an agent receives a 'red' signal from a neighbor, he acts red; if he receives a 'magenta' signal, he acts magenta, and so forth. If you as an agent happen to be in the 10% that can see your immediate environment on a given round, then, you have all the information you need to do well. If you have a communicating neighbor who is in that 10%, you derive the same benefit.

It is here that the Gricean maxims come in. Suppose that you as an agent have an immediate neighbor who is among the 10% who can see what's happening. We'll call him 'Seymour'. Because a system of communication is already in play between you and Seymour, and because Seymour can see what's happening, you are going to do fine. You are going to do fine, that is, *as long as he's following the appropriate pragmatic maxims:*

You will have the information you need on the condition that Seymour observes something like a maxim of *quantity*: as long as he gives you enough information. Agents start out with a set probability of transferring important information—a probability, if they see 'red', of signaling red to their neighbors. We labeled our probability intervals 1 to 10, corresponding to .1 intervals between 0 and 1.

You will have the information you need on the condition that Seymour observes a maxim of *quality*: as long as he is telling you the truth. Agents start out with a probability interval between 1 and 10 of sending a correct message if they send a message at all: a probability of sending 'red' in a red environment rather than 'magenta' or 'yellow', for example.

You will have the information you need on the condition that Seymour observes a maxim of *relation*: on the condition that he is giving you relevant information. I will detail the modeling of relevance in a moment.

The general question at issue is whether pragmatic behavior in accord with Gricean maxims will emerge in a communicative network of this kind. Here as in the semantic case, it should be emphasized, the model works entirely in terms of (1) individual gains and losses—gains from acting appropriately in one's environment and avoiding inappropriate action, and (2) local action in the distributed network itself—the local action of sending and receiving signals with immediate neighbors and of updating on the behavior of successful neighbors.

In the pragmatic studies that follow, our behavioral updating is simply by imitation. One of the lessons from the semantic studies seems to be that it is patterns of mutual advantage in an environment that are of importance, regardless of how agents learn from their neighbors. We expect much the same to hold in the case of pragmatics, but more complicated mechanisms of localized genetic algorithm and neural net training are left to further work.

In the following sections communicative behavior in accord with each of the maxims is developed successively; we then show that these behaviors can co-evolve as well.

1. Quantity

We begin with the pragmatics of quantity. What our model gives us is an environment of changing but spatialized events. On each round a random 10% of agents can detect those events. All individuals in the array are primed with appropriate signals—'red' for a red environment, 'yellow' for a yellow, and the like—but our agents start with different probabilities of signaling at all. Some transfer information whenever they have it. Some signal with only 60% probability, or 10%, or never signal at all. On each round, all agents gain and lose points on an individual basis. After 100 rounds, they look to see if any neighbor has gained more points—individual points gained on the basis of information *received*, of course. If so, they imitate the quantity probability in *sending* of their most successful neighbor—the probability of that neighbor signaling to others.

What happens if we start with an array of agents with randomized 'quantities'? In Figure 8 we use the darkest shades to code those with the least probability of signaling, using the brightest shades to code those with the highest probability (in the url, coding is done with darker and lighter shades of blue). Over the course of a few hundred generations the array comes to be increasingly dominated by those following a Gricean maxim of quantity: by those with the highest probability of signaling. Here signaling does not have to be free: the results hold even with a cost for signaling. Figure 9 shows a graph for a typical run.

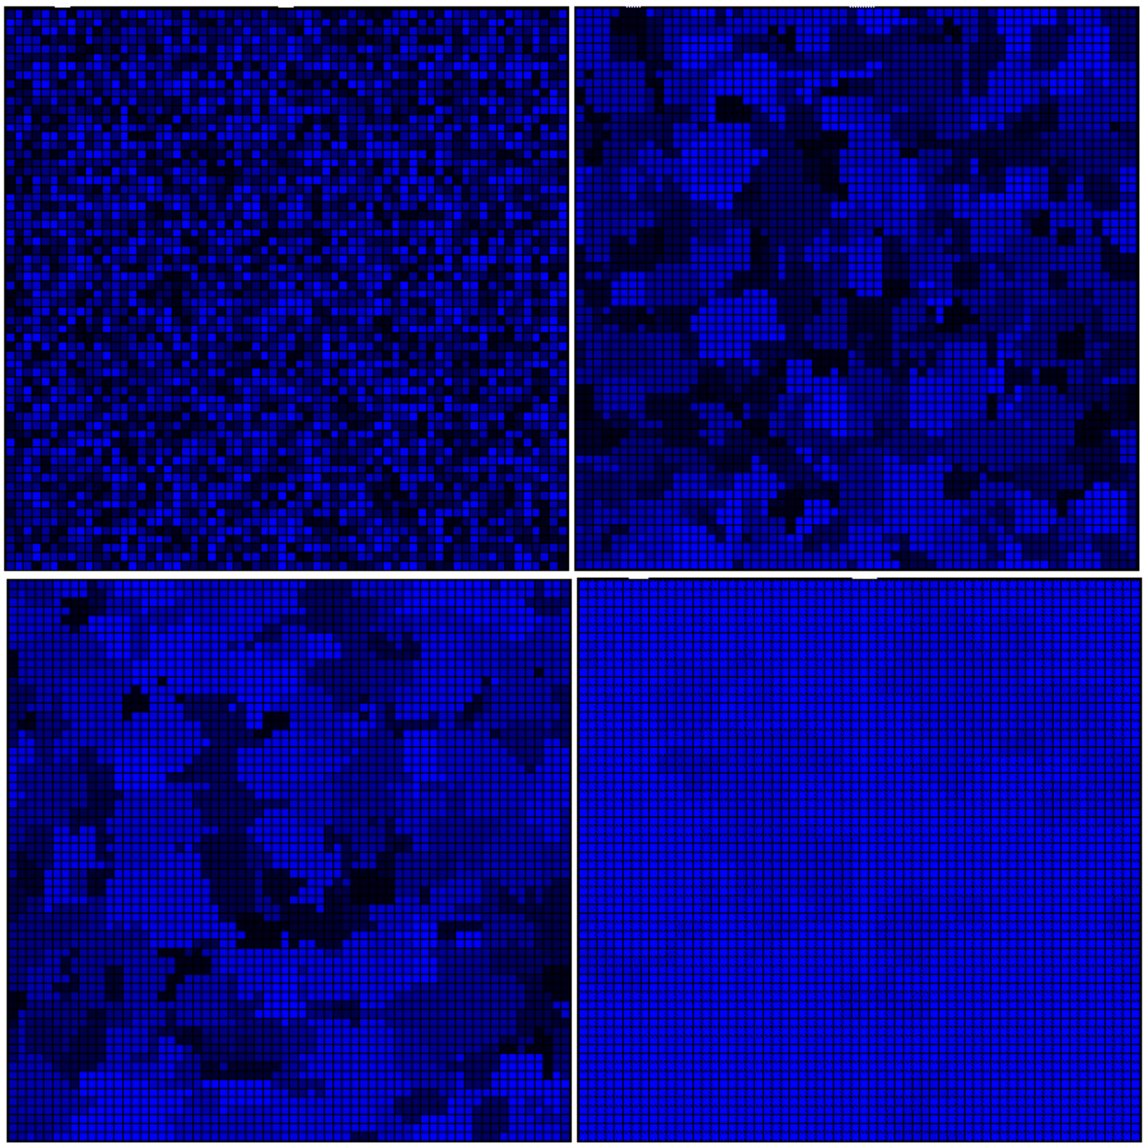

Figure 8 The Emergence of Quantity. Agents with lower probabilities of signaling are coded in darker shades. Evolution to pragmatic maxim of quantity by imitation of successful neighbors over 100 generations. Initial randomization and generations 5, 10, and 100 shown.

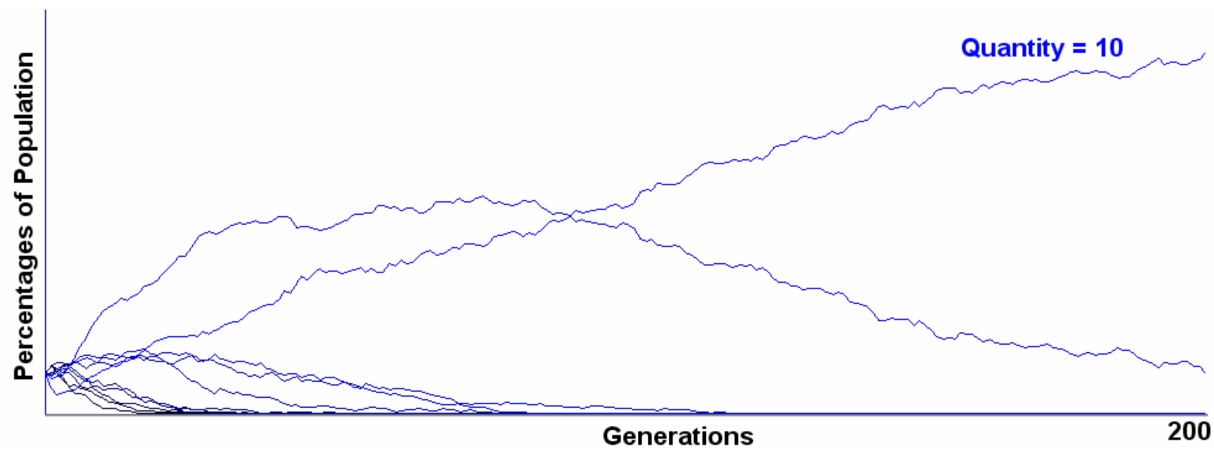

Figure 9 The Emergence of Quantity. Percentages of population shown over 200 generations.

2. Quality

In order to run the model for quality we assume that quantity is already in place: all those who detect a particular environmental condition signal it to their neighbors. They may or may not signal *accurately*, however. Our agents come with a probability between 1 and 10 of accurate signaling, otherwise sending a decidedly incorrect signal—'red' instead of 'yellow', for example, or 'yellow' instead of 'magenta'.

We start with a randomization in terms of quality. Those with the lowest quality are coded with the darkest color; those with the highest are coded with the brightest (Figure 10; in the url, coding is done with darker and lighter shades of green). Over the course of 100 generations the array comes to be dominated by those with following a Gricean maxim of quality: The population converges on truth-telling (Figure 11).

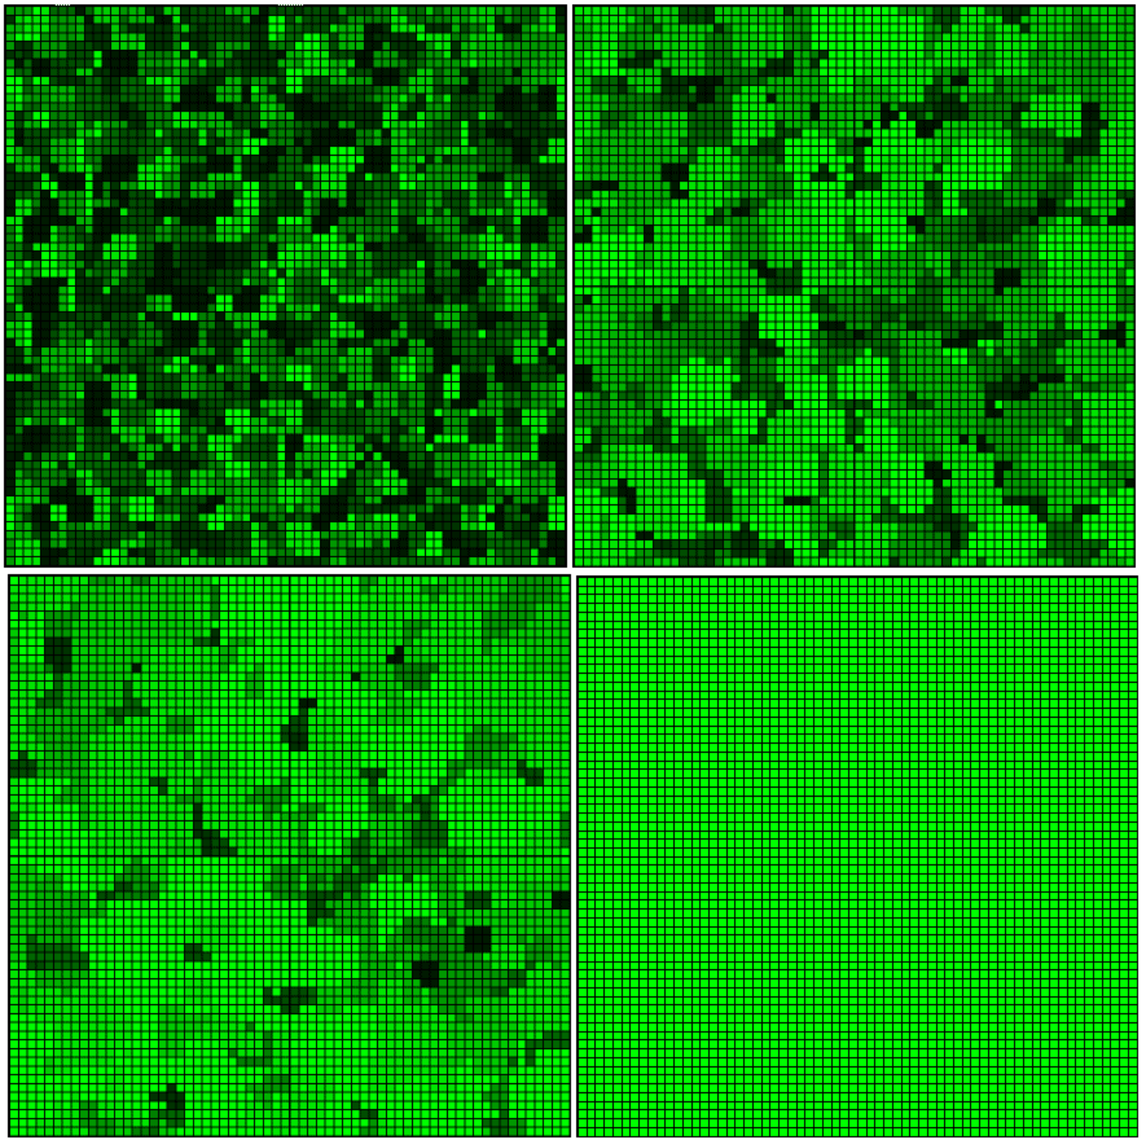

Figure 10   The Emergence of Quality. Agents with lower probabilities of accuracy in signaling are coded in darker shades. Evolution to pragmatic maxim of quantity by imitation of successful neighbors over 100 generations. Generations 1, 5, 10 and 90 shown.

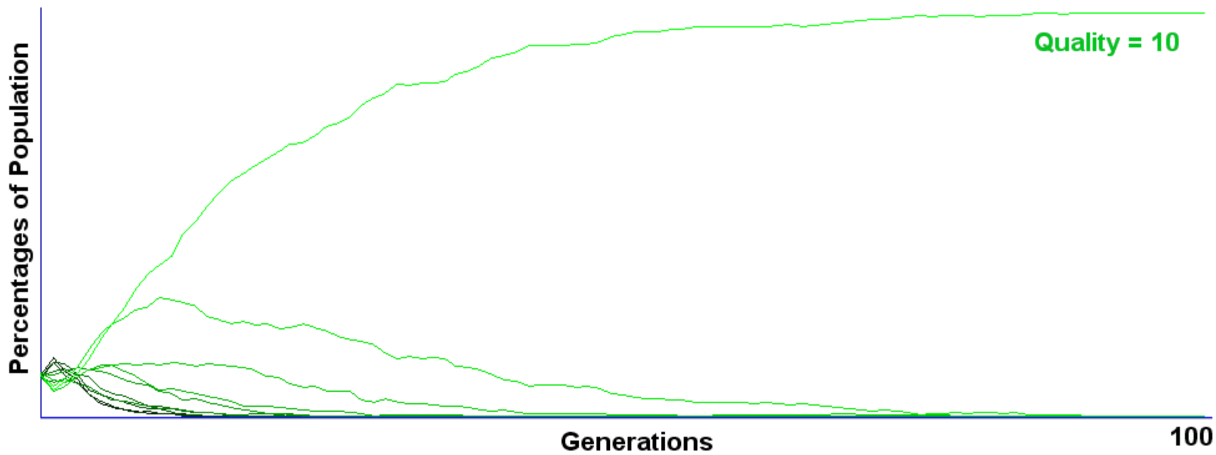

Figure 11 The Emergence of Quality. Percentages of population shown over 100 generations.

3. Relevance

Information regarding environmental conditions is relevant, on this model, to those who are actually *in* an event environment. Environments have fringes, however, and agents come with a 'degree of relevance' in their reports (Figure 12). Some agents signal precisely when an event is happening in their immediate neighborhood. Others report when something is happening in a slightly larger area, and which therefore has a lower probability of actually occurring where they are. Others report on a still wider area, offering information that is still less relevant.

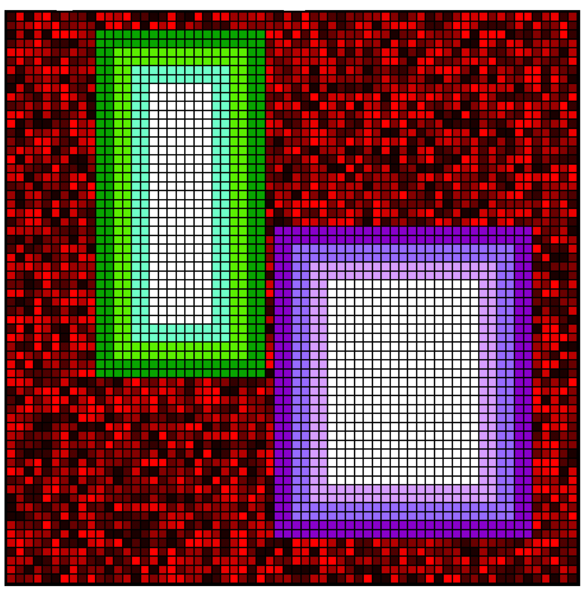

Figure 12 Relevance neighborhoods

Here we assume both quantity and quality to be in place: reports are forthcoming from all agents, truthful as to the character of the event reported, but of varying relevance. We start with an array randomized as to relevance, coloring those with least relevance in darker shades, those with most relevance in brighter shades (in the url coding is done in shades of red). The result, as shown in Figures 13 and 14, is convergence to the highest degree of relevance across the array. The population fixates on maximized relevance.

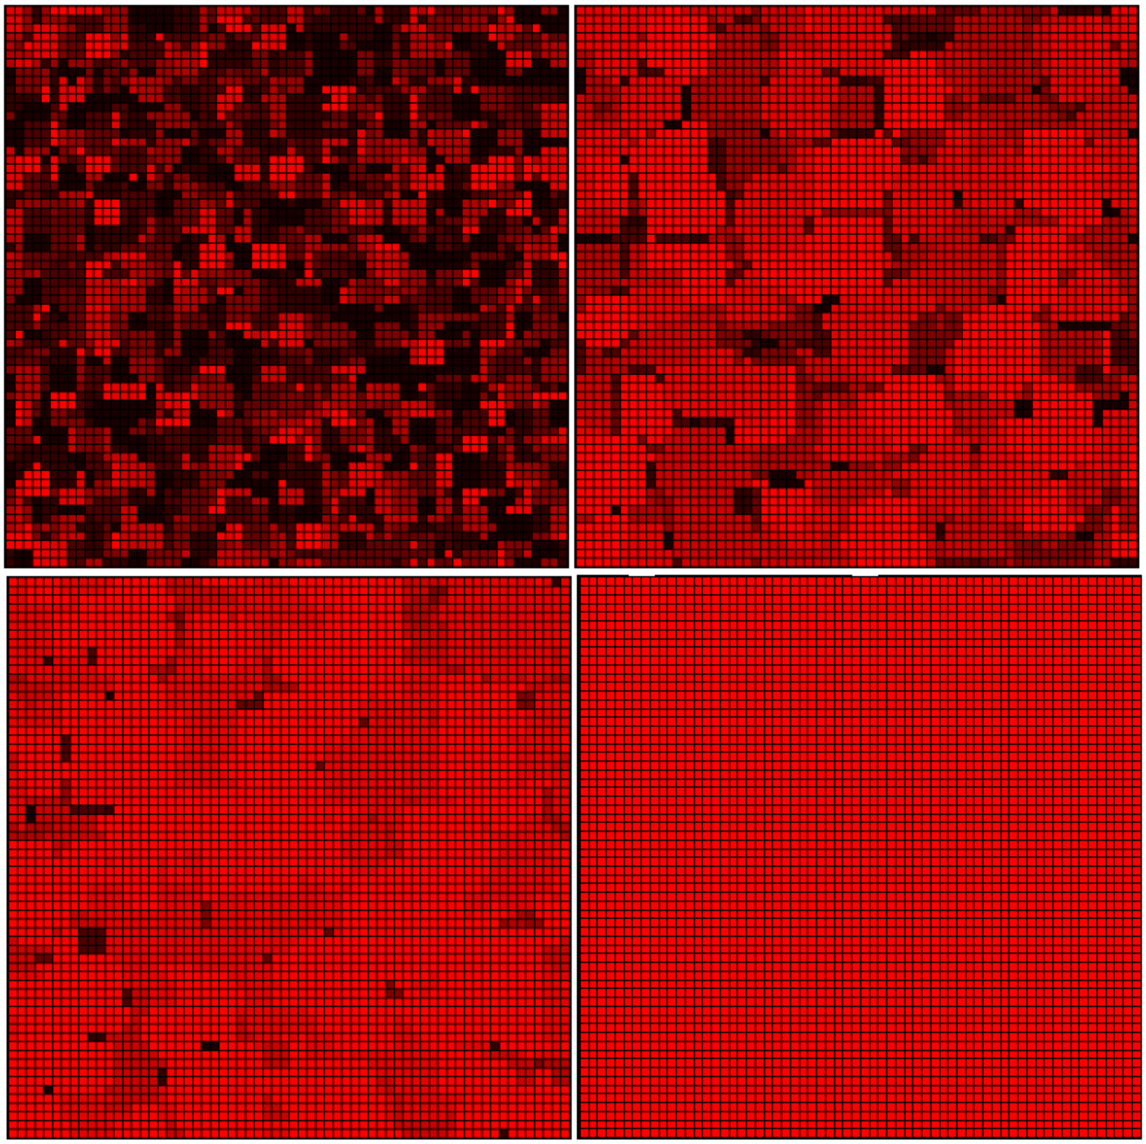

Figure 13 The Emergence of Relevance. Agents with lower relevance in signaling are coded in darker shades. Evolution to pragmatic maxim of relevance by imitation of successful neighbors over 100 generations. Generations 1, 5, 10 and 100 shown.

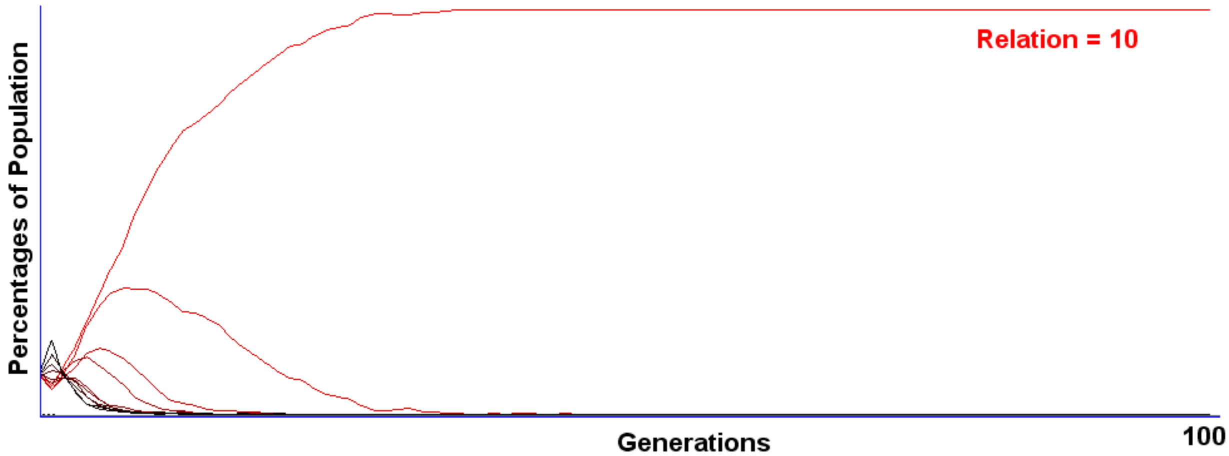

Figure 14 The Emergence of Relevance. Percentages of population shown over 100 generations.

D. Co-Emergence of the Maxims

What if we combine all three of these pragmatic parameters? Here we start with an array randomized for all pragmatic virtues: randomized for quality, quantity, and relevance. Shadings are combined: black or near-black indicates a combination of poor pragmatic values, a light grey indicates an individual virtuous in terms of all three pragmatic maxims. In the animation accessible at www.pgrim.org/pragmatics, we have simply combined the blue, green, and red codings of the pragmatic virtues above.

On each turn, each agent gains and loses points by acting on the information he receives. But of course the quality of that may information may vary—he may *not* receive important information from a neighbor, he may receive *incorrect* information, or he may receive *irrelevant* information. At the end of 100 rounds, if a neighbor has garnered more points from acting on *received* information in the environment, each agent updates on just one aspect of that neighbor's pragmatic character in *sending* information: He randomly adopts the neighbor's quality rating, quantity rating, or relevance rating. That means, of course, that he may update on the wrong thing—the neighbor's success may be due to something other than the trait he imitates.

The fact that only one character will change each time, and that they are all interacting, makes the result significantly slower in this simulation. Nonetheless, over 600 generations or so, convergence is to the highest value in all pragmatic parameters (Figures 15 and 16).

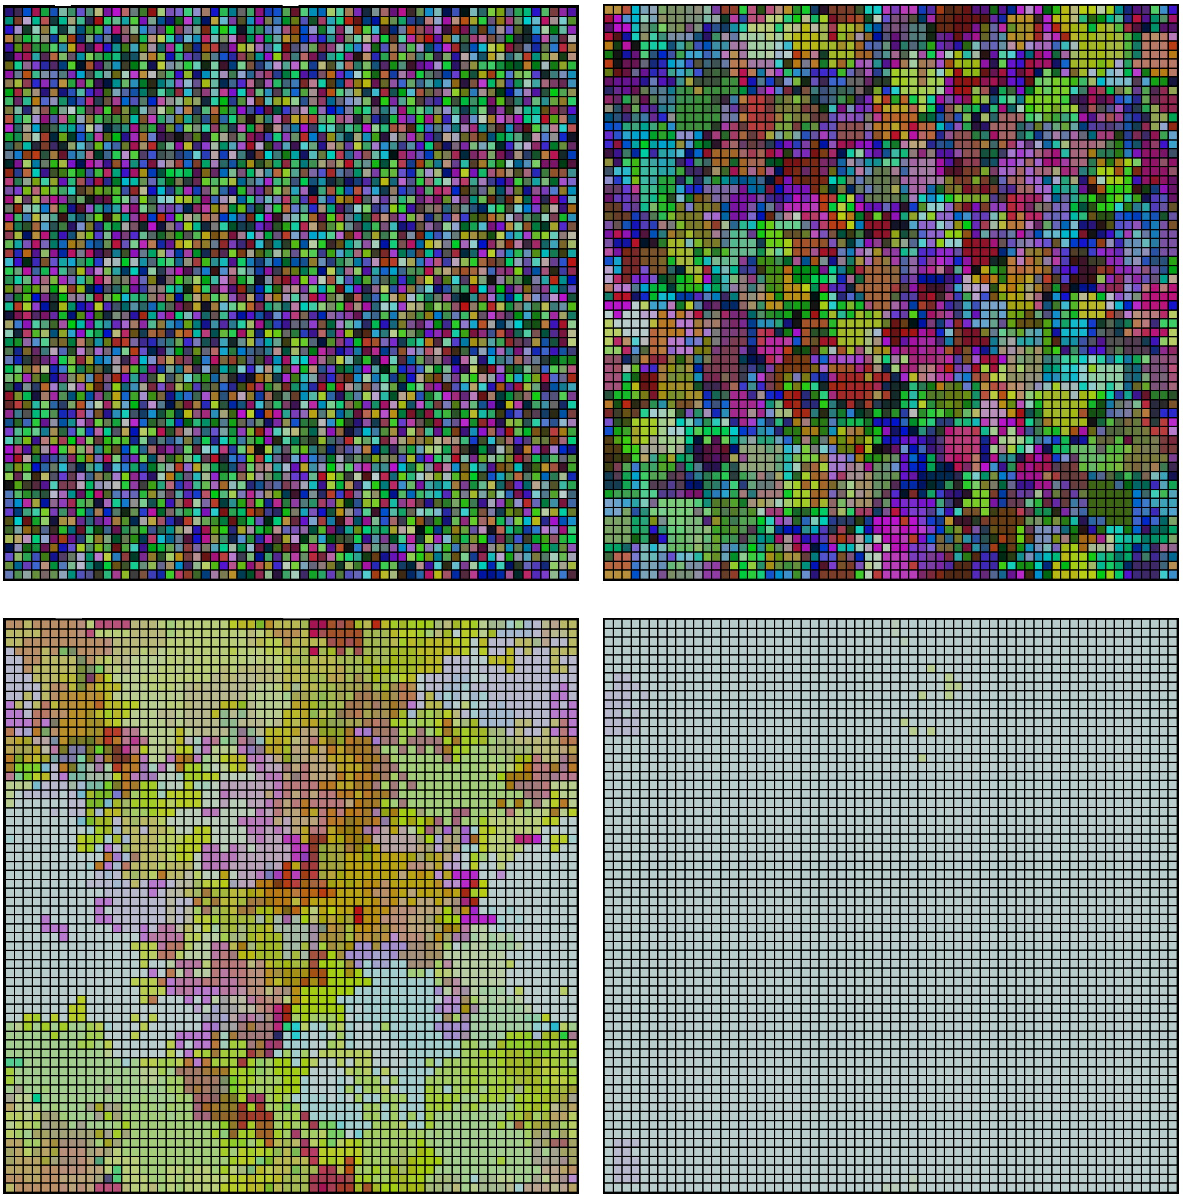

Figure 15 Co-Emergence of the Maxims together over 600 generations. Quality, quantity and relevance are coded from dark to light in three different shades. Initial randomization and generations 10, 100, and 600 shown.

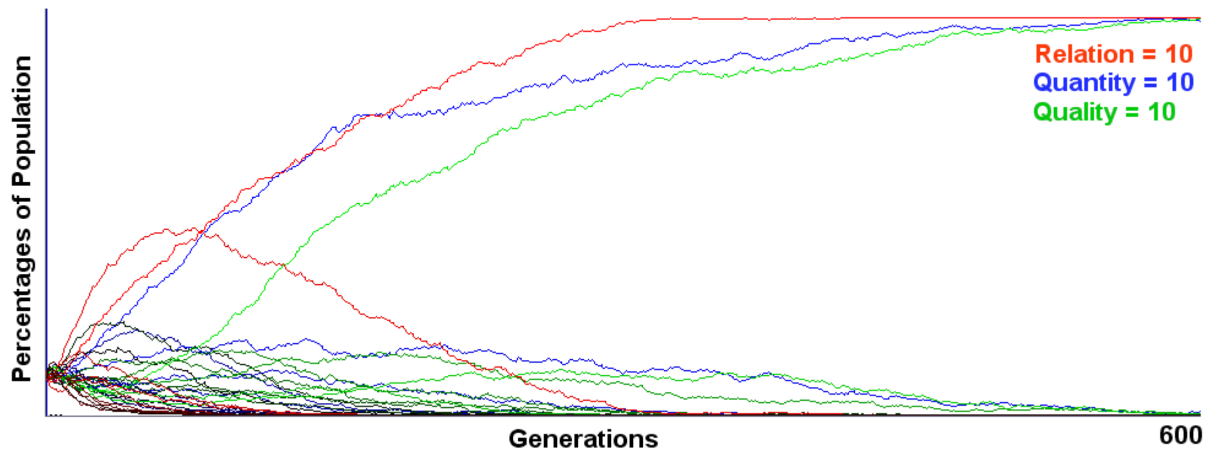

Figure 16 Co-Emergence Quantity, Quality, and Relation over 600 generations.

In the models for pragmatics as in those for signaling and cooperation, it should be emphasized, it is local spatialized action that is the key. In correlate work in evolutionary game theory, using assumptions of global interaction, Nicholas Asher was able to produce pragmatic conventions of sincerity only with very strong assumptions regarding the immediate advantage to a speaker of telling truth (Asher, Sher & Williams 2002; private correspondence). In the dynamics of localized action and imitation, on the other hand, a full range of pragmatic conventions emerge in terms of environmentally significant events, assuming no immediate advantage to the speaker, and despite signaling costs.

Although simulations should never be treated as conclusive, they can be wonderfully suggestive. What these results suggest is that fundamental phenomena of pragmatics emerge spontaneously in communicative networks. Those fundamental phenomena don't need to be added, they don't need to be cultivated, and they don't need to be assumed. In any informational network of a certain sort—one in which individuals are striving merely for individual informational advantage—behavior in accord with Gricean maxims of quality, quantity, and relation will come for free.

E. Simulating Grice

These results seem to go some way in fulfilling what Grice himself wanted:

"…I would like to be able to show that observance of the Cooperative Principle and maxims is reasonable (rational) along the following lines: that anyone who cares about the goals that are central to conversation/communication (such as giving and receiving information, influencing and being influenced by others) must be expected to have an interest, given suitable circumstances, in participation in talk exchanges that will be profitable only on the assumption that they are conducted in general accordance with the Cooperative Principle and the maxims. Whether any such conclusion can be reached, I am uncertain…" (Grice 1989, 28-29)

On the other hand, there are aspects of the simulational work that run directly counter to Grice. Grice's approach to meaning is heavily intentional, and his work in pragmatics follows suit. Following Grice's lead, a great deal of good and solid work in pragmatics assumes highly reflective and highly cognitive agents—agents capable of identifying the limits of literal meaning, capable of logical deduction, and capable of building cognitive models of their interlocutors.

What these simulations suggest is that something much simpler and much less cognitive may be going on far beneath that level. The agents in these simulations are too simple to form or recognize intentions. They have no explicit capacities for logical inference, and are far too simple to build cognitive models of their interlocutors. Despite that, something develops in networks of those simple agents that looks a lot like signaling. Behaviors develop in networks of those simple agents that also seem to be in accord with the Gricean maxims.

With more complicated agents, of course, these same dynamics may play out in more complicated forms involving intentions, logical inference, and cognitive modeling. But what the simulations suggest is that cognitive complexity of that sort isn't strictly necessary: information maximization in communication networks produces fundamental phenomena of pragmatics even with far simpler cognitive agents.

These results also offer some suggestions regarding the conceptual categories in play. Grice followed Kant in speaking of 'maxims'. But the term 'maxims' itself may suggest too high a cognitive level. Grice also followed Kant in dividing his maxims into quantity, quality, relation and manner, but even Grice himself seemed to recognize the artificiality of these categories. There are undoubtedly important distinctions to be drawn between forms of linguistic convention, but the lines should perhaps not be drawn where Grice drew them. In terms of dynamics within information systems, even the distinction between semantics and pragmatics may not be as hard and fast as it has sometimes been assumed to be.

The background work summarized in sections I and II concerns the emergence of signaling systems, read as semantic. The newer work outlined above assumes a semantic system in place, and reads further results as an emergence of pragmatics. What if we combined the simulations; what if we did the two together? We haven't yet done that work, but my prediction is that the results will be much the same. My prediction is that information maximization will emerge *both* in terms of coordinated use of signals and in pragmatics of appropriate, relevant, and accurate signaling.

Would that be the emergence of two things simultaneously—semantics and pragmatics? One could look at it that way. But I also think that one could take the result as an indication that there is a single core phenomenon at issue, of which semantics and pragmatics are somewhat artificially separated 'aspects'. The core phenomenon, and the engine that drives both, is environmental information maximization in communication networks.

V. Pragmatic Implicature and Inference

What the simulational results offer so far is the emergence of something like Gricean pragmatic conventions. Assuming a semantics of basic signals, the dynamics of local interaction, self-advantage, and imitation of successful neighbors is sufficient to produce conventions of quality, quantity, and relevance. But there is something important that this does not yet give us. It does not yet give us Gricean inference on the part of the hearer—reliance on conventions in play in order to draw particular inferences beyond what is said. It also does not yet give us Gricean implicature on the part of the speaker—the use of conventions in play in order to encourage particular inferences beyond what is said. Can the modeling techniques outlined here take us that next step? The answer appears to be ‘Yes and No,’ but an interesting ‘Yes and No’.

The reason for the ‘No’ answer is that most Gricean inferences and implicatures are *characterized* in highly cognitive terms. Some of Grice's examples seem to depend crucially on cognitive modeling of the other agent—an aspect rightly emphasized by Asher and his collaborators (Asher, Sher, & Williams 2002; Asher & Lascarides 2003). Some of the examples depend on an agent’s explicit recognition of the limits of literal meaning, or of conventions as conventions, or explicit recognition of patterns of logical entailment. Can the simulation techniques above be used to model those? No, because we are working with agents at far too low a cognitive level. In order to extend network techniques to most cases of Gricean inference and implicature we would need networks of more complex agents.

Gricean inference and implicatures are a mixed bag, however, and there are some phenomena in at least the general area that *do* show up even in simulations of agents this simple. One of those phenomena is a close relative of scalar inference and implicature—the inference and implicature from ‘some students came to the party’ to ‘not all came’, for example, or ‘there are five apples in the basket’ to ‘there aren’t seven.’

What is at issue in such cases is a particular kind of coordination between speaker and hearer. The speaker’s ‘five applies’ implies ‘not more.’ He doesn’t say ‘there are five apples in the basket’ when there are ten, even though it’s true that there are five apples in the basket when there are ten. The hearer, on the other end of the coordination, hears ‘five apples’ and infers ‘and not more’. In simplest terms, the speaker doesn’t use a scalar lower than the highest justified. He doesn’t say ‘three apples’ when there are five, for example. The hearer, on his side, doesn’t think ‘he said five apples, so there might be ten;’ he doesn’t act on a scalar higher than that conveyed.

Whether or not it fully qualifies as implicature, this kind of coordination between speaker and hearer *can* emerge even with agents as simple as those used here. That is the ‘Yes’ of the ‘yes and no’. This aspect of scalar implicature seems to be simpler than many of the things in the Gricean mixed bag; interestingly, it was also scalar implicature that Gerhard Jäger explored using the very different (and non-spatialized) ‘best response’ dynamic (Jäger 2007).

A further simulation shows the co-emergence of coordinated speaker and hearer behavior along the lines of scalar implicature. The environment is again one of randomized spatialized events, but the events differ in severity: there are storms of grades 1, 2, 3, and 4, for example. If you as an agent are prepared for a storm of grade 2, you are prepared for a storm of either grade 1 or 2, but you are still open to damage if hit with a 3 or 4. If you are prepared for a storm grade 3 and you only get hit with a 1, on the other hand, you have wasted precious resources—the effort expended in needless preparation for levels 2 and 3 (Figure 17).

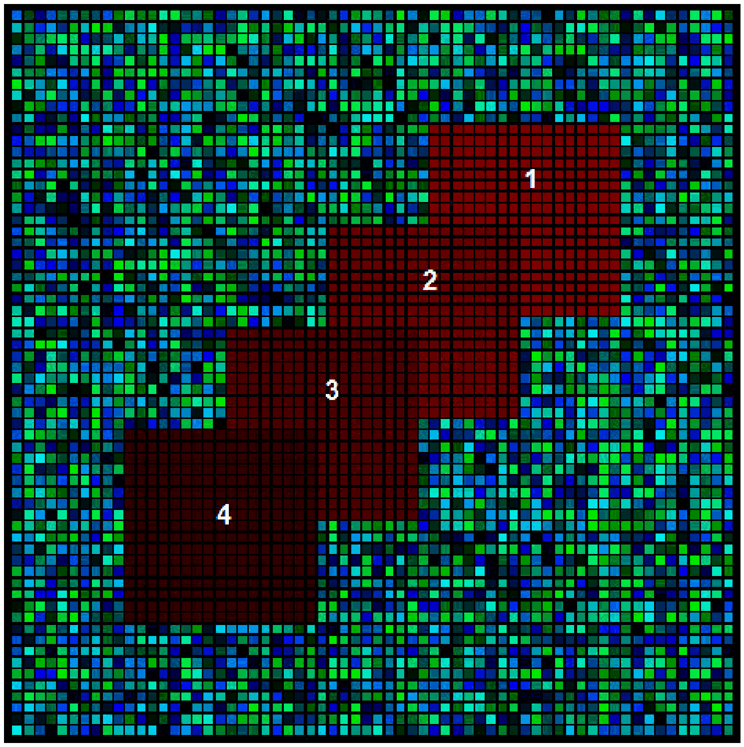

Fig. 17 Spatialized events of scalar ordered severity: storms of severity 1, 2, 3, and 4, for example.

The ideal in such an environment is to act in a way precisely appropriate to the storm: to act 1 and 2 in a 2 environment; 1, 2 and 3 in a 3 environment. In this simulation you as an agent get a positive point for every degree you are correct about—a positive point for both 1 and 2 if you act '2' in a two environment, for example. You get a negative point for every degree you are wrong about—a negative point for acting 3 if the storm is only a 2, for example. If the storm is a 4 and you only act 1, you get a positive point for your '1' but negative points for the 2, 3, and 4 that you missed. Here as before, only 10% of the population on any round can directly observe the relevant environment—hence the need for communication.

In modeling the emergence of pragmatic conventions, we started with an established semantics for red, magenta, and the like. Individuals who heard 'red' automatically acted red. Here we start with an established scalar semantics: If you hear '2' you act both 2 and 1. If you hear 3 you will act 3, 2, and 1. We also start with established pragmatic conventions of the forms evolved above. Speakers who observe a storm always send a signal, they send a signal regarding the immediate environment, and—most important for our purposes—they never send a signal that isn't true.

Those specifications still leave something important open—something about scalars. Our agents are both speakers and hearers. As speakers they never speak anything but the truth, but we start with an initial population in which speakers *may* say something that is true but does not convey the full scalar height. Storms of degree 4 are also storms of degree 3, 2, and 1. As speakers, therefore, our agents start with a probability interval between 1 and 10 of giving a signal less than the full observed scalar. As hearers, our agents always act 1 and 2 on hearing the signal '2'. But they may also act on a scalar higher than that heard: they may act not only 1 and

2 but 3, for example. As hearers, our agents start with a probability interval between 1 and 10 of acting not only on what they hear but on a scalar higher than that they hear.

In our initial array of randomized probabilities, then, a scalar signal does not necessarily indicate that the truth isn't higher. A heard scalar doesn't entail that you shouldn't act higher. There is as yet no scalar coordination between speakers and hearers. Indeed even within each agent there need be no coordination between speaking behavior and hearing behavior. We start with a situation in which the coordination characteristic of scalar implicature and inference doesn't yet exist.

Our agents code speaker probability in one color scale (blue, in the color version at the url for this chapter), with lighter shades indicating a lower probability of signaling below the observed scalar. Agents' hearer probabilities are coded in another color scale (green), with lighter shades indicating a lower probability of acting above the scalar message received. As before, all gains are purely individual. After 100 rounds of gains and losses, individuals see if any neighbor has done better in terms of over-all score. If so, they imitate just one of that agent's features, at random: that agent's probability of giving a scalar less than the observed reality, or that agent's probability of acting on a scalar higher than that heard.

Figures 18 and 19 show the evolution of the array. In such an environment, speakers *could* routinely understate the severity of observed storms. Hearers *could* routinely act above a signal that was likely to be low, following a general strategy of over-reacting to assumed under-statement. But that is not the direction in which this kind of array evolves. It evolves systematically to a coordination between speakers and hearers in which speakers have the lowest probability of sending a signal less than the observed scalar height, and in which hearers have the lowest probability of acting as if what they heard might be less than the full scalar height.

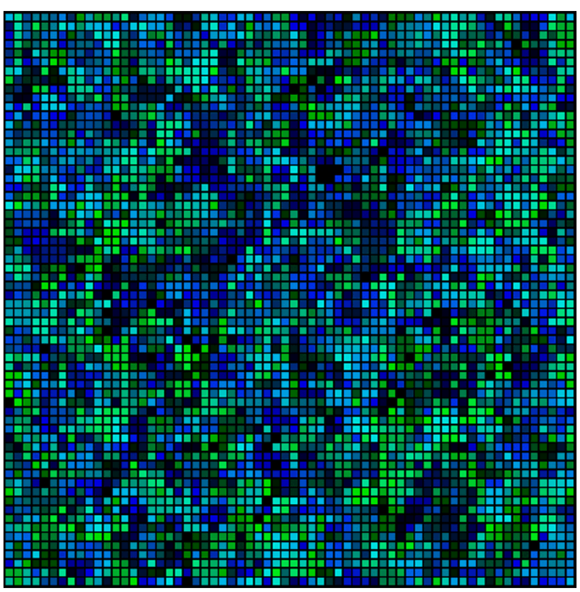

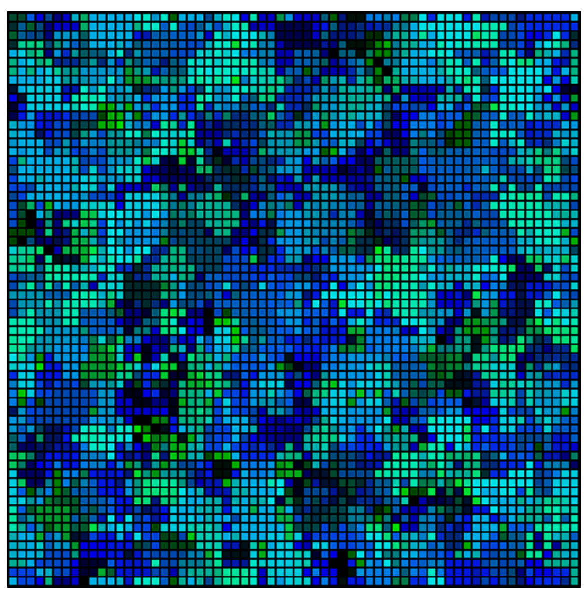

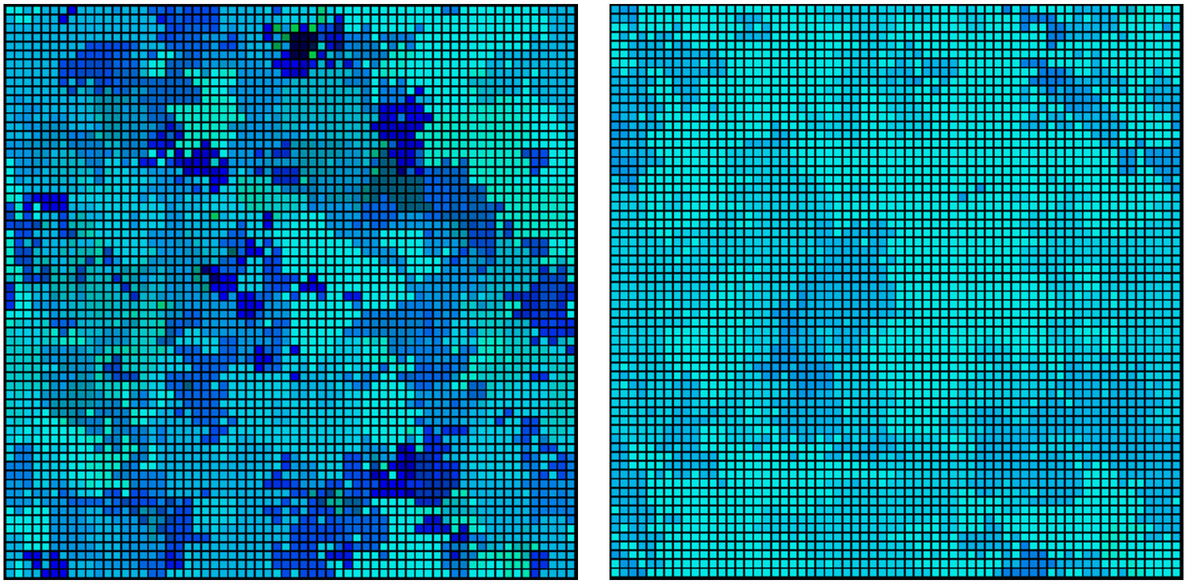

Figure 18 Emergence of scalar coordination between speakers and hearers over 200 generations. Speaker's tendency to understate scalars is in darker green, hearer's tendency to act on higher scalar than announced is in darker blue, combined here as darker gray. Generations 2, 11, 50 and 200 shown.

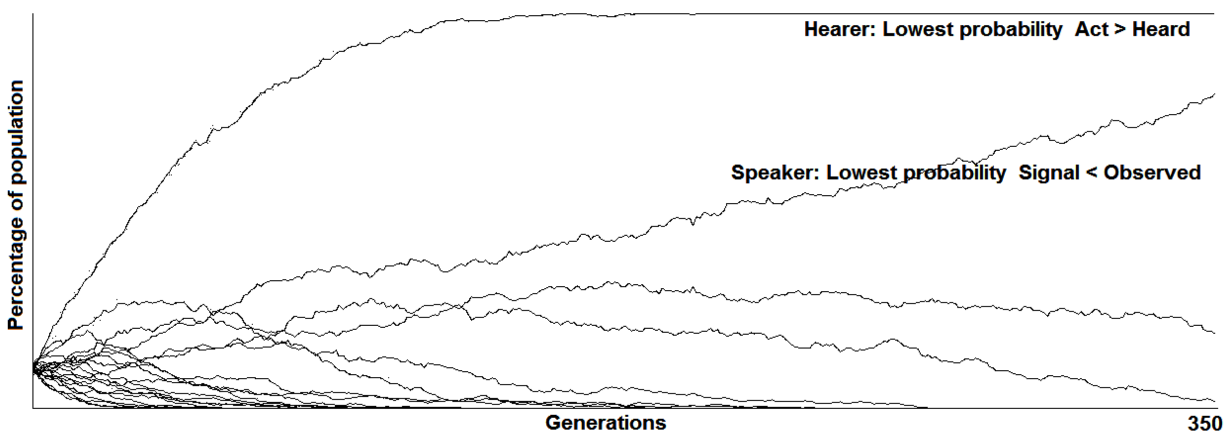

Figure 19 Emergence of scalar coordination between speakers and hearers over 350 generations.

As indicated above, it is unclear whether this qualifies as Gricean implicature in the full sense. It does show, however, that extremely simple speakers and hearers in information networks can develop precisely the communicative coordination that is characteristic of scalar implicature and inference. It is interesting to note, by the way, that speakers and hearers do not converge to this coordination at the same rate. It is the hearers who converge first, learning not to over-react. In large part because of the dynamics of payoffs, speakers are significantly slower in learning not to under-state.

VI. Conclusion

What simulations suggest is that maximization in spatialized information networks, even networks of very simple agents, is sufficient to produce aspects of semantics, of pragmatics, and even coordinated communication behaviors suggestive of scalar implicature and inference. Some of the core phenomena at issue in semantics, pragmatics, and implicature appear in the fundamental dynamics of information maximization, even for agents far below the cognitive level of those that appear in Grice.

Even if this captures some of the core phenomena at issue in pragmatics, it cannot be said to capture it all. There are things in the mixed bag of Gricean pragmatics that seem to *demand* a higher cognitive level—the level of agents that can explicitly recognize logical inferences, that can explicitly reflect on conventions as conventions, or that can cognitively model other agents. Those things will not appear in networks of agents this simple

Those other aspects might appear, however, with a similar dynamics, in spatialized networks of more complex agents. All of the phenomena illustrated here operate in terms of individual information maximization across spatialized arrays of communicative agents. It is to be expected, I think, that information maximization of that kind will use *all* the tools available to it. Information maximization in arrays of simple agents exploits the full cognitive abilities of those agents, even where—as here—those cognitive abilities are severely limited. Information maximization in arrays of more complex agents will have a wider range of cognitive abilities to exploit, and it is predictable that it will do so.

The next step along these lines would therefore be an exploration of the dynamics of maximization in information networks involving more complex agents. That could tell us whether a similar dynamics in the context of more complex agents might be sufficient for further aspects of pragmatics. If particular aspects of language use will appear in networks of agents only where those agents have particular cognitive abilities, such a research trajectory would also give us a new typology and a new understanding of different communicative phenomena: an understanding in terms of the level of networked cognitive abilities that different communicative phenomena require.

References

Asher, N., Sher, I., & Williams M., 2002. "Game Theoretic Foundations for Gricean Constraints" in R. v. Rooij. And M. Stokhof, eds., *Proceedings of the 13th Amsterdam Colloquium*, Institute for Logic, Language and Computation, Universiteit van Amsterdam.

Asher, N., & Lascarides, A., 2003. *Logics of Conversation*. Cambridge: Cambridge University Press.

Cangelosi, A., & 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.

Di Chio, C., & Di Chio, P., 2007. "Evolutionary Models of Language" in A-V. Pietarinen, ed., Game Theory and Lingusitic Meaning, Amsterdam: Elsevier.

Fausett, L., 1994. *Fundamentals of Neural Networks*. Upper Saddle River N.J.: Prentice Hall.

Grice, P., 1998. *Studies in the Way of Words*. Cambridge, MA: Harvard University Press.

Grim, P.,1995. "The Greater Generosity of 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. 2009a. "Philosophical Implications of Interaction and Information Networks," Evolution, Game Theory & the Social Contract Conference, Beckman Center for the National Academies, University of California, Irvine, March 2009.

Grim, P. 2009b. "Network Simulations and Their Philosophical Implications: Models for Semantics, Pragmatics, and Epistemology," Models and Simulations 3, Charlottesville Virginia, March 2009.

Grim, P., Mar, G., and St. Denis, P., 1998. *The Philosophical Computer: Exploratory Essays in Philosophical Computer Modeling*. Cambridge, MA: MIT Press.

Grim, P., Kokalis, T., Alai-Tafti, A., & Kilb, N. 2000. "Evolution of Communication in Perfect and Imperfect Worlds, *World Futures: The Journal of General Evolution* 56, 179-197.

Grim, P., Kokalis, T., Alai-Tafti, A., & Kilb, N. 2001. "Evolution of Communication with a Spatialized Genetic Algorithm," *Evolution of Communication* 3, 105-134.

Grim, P., St. Denis, P., & Kokalis, T., 2002. "Learning to Communicate: The Emergence of Signaling in Spatialized Arrays of Neural Nets," *Adaptive Behavior* 10, 45-70.

Grim, P., Kokalis, T., Alai-Tafti, A., & Kilb, N, and St. Denis, P., 2004. "Making Meaning Happen," *Journal for Experiemntal and Theoretical Artificial Intelligence* 16, 209-244.

Grim,P., Wardach, S., and Beltrani, V., "Location, Location, Location: The Importance of Spatialization in Modeling Cooperation andCommunication," *Interaction Studies: Social Behavior and Communication in Biological and Artificial Systems* 7, 43-78.

Jäger, Gerhard, 2007. "Game Dynamics Connects Semantics and Pragmatics." In A-V. Pietarinen, ed., Game Theory and Lingusitic Meaning, Amsterdam: Elsevier.

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.

Nowak, M., Plotkin, J., & 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.

Nowak, M., & Sigmund, K., 1990. "The evolution of stochastic strategies in the prisoner's dilemma." *Acta Applicandae Mathematicae* 20, 247-265.

Nowak, M., & Sigmund, K., 1992. "Tit for tat in heterogeneous populations." *Nature* 355, 250-252.

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.

Parikh, P., 2001. *The Use of Language*. Stanford: CSLI Publications.

Parikh, P., 2006. "Pragmatics and Games of Partial Information," in A., Benz, G. Jäger, and R. v. R., *Game Theory and Pragmatics*, New York: Palgrave Macmillan.

Parikh, P., & Clark, R., "An Introduction to Equilibrium Semantics for Natural Language," in A-V. Pietarinen, ed., *Game Theory and Lingusitic Meaning*, Amsterdam: Elsevier.

Notes

[1] My research collaborators in this work include Peter Ludlow, Dustin Locke, and Aaron Bramson at the University of Michigan and Brian Anderson at Stony Brook.

[2] This first example employs synchronous updating. Each of our agents simultaneously calculates whether any neighbor has done better, if so updating to the strategy of the most successful neighbor. Results are much the same, however, if we randomly update only a small proportion of cells at each turn. See Grim, Wardach, and Beltrani 2006.

[3] Although spatialized lattices of this form are my main tool throughout, it is clear that a study of the dynamics of interaction across the entire range of possible network structures is needed. With regard to cooperation in particular, it turns out that the results outlined here generalize nicely across network structures. See Grim 2009a, 2009b.

[4] Here again the ultimate goal would be an understanding of communication across various networks.

[5] Accessible but technically complete introductions to neural nets are unfortunately difficult to find. Perhaps the best is Fausett, 1994. We are grateful to Laurene Fausett for personal correspondence regarding construction of the model.

[6] A 'two-lobe' structure for signaling, it turns out, has been invented or reinvented numerous times in the literature. See de Saussure 1916; MacLennan 1991; Oliphant & Batali 1997; Cangelosi & Parisi 1998; Nowak, Plotkin & Krakauer, 1999; Nowak, Plotkin, and Jansen 2000.

[7] In a second series of studies we used more complex neural nets, capable of handling the full range of Booleans, with partial training by backpropagation. The results are just as strong. See Grim, St. Denis, Kokalis, 2002 and Grim, Kokalis, Alai-Tafti, Kilb & St. Denis, 2004.