Interaction Studies: Social Behavior and Communication in Artificial Systems
7 (2006), 43-78

Location, Location, Location:

The Importance of Spatialization in Modeling Cooperation and Communication

Patrick Grim, Stephanie Wardach, and Vincent Beltrani
Group for Logic & Formal Semantics
Dept. of Philosophy
SUNY at Stony Brook

Abstract

Most current modeling for evolution of communication still underplays or ignores the role of local action in spatialized environments: the fact that it is immediate neighbors with which one tends to communicate, and from whom one learns strategies or conventions of communication. Only now are the lessons of spatialization being learned in a related field: game-theoretic models for cooperation. In work on altruism, on the other hand, the role of spatial organization has long been recognized under the term 'viscosity'.

Here we offer some simple simulations that dramatize the importance of spatialization for studies of both cooperation and communication, in each case contrasting (a) a model dynamics in which strategy change proceeds globally, and (b) a spatialized model dynamics in which interaction and strategy change both operate purely locally. Local action in a spatialized model clearly favors the emergence of cooperation. In the case of communication, spatialized models allow communication to arise and flourish where the global dynamics more typical in the literature make it impossible.

Simulations make a dramatic case for spatialized modeling, but analysis proves difficult.

In a final section we outline some of the surprises of spatial dynamics but also some of the complexity facing attempts at deeper analysis.

Keywords: evolution of communication, cooperation, altruism, spatialization

1. Introduction

Space plays a major role in natural communication as we know it. It is spatial neighbors with which one tends to communicate, and from whom one learns strategies and conventions of communication. Most current computational models regarding the evolution or emergence of communication, however, ignore or underplay the significance of local action in spatialized environments.

Early models in related areas—altruism and cooperation—also neglected the role of space. The models from population genetics that were initially applied to the evolutionary question of altruism were non-spatial, though with time spatial effects have been firmly established under the term 'viscosity.' Classical game-theoretic models for cooperation are non-spatial, though there is now increasing recognition that spatial organization has a major role to play here as well. In studies of communication, on the other hand, the importance of spatialization has yet to be fully recognized.

In what follows we want to underscore the importance of spatialization in studies of altruism and cooperation and want to extend that lesson to modeling for communication. There are three primary features of importance, we want to argue, in understanding evolution or emergence of both cooperation and communication. Those features are location, location, and location.

We first offer some simple but graphic results regarding the emergence and spread of game-theoretic cooperation under two different conditions: (a) a model dynamic in which strategy change proceeds globally, and (b) a rival model dynamic in which interaction and strategy change both operate purely locally. In this comparison, local or spatialized dynamics clearly favor the emergence of cooperation. We then carry the lesson over to models for communication by showing a similar difference regarding the emergence and spread of simple signaling in (a) global and (b) spatialized models. Localizing action and strategy change allows communication to arise and flourish where the global dynamics more typical in the literature make it impossible.

Our primary goal, then, is simply to call attention to space as a major but neglected factor in modeling both cooperation and communication. Beyond simulation, it is clear that we will want an analytic understanding of the spatial dynamics at issue. In a final section we display some of the surprising phenomena of spatial interaction in general and introduce a few tools for analysis, but also illustrate the complexity that quickly swamps any such attempt.

2. 'Altruism' and 'Cooperation': The Tangled Background of Two Traditions

'Altruism' and 'cooperation' seem very close in meaning, and thus the uninitiated might expect a single tradition of modeling regarding the two. The fact of the matter, however, is that 'altruism' and 'cooperation' are better seen as code-words for two content-related but often methodologically distinct modeling traditions.

2.1 Altruism

'Altruism' carries a well-established biological definition: altruistic traits are those that decrease an individual's fitness and increase the fitness of others (Mitteldorf and Wilson 2000, 481; Sober and Wilson 1998, 17; Van Baalen and Rand 1998, 631; Brinkers & den Dulk, 1999, 499).[1] The problem for evolutionary theory is how any such traits could evolve.

The classical tools for evolutionary theory were those of population genetics, pioneered by Fisher (1930). The basic model, for differential growth of genetic types within a population, brings with it assumptions of large population size, global interaction, and global reproduction. Against this background, one mechanism proposed for evolution of altruism was 'inclusive fitness' (Hamilton 1963, 1964), later dubbed 'kin selection' by Maynard Smith (1964). Altruistic traits can evolve if they further the survival or fecundity of genetically related individuals, even if those traits reduce the survival or fecundity of the altruistic individual. A second explanation proposed for evolution of altruism was 'group selection,' which continues today under the name 'multi-level selection' and in the specific sense of a selective advantage for members of spatially separated subpopulations that contain a higher percentage of altruists (Sober and Wilson 1998).

Controversy remains as to the relative importance of kin selection and group selection. What is clear is that both theories work by compromising the global continuity of the classical model. On both theories, it is precisely because interaction and reproduction are *not* uniformly global that altruism can arise. The basic mechanism of each theory, moreover, tends to be conceived in spatial terms. The crucial mechanism of group selection is a separation of groups, usually envisaged as a spatial separation. (Also important, of course, is some mode of dispersal or territorial expansion.) Hamilton (1964) proposed a similar mechanism for kin selection in terms of a ‘viscous’ population—one with limited dispersal. If action is spatially restricted and offspring remain close to parents and kin, he reasoned, spatially restricted altruism can evolve because it will tend to benefit relatives.

It is now recognized on the basis of a broad range of simulation studies that ‘viscosity’ can play a strong positive role in the evolution of altruism (Wilson, Pollack, and Dugatkin, 1992; Mitteldorf and Wilson, 2000; van Baalen and Rand, 1998; Irwin and Taylor, 2001; Taylor and Irwin, 2000; Di Paolo 2000; Goodnight 1992; Durrett and Levin, 1994a, 1994b; Krakauer and Pagel 1995; Nakamaru, Matsuda, and Iwasa, 1997; Nakamaru, Nogami and Iwasa, 1998). There is a catch, however: in these models, clusters of altruists also prove more vulnerable to stiff local competition. Most of these models incorporate features which allow the benefits of spatialized altruism while compensating for this vulnerability. The clear role established for viscosity in general may indeed suggest that the dispute between group and kin selection is ultimately an empty one; that it is viscosity that is doing the real work, whether that dynamic is conceived in terms of kin or group (van Baalen and Rand, 1998; Grafen 1984; Quellar 1994; Di Paolo 1999; Pepper and Smuts 2000a, 2000b). Biological instances of some of the theoretical mechanisms at issue have been claimed for breeding habits of birds, levels of aggression in male insects competing for mates, virulence in diseases, and sex ratios in insects (Sober and Wilson, 1998; West, Pen and Griffin, 2002; Griffin and West, 2002).

Under the title of ‘viscosity’, then, the importance of spatial effects in evolution of altruism seems to be well established. This victory in understanding comes with an analytic cost, however; it turns out that the dynamics of viscous populations are much more difficult to analyze than is the continuous development of a globally homogenous population. As long as one is dealing with population genetics for a 'randomly mixed' population, the mathematics are linear and easy. When one builds a model realistic enough to capture the spatial effects of altruism, on the other hand, the math becomes difficult or we find that we lack the necessary analytic tools entirely. Reliance on computer simulation increases.

2.2 Cooperation

'Altruism' as a key word tends to signal the tradition outlined above. 'Cooperation', on the other hand, tends to bring up an importantly different set of models. For 'cooperation', the the ur-reference—that to which all others trace—is Axelrod (1981).

Here, interestingly enough, the tradition *begins* with computer simulations. In 1980, Robert Axelrod announced a computerized tournament for the Iterated Prisoner's Dilemma (Axelrod 1980a, 1980b). Participants were invited to submit any strategy they wished, no matter how complicated, as long as it could be coded in Fortran. Submitted strategies were pitted in a 'round robin' competition against each other, a random strategy, and themselves, with that strategy which amassed the most points over all declared as the 'winner'. In Axelrod's first tournament, the winner was Tit for Tat (TFT). Axelrod announced the results and called a second tournament. The winner was again TFT. From that humble beginning flows a river of literature on 'cooperation'.

This tradition, unlike the 'altruism' tradition, has always been explicitly game-theoretic. The payoff matrix is that of the Prisoner's Dilemma, with a rich variety of strategies at stake in the iterated game. Although TFT is taken as a paradigmatic 'cooperative' strategy, moreover, there is no strict definition of 'cooperation' akin to that offered by biologists for 'altruism'.

There is, in particular, no demand that a 'cooperative' strategy be one which promotes the fitness of its competitors at its own expense.[2] In general, the 'cooperation' tradition has conceived of its subject matter as a range of behaviors more contextual and more complex than any simple dichotomy between altruism and selfishness. Those subtleties, it has been thought, are appropriate for the exploration of a social realm of cooperation and competition in ways that the simpler categories of the biological model might not be.[3] It is thus the game-theoretic 'cooperation' tradition that has been pursued in experimental economics and psychology (Cooper 1996, Segal and Hershberger, 1999, Clark and Sefton, 2001; Wedekind and Milinski 1996).

Like the initial models in the 'altruism' tradition, the classical models in the 'cooperation' tradition are global. Axelrod's initial tournaments are 'round-robin' competitions in which each player plays all others in the strategy pool. Axelrod and Hamilton's 'ecological' model uses replicator dynamics; each player plays against all others, weighted by their proportional representation in the pool, and strategies then reproduce selectively as a function of their comparative success (Axelrod and Hamilton, 1981; Axelrod 1984, 1985). In a more recent variation using genetic algorithms, Axelrod continues to play all strategies against all others in the pool, with relatively successful individuals selected to have more offspring (Axelrod 1997.[4] Martin Nowak and Karl Sigmund have attempted to increase the realism of the Axelrod-Hamilton model by adding stochastic 'noise' or 'grit' (Nowak 1990, Nowak and Sigmund, 1992). Their algorithms for action and strategy-change, however, still follow the Axelrod and Hamilton model and thus remain thoroughly global.

Occasionally there have been attempts to make contact across the 'altruism' and 'cooperation' traditions. Given the definitional strictures on 'altruism', however, this has generally occurred by eliminating most of the game-theoretical interest, cutting the Prisoner's Dilemma strategies at issue to only two (Nakamaru, Matsuda, and Iwasa 1997, Nakamaru, Nogami and Iwasa, 1998, Epstein 1998, Sella and Lachmann 2000; Vainstein and Arenzon, 2001). A population composed purely of All-C and All-D, for example, can be used to model a population of pure 'altruists' and the purely 'selfish'.[5]

The first use of spatialization in the 'cooperation' tradition employs precisely this kind of simplification. Nowak and May (1992, 1993) introduced a two-dimensional cellular automata model using only individuals who simply cooperate C or who simply defect D. [6] On each round each cell plays its eight neighbors (and, somewhat peculiarly, itself) and totals its scores. Each cell then adopts the strategy of its most successful neighbor. Nowak and May reported that "spatial effects can change the outcome of frequency dependent selection" (35); in particular, C is able to evolve as a significant player in a spatialized environment where a global algorithm leads to quick dominance by D. Nowak and May also offered a gallery of 'dynamical fractals' and 'evolutionary kaleidoscopes,' showing the complexities of spatial evolution for arrays with a variety of starting configurations and relative payoffs.

This first spatialized model was given a cold reception in Huberman and Glance (1993). Nowak and May=s algorithm employed an assumption of simultaneous updating for all cells in the array, common in work using cellular automata. It is clear that the symmetry of Nowak and May=s 'kaleidoscopes' depends crucially on that assumption; without simultaneous updating the symmetry is quickly broken. What Huberman and Glance suggested was that greater prospects for cooperation in a spatialized array might be artifactual consequences of simultaneous updating as well. Repeating Nowak and May's computer experiments but without synchronous updating, Huberman and Glance found results "which do not correspond at all to the behavior found for synchronous updating," and which thus "cast into doubt the conclusions recently obtained concerning territoriality and the universality of long-term averages of cooperation" (7716). An early move toward spatialization was stalled.

Oliphant (1994) countered Huberman and Glance's criticism, showing that any claim that Nowak and May's general results depended on synchronized updating were overstated. Oliphant's model was again restricted to simple strategies of C and D, this time organized in a one-dimensional wrap-around ring, but with asynchronous updating in Huberman and Glance's preferred sense.[7] Oliphant finds a wide range of behavior in such a model, but notes that spatialization favors cooperation even without simultaneous updating: "While non-spatial populations quickly fall into defection, spatial populations are able to evolve and maintain cooperative behavior" (351).

Lindgren and Nordahl (1994) take spatialization further in considering a lattice of strategies with a range of memory depths. They conclude that spatial dynamics allows for diverse communities using either synchronous or asynchronous updating, including forms of coexistence between cooperative and non-cooperative strategies. With a brief treatment of the updating issue, our own previous work (Grim 1995, 1996) uses a spatialization of stochastic strategies, showing success of more generous strategies in a spatialized environment. Brauchli, Killingback, and Doebeli (1999) use an extended set of stochastic strategies, showing greater success for generous forms of Pavlov in a spatialized environment.

There have also been more recent indications that spatialization has an important role to play in game theory. Most of the modeling in Danielson (2001), regarding reciprocity and cooperation, is non-spatial. In a final section, however, he suggests that spatial effects—specifically, the mechanism of imitating one=s most successful neighbor in a one-dimensional cellular automata—may be important for the spread of cooperation.[8] William Harms (2000) introduces an 'agent-patch' model, which distributes a population of cooperator and defector agents across a spatialized grid of 'patches'. Agents interact randomly with other agents on the same patch, reproduce when they accumulate enough 'fitness points,' and die when their points fall to zero. Harms identifies population viscosity—the fact that kin are born and remain in the same basic location—as one factor that can result in increased gains and thus population growth for cooperators; viscosity allows cooperators to benefit from interaction with their own kind and to avoid interacting with defectors.[9] Although the models in Skyrms (1996) tend not to be spatial, he occasionally mentions viscosity as a feature of potential interest.[10] More recently, Skyrms has argued that the Stag Hunt may be a more appropriate game-theoretic model for studying cooperation than the standard Prisoner=s Dilemma,[11] but finds that local action alone is insufficient to generate cooperation in the Stag Hunt (Skyrms 2001; see also Skyrms and Pemantle, 2000). Cooperators do prosper, however, if the model also includes the possibility of reinforced patterns of interaction, allowing stag hunters to seek out other stag hunters with which to interact, and if it incorporates localized strategy change in which an agent imitates successful neighbors. "Here, we finally have a model that can explain the institution of a modest social contract" (Skyrms 2001, 38).

Work on 'altruism' and 'cooperation', then, falls into two related but distinct research traditions. Within the biological 'altruism' tradition the importance of spatialization is now clearly recognized under the category of 'viscosity'. Despite occasional indications, the importance of spatialization is not yet as clearly recognized in the socially more suggestive work on game-theoretic cooperation. In what follows we want to offer some graphic examples that underscore the importance of space.

3. Spatialized Models for Cooperation

The classical models of Axelrod, Axelrod and Hamilton, and Nowak and Sigmund, though global with respect to both action and strategy change, do show emergence of cooperation. The more recent suggestions in Nowak and May, Harms, Danielson, and Skyrms are that spatialization further favors cooperation: that only with spatialization does cooperation appear in certain contexts (Skyrms 2001; Nowak and May 1992, 1993), or that cooperation appears in a different and more robust way in spatialized environments (Danielson 2001, Harms 2000). Our results further support these more recent suggestions. Here we offer a simple illustration of the dramatic difference that simple spatialization—in particular, spatialized reproduction or strategy change—can make for the appearance of cooperation.

Consider first an array of the 8 elementary ‘reactive’ strategies in a Prisoner=s Dilemma: those which use only the opponent=s play on the previous round as their input for choice on the current round. Unlike Nowak and May (1992, 1993) and Harms (2000), it should be noted, we employ a range of strategies beyond All-C and All-D. We instantiate our 8 strategies at random in cells of a two-dimensional wrap-around or torroidal array. Each cell plays an iterated Prisoner’s Dilemma game of 200 rounds with only its immediate neighbors—those eight cells touching it at sides or diagonally. At that point each cell totals its points, and the most successful cell on the board then ‘reproduces’ into random sites across the display. Here 5% of sites are replaced across the board with that strategy that has proven most successful in its local action. In this first model, then, competitive action is local—each cell plays an iterated Prisoner=s Dilemma with only its immediate neighbors—but reproduction is global. Cells again play an iterated Prisoner's Dilemma with their new neighbors, and the cycle is repeated.

Global reproduction of this sort shows a clear and complete conquest by All-D. Cooperation, in the form of TFT, makes no showing at all. Typical evolution of a randomized array is shown in Figure 1.

Let us now localize strategy replacement as well as strategy interaction. In this variation of the model, each cell identifies its most successful neighbor. If that neighbor has a higher score than its own, the cell adopts that neighbor's strategy.[12] There is no guarantee, of course, that the most successful neighbor of a given cell is the most successful strategy on the board as a whole: this is a model in which strategy replacement, like strategy interaction, proceeds purely locally.[13]

Here as in the global model All-D shows early gains. Defectors initially gain by exploiting neighboring 'sucker strategies.' Although All-D does well in that environment, however, it does very poorly in play against itself, gaining only 1 point in each exchange. Tit for Tat, on the other hand, does well against itself, chalking up 3 points per exchange. Once All-D occupies the bulk of the field, then, clusters of Tit for Tat start to grow, thriving in groups in which members benefit from mutual cooperation. Neighboring cells convert to the higher-scoring Tit for Tat, which eventually conquers the full array (Figure 2). Together, localized action and localized reproduction in a spatialized model favor the growth of cooperation.

As noted above, Huberman and Glance (1993) suggested that the spatialization results favoring altruism in Nowak and May (1993) might be an artifact of simultaneous updating. It is thus interesting to repeat this exercise without simultaneous updating. Figure 3 shows typical evolution of a spatialized array in which only 1% of cells in the array, chosen randomly, update in any generation. It is clear that here at least it is the use of local action and reproduction in a spatialized model, rather than simultaneous as opposed to non-simultaneous updating, that is responsible for the triumph of cooperation.

Still working in a global model, Nowak and Sigmund showed that changing to a range of imperfect game-theoretic strategies increases the level of generosity (Nowak 1990, Nowak and Sigmund, 1992). In a world of stochastic noise, the ultimate winner is 'Generous Tit for Tat,' more willing to forgive defection against it than standard TFT. If you cooperate with GTFT, its probability of cooperating with you is close to 1; if you defect against GTFT, however, it will forgive that defection with a probability of 1/3. In previous work, we have found that spatialization increases the role of generosity still further: in a spatialized version of Nowak and Sigmund's stochastic model, the optimal strategy turns out to be one that forgives defection with a probability of 2/3, twice that of Nowak and Sigmund's GTFT (Grim 1995, 1996).

4. Spatialization and the Evolution of Communication

Although the classical models for emergence of cooperation are global, emergence of cooperation is strongly favored in a spatialized environment. We want to suggest that the same is true for communication.

Previous work in modeling communication can be divided into (1) entirely non-spatialized models and (2) partially spatialized models—usually with a spatialized task but without spatialized reproduction. The models we offer here use a more thorough spatialization of both task and reproduction.

4.1 Non-spatialized models

Entirely non-spatialized models appear in MacLennan (1991), MacLennan and Burghart (1994), Noble and Cliff (1996), Levin (1995), Batali (1995), and Oliphant (1999).[14] MacLennan and Burghardt's 'synthetic ethology' uses finite-state 'simorgs' which emit signals into a shared environment. Communication is taken to occur when one organism reacts appropriately to the signal of the simorg who last emitted; at the end of 10 cycles, the two organisms with the highest accumulated fitness breed a single simorg. Neither tasks nor strategy change are intended to be spatialized. Noble and Cliff (1996) offer a replication of the MacLennan and Burghart model in which they attempt to remove any remaining aspects of spatialization, and claim an increase in rate of fitness.[15] Levin (1995) simulates a population of agents with internal states and external observable qualities, where each individual's genome determines the mapping of internal to external observables. The top percentage of individuals successful in decoding other agents' external observables are bred by genetic recombination into the global population. In Batali (1995), the 'receive maps' of each individual are compared to the 'send maps' of the rest of the population. The fittest of the population is chosen to reproduce and a new individual is created with send and receive maps identical to its parent. The model in Oliphant (1999) is also entirely non-spatial, despite his earlier advocacy for spatialized models in game theory. Here Oliphant uses a form of Hebbian learning to produce meaning-signaling correlations across a population of associative nets. Single 'learners' train globally on the behavior of all members of a population of 'behavers.' They are then added to the global pool of 'behavers', another member is removed at random, and the cycle is repeated with a new ‘learner’.

4.2 Partially spatialized models

Some studies employ a spatialized task, but without spatialized reproduction or strategy change.

Hutchins and Hazelhurst (1995) is a borderline case, showing just a hint of task spatialization. At each time interval, two autoassociator networks—a ‘speaker’ and a ‘listener’—are randomly chosen from the population. Exposed to a random choice of 12 phases of the moon, the speaker produces a verbal output and the listener ‘corrects’ what it would say about the scene. Hutchins and Hazelhurst show that such a population can converge on a shared lexicon. In the course of their simulations, however, they attempt to make communication easier by biasing individuals toward choosing a listener with a similar speaking history, thereby reducing the randomness of interactions and offering a glimmer of spatialization.

The classical model for evolution of communication involving a spatialized task is that of Werner and Dyer (1991). The females in their population cannot move, but have an ability to ‘see’ males a short distance away. When a female senses a nearby male, she emits a signal; though males are blind, they respond to signals ‘heard’ from females. On mating, two offspring (one male and one female) are produced whose genomes are formed through genetic crossover and mutation of their parents’ genomes. To this point Werner and Dyer’s model is nicely spatialized. After reproduction, however, the parents are moved to a random placement elsewhere in the grid. Although the mating task in the model is conceived spatially, therefore, reproduction results in random relocation in the global array: action but not reproduction is spatialized, and Werner and Dyer admit that this global randomization may have hindered speciation in the population.

Use of a spatialized task without spatialized reproduction also characterizes Saunders and Pollack (1996), Cangelosi and Parisi (1998), and Wagner (2000). Saunders and Pollack (1996) use a population of neural network agents and a cooperative task of resource consumption. No agent can consume the resource by itself; for that it must recruit others. Fitness is rated on how much of the resource was consumed; those with high fitness are not mutated in the next generation, while those with lose fitness are. Cangelosi and Parisi (1998) have members of a community of neural nets communicate in the process of mushroom-hunting.[16] Despite a clearly spatial task, reproduction in the model remains global: those agents with the highest energy levels each give birth to 5 randomly distributed offspring. Wagner (2000) expands on MacLennan and Burghardt (1994) by adding a spatialized task. Each of Wagner's 'orgs,' coded by a genome, acquires fitness by gaining resources, but at least one other organism must be present in a sector containing resources; at that point both benefit. Mate selection, however, remains non-spatial; mates are chosen based on fitness levels across the population as a whole, and the population tends to converge to a single genotype as a result.

4.3 Spatialization and reward structure

The most thoroughly spatialized study to date, other than our own previous models, is perhaps that of Ackley and Littman (1994). Theirs is a population of neural networks conceived hierarchically on individual, local, and and global levels. The task is spatial in that individuals move horizontally on their own personal tracks, either toward or away from food and predators placed at the ends. Food and predator placement is the same for all members of a local group, and agents are able to emit signals heard by other members of the group. At one level, at least, reproduction is also spatialized, since a genetic algorithm is applied first to those fittest individuals within each ‘quad’. Unfortunately, the model becomes increasingly complex, with reproduction mechanisms at the global level involving winds and festivals. It is perhaps these complicating factors that lead Batali to conclude that “the details of such simulations make it difficult to assess the significance of the results” (5).[17]

With the exceptionof Ackley and Littman (1994) and our own previous work (Grim, Kokalis, Tafti, and Kilb, 2001; Grim, Kokalis, and St. Denis 2002), thorough spatialization with regard to both task and reproduction has been rare in the communication literature.[18] It should also be noted that in almost all models communication has been 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. As Ackley and Littman note regarding MacLennan (1991), the models that result reflect an artificial environment “where ‘truthful speech= by a speaker and ‘right action= by a listener cause food to rain down on both” (40). In later work, though they continue to use symmetrical reward models, MacLennan and Burghardt (1994) explicitly restrict their claims to the special case of communication regarding cooperative tasks. Here they use a story of communication for cooperation in bringing down a large animal, for example. That restriction of focus also limits any attempt to generalize their results to evolution of communication in general, however. Through explicit assumption or shared tasks, symmetrical rewards for communication also characterize the work of Werner and Dyer (1991), Hutchins and Hazlehurst (1995), Levin (1995), Saunders and Pollack (1996), Noble and Cliff (1996), Cangelosi and Parisi (1998), Oliphant (1999), and Wagner (2000).

The need for a model of how communication regarding *non*-shared tasks might originate is noted explicitly by a number of authors(Dyer (1995), Ackley and Littman (1994), Noble and Cliff(1996), Cangelosi and Parisi (1998). Batali (1995) 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. (2)

In an overview of approaches, Parisi (1997) 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=? (129)

The two studies of the type called for, with no instant reward for the ‘sender,’ are Batali’s own and Ackley and Littman (1994). Batali’s model is entirely non-spatialized, and continues to reward communicative abilities *per se* even if not symmetrically: agents are rewarded and selectively reproduce to the extent that their ‘receiving map’ matches the average ‘sending map’ of the population as a whole. Ackley and Littman do reward functional communication non-symmetrically, but this comes with undue complications in the model. Their results are also plagued by persistent ‘free-riding’ parasites. In a thoroughly spatialized model, we want to suggest, emergence of a naturally functional and non-symmetrically rewarded communication appears easily and without similar threats by parasites.

5. Spatialized Models for Communication

Here as with the case of cooperation we want to show an important contrast between spatialized and non-spatialized models.

We again use a randomized 64 x 64 cellular automata array of individuals carrying different strategies. The individual cells themselves are thought of as stationary, perhaps like coral in a reef. Although each cell carries a strategy, our communication model leaves the Prisoner’s Dilemma behind. Each cell’s strategy dictates those conditions in which it opens its mouth, those conditions in which it hides, and when it makes one of two arbitrary sounds. Our cells are eating-hiding-and-sounding machines, each with a strategy that governs how it manages its behavioral repertoire.

Our cells do not move, but both food sources and predators do, migrating in a random walk across the array.[19] If a cell has its mouth open when a food source lands on it, that individual 'eats' and thus gains a point. The food source does not then disappear, which is why we think of these as food *sources* rather than individual food items. If a predator lands on a cell that is not 'hiding', that individual is 'harmed' and loses one point.

We encode our strategies as four-tuples <*f*, *h*, *s1*, *s2*>, with three possibilities for each variable: an individual's strategy specifies what it 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). In other work we have used more complex strategies (Grim, Kokalis, Tafti and Kilb 2001; Grim, St. Denis, and Kokalis 2002), but here we restrict them to four simple variables. What these give us is a total of 81 strategies, which we can envisage in base 3 using 0, 1, and 2 for 'make no sound', 'make sound s1', and 'make sound s2' for places *f* and *h*, and for 'random selection', 'open mouth', and 'hide' for places *s1* and *s2*. Strategy <0,1,2,1>, for example, makes no sound when fed, sounds S1 when hurt, hides on hearing S1 but opens its mouth on hearing S2.

Among these 81 strategies are precisely two that we consider 'perfect communicators': strategies <1,2,1,2> and <2,1,2,1>. Strategy <1,2,1,2> makes sound S1 when fed and sound S2 when hurt, responding symmetrically to S1 by opening its mouth and to S2 by hiding. Strategy <2,1,2,1> follows the same pattern with sounds S1 and S2 interchanged.

Consider a uniform field of ‘perfect communicators’ <1,2,1,2>. On successfully feeding, a cell of this strategy will make sound S1. Its neighbors, on hearing that sound, will open their mouths. Since food sources migrate cell by cell in a random walk, each neighboring cell will thereby increase its probability of feeding. A similar but slightly more complicated pattern can be expected if neighbors hide in response to sound S2 when a cell is ‘hurt’ by a predator.[20] A uniform field of the other ‘perfect communicator’ strategy <2,1,2,1> would show the same advantages, though with the role of signals S1 and S2 reversed.

Let us begin with an array randomized across 6 different ‘Adam and Eve’ strategies, chosen to avoid any ‘perfect communicators’ but to represent each option (0, 1, or 2) in each of our four variable places. Our environment includes 50 food sources and 100 predators moving in random walks across the array, and our agents gather and lose points by feeding and predation over the course of 100 rounds. At that point we select our two highest-scoring distinct strategies to use in a genetic algorithm. We replace a random 5% of our population with 2-point crossover hybrids of our two top strategies, using a different random choice of crossover points in our parent strategies for each cell replaced. In each generation, therefore, 5% of the array is replaced with (variant) offspring of our two highest strategies. The genetic variability of the array is increased, new strategies may emerge as the highest-scoring, and hybrids from these replace 5% of the population in turn. The result is a fairly standard global genetic algorithm, similar to many used in work regarding communication.

What happens with communication under this form of global algorithm? The short answer is that patterns of communication do *not* emerge. Figure 4 shows percentages for particular strategies over the course of 10,000 generations, beginning with our Adams and Eves.

Given the character of our genetic algorithm (breeding always the two highest-scoring distinct strategies), no strategy can ever go to complete fixation in the array. In fact, competition continues through 10,000 generations primarily between strategies <1,1,1,1>, <1,0,1,1>, <2, 1,1,1> and <2,0,1,1>. None is even close to a perfect communicator. Figure 5 shows an enlargement of the last 2500 generations of the run. In none of our repeated runs, using different randomized seeds and different initial Adams and Eves, did either of our perfect communicators play a significant role.

We also varied the global genetic algorithm so as to 'breed' the two highest-scoring strategies in the array even when these happened to be the same. When they are the same,of course, it is not hybrids but a pure parental strategy that is scattered into 5% of the array. In this variation we were surprised to see that a 'perfect communicator' did significantly better, playing a subsidiary role through approximately the 1400th generation. In the end, though, it could not compete: it is <2,1,1,1> and <2,2,1,1> that dominate (Figure 6), with success to the former in 3221 generations.

Replacing a global genetic algorithm with a local form of reproduction makes a dramatic difference. In the local variation, at the end of every 100 rounds, each cell looks around to see if it has an immediate neighbor that has proven more successful. If so, it 'breeds' with that neighbor and is replaced by a hybrid. The two-point crossover mechanism of genetic algorithm remains the same, but here all reproduction has become purely local.

Figure 7 shows the clear emergence of a perfect communicator using a localized rather than a global genetic algorithm. All other parameters of the model, including the initial randomization pattern and initial set of Adams and Eves, are the same. By 1039 generations, the perfect communicator <2,1,2,1> occupies the entire array. In the spirit of Batali and Parisi quoted above, it should be emphasized that our perfect communicators succeed despite the fact that the model gives no reward for communication *per se*. Our communicators flourish without free riders despite the fact that there no immediate reward for the 'sender' in an exchange. All gains at issue are the 'natural' individual gains of successfully feeding or avoiding predation.

This local genetic algorithm, like our initial cooperation example, employs synchronous updating: All cells simultaneously gauge the success of neighbors and reproduce. In Figure 8 we show a non-synchronous variation as well. Here the localized genetic algorithm is applied only to a random 1% of the array at each generation. The result is slower, but it is clear that success still goes to our perfect communicator.

In the case of communication as in cooperation, then, localized action makes a major and important difference. This similarly holds despite the fact that our two models are otherwise very different. The cooperation model above involves Prisoner's Dilemma games between immediate neighbors and an 'imitate the most successful neighbor' mode of strategy change. Our communication model involves capture of food and avoidance of predators, response to arbitrary signals from immediate neighbors, and reproduction by localized genetic algorithm.

6. Spatial Dynamics

We take the work of the previous sections to establish a simple point: that localized action in space can be of major importance in the emergence of both cooperation and communication. Models that are not spatialized in such a way are ruling out what appears to be a major factor in the emergence of cooperative phenomena in general and communication in particular.

Here we want to offer illustrations of some of the surprises of spatial dynamics, with some first moves toward an analytic understanding.

We simplify by dealing with only two strategies at a time, using only four variables for each pair of strategies: what strategy A gains in interaction with itself, what B gains in interaction with itself, what A gains in interaction with B, and what B gains in interaction with A. These are treated as determinate amounts: neither the environment nor either of our strategies has a stochastic element. This sketch abstracts, moreover, from what the model might be about: A and B might be pitted against each other in the Prisoner's Dilemma or some other familiar game, might or might not be in competition for resources, and might or might not employ communication. Our attempt here is to abstract from some of the particulars of models of cooperation and communication in order to concentrate on spatial dynamics.

We use the 'most successful neighbor' aspect of strategy change employed in both cooperation and communication studies above, here in the form of pure imitation. After competition with each of its immediate neighbors, each cell looks around to see if any neighbor has a higher total score. If so, it copies the strategy of its most successful neighbor.

The first surprise is that an 'inferior' strategy—one that does more poorly playing against itself—can conquer a 'superior' strategy in a spatialized array of this type. Crucial here is the gain each strategy makes in competition with the other at an interface. Consider, for example, a simple array divided between two strategies, as in Figure 9a. Although strategy A may score only one point in play against itself, while strategy B scores 2 points against itself, strategy A can still invade B. Using AB to represent A's score when playing B, this will for example happen when gains <AA, AB, BA, BB> are <1, 4, 2, 2>. Here spatialization alone seems to counter what one might expect in terms of evolution toward fixation. Although a uniform population of B's would have twice as high an average score as a uniform population of A's, it is A's that invade to conquer.

In this case A invades B, despite an inferior score against itself, just because A's score against B at the border is so much higher than B's against A. But it is not the case that relative AB and BA scores at the interface are sufficient for success or failure. If we leave scores otherwise the same but raise B's gain against itself from 2 points to 3, the invasion will be reversed. For gains <AA, AB, BA, BB> at <1, 4, 2, 3>, it is the B that will invade A, as in Figure 9b. The set of scores <1,3,1,2> gives us a standoff, in which neither strategy is able to invade the other (Figure 9c).

For elementary boundary cases of this sort, the algebra is simple. Where AA is A's score against itself and AB its score against B, A will invade B just in case

$$5AA + 3AB > \max[5BB + 3BA, 8\,BB]\ .$$

B will invade A just in case

$$5BB + 3BA > \max[5AA + 3AB, 8\,AA].$$

We will have a standoff, with no invasion, when neither inequality is satisfied. It should be notd, however, that two importantly different strategy relationships can produce a standoff. In one case, which might be thought of as a static standoff, the scores of all bordering cells are equal. Each cell thus chooses to retain its current strategy. In another case, which might be thought of as a dynamic standoff, a cell A on the border may have B neighbors with a higher score, but may retain strategy A because its non-border A neighbors get a higher score still. Here a boundary remains in place not because all scores are equal, but because the cells with a strategy losing at the interface are replaced from their own ranks.

Unfortunately, analysis for spatial dynamics quickly becomes more complex. The inequalities above do not hold, for example, when the configuration of the border is changed by even a single cell (Figure 9d).

In this case we can still compute the score balances required for invasion and standoff, but the required inequalities become clumsier. In the array shown in Figure 9d, for example, A will create companion invaders to the sides of the single cell on the next round if

$$\max[3AA + 5AB, 6AA + 2AB, 5AA + 3AB) > \max [4BB + 4BA, 5BB + 3BA, 7BB + 3BA, 8BB].$$

It will invade forward and to the side if

$$3AA + 5AB > \max[7BB + BA, 4BB + 4BA, 8BB].$$

That single cell will shrink back if

$$\max[4BB + 4BA, 7BB + 1BA] > \max[3AA + 5AB, 6AA + 2BA].$$

And there will be an invasion to the left adjacent to the protuberance if:

$$\max[5B + 3BA, 4BB + 4BA] > \max[6AA + 2AB, 5AA + 3AB, 8AA].$$

This is the calculation required for a single generation at a single site. We can of course crunch the numbers in each such case, but any general patterns dictating invasion patterns seem to be lost in mere arithmetic.

Similar analytical complexities quickly arise regarding other simple dynamics. Consider invasion by a single cell of strategy A in a field of strategy B, growing to occupy a 3 x 3 block of cells (Figure 10). In this case the algebra is simple: Strategy A can invade only if 8AB > max

(8BB, 7BB + BA). It is also possible to prove that invasion will occur with a limit of precisely 50% for the range of possible values that might be put in for AB, BB, and BA (see Appendix 1). AA can be left out of consideration because our initial cell does not play itself.

Algebra will take us this far. But what occurs at the next step, in those 50% of cases in which a block of 9 cells has been established?

At this second stage there are at least four possibilities for growth, shown in Figure 11. Values AA, AB, BA, and BB all prove important, and here we were forced to resort to surveys. In a first survey we used all integer values between 0 and 1000 for each of AA, AB, BA, BB, checked in a second variation using values at .01 intervals between 0 and 1. In a third survey we used random values between 0 and 1000000, with the same results. In the pattern of growth shown in Fig. 11a, the block of 9 shrinks again to 1. In our surveys we found this to occur in approximately 11% of the nine-cell cases. In the pattern shown in Fig. 11b, the block expands from 9 cells to 25, a result we have found for approximately 58% of cases. In the pattern of Fig. 11c the block of 9 shrinks to a static cross formation. This we found in only about 1.5% of our sample. Finally, in the pattern shown in Fig. 11d, the block of 9 expands to a ‘castellated’ configuration with a nick on each side. This result shows up with a frequency of approximately 30% in our surveys.

The illustrations we have offered here consider only two strategies and invasion from a single cell or along a linear border. The increase in complexity even from a first step to a second is intimidating. Our experience thus parallels that of other investigators: simulation is easy but analysis is hard. Van Baalen and Rand (1998) note that “Invasion in viscous populations is a process that is hard to analyze” (632).

> Thus, for the evolution of altruism there must be discreteness and associated stochasticity (Goodnight, 1992). This means that we should analyze models that are individual-based as well as spatial. Probabilistic cellular automaton (PCA) models ... satisfy these criteria. However, even when we assume haploid reproduction (and thus ignore genetics) such PCA models are easy to stimulate but very hard to analyze. (633)

It should also be noted that the complexities of spatial dynamics make 'invasion' too simple a category. Which of the typical dynamics shown in Figure 12, for example, should count as 'invasion'? Some clearly do, in that any finite area is eventually occupied by the invading strategy. But in some the introduced strategy grows or expands only to shrink in a later generation. In others the introduced strategy expands without ever taking over, always leaving islands of the other strategy. What happens in the spatial dynamics of a more diverse population will depend on all of these dynamics, and a simple single category of 'invasion' seems inadequate to do them justice.[21] Figure 13 shows a particularly elaborate form of invasion, here using <3.49, 1, 5, 1> as our values for <AA, AB, BA, BB>.

On the basis of simulations, it is clear that spatialization is an important factor in emergence of both cooperation and communication. Difficulties of complexity, however, quickly swamp the search for analytic understanding. It is possible and perhaps even likely that there can be no algorithm adequate for predicting invasion patterns in general—it is known that similar questions are undecidable for the general case in the spatialized prisoner's dilemma, for example (Grim 1997, Grim, Mar, and St. Denis, 1998). At this point resort to approximation techniques seems a necessary alternative (Durrett and Levin 1994a, 1994b; Durrett 1999; Van Baalen and Rand 1998; Snyder and Nisbet 2000; Filipe and Maule 2003).

7. Conclusion

There are three neglected features of primary importance, we want to argue, in understanding evolution of both cooperation and communication. Those features are location, location, and location.

Under the category of 'viscosity', the importance of spatialization has been recognized in work in evolution theory and simulations regarding 'altruism'. For some reason it has been given significantly less attention in the corresponding game-theoretic tradition of work on 'cooperation'. In studies of communication, spatialized tasks have sometimes been part of the model. The number of communication studies that incorporate spatialization across both task and mode of reproduction, on the other hand, is very small indeed. What we have tried to argue, contrasting global with spatialized simulations for both cooperation and communication, is that this is a mistake.

Emergence of both cooperation and communication can depend crucially on local as opposed to global organization, and we ignore the role of spatialization at the risk of missing major aspects of the phenomenon under study. Contemporary resources for simulation make development of spatial models in both cases relatively easy. A full analytic understanding of the dynamics of spatial organization, on the other hand, is another matter.

References

Ackley, D.H., & Littman, M.L. (1994). Altruism in the evolution of communication. In R. Brooks & P. Maes (Eds.), *Proceedings of the Fourth Artificial Life Workshop*, pp. 40-48. Cambridge, MA: The MIT Press.

Alexander, R. D. (1987). *The Biology of Moral Systems*. New York: Aldine de Gruyter.

Axelrod, R. (1980a). Effective choice in the prisoner's dilemma. *Journal of Conflict Resolution* 24, 3-25.

Axelrod, R. (1980b). More effective choice in the prisoner's dilemma. *Journal of Conflict Resolution* 24, 379-403.

Axelrod, R. (1984). *The Evolution of Cooperation*. New York: Basic Books.

Axelrod, R. (1985). The emergence of cooperation among egoists. In T. Campbell & L. Sowden (Eds.), *Paradoxes of Rationality and Cooperation*, pp. 320-339. Vancouver: University of British Columbia Press.

Axelrod, Robert (1997). The evolution of strategies in the iterated prisoner's dilemma. In C. Bicchieri, R. Jeffrey, & B. Skyrms (Eds.), *The Dynamics of Norms*, pp. 1-16. Cambridge: Cambridge University Press.

Axelrod, R. & Hamilton, W. (1981). The evolution of cooperation. *Science* 211, 1390-1396.

Batali, John (1995). Small signaling systems can evolve in the absence of benefit to the information sender. *http://cogsci.ucsd.edu/~batali/cv.html*.

Bergstrom, C. T., and Lachmann, M. (2001). *Animal Behavior* 61, 535-543.

Brauchli, K., Killingback, T., and Doebeli, M (1999). Evolution of cooperation in spatially structured populations. *Journal of Theoretical Biology* 200, 405-417.

Brinkers, M., & P. den Dulk (1999). The evolution of non-reciprocal altruism. *Lecture Notes in Artificial Intelligence* 1674, 499-503.

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

Clark, K., & Sefton, M. (2001). The sequential prisoner's dilemma: Evidence on reciprocation. *The Economic Journal* 111, 51-68.

Cooper, R. (1996). Cooperation without reputation: Experimental evidence from prisoner's dilemma games. *Games and Economic Behavior* 12, 187-218.

Danielson, P. (2001). Competition among cooperators: Altruism and reciprocity. Sackler Colloquium on Agent-Based Modeling. *http:/www.ethics.ubc.ca/pad/*

Di Paolo, E. A. (1999). A little more than kind and less than kin: The unwarranted use of kin selection in spatial models of communication. *Lecture Notes in Artificial Intelligence* 1674, 504-513.

Di Paolo, E. A. (2000). Ecological symmetry breaking can favour the evolution of altruism in an action-response game. *Journal of Theoretical Biology* 203, 135-152.

Durrett, R. (1999). Stochastic Spatial Models. SIAM Review 41, 677-718.

Durrett, R., & Levin, S. (1994a). The importance of being discrete (and spatial). *Theoretical Population Biology* 46, 363-394.

Durrett, R., & Levin, S. (1994b). Stochastic spatial models: A user's guide to ecological applications. *Philosophical Transactions of the Royal Society of London B* 343, 329-350.

Epstein, J. (1998). Zones of cooperation in demographic prisoner's dilemma. *Complexity* 4, 36-48.

Filipe, J. A., & Maule, M. M. (2003). Analytic methods for predicting the behavior of population models with general spatial ineractions. *Mathematical Bioscience*s 183, 15-35.

Fisher, R. A. (1930). *The Genetical Theory of Natural Selection*. New York: Clarendon Press.

Gintis, H., Smith, E., & Bowles, S. (2001). Costly signaling and cooperation. *Journal of Theoretical Biology* 213, 103-119.

Godfray, H. J. (1991). Signalling of need by offspring to their parents. *Nature* 352, 328-330.

Goodnight, K. F. (1992). The effect of stochastic variation of kin selection in a budding-viscous population. *American Naturalist* 140, 1028-1040.

Grafen, A. (1984). Natural selection, group selection, and kin selection. In Krebs, J. & Davies, N. (Eds.), *Behavioral Ecology: An Evolutionary Approach 2nd Ed.*, pp. 62-84. Oxford: Blackwell.

Grafen, A. (1990). Biological signals as handicaps. *Journal of Theoretical Biology* 144, 517-546.

Griffin, A. S. & West, S. A. (2002). Kin selection: Fact and fiction. *Trends in Ecology and Evolution* 17, 15-21.

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. (1997). The undecidability of the spatialized prisoner's dilemma. *Theory and Decision* 42, 53-80.

Grim, P. (2002). The basic question: What is reinforced? What is selected?, commentary on H. Rachlin, Altruism and selfishness. *Behavioral and Brain Sciences* 25, 261.

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

Grim, P., Kokalis, T., Tafti, A., and Kilb, N. (2001). Evolution of communication with a spatialized genetic algorithm. *Evolution of Communication* 3, 105-134.

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.

Hamilton, W. D. (1963). The evolution of altruistic behavior. *American Naturalist* 97, 354-356.

Hamilton, W. D. (1964). The genetical evolution of social behavior, i and ii. *Journal of Theoretical Biology* 7, 1-16, 17-52.

Harms, William (2001). Cooperative boundary populations: The evolution of cooperation on morality risk gradients. *Journal of Theoretical Biology* 213, 299-313.

Harms, William (2000). The evolution of cooperation in hostile environments. *Journal of Consciousness Studies* 7, 308-313.

Hauert, C., & Schuster, H. G. (1996). Effects of increasing the number of players and memory size in the iterated prisoner's dilemma: A numerical approach. *Proceedings of the Royal Society London B* 264, 513-519.

Huberman, B., & Glance, N. (1993). Evolutionary games and computer simulations. *Proceedings of the National Academy of Sciences USA* 90, 7715-7718.

Hutchins, E., & Hazlehurst, B. (1991). Learning in the cultural process. In C. G. Langton, C. Taylor, J. D. Farmer, & S. Rasmussen (Eds.), *Artificial Life II: SFI Studies in the Sciences of Complexity*, pp. 689-708. Addison-Wesley.

Hutchins, E., & Hazlehurst, B. (1995). How to invent a lexicon: The development of shared symbols in interaction. In N. Gilbert & R. Conte (Eds.), *Artificial Societies: The Computer Simulation of Social Life*, pp. 157-189. London: UCL Press.

Irwin, A. J., & Taylor, P. D. (2001). Evolution of altruism in stepping-stone populations with overlapping generations. *Theoretical Population Biology* 601, 315-325.

Jefferson, D., Collins, R., Cooper, C., Dyer, M., Flowers, M., Korf, R, Taylor C., & Wang, A. (1992). Evolution as a theme in artificial life: the Genesys/Tracker system. In C. G. Langton, C. Taylor, J. D. Farmer, & S. Rasmussen (Eds.), *Artificial Life II: SFI Studies in the Sciences of Complexity*, pp. 549-577. Addison-Wesley.

Kennedy, J., Eberhart, R., & Shi, Y. (2001). *Swarm Intelligence*. San Francisco: Morgan Kaufman.

Killingback, T., and Doebeli, M. (1996). Spatial evolutionary game theory: hawks and doves revisited. *Proceedings of the Royal Society London B* 263, 1135-1144.

Krakauer, D., & Pagel, M. (1995). Spatial structure and the evolution of honest cost-free signaling. *Proceedings of the Royal Society London B* 260, 365-372.

Leimar, O., & Hammerstein, P. (2001). Evolution of cooperation through indirect reciprocity. *Proceedings of the Royal Society London B* 268, 745-753.

Levin, M. (1995). The evolution of understanding: A genetic algorithm model of the evolution of communication. *BioSystems* 36, 167-178.

Lindgren, K. and Nordahl, M. (1994). Evolutionary dynamics of spatial games. *Physica D* 75, 292-309.

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*, Society for the Study of Artificial Intelligence and Simulation of Behavior, AISB Press, pp. 139-146.

MacLennan, B. J. (1991). Synthetic ethology: An approach to the study of communication. In C. G. Langton, C. Taylor, J. D. Farmer, & S. Rasmussen (Eds.), *Artificial Life II: SFI Studies in the Sciences of Complexity*, pp.631-657. Addison-Wesley.

MacLennan, B. J., & Burghardt, G. M. (1994). Synthetic ethology and the evolution of cooperative communication. *Adaptive Behavior* 2, 161-888.

McAndrew, F. T. (2002). New evolutionary perspectives on altruism: Multilevel-selection and costly-signaling theories. *Current Directions in Psychological Science* 11, 79-82.

Miller, J. H., Butts, C. T., & Rode, D. (2000). Communication and cooperation. *Journal of Economic Behavior & Organization* 47, 179-195.

Mitteldorf, J., & Wilson, D. S. (2000). Population viscosity and the evolution of altruism. *Journal of Theoretical Biology* 204, 481-496.

Nakamaru, M., Matsuda, H., & Iwasa, Y. (1997). The evolution of cooperation in a lattice-structured population. *Journal of Theoretical Biology* 184, 65-81.

Nakamaru, M., Nogami, H., & Iwasa, Y. (1998). Score-dependent fertility model for the evolution of cooperation in a lattice. *Journal of Theoretical Biology* 194, 101-124.

Noble, J., & Cliff, D. (1996). On simulating the evolution of communication. In P. Maes, M. Mataric, J. A. Meyer, J. Pollack, & S. W. Wilson (Eds.), *From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*, pp. 608-617. Cambridge, MA: MIT Press.

Nowak, M. (1990). Stochastic strategies in the prisoner's dilemma. *Theoretical Population*

Biology 38, 93-112.

Nowak, M., & May, R. (1992). Evolutionary games and spatial chaos. *Nature* 359, 826-869.

Nowak, M., & May, R. (1993). The spatial dimensions of evolution. *International Journal of Bifurcation and Chaos* 3, 35-78.

Nowak, M., & Sigmund, K. (1992). Tit for tat in heterogenous populations. *Nature* 355, 250-252.

Nowak, M., & Sigmund, K. (1993a). Chaos and the evolution of cooperation. *Proceedings of the National Academy of Sciences* 90, 5091-5094.

Nowak, M., & Sigmund, K. (1993b). A strategy of win-stay, lose-shift that outperforms tit-for-tat in the prisoner's dilemma game. *Nature* 364, 56-58.

Nowak, M., & Sigmund, K. (1995). Invasion dynamics of the finitely repeated prisoner's dilemma. *Games and Economic Behavior* 11, 364-390.

Nowak, M., & Sigmund, K. (1998). The dynamics of indirect reciprocity. *Journal of Theoretical Biology* 194, 561-574.

Oliphant, M. (1994). Evolving cooperation in the non-iterated prisoner's dilemma: the importance of spatial organization. In R. A. Brooks & P. Maes (Eds.), *Artificial Life IV*, pp. 349-352. Cambridge: MIT Press.

Oliphant, M. (1999). The learning barrier: Moving from innate to learned systems of communication. *Adaptive Behavior* 7, 371-384.

Parisi, D. (1997). An artificial life approach to language. *Brain and Language* 59, 121-146.

Pepper, J. W., & Smuts, B. B. (2000a). The evolution of cooperation in an ecological context: An agent-based model. In T. A. Kohler & G. J. Gumerman (Eds.), *Dynamics of Human* *and Primate Societies: Agent-Based Modeling of Social and Spatial Processes*, pp. 45-76. Oxford: Oxford University Press.

Pepper, J. W., & Smuts, B. B. (2000b). Agent-based modeling of multilevel selection: The evolution of feeding restraint as a case study. Swarm User Group Meeting, Santa Fe Institute, NM.

Queller, D. C. (1994). Genetic relatedness in viscous populations. *Evolutionary Ecology* 8, 70-73.

Rachlin, H. (2003). Altruism and Selfishness. *Behavioral and Brain Sciences*.

Saunders, G. M., & Pollack, J. B. (1996). The evolution of communication schemes over continuous channels. In P. Maes, M. Mataric, J. A. Meyer, J. Pollack, & S. W. Wilson (Eds.), *From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*, pp. 580-589. Cambridge, MA: MIT Press.

Segal, N. & Hershberger, S. (1999). Cooperation and competition between twins: Findings from a prisoner's dilemma game. *Evolution and Human Behavior* 20, 29-51.

Sella, G., & Lachmann, M. (2000). On the dynamic persistence of cooperation: How lower individual fitness induces higher survivability. *Journal of Theoretical Biology* 206, 465-485.

Sigmund, K. (1998). Complex adaptive systems and the evolution of reciprocation. *Ecosystems* 1, 444-448.

Sigmund, K., & Hauert, C. (2002). Altruism. *Current Biology* 12, R270-R272.

Sober, E., and Wilson, (1998). *Unto Others: The Evolution and Psychology of Unselfish Behavior*. Cambridge, MA : Harvard University Press.

Skyrms, B. (1996). *Evolution of the Social Contract*. Cambridge: Cambridge University Press.

Skyrms, B. (2000). Game theory, rationality, and evolution of the social contract. *Journal of Consciousness Studies* 7, 269-284.

Skyrms, B. (2001). The stag hunt. *Proceedings and Addresses of the American Philosophical Association* 75, 31-41.

Skyrms, B., & Pemantle, R. (2000). A dynamic model of social network formation. *Proceedings of the National Academy of Sciences USA* 97, 9340-9346.

Smith, E. A. & Bleige Bird, R. L. (2000). Turtle hunting and tombstone opening: public generosity as costly signaling. *Evolution of Human Behavior* 21, 245-261.

Smith, J. M. (1964). Group selection and kin selection. *Nature* 201, 1145-1147.

Smith, J. M. (1991). Honest signalling: The Sir Philip Sidney Game. *Animal Behavior* 42, 1034-1035.

Snyder, R. E., and Nisbet, R. M. (2000). Spatial structure and fluctuations in the contact process and related models. *Bulletin of Mathematical Biology* 62, 959-975.

Taylor, P. D., & Irwin, A. J. (2000). Overlapping generations can promote altruistic behavior. *Evolution* 54, 1135-1141.

Trivers, R. L (1971). The evolution of reciprocal altruism. *Quarterly Review of Biology* 46, 35-57.

van Baalen, M., & Rand, D. (1998). The unit of selection in viscous populations and the evolution of altruism. *Journal of Theoretical Biology* 192, 561-565.

Vainstein, M. H, & Arenzon, J. J. (2001). Disordered environments in spatial games. *Physical Review E* 64, 051905-1 – 051905-6.

Wagner, K. (2000). Cooperative strategies and the evolution of communication. *Artificial Life* 6, 149-179.

Wedekind, C., & Milinski, M. (1996). Human cooperation in the simultaneous and the alternating prisoner's dilemma: Pavlov versus generous tit-for-tat. *Proceedings of the National Academy of Sciences USA* 93, 2686-2689.

Werner, G. M., & Dyer, M. G. (1992). Evolution of communication in artificial organisms. In C. G. Langton, C. Taylor, J. D. Farmer, & S. Rasmussen (Eds.), *Artificial Life II: SFI Studies in the Sciences of Complexity*, pp. 659-587. Addison-Wesley.

West, S., Pen, I., & Griffin, A. (2002). Cooperation and competition between relatives. *Science* 296, 72-75.

Wilson, D. S., Pollock, G. B., & Dugatkin, L. A. (1992). Can altruism evolve in purely viscous populations? *Evolutionary Ecology* 6, 331-341.

Zahavi, A. (1977). Reliability in communication systems and the evolution of altruism. In B. Stonehouse & C. M. Perrins (Eds.), *Evolutionary Ecology*, pp. 253-259. London: MacMillan.

Appendix 1

In growth from a single cell to a block of 9, as illustrated in Fig. 10, it is possible to prove that invasion will occur with a limit of precisely half for the range of possible values that might be put in for AB, BB, and BA. AA can be left out of consideration because our initial cell does not play itself. Using the inequality for expansion $8AB > \max(8BB, 7BB+BA)$, we proceed by cases on the right side:

Case 1. Max (8BB, 7BB + BA) = 8BB just in case BB > BA (1). Where values for BB > BA are chosen randomly and independently, this condition will hold in 50% of cases. Expansion from 1 cells to 9 occurs in this first case just in case 8AB > 8BB and thus just in case AB > BB. Where values for AB and BB are chosen independently, this condition will also hold in 50% of cases. Choice of values for AB, BA, and BB will thus give us expansion under case 1 in 25% of all cases.

Case 2. Max (8BB, 7BB + BA) = 7BB + BA just in case BA > BB (2). Where values for BB and BA are chosen randomly and independently, this condition will hold in 50% of cases. Expansion from 1 to 9 cells occurs in this second case just in case 8AB > 7BB + BA. But BA > BB by (2). So BA = BB + , for some , > 0. Since BB = BA - ,,

7BB + BA = 7(BA - ,) + BA

= 7BA - 7, + BA

= 8BA - 7,

Since the lowest bound for , is 0, the upper bound for 7BB + BA = 8BA - 7, is 8BA. Expansion occurs in this second case just in case 8AB is greater than that upper bound; just in case 8AB > 8BA, or AB > BA. Where values for AB and BA are chosen independently, this condition will also hold in 50% of cases. Choice of values for AB, BA, and BB will thus give us expansion under case 2 in 25% of all cases.

In total, we can thus expect expansion from 1 cell to a 3 x 3 block of 9 for precisely ½ of assignable values.

Notes

---

[1] Questions have been raised about altruism as a basic category, and seem particularly appropriate where complex and flexible behaviors are at issue. Di Paolo (2000) cites the standard definition, but notes that given such a definition it will depend on context whether a behavior counts as altruistic or not. "...[W]hether a behavior is altruistic or not depends on the current pool of behavioral strategies in the neighborhood. This form of contextual altruism is more appropriate for describing social behaviors where benefits and costs depend both on the initiating behaviour and on the response of other individuals" (136).

[2] The Axelrod tradition is sometimes spoken of as instantiating a theory of 'reciprocal altruism' in the sense of Trivers (1971): with iterated exchanges over time and benefits that outweigh costs, a pattern of individual acts each of which may count as 'altruistic' can produce long-range benefits for both parties. But it is far from clear that 'reciprocal altruism' in this sense qualifies as 'altruism' in the biologist's sense at all. "Reciprocal altruism is defined as cooperative behavior among unrelated individuals that benefits everyone involved" (McAndrew 2002), but altruistic traits are explicitly defined in the biological tradition as traits which do *not* benefit both parties. It may also be the case that who an act benefits and thus whether it counts as 'altruism' may depend both on context (Di Paolo 2000) and on the level of description (Rachlin 2002, Grim 2002).

[3] Further bodies of theory relevant to altruism, incorporating still more variables, are 'costly-signaling theory' (McAndrew 2001, Gintis, Smith, and Bowles 2001; Smith and Bliege Bird, 2000) and 'indirect reciprocation' (Sigmund 1998, Sigmund and Hauert 2002; Leimar and Hammerstein, 2001; Nowak and Sigmund 1998). Status is the core of some theories of both types: that altruistic acts increase an individual's perceived value, and thus increase contributions

---

to his fitness from the community (Alexander 1987, Grafen, 1990; Zahavi, 1977). On this pattern, an individual might benefit from his altruistic acts even if he never again interacts with the beneficiaries of those acts. Here again, however, the fact that an individual benefits from such a pattern of behavior makes it doubtful that it qualifies as 'altruistic' in the strict biological sense.

[4] In Axelrod 1997, an average individual is given one mating, while an individual one standard deviation more effective than the average is given two. Successful individuals are randomly paired for crossover, and the process is repeated for the new pool. This is actually the more interesting of two models Axelrod explores. In the other, a pool of strategies changes over generations, but the fitness measure remains the same: success against 8 strategies representative of those submitted in the second of Axelrod's tournaments.

[5] Use of TFT in place of All-C gives a similarly simple but different result. For All-C vs All-D in a sufficiently iterated game, where CD indicates the gain of All-C against All-D, DC > CC > DD > CD. Replacing All-C with TFT, on the other hand, gives us TT > DT > DD > TD. Some more recent attempts have at least broadened the field to 3 strategies, using All-C, All-D, and TFT as a 'discriminating' strategy (Harms 2001).

[6] Neither strategy gears its play to that of its opponent, and thus C and D might for different values be considered either All-C and All-D in an iterated game or a simple 'cooperate' and 'defect' in a non-iterated game. Nowak and May seem to think of them in the former way, but use a simplified matrix in which cooperation against defection and defection against defection both get a value of zero. Oliphant 1994 uses a more standard Prisoner's Dilemma matrix, and seems to think in terms of a non-iterated game.

---

[7] Individuals select opponents to play based on a gaussian or bell curve distribution around them. In each cycle of approximately 32 games per player, a single individual is chosen randomly (1) to reproduce, but with a bias in favor of successful strategies, and (2) to be replaced, though with a bias toward the less successful. The model also employs an unspecified level of mutation.

[8] Danielson's strategies are also slightly more sophisticated: the array includes straight Defectors, R1 reciprocators which cooperate when at least one immediate neighbor has cooperated, and R2 cooperators who cooperate only when both immediate neighbors have cooperated.

[9] In Harms's initial simulations successful populations grow and, with no boundary separating the communities, eventually collide. In that instance defectors dominate over cooperators. Here Harms proposes variable extinction rates, and finds an intermediate extinction rate for each population size—a rate which allows cooperators, because of their superior fitness levels from auto-associating, to outlive defectors and thus to succeed.

[10] William Harms (2000) also notes that Skyrms (2000) leaves out spatialization.

[11] Stag Hunt is defined by the ordering CC>DC>DD>CD, where for example DC indicates the gain for a player defecting against an opponent who cooperates. The Prisoner's Dilemma is defined by DC>CC>DD>CD, with the additional condition that CC>(DC+CD)/2.

[12] See also Killingback and Doebeli, 1996.

[13] Although developed independently (Nowak and May 1992, 1993; Grim 1995; Grim 1996; Grim, Mar, and St. Denis, 1998), we note that spatialization of this form fits the main outlines of the particle swarm paradigm elaborated in Kennedy, Eberhart, and Shi (2001).

[14] Theoretical work on biological questions of 'costly signaling' has also generally been non-spatial. See Zahavi (1977), Smith (1991), Godfray (1995), Bergstrom and Lachmann (2001).

---

[15] Noble and Cliff claim to find "at least a topology, if not a geometry" in the MacLennan and Burghardt study: "simorgs will tend to receive signals from their immediate neighbors in one direction, and send signals to their neighbors in the other direction" (613). We take this to be a misunderstanding of the MacLennan and Burghardt study, since all simorgs receive any signals and are judged in terms of 'internal' matches to those signals. For our purposes, however, since we are arguing that spatialization should be included in models for communication, the degree of purity in excluding spatialization seems moot.

[16] In Cangelosi and Parisi (1998) the task is that of identifying nearby mushrooms as edible or poisonous and moving towards edible or away from poisonous mushrooms. Communication is somewhat artificially encouraged by limits on perception. A mushroom-hunting organism can see mushrooms at some distance, but can only observe their perceptual properties (whether edible or poisonous) if they are one square away. To aid agents in identifying mushrooms from further away, they are offered a far-sighted conspecific, randomly chosen from the population, which emits signals to aid in identification of the mushrooms. The near-sighted organisms must then learn to encode and interpret these signals correctly.

[17] Di Paolo (1999) notes that spatialization of models can introduce very different results in evolution of communication. But he criticizes some, particularly Ackley and Littman (1994) and Oliphant (1996), for attributing these differences to 'kin selection' while doing little to confirm the theory's applicability. Di Paolo argues that "[t]his careless appeal to kin-selective arguments is dangerous" (505) and offers a model which concludes that the emergence of cooperation in spatial models should not be necessarily attributed to kin selection. Di Paolo's charges may be unduly barbed: as far as we can tell, the authors criticized mention kin selection only as a

---

possible explanation.

[18] In a very different study regarding linguistic change and formation of dialects, Livingstone and Fyfe (1999) do explicitly compare global and local models, as we do here. In that context they show that spatialization produces greater linguistic diversity.

[19] On each round, each food source and each predator is assigned a random number between 1 and 9; depending on that number it will move into one of the neighboring 8 cells or remain where it is. A cell adjacent to one with a food source thus has a 1 in 9 chance of having that food source on it in the next generation; similarly for predators. The 'payoff' for received communication in a community of perfect communicators is thus an increase in the probability of successful feeding and of harm avoidance.

[20] We use a larger number of predators because of different dynamics in response to a food call and a predator warning. In the case of a food call a chain reaction of eating and signaling is set up across a community of communicators, with exploitation of the food source on every round. In response to a predation warning, on the other hand, neighbors 'hide', thereby preventing a predator hit the next time and preventing a further warning as well. As a result, even perfect communicators can avoid a predator hit only every other time. The number of food sources and predators is adjusted accordingly; it is for this reason that we use 50 food sources and 100 predators. For details see Grim, Kokalis, Tafti, and Kilb (2001).

[21] See also Grim (1995, 1996).