The emergence of differentiating signals within groups

Edward Augenblick
Stanford University
Stanford, CA, USA

Annerieke Heuvelink
Vrije Universiteit
Amsterdam,
The Netherlands

Radu Iovita
University of
Pennsylvania
Philadelphia, PA, USA

Graham Ritchie
University of
Edinburgh
Edinburgh, UK

Abstract

Groups of people seem to naturally develop auditory and visual "signals" that are unique to that group. Furthermore, people seem to preferentially interact with others that share their set of "signals." This paper proposes an explanation of this behavior that focuses on social norms, which are "rules of thumb" that allow successful coordination in human interaction. The basic hypothesis is that matching signals is an extremely efficient way of determining if two people share the same social norms. In order to support this hypothesis and provide some sense of how it might have evolved, we have developed two computer models. These models are reviewed and the main results are presented.

1. Introduction

This paper is intended to provide a relatively simple explanation for a set of human behaviors that are very common, but not particularly well understood. We will introduce these behaviors through three related and relatively uncontroversial observations.

First, groups of people seem to naturally develop auditory and visual "signals" that are unique to that group. For example, subgroups of speakers of a language often develop slightly different vocabularies, methods of annunciation, and grammatical structure that allow them to communicate with the wider group but still recognize members of the subgroup. Visually, subgroups of a population differentiate themselves through the use of different clothing, jewelry, or body art. More subtly, groups of people are often able to differentiate themselves by knowing certain information or facts that other groups do not know. After some reflection, it becomes clear that there are a large number of similar examples of this behavior.

The next observation is that people seem to naturally imitate the signals of the group that they spend time with. On the most basic level, people naturally emulate the signals of the culture that they are raised in. A person raised in a specific tribe in Africa will obviously take on the verbal and visual signals of that tribe, rather than those of a different tribe. However, if a person spends time in a new culture, they will often "pick up" some of the fashion or verbal tendencies of the new group. This suggests that the signals that people choose to display are largely dependent on the groups that they spend time with.

The final observation is that people seem to preferentially interact with people that have similar "signals." It is fairly clear that people, subconsciously or not, gravitate towards "similar" people, while they shun people that are "different", where these similarities and differences lie in the outward signals that people present to others.

It would seem that there must be some evolutionary advantage to this behavior as it is so common, but it is not easy to find a clear explanation. As the behavior involves preferentially interacting with others, one explanation might be that the behavior is used to signal kin membership. At first, this makes some sense, as people copy the signals of the people they spend time with, which is commonly their kin group. However, people use the signals to determine if they should preferentially interact with people that they *don't* know, who are most commonly *not* close kin. Basically, people who grow up around their close kin can recognize them by sight, and they don't need these signals to differentiate them. Therefore, it would seem that there must be better explanation.

A more satisfying explanation revolves around the idea of social norms and social interaction. To understand this explanation, it is necessary to recognize that successful social interaction requires that people "coordinate" their actions. For example, if a group of people is hunting an animal, they must all coordinate on many actions, such as when to attack the animal or what to do if a member of the group is injured. Note that there is not just one way to coordinate these actions, but in order for the hunt to be successful, the actions must be coordinated in one way or another. Another example involves punishment within groups. Many people have pointed out that good punishment systems are essential to successful social interaction. However, a good punishment system necessarily requires that people coordinate their actions: every person must know which actions can lead to punishment and their role in punishing others in order for the punishment system to work. Again, there is not only one good punishment system, but any punishment system must involve coordination to be successful. If fact, with some reflection, it seems that a large amount of social interaction requires some coordination in order to be successful.

Social norms are the rules that govern how people behave in these situations. Different cultures have different ways of hunting or punishing people, and consequently have different social norms. When people who share the same social norm interact, their behavior is coordinated well. Conversely, when people who share different social norms interact, their interaction can be "mixed-up" and lead to poor results. This suggests that interacting preferentially with people that share the same social norms can lead to much higher payoffs than interacting with others. But, how can a person tell if another person, who they have never met, shares their social norms? It would be impossible for a person be able to explain their social norms to another person as these norms are numerous, complex and often exist in the subconscious.

However, the visual and auditory signals mentioned above provide a fantastic solution. Social norms are learned from the group that one spends time with. If an easily accessible visual or auditory signal was learned alongside the social norm, then two people who meet would know that they share the same social norms if and only if they share the same signal. Notice how the signal compresses and simplifies the incredibly complex problem of matching social norms. When you meet a person, it is not necessary to know their exact social norms, but only if their social norms match your own. The signal perfectly and efficiently achieves this purpose. If you meet someone with a different signal, you are not given information about their social norms, but that is irrelevant: all you care is if their signal matches your own. If you meet someone with the same signal, you do not need them to express the social norm explicitly as you share the same strategy. The signal circumvents the problem that social norms themselves are so difficult to communicate.

In order to support this hypothesis about the connection between the signals and social norms, as well as to provide some sense of how it might have evolved over time, we have developed two computer models. In the sections 3 and 4 of this paper, these models are reviewed and the main results are presented. The next section will first introduce some research related to the topic introduced.

2. Related Research

The evolution of the social basis for separation into groups is one of the most complex questions in human evolution. Human group size and organisation are clearly limited by biological and environmental factors, but are also subject to significant alteration due to internal fractionation based on cultural rules of engagement, i.e., social norms. One view of this process assumes imitation as a prime mover. From an evolutionary point of view, imitation is selected for whenever individual learning is difficult or costly and the environment is changing slowly enough to make reliance on past experience more beneficial than individual learning (Boyd and Richerson 1994). Proximity to the source of knowledge (in the case of social knowledge this is simply the collection of practitioners) is considered enough to stimulate the perpetuation of a 'tradition'. Boyd and Richerson cite the example of Salomon (1985) of German and Yankee farmers in Illinois, whose different ethnic and religious heritage led to different adaptations in the face of modernity.

The other view, espoused by many economists, is that semiotic activity related to group fission is used in coordination games. In reality, equilibria are unstable and ideas get lost if there is no mechanism of maintenance and replication. More recently, Boyd et al (2003) have proposed such a model based on empirical data and theoretical elaboration of the concept of altruistic punishment as a corrective factor in group selection.

Taking human language as the pinnacle of human semiotic achievement, the best observed and studied process of fission is the formation of registers. According to Reid (1956), linguistic registers are functionally different uses of a language depending on social situations. Examples include scientific and sports jargons (Halliday 1988; Hoyle 1993), criminal registers (Maurer 1955), and honorifics (Agha 2002; Morford 1997). In a technical sense, these involve metapragmatic models of social action (Agha 2003) and are often characterised by differences in repertoire (size, grammatical and semiotic range), stereotypes of users and usage, and categories of persons competent in performance *and* recognition. Together, these factors contribute to recursively define the semiotic criteria for subdivision *and* the set of persons that fit these criteria.

In the case of linguistic registers, replication is often done through prescriptive socialisation within a group (e.g. honorifics as learnt in childhood, but also scientific or legal jargon socialised in the workplace). Many registers are also dispersed through the circulation of more perduring semiotic artifacts such as books and electronic media. However, prescriptive means of replication and maintenance form only a small part of the mechanisms used by a social group to preserve linguistic registers. Oftentimes this is achieved through implicit metalinguistic activities such as jocular accounts of defective speech (Agha 1998; 2003). These metalinguistic actions often serve the double purpose of defining the speaker as a member of two groups (such as a scientist poking fun at scientific jargon, defining himself as a good scientist, and also a relaxed, informal person at the same time). Multiple group membership is often seen in primate groups as well as in humans and any realistic model of these processes should incorporate the fluidity of between group relations.

3. NetLogo Model

Before we start questioning how a system of social entities copying signals from their neighborhood and using these signals in a later stage to coordinate their actions can have evolved, we first want to investigate whether, assuming this system, it will be the case that various groups of entities with specific signals and equal underlying strategies will emerge. In order to do this we develop a proof-of-concept model in NetLogo, which is a multi-agent based programming environment developed by Northwestern University (Wilensky 1999). In the next section we will first introduce the details of our model followed by some initial results.

3.1. Implementation details

The model consists of an environment of 15 x 15 patches, with each patch containing one agent. The neighborhood of each agent is said to be the eight patches surrounding its own patch, with the environment being wrapped, i.e. the right border is linked to the left border, as is the top to the bottom.

At the start of the simulation all the agents are initialized with two features. First, they all hold a list of ten values that, when normalized, represent the chances for that agent of selecting one of the ten possible strategies. These chances are equal at the beginning, i.e., 10 % each, as represented by this list [1 1 1 1 1 1 1 1 1 1]. Second, they all receive a random signal. Signals are made up of a list of three digits, with the digits ranging from 0 – 6 and being wrapped, i.e., the 6 is just as close to the 5 as to the 0.

After the initialization the model starts looping through the following sequence. At the beginning of every loop all the agents check the agents in their neighborhood and select one agent to play a coordination game with. The selection is based on the agents' signals with the agent being picked having the most similar signal. To illustrate: when agent A has signal [1 0 3] it will select agent B with signal [1 6 2] - so with a distance of (1 – 1) mod 7 + (0 – 6) mod 7 + (3 – 2) mod 7 = 2 - over agent C with signal [2 4 5] - so with distance 6. After selecting an agent to coordinate with, both the agents select the strategy they will play. The chance for a strategy to get selected is proportional to its value in the strategy list, with the value of strategy 1 found on position 1, 2 on position 2 and so on. When they have selected the same strategy, the agent that initialized the game will get positively reinforced both in selecting that strategy and in picking that agent. The former is done by adding the reinforcement value to the value of that strategy in the strategy list. The latter is done by decreasing the distance of its signal with the selected agent's signal by 1, e.g., after successful coordination between agents A and B agent A's signal will become [1 6 3] or [1 0 2].

However, when the agents select different strategies, agent A will get negatively reinforced in selected that strategy and agent again. The chance of selecting that strategy gets lowered by subtracting 1 from the value the selected strategy holds in agent A's list (except when it is already 0), and its signal will be adapted to increase the distance with the signal of agent B by 1.[1]

So in this model the agents do not know what the chances of other agents playing certain strategies are. The only aspects they can pay attention to are the agents' signals.

3.2. Results

While setting up this model we wonder several things. Will the agent cluster in groups displaying the same kind of signals and playing the same strategy? How much positive reinforcement will be needed to get the agents coordinated? What will change with changing reinforcement levels?

In order to answer these questions, the model is run several times with various levels of reinforcement. The reinforcement level influence how much the chance of selecting a strategy that was last selected is changed based on the outcome of the coordination game.

We find that when we run the model with just a small amount of reinforcement, nothing happens in the sense of coordination of strategies or signals. However, when we increase the reinforcement a bit, from 1 to 2, we see that after many time steps slowly a group of neighboring agents emerge that play the same strategy most of the time and display similar signals. After even many more time steps the entire world gets coordinated and starts playing that same strategy and displaying similar signals, see figure 1a, b and c.

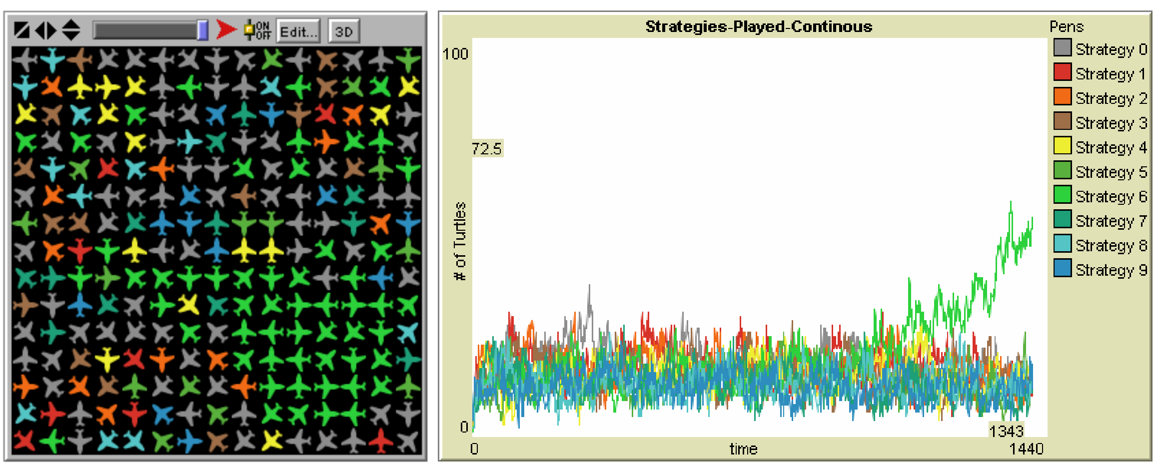

---

[1] We also run the model using not a fixed value of 1 as negative feedback, but using the reinforcement-level set in the interface for the positive feedback for the negative feedback as well. Under these circumstances we still find clustering of groups of agents on signal and strategy. However, the reinforcement-level needed to get this coordination kicked off is found to be higher; it only starts when the reinforcement-level is minimally 4. Using that value one strategy slowly takes over, resembling the results of the model using the fixed negative reinforcement with a positive reinforcement of 2, see figure 1.

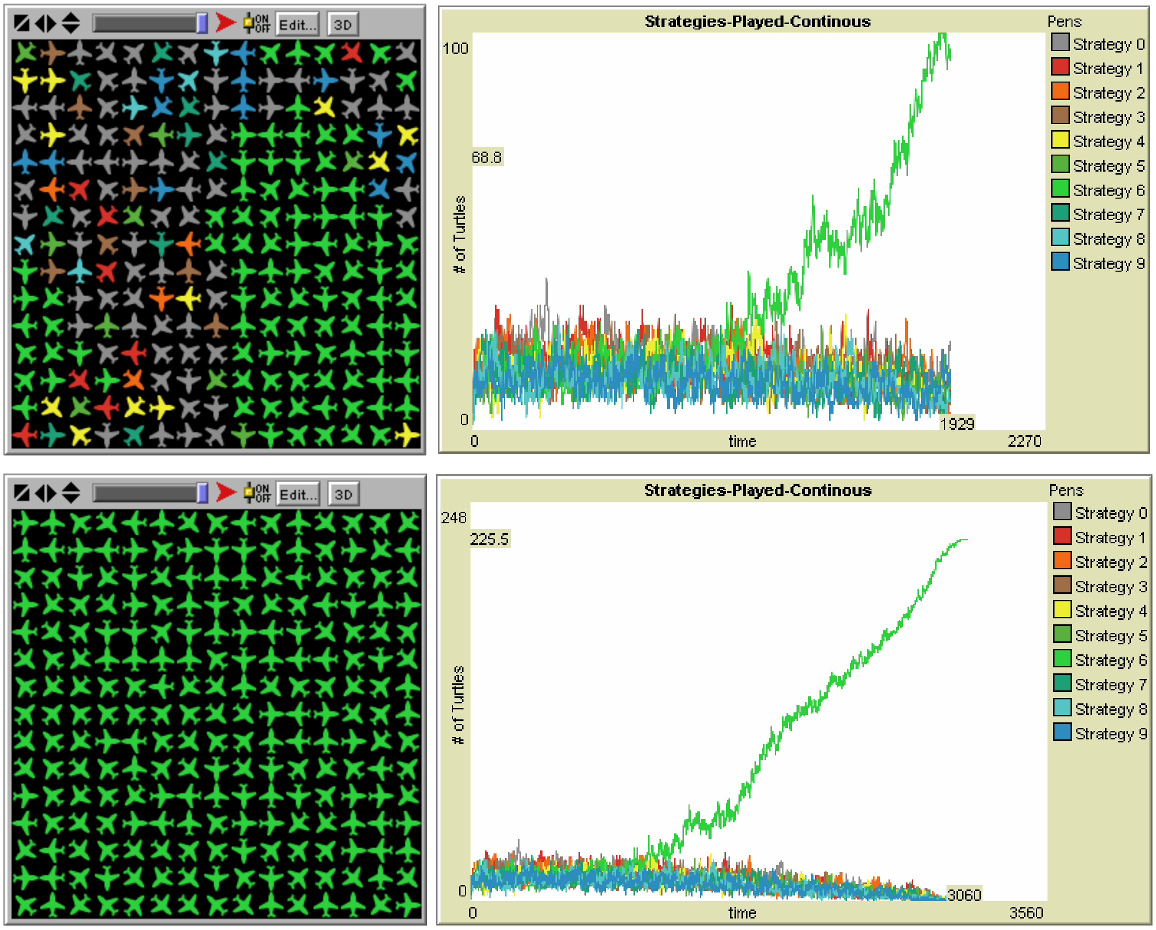

Figure 1a, b, and c: The model with a reinforcement of 2 at three different time steps

When we increase the reinforcement level further, we find again coordination among the agents. Neighboring agents end up with similar signals, which are truthful labels for the underlying strategy. However, not just one signal and strategy prevails, but various groups with unique signaling and played strategies emerge, see figure 2.

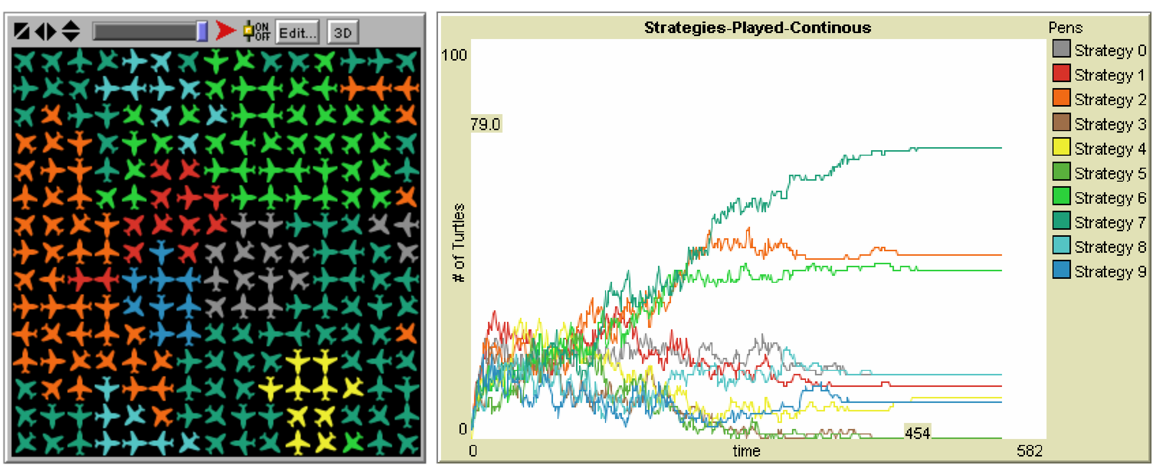

Figure 2: The model with a reinforcement of 3

When we increase the reinforcement level even further to 4, we find that many persistent small groups are formed very rapidly, see figure 3. The explanation behind this is as follows: when in the beginning two agents play the coordination game and by chance select the same strategy, they get greatly reinforced and the chance that they will select that strategy again almost triples, from 1/ 10 to 5 / 15. (The list holding the chances for selecting the various strategies changes from [1 1 1 1 1 1 1 1 1 1] to [1 5 1 1 1 1 1 1 1 1] when strategy 2 is successfully played). Since they will also adapt their signals, the chance that they will select each other again is big as is the chance that they will then select the same strategy again. As soon as this happens they will get reinforced again and very soon they lock into selecting certain agents and a certain strategy.

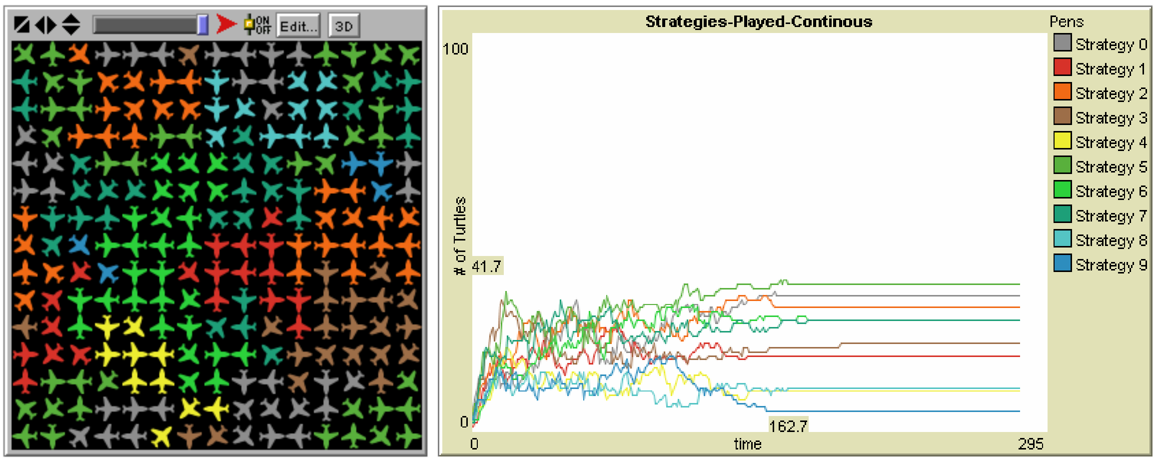

Figure 3: The model with a reinforcement of 4

Now we see that our model works in forming groups of agent displaying the same kind of signals and playing the same kind of strategy - therefore making the label truthful - we wonder in how far these results are plausible compared to the real world. The main difference we see is that the real world is much more dynamic. Groups emerge but also disappear again, people get into and out of groups, signals change etc. However, in the model and results above we see that after an initial "search-period", quite static groups of agents emerge, signaling the same signals and playing the same strategies over and over again. We wonder: what will happen when we introduce a bit more dynamics in the model? We decide to extend the model with death and birth to research this question. Through the introduction of a chance-of-death-parameter we can set the chance for an agent to die at each time step. When an agent dies, a new agent is initiated with a signal randomly selected out of his environment and with a list denoting equal chances for selecting each strategy.

3.3. Results of the model with death and birth

We run the new model with death and birth, with the chance of dying set at 1%, under the same conditions under which we have run the original model. When the reinforcement is set at 2, we find that, unlike in the original model, no clustering emerges. It seems that the reinforcement level needs to be higher under these more dynamic circumstances to take the threshold and to kick off the coordination. When we increase the reinforcement level to 3, we indeed see groups of agents emerge again signaling the same signals, interacting with each other and playing the same strategies. Interesting to see is that under these circumstances the groups are much more dynamic, e.g., see figure 4. In this figure we see that although the group playing strategy 1 starts off very strong and keeps growing, at a certain moment in time it stops growing and slinks again, giving space to another group with another strategy to become most prevalent. These results resemble much more the real-world dynamics in the formation and signaling of groups of entities.[2]

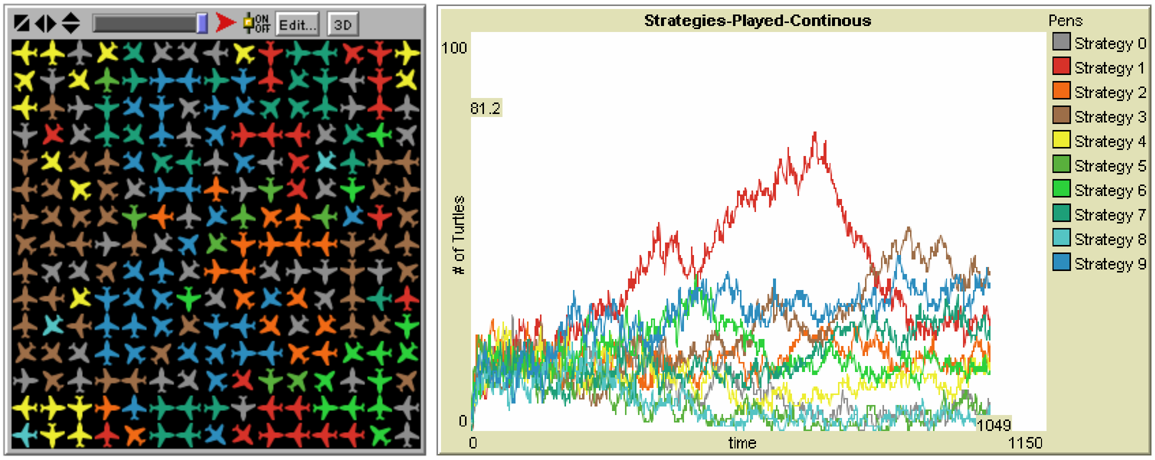

Figure 4: Model with death chance of 1 percent, with a reinforcement of 3

4. Repast Model

The NetLogo model demonstrates that agents can form groups with specific signals denoting specific strategies when it is the case that they pay attention to each other signals and make them more similar after successful coordination. We now want to continue our research with the more fundamental question: how is it that these two traits could have evolved? In order to research this question we have developed a model in Repast in which agents are equipped with genes that code for these traits. This section introduces the model and its results.

4.1. Implementation details

A population of p agents (p was set to 100 for all results presented here) lives in a 2D toroidal grid of 400 units and passes through the following ‘life stages’:

Birth: The agent’s ‘genomes’ consist of two genes, which serve the following functions:

Gene 1: a ‘learn local signal’ gene (denoted as a below) which determines the extent to which an agent is influenced by the signal of the previous generation when selecting their own signal.

[2] We also ran the model with a higher chance-of-death value and found that the same behavior emerged. However, there was a limit to how high the death percentage could be in order for coordination to happen, due to the increasing randomness. Below that limit the general rule: the higher the death-chance, the higher the reinforcement-level needs to be to kick coordination off, holds.

Gene 2: a 'pay attention to signal' gene (denoted as *b* below) which determines the extent to which the agent will pay any attention to another agent's signal when picking opponents for the coordination game.

These genes are simply real numbers constrained to be in the range 0-1. Every agent's strategy and signal are set after birth as described below. Both of these values are always in the range 0-99.

The strategy of an agent *i*, $strat_i$, is determined by the strategy of *n* nearby agents from the previous generation known as the agent's 'teachers' (*n* was set to 10 for all results presented here). Each teacher is selected in the following way; firstly a number is drawn from a normal distribution with mean 0 and standard deviation *t* (set to 1 for all results presented here). The absolute value of this number is computed, the number is incremented by 1 (to ensure the agent does not interact with the agent in the same location as itself) and the floor of the resulting value is used to determine the distance from the current agent the teacher will be chosen from. If there is more than one agent occupying a square the same distance from the agent one is selected at random. This process is repeated to give a total of *n* teachers. All the teacher's strategies are then sampled and the most common of these (i.e. the mode) is simply set to be the agent's strategy. This means that a new agent's strategy is directly set to be the same as the majority of his neighbors.[3]

The signal of agent *i*, sig_i, is also influenced by the signals of the same teachers used above. All of agent *i*'s teachers' signals are sampled and the most common (i.e. the mode), m_i, of these signals is recorded. Agent *i*'s signal is then given by the following formula, where a_i is the value of agents *i*'s gene 1 and *rnd* is a uniform random number in the range 0-99:

$$sig_i = (a_i \times m_i) + ((1 - a_i) \times rnd)$$

This means that an agent with an *a* value of 0 has his signal set entirely at random, while a value of 1 means the signal is set to the most common of those around him. This essentially gives the agent a 'choice' as to the weight of attention paid to the signals of those around him when selecting his own signal.

Adulthood: An agent's fitness is evaluated in *f* trials (set to 10 for all results presented here) in which they play a coordination game with other agents. In each trial an agent selects a group of *m* nearby agents as potential partners (*m* was set to 30 for all results presented here). This group is selected in the same way as described above for selecting the group of teachers; however, in this case the standard deviation of the normal distribution is set to be *u* (set to 4 for all results described here). In general we set $u > t$ to model an asymmetry between the group of agents from whom signals and strategies are learned and the group with whom the coordination game is played, as it seems reasonable

[3] This is a very simple method of assigning an agent's strategy and we would like to investigate alternative strategy selection schemes, but as it stands it seems a reasonable assumption that an individual's social norms (as the strategy value is intended to model) are simply inherited from those with whom they are in close contact as children.

that the agents will interact with some agents from further afield in their adult life than in childhood. From this group the partner is selected according to the value of the agent's gene 2, b_i, in the following way. A random real number in the range 0-1 is picked from a uniform distribution, if this number is less than b_i the agent with the signal that is closest to the agent's own signal, i.e. $| sig_i - sig_j |$ is lowest, is selected. If the random number is greater than b_i a partner is selected from the group at random. The result of the game is simply the absolute difference between the two agent's strategies, i.e. $| strat_i - strat_j |$. This value is subtracted from the agent's total fitness score, this means that the highest obtainable fitness is 0, if in all f trials the agent selects a partner with an identical strategy to his own.

Reproduction: After a fitness value has been calculated for each agent, the genes of each agent are subject to recombination and mutation implemented by a standard genetic algorithm (GA) using the fitness value calculated as above to determine the relative reproductive success of each agent. We used standard tournament selection with a tournament size of 4, a crossover rate of 0.6, single point crossover and a mutation rate of 0.01. The mutation operator used was the "reflect" operator described in (Bullock 1999). The spatial nature of the agents' world is taken into account by only allowing tournaments to be played with agents located within three grid squares of each other.

Death: Once all the previously occupied positions in the grid have been filled up by new child agents, the signals and strategies of all previous agents are recorded to be used to teach the next generation. The previous generation is then removed.

4.2. Results

For the results shown here we used the default parameter values as described in the text, and we show results for one typical run of the simulation. We added agents to random points in the toroid, ensuring that only one agent existed in one square. The strategies, signals and gene values were all randomly initialised to a value within the allowed constraints (i.e. signals and strategies were set to within 0-99 and genes were initialised to within 0-1) and the simulation was run for 200 generations.

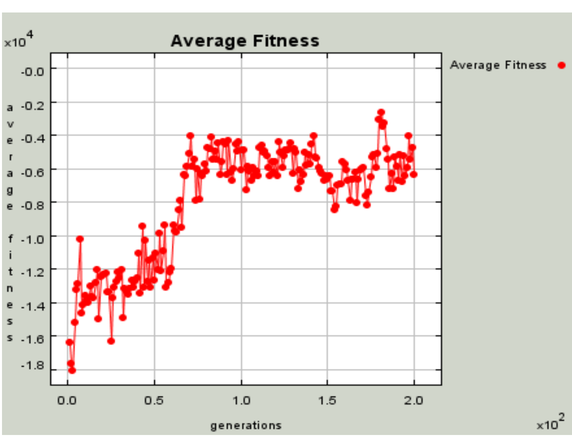

Figure 5: Average fitness of the population

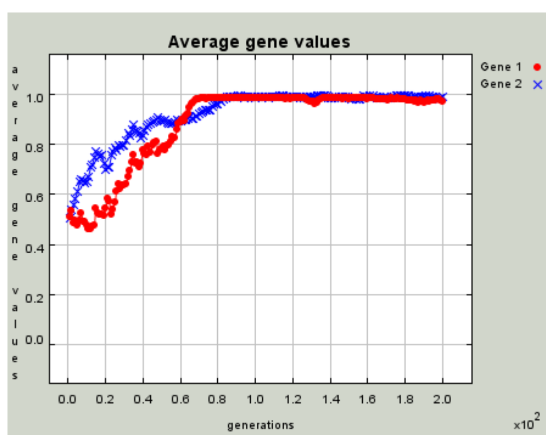

Figure 6: Average values of genes 1 & 2

Figure 5 shows the average fitness of the population at each generation. This graph shows that fitness rapidly increases from the start of the run and then stabilizes after around 75 generations.

Figure 6 shows the average values of genes 1 and 2 across the whole population for each generation. We see that the average value of both genes quickly increase and by around generation 75 the values are very close to 1. The fact that the increase in these values is closely correlated with an increase in fitness demonstrates that larger values have been the subject of strong selection pressure.

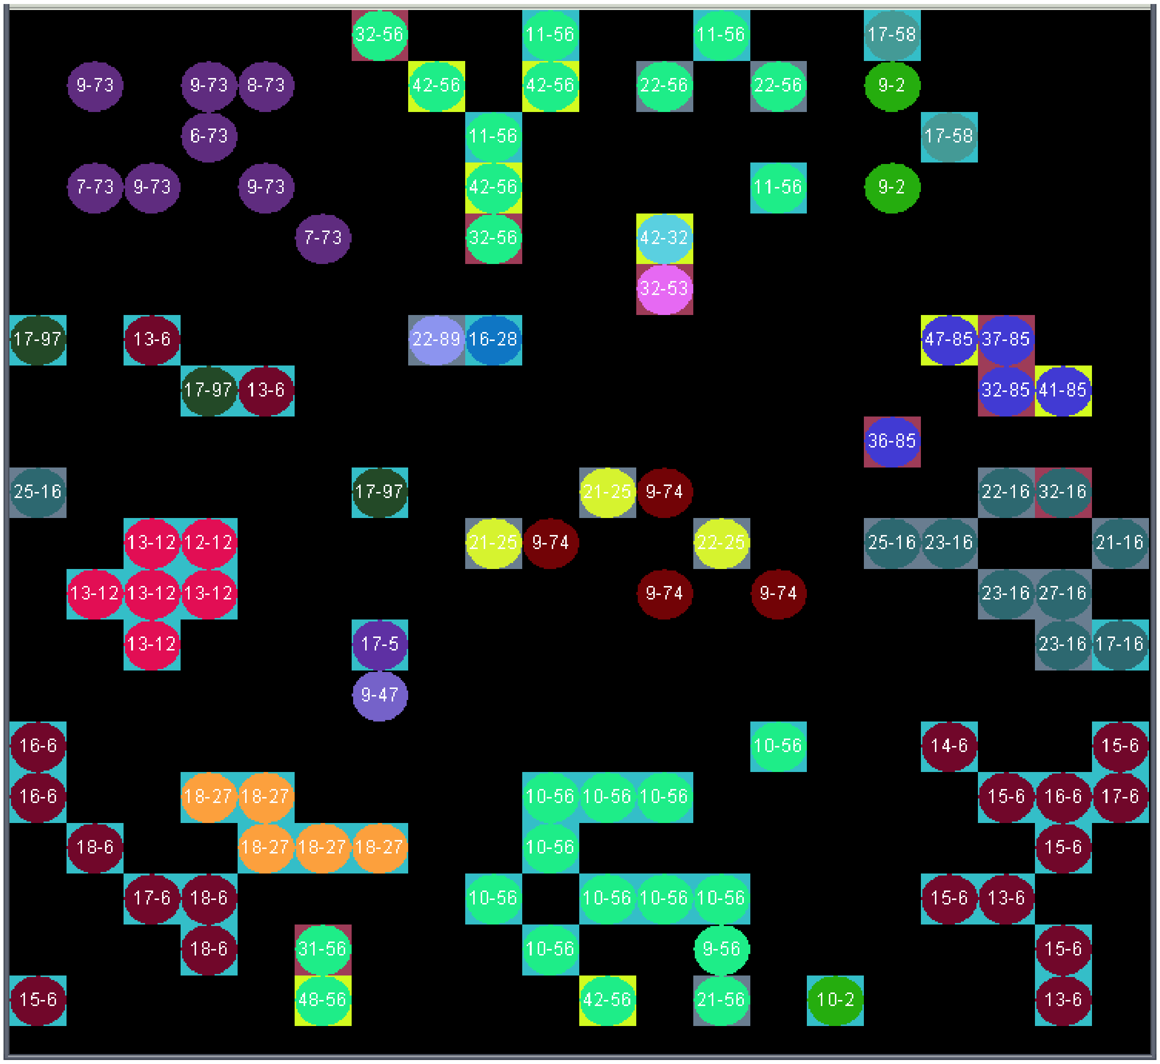

Figure 7: The agents' world at the end of the simulation run

Figure 7 shows an image of the agents' world at the end of the simulation run. Each agent is represented by a circle inside a square, the values shown are each of the agents signals and strategies respectively (The colors of the squares also correspond to the signal and the color of the circle corresponds to the strategy). We can see that most of the agents have formed clearly defined clusters, which share signals and strategies.

While the model does provide some support for our hypothesis, there are several ways in which it could be improved and extended. The parameter values chosen are rather arbitrary and while they seem reasonable, more experiments should be run to test the effect of several parameters on the behavior of the model. In particular the standard deviation of the distributions used to select teachers and coordination game partners seems to have a large effect on the results. We'd also like to investigate different ways of selecting an agent's strategy, including perhaps having some genetic component. Finally, we'd like to investigate the effect of changing the coordination game over time.

5. Conclusions and Further Work

The NetLogo model showed us that, when agents pay attention to each other signals and when these signals coevolve with strategies, locally groups of agents with similar signals as well as equal hidden strategies can evolve. The Repast model consequently demonstrated that, under some conditions and assumptions, agents can evolve to pay attention to signals as an indicator of other agent's strategies, without direct selection on the signal. As a result of this behavior, the agents form local groups that share strategies and signals and interact preferentially with members of the group.

The models above provide support for the hypothesis that the visual and auditory signals that people develop and use to identify group membership have evolved to signal social norms. Furthermore, the models suggest how this social behavior might have arisen through natural selection that works on an individual level. We are happy with these results, but more work can be done to improve the models besides the already mentioned directions. For example, many of the strategies in the models, such as choosing to copy the strategies of one's surrounding networks, are preprogrammed. It would be more convincing if these behaviors "emerged" without being preset. Once these modifications have been made, it would also be interesting to explore more complex questions related to this behavior. For example, it seems that people not only copy the signals of the people around them, but purposely mutate the signals so that the people that they spend more time with will have a more similar signal to them than others. It would be interesting to incorporate this, and other behavior, into the model in the future.

Acknowledgements

We would like to thank the Santa Fe Institute and especially Melanie Mitchell and Tom Carter for organizing the Complex Systems Summer School 2005, Santa Fe, New Mexico. Furthermore, thanks to all the lecturers and other participants for creating a very inspiring environment and for simply having a great month with us.

References

Agha, A. 1998. Stereotypes and registers of honorific language. *Language in Society* 27 (2): pp. 151-194.

Agha, A. 2002. Honorific registers. In: *Culture, Interaction and Language*, K. Kataoka and S. Ide, Eds., pp. 21-63. Tokyo: Hituzisyobo.

Agha, A. 2003. Registers of language. In: *A companion to linguistic anthropology,* A. Duranti, Ed., pp. 23-45. Oxford: Blackwell.

Boyd, R., H. Gintis, S. Bowles, and P. Richerson 2003. The evolution of altruistic punishment. *PNAS* 100: pp. 3531-3535.

Boyd, R. and P. Richerson 1994. The evolution of norms: an anthropological view. *Journal of Institutional and Theoretical Economics* 50(1): pp. 72-87.

Bullock, S. 1999 Are artificial mutation biases unnatural? In: *Fifth European Conference on Artifical Life* (ECAL99), D. Floreano, J-D Nicoud, and F. Mondad, Eds., pp. 64-73. Springer-Verlag.

Halliday, M.A.K. 1988. On the language of physical science. In: *Registers of Written English*, M. Ghadessy, Ed., pp. 162-178. London: Pinter.

Hoyle, S. M. 1993. Participation frameworks in sportscasting play: imaginary and literal footings. In: *Framing in Discourse*, D. Tannen, Ed., pp. 114-145. New York: Oxford.

Maurer, D. 1955. Whiz Mob: A Correlation of the Technical Argot of Pickpockets With Their Behavior Pattern. Gainesville, FL.: American Dialect Society.

Morford, J. 1997. Social indexicality and French pronominal address. *Journal of Linguistic Anthropology* 7(1): pp. 3-37.

Reid, T. B. W. 1956. Linguistics, structuralism and philology. *Archivum Linguisticum* 8(1): pp. 28-37

Salomon, S. 1985. Ethnic communities and the structure of agriculture. *Rural Sociology* 50: pp. 323-340.

Wilensky, U. 1999. NetLogo. *http://ccl.northwestern.edu/netlogo*. Center for Connected Learning and Computer-Based Modeling. Northwestern University, Evanston, IL.