

Emergence of Cooperative Strategies from Elementary Actions in Agents with Neural Nets

Mikhail Burtsev

Peter Turchin

Department of Ecology and Evolutionary Biology

University of Connecticut

Storrs, CT 06269-3042, USA

e-mail: turchin@uconn.edu

January 2005

One of the greatest challenges in the modern biological and social sciences has been to understand the evolution of altruistic and cooperative behaviors. General outlines of the answer to this puzzle are currently emerging as a result of developments in the theories of multilevel selection (1), cultural group selection (2), and strong reciprocity (3). The main conceptual tool used in probing the logical coherence of proposed explanations has been game theory, including both analytical models and agent-based simulations (4-7). The game-theoretic approach yields clearcut results, but the simple structure of payoffs and a small fixed set of strategies, imposed by investigators, in some cases may be an unrealistic assumption. We describe a much more stringent test of the theory by developing a computer model in which agents are endowed with a limited set of receptors, a set of elementary actions, and a neural net in between. Behavioral strategies are not predetermined, but allowed to emerge in the process of evolution. Our results indicate that cooperative strategies can evolve even under such minimalist assumptions, provided that agents are capable of perceiving heritable external markers of other agents.

The artificial world in our model is divided into cells, which either contain a resource bundle, or are empty. An empty cell can acquire resource with a certain probability per time step. Agents are characterized by a set of receptors, a set of effectors, and a neural net connecting receptors to effectors. Each effector is responsible for a particular action. Agents can do nothing (rest), consume resource if it is present (eat), produce offspring (divide), go forward to a neighbor cell (move), make a turn to left or right (turn), and attack another agent if present in the same cell (attack). All actions spend energy taken from the agent's internal store. If internal energy is completely depleted, the agent dies. The least energetically demanding action is rest, the most—attack. Consumption of resource increases internal store of energy subject to an upper limit (the maximum energy that can be stored). When one agent hits another, the victim loses a certain amount of energy, part of which is gained by the attacker.

Sensory inputs of agents include its internal store of energy, whether there is resource in the agent's field of vision (the cell it is in, the neighbor cell in front of the agent, and the cells on the right and left), and how many other agents are in the field of vision. Each agent has external phenotype that is coded by a vector of integer values (markers). The Euclidian distance between an agent's markers and the markers of another agent in the cell (a potential subject for attack) is also a sensory input. Behavior of an agent is controlled by a simple one-layer neural net. The neural net is a matrix of weights that transform the values of sensory inputs into the strengths of output signals, or "motivations." The agent selects the action with the strongest motivation. A specific neural net with a certain combination of weights, therefore, corresponds to a specific behavior strategy. Both weights of the neural net and external markers are inherited by the offspring when an agent divides, subject to a certain rate of mutation. Detailed description of the model could be found in (8).

Each agent in our model has 11 receptors and 7 effectors. The number of potential behavioral strategies in our model is literally astronomical (greater than 10^{1000}). Because of the huge number of potential strategies, it is somewhat difficult to understand just what are the behaviors that evolve in the simulation—each matrix of neural weights is a “black box.” To make sense of our results we confronted agents, which evolved in our simulations, with a discrete set of stimuli and noted the action taken. Thus, we placed each agent in a cell with a phenotypically similar (“in-group”) versus dissimilar (“out-group”) other agent, and took note of the action that the agent took: escape, attack, or no interaction (rest, eat, etc). This approach allowed us to classify strategies into aggressive or not, and those discriminating, or not, in-group versus out-group members.

Our simulation study examined the spectrum of strategies evolving in the full model versus the simplified model in which agents could not detect external markers. We also determined how the carrying capacity of the environment (varied by increasing the size of food bundles) affects the strategies adopted by the agents.

Analysis of the model without markers showed that agents adopted strategies corresponding to those in the well-known game of dove–hawk–bourgeois (9). Doves never attack other agents and attempt to escape when attacked, while hawks make a living by predation on other agents. The bourgeois strategy is to stay in the same cell and immediately attack any invader, while ignoring agents in neighboring cells (unlike the hawks). In the model without markers the dominant strategy is bourgeois, provided that the carrying capacity of a single cell is sufficient for supporting a sedentary agent. Below this threshold, C_1 , the bourgeois strategy is impossible, because agents are forced to keep moving in order to get enough food to survive, and the population is divided between doves and hawks. The long-term population density increases linearly with carrying capacity until it reaches C_1 and then becomes flat (Figure 1). What happens is that once the bourgeois strategy takes over, each cell can be occupied only by a single agent. Even when resources are sufficient to support more than one agent per cell, they fight until only one survives.

In the full model, in which agents could detect phenotypic similarity, three kinds of cooperative strategies emerged. The first one was simply the cooperative version of the dove strategy. Cooperative doves ignored out-group (phenotypic distance large) members, but left cells with in-group (phenotypic distance is small) members in order to avoid competing with them. In the second strategy, agents also left cells with in-group members, but when they encountered out-group members they attacked them. We term this strategy “raven,” because, according to the Russian proverb, “a raven will not peck out the eye of another raven.” The third cooperative strategy was to stay in the same cell with in-group members and collectively fight with any out-group invader. Having to share limited resources of the cell meant that agents utilizing this cooperative defense strategy were small, but they still had a good chance of defeating a large invader because of their advantage in numbers. This strategy resembles the “mobbing” behavior that many species of small birds, such as starlings, use to drive away large predators. For this reason, we call it the “starling strategy.”

The emergence of the starling strategy has a dramatic effect on the relationship between carrying capacity and the long-term population density. For the lower values of carrying capacity, the curve in the full model follows that of the model without markers. But once it exceeds the threshold $C_2 = 2C_1$, the density curve again begins to rise (Figure 1). The implication is that cooperating agents develop the capacity to utilize the resource more efficiently than the noncooperative ones. Analysis of the effect of carrying capacity on the prevalent strategies, evolved by agents, shows the mechanism of this rise. While carrying capacity is insufficient to support survival of at least two agents in a cell ($C < C_2$) the starling strategy cannot invade the

population. Instead, the dominant strategy is raven, whose frequency increases linearly with C for $C < C_1$, and saturates at a high level for $C_1 < C < C_2$ (Figure 2). So far the only difference resulting from the agents ability to use markers is the switch from the hawk to the raven strategy (predators recognize and do not attack in-group members). Once carrying capacity goes beyond C_2 the starling strategy becomes viable. Starlings, however, do not drive ravens to extinction. Instead we observe sustained oscillations in the numbers of starlings, ravens, and doves (Figure 3). As a result, all three strategies manage to coexist in the long term (Figure 2).

In conclusion, our study shows that within the artificial evolution framework it is possible to model not only how one strategy displaces another (or not), but the very process by which novel strategies emerge. Our model did not endow agents with a set of preconceived strategies—all that we have assumed about our agents was that they have a set of elementary sensory inputs and a set of actions. How they connect inputs to actions was molded by the process of evolution. It is remarkable that the agents in our simulations evolved many of the strategies that were postulated by previous researchers. Thus, in the absence of phenotypic markers, three distinct strategies emerged, corresponding to the dove, the hawk, and the bourgeois. In the presence of markers, the evolution resulted in some predictable modifications of these basic strategies, but also in the emergence of a novel one. Cooperative doves avoided competition with in-group members, while cooperative hawks—ravens—avoided attack on phenotypically similar agents. The novel strategy was the starlings, who lived in groups and defended themselves cooperatively against predation. The starling strategy allowed agents to utilize the resources much more efficiently than any other strategy.

While we saw the emergence of cooperative defense, there was no corresponding cooperative attack strategy. This is probably because agents did not have an effector that would allow them to travel in groups in search of prey. In future work we plan to investigate whether adding such action as “follow another agent” could enable evolution of cooperative predation (“wolves”). Another limitation of our study was that agents could transmit traits (including phenotypic ones) only vertically, from parent to offspring. This means that our “in-group members” were also relatives. But one of the greatest puzzles about human ultrasociality is how cooperation between *unrelated* individuals could arise in the process of evolution. This issue can be addressed (and we plan to do so) by allowing cultural transmission of traits between group members.

References

1. E. Sober, D. S. Wilson, *Unto Others: the evolution and psychology of unselfish behavior* (Harvard University Press, Cambridge, MA, 1991).
2. P. J. Richerson, R. Boyd, in *Ethnic Conflict and Indoctrination* I. Eibl-Eibesfeldt, F. K. Salter, Eds. (Berghahn Books, Oxford, 1998).
3. S. Bowles, H. Gintis, *Theoretical Population Biology* **65**, 17-28 (2004).
4. R. Axelrod, *The complexity of cooperation: agent-based models of competition and collaboration* (Princeton University Press, Princeton, NJ, 1997).
5. J. M. Epstein, R. Axtell, *Growing artificial societies: social science from the bottom up*. R. A. Joshua M. Epstein, Ed., Complex adaptive systems (Brookings Institution Press, Washington D C, 1996).
6. M. Nowak, A. Sasaki, C. Taylor, D. Fudenberg, *Nature* **428**, 646-650 (2004).
7. H. Gintis, *Manuscript* (2004).
8. M. S. Burtsev, *Artificial Life* **10**, 397-411 (2004).
9. J. Maynard Smith, *Journal of Theoretical Biology* **47**, 209-222 (1974).

Figure Captions

Figure 1. Population density (average number of agents per cell) as a function of the abundance of resources (the size of resource bundle randomly appearing in a cell with a fixed probability). Broken curve: the model version without external markers; solid curve: the model with markers.

Figure 2. Average proportion of agents employing the raven (solid curve), the cooperative dove (broken curve), and the starling (dotted curve) strategies in the full model with markers. Proportions do not add up to one, because there are other strategies, including the non-cooperative ones.

Figure 3. Numbers of agents employing one of the three cooperative strategies (starling, raven, or dove) as a function of time in one realization of the model.

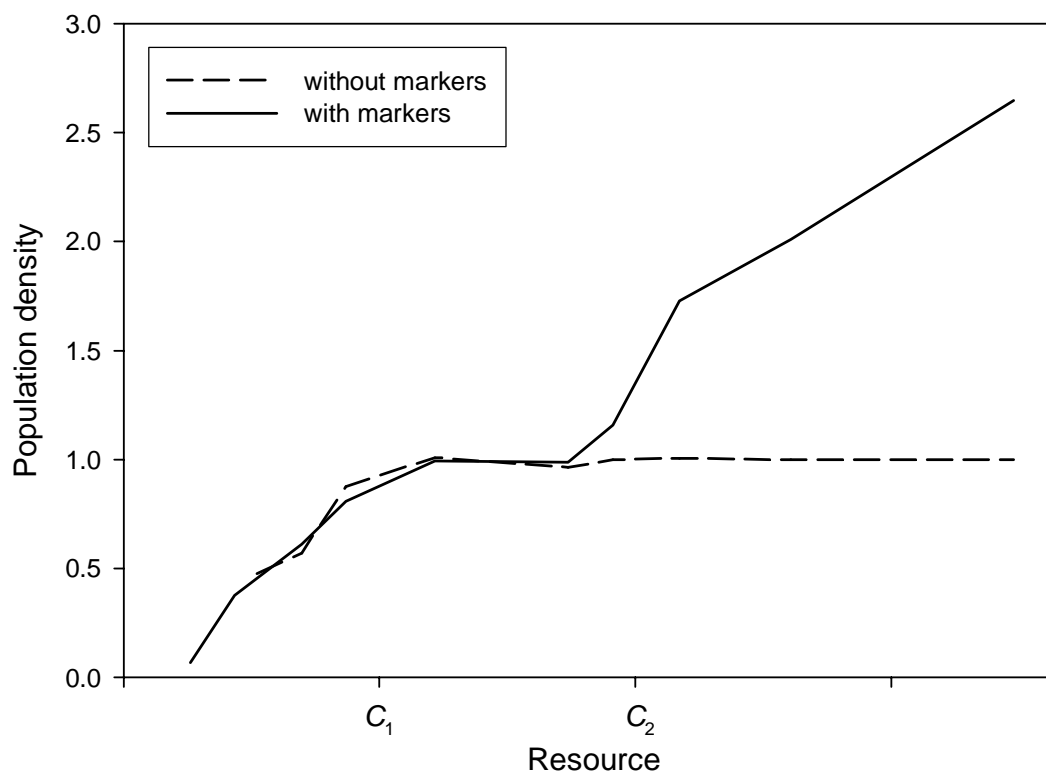


Figure 1




Figure 2

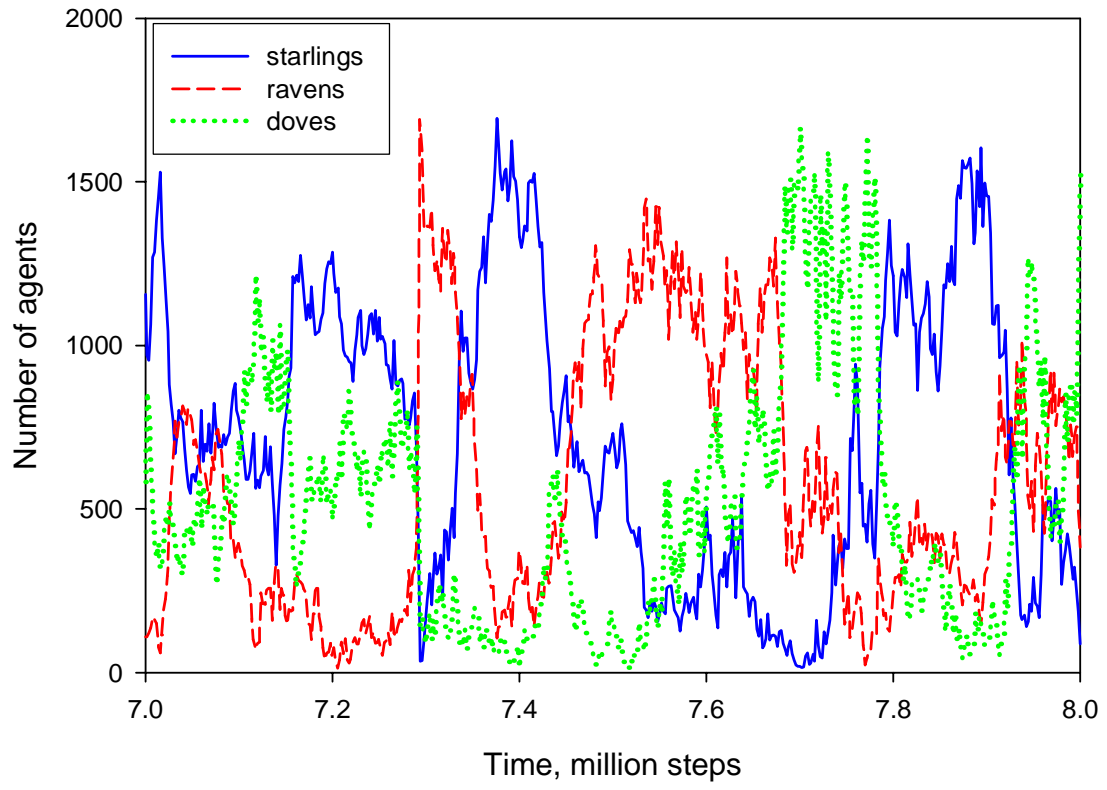


Figure 3