

Language, Altruism and Docility: How Cultural Learning Can Favour Language Evolution

Marco Mirolli^{1,2}, Domenico Parisi¹

¹ Institute of Cognitive Sciences and Technologies, National Research Council, 15 Viale Marx, 00137 Rome, Italy

² Department of Philosophy and Social Sciences, University of Siena, 47 Via Roma, 53100 Siena, Italy
mirolli2@unisi.it

Abstract

Human language serves a number of different functions, one of the most prominent being communicating about relevant features of the environment. From the point of view of the speaker, if the communicated information is advantageous for the hearer but not for the speaker, this is an altruistic use of language, and, as such, it requires an explanation of its evolution. Simon 1990 proposed an explanation of altruism in humans based on the genetically inherited 'docility' of our species. In this paper we present artificial life simulations that apply Simon's ideas to the problem of the emergence of the altruistic use of language described above. From the point of view of evolutionary theory, the present work represents the first attempt to test Simon's 'docility' theory of altruism with agent-based computer simulations. From the point of view of language evolution, our simulations give an original explanation of (the altruistic aspect of) human language based on one of its most peculiar characteristic, namely, the fact that it is culturally transmitted.

Introduction

Among the many functions carried out by human language, one of the most important is its use to inform another individual about some significant feature of the environment. It is an open problem whether this use of language was the principal function for which human language evolved (Bickerton 2002) or language evolution started for more social reasons such as facilitating social interaction and social coordination (Knight et al. 2000). In any case, the use of language for communicating about the environment posits a problem for an adaptationist account of language evolution. In fact, if what is communicated about the environment is useful for the hearer but the act of communicating has no advantage for the speaker, sending appropriate messages about the environment is an altruistic behavior on the part of the speaker.

In this section we describe various evolutionary explanations of altruism, then we briefly review agent-based simulations that have been used to solve the problem of altruistic communication, and finally we describe the rationale of the present work. In the next section we describe our simulations and their results and, finally, in the last section we discuss these results and make some conclusive remarks.

Evolutionary explanations of altruism

To solve the puzzle of the presence of altruistic behaviors in the animal kingdom evolutionary theories have usually adopted one (or more) of the following four kinds of arguments: reciprocity, kin selection, group selection, and cultural selection.

Arguments based on the concept of reciprocity state that if individuals interact repeatedly with each other, then altruistic behavior can evolve because the individuals can adapt their strategies according to the results of previous interactions (Axelrod and Hamilton 1981; Trivers 1971).

Kin selection theory (Hamilton 1964) constitutes the most unchallenged explanation for the evolution of altruism (but see Henrich 2003 and Queller 1992 for some interesting discussions). According to this theory, altruistic behavior can evolve if the product between the benefit the behavior gives to the receiver (B) and the coefficient of kin-relatedness between the emitter and the receiver (r) is greater than the cost for the emitter (C), as stated by Hamilton's Rule: $-C + rB > 0$, that is, $rB > C$.

Another classical - but more controversial - mechanism which has been used for the explanation of the evolution of altruistic behaviors is group selection. According to group selection theory, if a population is divided into groups competing with each other, the total number of altruists in the population can increase even though the number of altruists inside each group is bound to decrease: the reason is that groups with small percentages of altruists will tend to disappear in favour of groups with a larger number of altruists (Sober and Wilson 1998).

Finally, some form of cultural evolution (Cavalli-Sforza and Feldman 1981; Boyd and Richerson 1985) has been suggested as a possible explanation of altruistic behavior (Richerson et al. 2002; Simon 1990). In particular, Simon's explanation of altruism runs as follows: if cultural learning is advantageous for individuals, genetic selection will favour docile individuals, that is, individuals that tend to learn from others¹. Therefore, altruistic behavior can

¹ This is not in contradiction with the fact that cultural learning is unfrequent in the animal kingdom apart from the human species. The reason is that cultural learning requires a number of social, cognitive and neural pre-adaptations which happened to be present in early hominids but not in any other species.

emerge given that its cost is lower than the benefit of being docile and that individuals are not able to distinguish between selfish and altruistic behaviors.

Previous Simulations

In this section we review very briefly some simulations that have focused on the problem we are discussing in this paper, that is, the altruistic character of the use of language for informing conspecifics about some feature of the environment if the information benefits the hearer but not the speaker.

As far as we know, no simulations have been done for testing the plausibility of an account of language evolution which relies on reciprocal altruism. Simulations by Ackley and Littman 1994 and Oliphant 1996 have explored the possibility that kin selection could have played a role in the evolution of language, but Di Paolo 1999 has criticized this work for its improper use of kin-selection. Noble et al. 2001 explore some other adaptive factors that may be relevant for the emergence of a simple signalling system, including group selection, sexual selection, and the handicap principle (Zahavi 1975).

In the present work we test yet another possible mechanism for the evolutionary explanation of the altruistic use of human language which we are dealing with, a mechanism which relies on one of the most peculiar characteristics of human language, that is, the fact that human language is culturally transmitted.

Docility and the evolution of language

In Simon's model the conditions for the evolution of an altruistic behavior are the following:

- 1) there is some advantage d in being disposed to learn from others, i.e., in being docile;
- 2) organisms are not able to evaluate the contribution of each behavior to their own fitness;
- 3) the advantage d of being docile is greater than the cost c of the altruistic behavior.

With his argument Simon intended to show how the presence of altruism among human beings should be considered differently from the presence of altruism in other species in that cultural learning plays a crucial role in humans and this makes possible the emergence of altruism with a mechanism peculiar to our own species. As the use of language that we are dealing with is altruistic and as human language is the only communicative system in the animal kingdom to be culturally transmitted, Simon's explanation of altruism seems to be applicable to language evolution.

The present work represents the first attempt (to our knowledge) to test Simon's docility theory with agent-based computer simulations and, at the same time, gives an original explanation for the evolution of (the altruistic aspect of) human language based on its learned character.

Simulations

In order to test the soundness of Simon's theory with respect to the evolution of human language, we ran two simulations based on a simplification of the simulative scenario used by Cangelosi and Parisi 1998 and discussed in Parisi 1997. First, we describe the simulative set-up common to both simulations, then we indicate what makes the two simulations different, and finally we describe the main results of the simulations.

The simulative set-up

The population is a succession of 500 generations of 100 individuals each. In each generation the 20 individuals which have the most energy at the end of life are selected for reproduction. Each individual generates 5 offspring and the $20 \times 5 = 100$ offspring constitute the next generation². All individuals have the same network architecture but connection weights can vary among individuals (see below). Each individual lives in its own copy of the environment which is a linear succession of 11 cells. At the beginning of each 'epoch' an individual is placed in the first cell while the last cell contains a 'mushroom'.

There are 210 edible mushrooms and 210 poisonous ones, each different from all the others. The perceptual properties of the mushrooms are represented as variations from two prototype sequences of ten bits (+1 and -1) each, one for the edible mushrooms and one for the poisonous mushrooms. The 210 edible mushrooms are the 210 bipolar patterns of ten bits which have four +1s (that is, all and only the patterns that differ by 4 bits from the prototype which has all -1s). The perceptual properties of poisonous mushrooms are the 210 bipolar patterns which have six +1 (those that differ by 4 bits from the prototype which has all +1s). If the individual enters the mushroom cell located at the end, the individual eats the mushroom. If the mushroom is one of the 10 edible mushrooms, the individual's energy is increased by 30 energy units. However, if the mushroom is poisonous, its energy is decreased by 5 units. Furthermore, an individual's energy is decreased by 1 unit for each step from the initial to the final cell. This is the fitness formula:

$$f(x) = 30 \times \text{number of edible mushrooms eaten} - 5 \times \text{number of poisonous mushrooms eaten} - \text{number of steps}.$$

As the life of each individual last 420 'epochs', one for each possible mushroom, the maximum possible fitness will be:

² The results of our simulations are robust with respect to changes in most of the simulation's parameters, including: selection algorithm (range-based vs. proportional to fitness); use vs. non-use of cross-over; probability of mutations; type of mutation (substitution of a weight vs. gradual change); presence vs. absence of limits to weights' values; learning rate and momentum values (for the simulation with learning).

$F_{\text{Max}} = 30 \times 210$ (edible mushrooms eaten) $- 5 \times 0$ (poisonous mushrooms eaten) $- 210 \times 10$ (steps necessary for eating a mushroom) $= 6300 - 2100 = 4200$

The organism's behavior is controlled by a neural network with 12 input units, 2 hidden units, and 3 output units. Ten input units are 'visual' units. If the organism is sufficiently near the mushroom, that is, it is in the last but one cell of the corridor, the perceptual properties of the mushroom are encoded in this set of input units; otherwise, the activation of all the visual units is set to zero. The other two input units are the linguistic ones: their activity depends on the linguistic signal which is produced by another individual randomly chosen from the rest of the population, the speaker. All the input units are fully connected with the two hidden units which in turn are fully connected with all three output units. The activation of the first output unit is thresholded to a value of either 1, in which case the organism moves one step forward, or -1 , in which case the organism stays still. The continuous activations of the other two output units, in the interval $[-1; 1]$, constitutes the signal produced which is copied, when the organism acts as a speaker, in the two linguistic input units of the hearer.

We ran two simulations: in the first one, the genome of our organisms contains their connection weights, while in the second one, the genome is constituted by a single gene, the docility gene, which stands for the disposition of the organisms to learn from their parents.

The two conditions: genetic selection and cultural selection through docility

In the first simulation – the “genetic simulation” – the values of the connection weights of the initial generation are randomly chosen in the range $[-0.5; 0.5]$, but the connection weights of each successive generation are inherited from parents to offspring with each weight having a probability of 2% to have its value changed by a random number in the interval $[-0.5; 0.5]$.

In the second simulation – the “docility simulation” – the connection weights of all individuals are always random at birth and they are not inherited from parents. Instead, the genome of these organisms is constituted by one only gene, encoded as an integer number, which stands for an individual's 'docility', i.e., the number of learning trials for that particular individual. In the first generation each individual is assigned a random value in the interval $[0; 200]$ for this gene and this value is genetically transmitted with a 2% probability of being changed by adding or subtracting a random number in the range $[-100; 100]$. In any case, docility is forced to stay in the interval $[0; 500]$.

The life of organisms in this second simulation is divided into two periods: infancy and adulthood. During infancy, the organism is supposed to follow its parent and learn from it how to behave in different situations. In short, the docility gene determines the number of back-propagation learning cycles to which the infant exposes itself: the learning is imitative in that the teaching input of the back-

propagation algorithm comes from the output of the infant's parent. Since there are three kinds of situations to which organisms are exposed during their life, there are also three different learning conditions: 1) 'comprehension learning', 2) 'decision learning' and 3) 'naming learning'.

Comprehension learning corresponds to the situation in which the tested organism is distant from the mushroom and has to decide whether to move or not to move according only to the signal it receives from another organism; decision learning corresponds to the situation in which the organism is near the mushroom and so its decision whether to move or not to move depends on both the visual input and the linguistic input; finally, naming learning corresponds to the situation in which the organism acts as a speaker: it receives only the perceptual properties of a mushroom and has to produce a linguistic signal.

In short, for each learning cycle determined by one's docility, this is what happens:

- 1) one of the three learning situations is randomly chosen together with one of the 420 mushrooms;
- 2) the appropriate input is given both to the learner and to its parent;
- 3) both the organism's output and its parent's output are calculated;
- 4) the output of the parent is given to the child as teaching input (a random value chosen in the interval $+0.25/-0.25$ is added to the teaching input³);
- 5) finally, the child's connection weights are changed according to the back-propagation algorithm (with a learning rate of 0.3 and a momentum of 0.8).

After infancy, individuals start adult life, which is identical to that of the genetic simulation.

Results

We describe the average results of 10 replications of both simulations using three measures: language quality, average fitness, and average value of the docility gene (in the pictures, all values are normalized with respect to their maximum possible value).

This is the way we calculate language quality. A linguistic signal is constituted by a vector of two continuous numbers in the range $[-1; 1]$ (the vector of activation of the linguistic output units of an organism). So, a signal can be considered as a point in a bi-dimensional Cartesian space. Let's call E and P the set of points ('clouds') which represent the signals produced by the organisms of one generation in presence of all edible and poisonous mushrooms, respectively. A good language is one in which the two clouds E and P are 1) small (all

³ The results are robust even with respect to the quantity of noise added, provided that this quantity is adequate for cultural evolution: if there is no noise, there is no room for improvement in behavioral capacity; on the other hand, if there is too much noise, good behaviors cannot be preserved. All noise values between 0.1 and 0.4 produce the same qualitative results.

mushrooms belonging to the same category are named in similar ways) and 2) distant from each other (mushrooms of different categories are named in different ways). A measure for 1) is the mean distance of the points of a cloud from its geometric center; a measure for 2) is the distance between the two centers. We normalize those two values in the range [0; 1] so that the maximum quality obtainable (1) is achieved when each cloud is a single point, in opposite corners of the space. In order to plot a single value for each simulation, we measure the language quality as the product of the two normalized values.⁴




Figure 1: Average language quality of the genetic (thin line) and docility (thick line) simulations as a function of generations.

As it turns out, the language quality of the genetic simulation (fig. 1, thin line) is very bad: in this condition a good communication system does not emerge. The reason is the altruistic character of the type of communication we are dealing with here. Since producing appropriate messages (good speaking) gives advantages only to hearers but not to speakers, egoists (bad speakers) are selected against altruists (good speakers); hence the very low quality of language. As a result, the average fitness of the organisms of this simulation is sub-optimal. In fact, the presence of a good communication system is a necessary condition for optimal behavior. Without being told by others an organism cannot know which kind of mushroom is present at the end of the corridor and has to waste energy to go and check by itself. We can also calculate the maximal possible fitness for organisms which are not helped by language at all:

⁴ The reason for using the product instead of the mean is that the product gives high values only if both values are high, and this seems quite appropriate. For example, the language quality of a communication system in which both clouds are collapsed in the same point (all mushrooms are named in exactly the same way) is certainly 0: this is in fact the value this system would reach using the product of the two measures while by using the mean the same system would reach the inappropriate value of 0.5.

$$F_{\text{MaxWithoutLanguage}} = F_{\text{Max}} - 210 \text{ (poisonous mushrooms)} \times 9 \text{ (steps to be made in order to see the perceptual properties of the mushroom)} = 4200 - 1890 = 2310$$




Figure 2: Average fitness of the genetic (thin line) and docility (thick line) simulations as a function of generations.

If we normalize this quantity, we get $2310 / 4200 = 0.55$, that is just a little lower than the average fitness reached by the genetic simulation, which fluctuates in between 0.6 and 0.65 (fig. 2, thin line).

The situation of the docility simulation is very different. In this case, there is a strong selective pressure for the evolution of docility, since organisms that do not learn culturally are bound to behave randomly. As a result, the value of the docility gene increases constantly until it reaches almost its maximum value (fig. 3). To this increase in docility corresponds a parallel increase in the quality of the language produced by those organisms, which reaches the quite high value of about 0.75 (fig. 1, thick line). As it turns out, the correlation between docility and language quality is very high: 0,988. As a result, the organisms of this simulation can exploit all the advantages given by a good communication system and consequently their average fitness reaches almost the maximal possible value (fig. 2, thick line).




Figure 3: Average value of the docility gene of organisms of the docility simulation as a function of generations.

Discussion and Conclusions

There are a number of possible solutions to the problem of the evolution of altruism which rely on reciprocity, kin selection, group selection or cultural evolution, and we think that more than one of these factors may have played a role in the evolution of the use of language which is taken into consideration in our simulations.

Pedone and Parisi 1997 have suggested that the crucial factor for the evolution of an altruistic behavior is the similarity of behavior between interacting individuals. Henrich 2003 presents a generalization of Hamilton's Rule which substantiates this suggestion: what really matters for the evolution of altruism is the probability for an altruist to encounter another one. Consequently, all the theoretical solutions to the problem of altruism consist in finding plausible mechanisms for maintaining this probability high. Our docility simulation strongly confirms Simon's theory according to which docility can be such a mechanism. In fact, if docility evolves due to the egoistic advantages it confers to organisms, then cultural transmission can guarantee the similarity of behaviors necessary for altruism to emerge.

Turning specifically to the evolution of human language, our simulations suggest that the use of language for informing others about the environment (unless it benefits both speaker and hearer in coordinating their behavior, e.g., in group hunting) might have emerged relatively lately during hominid evolution, namely, after hominids had become docile and cultural evolution had started.

In fact, from the point of view of the informer, this use is an altruistic behavior, present in the human species, which needs an evolutionary explanation. The results of the simulation in which behaviors are genetically inherited (through the neural networks' connection weights) confirm the theoretical prediction that in such a condition a good communication system does not emerge because of its altruistic character. Consider also that in our simulations communication is in a sense hardwired: our organisms are *forced* to produce signals for others. So, we are not actually dealing with the emergence of communication as such, but with the emergence of a *good* system for communicating about the environment. This is justified by the assumption that hominids' proto-language may have initially served other, more 'social', functions, such as strengthening social relationships (Dunbar 1996). This use of language does not pose the theoretical problems of altruism, since it is not altruistic.

So, this is the evolutionary scenario we are suggesting. First, a very complex social structure created the basis for the evolution of the first kind of hominids' communication system which in turn favoured the development of a still more complex social structure. Second, this more complex social structure constituted one of the preconditions for the evolution of cultural learning. Finally, (proto-)language started to be used also for informing others about the environment. Our simulations show that the homogenisation of behavior induced by cultural

transmission could have favoured the emergence of this kind of altruistic use of language.

Acknowledgments. The research presented in this paper is supported in the framework of the European Science Foundation EUROCORES programme "The Origin of Man, Language and Languages"

References

- Ackley, D. H.; and Littman, M. L. 1994. Altruism in the Evolution of Communication. In R. A. Brooks and P. Maes eds. *Artificial Life IV: Proceedings of the International Workshop on the Synthesis and Simulation of Living Systems*, Cambridge: MIT Press, 40-48
- Axelrod, R.; and Hamilton, W. D. 1981. The Evolution of Cooperation. *Science* 211: 1390-1396
- Bickerton, D. 2002. Foraging Versus Social Intelligence in the Evolution of Protolanguage. In A. Wray ed. *The Transition to Language*. Oxford: Oxford University Press
- Boyd, R.; and Richerson, P. J. 1985. *Culture and the Evolutionary Process*. Chicago: University of Chicago Press
- Cangelosi, A.; and Parisi, D. 1998. The emergence of a 'language' in an evolving population of neural networks. *Connection Science* 10(2): 83-97
- Cavalli-Sforza, L. L.; and Feldman, M. W. 1981. *Cultural transmission and evolution: a quantitative approach*. Monographs in Population Biology 16, Princeton: Princeton University Press
- 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. In D. Floreano, J.D. Nicoud, and F. Mondada eds. *Advances in Artificial life Proc. ECAL'99*, LNAI 1674, Lausanne: Springer-Verlag, 504-513
- Dunbar, R. I. M. 1996. *Grooming, gossip and the evolution of language*. London: Faber and Faber
- Hamilton, W. D. 1964. Genetic evolution of social behavior. *Journal of Theoretical Biology* 7(1): 1-52
- Henrich, J. 2003. Cultural group selection, coevolutionary processes and large-scale cooperation. *Journal of Economic Behavior and Organization*
- Knight, C.; Studdert-Kennedy, M.; and Hurford, J. eds. 2000. *The evolutionary emergence of language: social function and the origins of linguistic form*. Cambridge, Mass.: Cambridge University Press

Knudsen, T. 2003. Simon's selection theory: why docility evolves to breed successful altruism. *Journal of Economic Psychology* 24: 229-244

Noble, J.; Di Paolo, E. A.; and Bullock, S. 2001. Adaptive factors in the evolution of signalling Systems. In A. Cangelosi and D. Parisi eds. *Simulating the evolution of language*. London: Springer-Verlag, 53-78

Oliphant, M. 1996: The dilemma of Saussurean communication. *Biosystems* 37(1-2): 31-38

Parisi, D. 1997. An Artificial Life approach to language. *Brain and Language* 59: 121-146

Pedone, R.; and Parisi, D. 1997. In what kinds of social groups can altruistic behavior evolve?. In R. Conte, R. Hegselmann and P. Terno eds. *Simulating social phenomena*. Berlin: Springer-Verlang, 195-201

Queller, D. C. 1992. A general model for kin selection. *Evolution* 46: 376-380

Richerson, P.; Boyd, R.; and Henrich, J. 2003. The cultural evolution of cooperation. In P. Hammerstein ed. *Genetic and cultural evolution of cooperation*. Cambridge, Mass.: MIT Press

Simon, H. A. 1990. A mechanism for social selection and successful altruism. *Science* 250: 1665-1668

Sober, E. R.; and Wilson, D. S. 1998. *Unto others: The evolution and psychology of unselfish behavior*. Cambridge, Mass.: Harvard University Press

Trivers, R. L. 1971. The evolution of reciprocal altruism. *Quarterly Review of Biology* 46(1): 35-57

Zahavi, A. 1975. Mate selection. A selection for a handicap. *Journal of Theoretical Biology* 53: 205-214