

n



n



n




n

n



n



n



UNIVERSITY OF



communication evolves when it is beneficial to have one's behavioural intentions predicted (or falsely predicted). Some authors (Grafen, 1990) have suggested that this places too much emphasis on dishonest communication, and a typical definition from the biology literature (Lewis & Gower, 1980) is as follows:

... the transmission of a signal or signals between

1.3 Work in theoretical biology, ethology, and behavioural ecology

There exists an extensive empirical and theoretical literature on animal communication. Empirical work tends to describe a particular type of signalling within one or more animal groups.

then validate and extend that model using iterative, computational techniques. For example, de Bourcier and Wheeler (1994) look at aggressive signalling and territoriality. They state that their method of *synthetic behavioural ecology* “is pitched at an intermediate level between, on the one hand, abstract theories

2.2 Method

MacLennan and Burghardt¹ used populations of simulated organisms that they refer to as “simorgs”. The simorgs all have access to a shared global environment, and each individual has access to a private local environment. The global environment provides a medium for communication, and the local environments are a source of significant information that the simorgs may evolve to communicate *about*. Each of the environments is as simple as possible, represented by a single variable that can take on a finite number of values. It is emphasised that “there are no geometrical relations among [the simorgs]... they are not in a rectangular grid, nor are some closer than others” (MacLennan & Burghardt, 1994, p. 166).

MacLennan and Burghardt suggest, by way of analogy, that the global environment can be thought of as the air, capable of transmitting only one sound at a time, and the local environments can be considered exclusive hunting grounds, into which different species of prey may wander. In other words, states of the global environment have the potential to be exploited as signals, and states of the local environment are particular circumstances that it will pay simorgs to signal about.

Simorgs have only two classes of behavioural choice open to them: they can *emit* a signal (into the global environment), or they can *act* in an attempt to respond to the signal of another. The state of the global environment can be changed by any of the simorgs if that simorg emits a signal when its turn comes; the states of the simorgs’ local environments are not under their control, and are periodically reset to random states.

In the synthetic world, simorgs achieve fitness by successfully cooperating with another simorg. Specifically, by responding to a signal with an action that matches the local environment state of the signaller. When this occurs, both the signaller and the respondent are rewarded with a point of fitness. Continuing their analogy, MacLennan and Burghardt suggest that this is to be regarded as two hunters bringing down a prey animal that neither could bag alone. Assuming that successful communication has taken place, note that the signal does not mean “I’ve got some prey here”, but “I’ve got prey of type λ here; would you mind helping out with action- λ ?” The state of another simorg’s local environment is not directly knowable, and successful cooperation can only come about through a lucky guess or the employment of communication.

In order to implement their ideas in a computer program, MacLennan and Burghardt had to make a number of somewhat arbitrary practical decisions. Thus, time in the synthetic world is discrete. Once each time step, the simorgs respond (i.e., act or emit) in a fixed order; effectively they are arranged in a ring. The program keeps track of the “owner” of the symbol currently occupying the global environment. It is possible, for example, for one simorg to emit and then earn several fitness points consecutively as a series of other simorgs act in response to the same persistent signal.

Every five time steps (one environment cycle) the local environments are reset to a random value, ensuring that the simorgs must react to changing circumstances if they are to succeed. Every fifty time steps there is a breeding

¹ MacLennan and Burghardt’s methodology is difficult to describe briefly, and the reader is referred to their work (MacLennan & Burghardt, 1994) and MacLennan’s earlier article (MacLennan, 1991) for a complete account.

cycle: two fit simorgs are stochastically selected as parents and, using two-point crossover with a small chance of

2.3 Results and conclusions

MacLennan and Burghardt report that communication did indeed evolve in the synthetic world. The results reported are for a single random initial population subjected to each of the three conditions; MacLennan and Burghardt assure us that these results are typical. In the $-L^-$ condition, there was only a very slight increase in fitness over the length of an experimental run, whereas in the $+L^-$ condition the rate of fitness increase was an order of magnitude greater. In the $+L^+$ condition, the rate of fitness increase was higher still. MacLennan and Burghardt conclude that, when it is not suppressed, communication is selected for and leads to higher levels of cooperation. The provision of the single case learning rule further increases the effectiveness of the communicative strategy.

Analyses of the denotation matrices showed that in the $-L^-$ condition, the pattern of symbol use was almost random. When communication was permitted the matrices were quite structured, as measured by the entropy statistic. Visual inspection of the denotation matrices made it clear that certain symbols had evolved to (almost uniquely) represent certain local states. There was ambiguity in two senses: sometimes a symbol would be used to represent two or more states, and sometimes a state was represented by two or more symbols. MacLennan and Burghardt suggest that the ambiguity is either due to two subpopulations using different symbol dialects, or to individual simorgs using one symbol to represent two different states.

That there should be any fitness increase at all in the $-L^-$ condition is not obvious. MacLennan and Burghardt refer to this phenomenon as “partial cooperation through co-adaptation”, and regard it as a “low-level effect” (1994, p. 185). They explain it by noting that simorgs can do better than chance if they emit a symbol only in a subset of their local situations, and guess actions within that same subset (see s40Td[(qC)TJ1Maceainls

	MacLennan & Burghardt	Our results		
		<i>Mean</i>	<i>SD</i>	<i>p</i>
Fitness increase				
$-L^-$	0.37	0.99	1.16	n.s.
$+L^-$	9.72	14.6	6.54	n.s.
$+L^+$	37.1	10.6	10.6	0.025
Final mean fitness				
$-L^-$	≈ 6.6	6.74	0.43	n.s.
$+L^-$	10.28	12.71	2.68	n.s.
$+L^+$	59.84	46.13	4.02	0.004

Table 1: Rate of fitness increase (determined by linear regression and measured in units $\times 10^{-4}$ breeding cycles) and final mean fitness scores. Note that mean fitness data was a moving average smoothed over 50 breeding cycles, and that final mean fitness in the $+L^+$ condition is much higher because the simorgs had four chances per environment cycle to respond after correction by the learning rule: fitness scores in this condition start at 40+ rather than the usual chance level of 6.25. Rates of increase are thus a better comparison across conditions.

hypothesis that MacLennan and Burghardt’s result could have come from the same distribution as our data (“n.s.” means not significant, i.e. $p > 0.05$).

The $-L^-$ and $+L^-$ conditions showed slightly higher rates of fitness increase in our own experiment. More importantly, the rate of fitness increase in the $+L^+$ condition was more than three times *smaller* in our data than in

	MacLennan &	Our results		
	Burghardt	<i>Mean</i>	<i>SD</i>	<i>p</i>
$-L^-$	5.66	4.96	0.15	< 0.001
$+L^-$	3.95	3.36	0.50	n.s.
$+L^+$	3.47	4.45	0.36	0.015

Table 2: Entropy statistics, calculated on the denotation matrix of the final 50 breeding cycles of the experiment. An entropy value of 6 would indicate a completely random matrix. A value of 3 indicates a perfectly structured matrix, with one symbol per situation.

4 Extension and critique

Having described the methods used by MacLennan and Burghardt, and noted the degree to which our results match theirs, we now wish to comment critically on certain aspects of their experiment. Several questions are raised as to what might be an appropriate methodology for studying the evolution of communication, and we hope to answer these questions in section 5.

4.1 No geometry?

MacLennan and Burghardt claim that there are “no geometrical relations” (1994, p. 166) among the simorgs. This is in keeping with their goal of constructing a synthetic world that is as simple as possible while still permitting communication to evolve. If the simorgs were arranged on a toroidal grid and could communicate only locally, for example, this would certainly complicate things.

However, in the current set-up, the simorgs are effectively arranged in a ring. As MacLennan and Burghardt (1994, p. 170) put it, “The simorgs react one at a time in a fixed order determined by their position in a table.” Thus there is at least a topology, if not a geometry: simorgs will tend to receive signals from their immediate neighbours in one direction, and send signals to their neighbours in the other direction.

The experiment could have been performed without this modest topological assumption if the simorgs were updated in a different random order at each time step. We modified our version of the program to use just such an updating procedure. Table 3 shows the rates of fitness increase and final fitness scores under this method.

There is a dramatic difference between the two updating methods. In the communication only ($+L^-$) and no communication ($-L^-$) conditions, similar performance is observed under both updating methods. The effect of the learning rule, on the other hand, depends very much on the updating method used: under random-order updating, the rate of fitness increase is much higher. Curiously, the rates of fitness increase under random-order updating come closer to the rates observed by MacLennan and Burghardt — perhaps this is a clue as to the cause of our differing findings.

Furthermore, random-order updating clears up an irksome feature of MacLennan and Burghardt’s results. Fitness in the learning condition commences close to the random level of 6.25 (see the notes to table 1), which makes mean

	<i>Mean</i>	<i>SD</i>	<i>Effect</i>
Fitness increase			
$-L^-$	0.94	1.52	-4.5%
$+L^-$	18.6	7.05	+27.4%
$+L^+$	33.7	13.8	+218%
Final mean fitness			
$-L^-$	6.76	0.53	+0.23%
$+L^-$	14.47	2.83	+13.9%
$+L^+$	22.24	5.21	-51.8%

Table 3: Effect of random-order updating. Rate of fitness increase $\times 10^{-4}$ breeding cycles (determined by linear regression), and final mean fitness scores are shown, with means and standard deviations across 20 runs. The “effect” column compares the random-order results with our standard updating results (see table 1); note that if the updating method was not influencing the results, we would expect this value to be close to zero.

fitness directly comparable with the other conditions — note the 51.8% drop in final fitness scores. Under standard updating, a simorg will often have its look-up table corrected on the first time step of an environment cycle, then find itself in exactly the same context on the next four time steps, and score up to four “free hits”. When simorgs are responding in a different random order each time step, it is no longer the case that a simorg will be communicating with the same near neighbours every time, and the learning rule loses this bonus property.

The most important point about the random updating procedure, however, is that it demonstrates that MacLennan and Burghardt’s results could be dependent upon such apparently minor assumptions built in to their procedure. Their goal is to uncover general laws that can be translated back into the realm of real biology, but if the effect of learning on the evolution of communication is dependent on the updating method used, it is difficult to know what biological conclusions should be drawn. Does learning facilitate the development of a communicative system, or doesn’t it?

4.2 Dialects or suboptimal look-up tables?

MacLennan and Burghardt, noting the ambiguous symbol use evident in the denotation matrices, comment that “we cannot tell from [the denotation matrix] whether this multiple use of symbols results from two subpopulations or from individual simorgs using the symbol to denote two situations” (1994, p. 179). The idea that there could be subpopulations using different dialects seems quite plausible, especially given that the topology of the simorgs’ environment (see section 4.1) ensures that simorgs will only be communicating with near neighbours. One can imagine a series of simorgs using variant *A* in one section of the ring, shading gradually into variant *B* in the opposite section, and back again.

MacLennan and Burghardt claim that the facts of the matter could easily be uncovered: given that the underlying finite state machines are available

in computer memory, “there need be no mystery about how the simorgs are communicating, because the process is completely transparent” (1994, p. 179). However, they make no clear statement as to whether they in fact believe there are two or more subpopulations using variants of the evolved “language”. MacLennan, in his earlier paper, is less conservative: “the differing use of symbols in various contexts makes it quite possible for every simorg to be using a different dialect of the ‘language’ manifest in the denotation matrix.” (1991, p. 653).

In an attempt to resolve this question, we used a convergence statistic in our experiments. We examined each position on the genome in turn, and calculated the mean percentage of identical entries across the population of simorgs. Thus, a convergence statistic of 100% would indicate a population of simorgs with identical genomes and, thus, identical FSMs.

In runs of 5000 breeding cycles duration, the final convergence statistic was typically between 75% and 85%. This is not conclusive: it means that up to 25% of the simorgs could have been different from the norm, or that 25% of the genetic material of each simorg could be unique, and so leaves plenty of room for the possibility of different dialects. However, when the runs were extended to 2×10^4 breeding cycles or more, final convergence statistics in the $+L-$ condition were approximately 99.5%, and denotation matrices were qualitatively similar, i.e. they still showed ambiguous communication. It is implausible to suggest that there might be different dialects when the simorgs in a population are 99.5% identical to each other. We conclude that the suggestive ambiguity in the denotation matrices is nothing more than the net effect of (more or less) the whole population using a single, inefficient “language” that sometimes represents a state by more than one symbol, or uses one symbol to denote more than one state.

Despite their stated wariness about adopting any sort of intentional stance towards the simorgs (1994, p. 163), we believe that MacLennan and Burghardt are not immune to the temptation to think of the simorgs as intentional agents, and that, in this case, that temptation has led them astray. Language without the scare quotes is undoubtedly the exclusive province of sophisticated intentional agents (argued in e.g. Bennett, 1976; Dennett, 1987), but having drawn the analogy between human language and simorg communication, MacLennan and Burghardt were too ready to suspect that, like real language users, simorgs might have dialects. We are not claiming that they were wrong to draw the analogy in the first place; indeed, a central premise of adopting the intentional stance is to take such analogies very seriously. However, in dealing with simple systems like the simorgs, lower level mechanistic explanations are potentially open to us. The intentional story (dialects) can be shown, via the convergence statistic, to be inferior to the mechanistic story (less than ideal structure in the FSM look-up tables).

4.3 Consequences of the FSM look up table approach

Imagine for a moment that you are a simorg. Disregarding the fact that simorg “decisions” are entirely determined by the look-up table, imagine that you have decided to emit a symbol. The only context that is important to you is the state of your local environment: you need to choose the right symbol to describe your situation, according to the “language” conventions that have developed. The identity of the symbol in the global environment is unimportant, because you’re

going to overwrite it anyway.

Similarly, if you're going to act, you don't care about the state of your local environment; you only want to interpret the global symbol in such a way as to correctly match the environment of the last emitter, and thereby score a point of fitness.

For the real simorgs of MacLennan and Burghardt, things are not this simple. There is no prior decision to emit or to act, only the consultation of a table with an entry for every possible combination of local and global environment states. As MacLennan and Burghardt put it, "finite-state machines have a rule for every possible condition" (1994, p. 168).

Surprisingly, this means that the choice of the FSM architecture makes evolving a communication system harder for the simorgs than it might be under some other control architectures. For example, if during a particular run it became ad-

the number of local states. To date we have only seen this phenomenon evolve spontaneously when $L \leq G \leq 4$, but the principle remains.

The trouble with this result is that one presumably does not want to call it an evolved communication system or “language”, even though the simorgs are ostensibly fitter than ever before. If the global environment is (almost) always in the same state, it is difficult to describe it as carrying any information. The simorgs in such a situation appear to be exploiting a loophole in the experimental design.

MacLennan and Burghardt were aware of this possibility (see section 2.3). They saw it as most relevant to the $-L-$ condition, in that it provided an explanation for the otherwise mysterious increase in fitness observed. MacLennan (1991, p. 653) felt that “in most cases [it] is a low level effect that is unintrusive and can be ignored”.

4.5 Fewer symbols: faster improvement

The point outlined in 4.4 has a number of implications. Given that the optimal strategy involves the utilisation of only one symbol, we hypothesised that giving the simorgs progressively *fewer* symbols to work with would steer them towards that strategy and thus *improve* their performance. This contrasts with the intuitive hypothesis that n local states will require simorgs to use n symbols to denote them. MacLennan and Burghardt seem to have assumed the truth of the intuitive hypothesis: they speak of the ideal denotation matrix as having one symbol to denote each situation, and refer to the fact that $L = G$ as meaning that “there were just enough possible sounds to describe the possible situations.” (1994, p. 175).

To test our hypothesis we used the $+L-$ condition, held the number of local environment states constant at $L = 8$, and varied the number of global environment states G , i.e. the number of possible symbols, from eight down to one.

to use ideas from biology, ethology, behavioural ecology, and signalling theory. In this context we believe that, despite its apparently simple nature, MacLennan and Burghardt’s work is overly complex and ambitious.

MacLennan and Burghardt are trying to do a number of things at once. Primarily, they are attempting to provide an existence proof for the synthetic evolution of communication, and they make no secret of having constructed the synthetic world so that the simorgs will be likely to reproduce only if they cooperate (i.e. communicate) in the specified way. They are also examining a process by which arbitrary symbols can evolve to denote something in a simple “language”. As they put it, “beyond merely detecting the presence of communication, we are also interested in studying its structure” (MacLennan & Burghardt, 1994, p. 173). Further, because the simorgs must come to know not only the correlations between symbols and local states, but also when to act and when to emit, MacLennan and Burghardt are effectively looking at the evolution of turn-taking. Finally, they are interested in the effect of learning on the evolution of communication.

With the possible exception of the basic existence proof, each of these phenomena are poorly understood, and each is worthy of a separate, narrowly-focused simulation experiment. When all of these questions of interest are thrown in together, they interfere with each other and make the extraction of general principles impossible. For instance, in trying to push the simulation towards communication, they choose to reward both the sender and the receiver of a message, and in an effort to leave things open-ended enough for spontaneous symbol-meanings to develop, they use the FSM architecture. But what is the relative importance of these factors in causing the observed results? MacLennan and Burghardt allow spontaneous strategies for emitting vs. acting to develop amongst the simorgs, presumably to leave them as unconstrained as possible, but this decision creates the loophole described in section 4.4. Would the same type of communication develop if the simorgs were constrained to be senders and then receivers in turn?

In principle, it may be that communication between simorgs is entirely dependent on their internal architecture, or on the fitness reward structure used, or some other quirk of the methodology — MacLennan and Burghardt themselves note that when the method for selecting parents was deterministic rather than stochastic, communication did not develop. It is not possible, from MacLennan and Burghardt’s results alone, to determine any necessary or sufficient conditions for the evolution of communication; they are doing the equivalent of commencing the study of gravitation with a four- or five-body problem.

Of course, we are not claiming that if only the various factors bearing upon the behaviour of MacLennan and Burghardt’s simorgs could be isolated, then the general principles governing naturally evolved communication would be laid bare. It is quite likely that there are complicated, non-linear interactions even in their small system. However, if we do not understand the effect of each factor alone (e.g. cost or benefit of communication, updating method, simorg architecture) then it would seem optimistic to hope to understand the complex case.

The difficulties with MacLennan and Burghardt’s experiment can be seen in another light: they compare synthetic ethology favourably with empirical ethology in that experiments in the former are repeatable, and full access to all variables is possible. However, this comes at a price. MacLennan and Burghardt

are forced to rigorously specify the environment and the internal nature of the simorgs, making several ad hoc decisions along the way. In a sense, they have to go down to the level of simorg genetics. This is interesting, because one of the great strengths of ethology comes from what Grafen calls the “phenotypic gambit” (1991, p. 6), in which genetics is almost entirely abstracted away, and broad behavioural strategies are considered at a functional level⁵. Most of the time, the conclusions so derived are borne out in the real world. The parallel to be drawn with MacLennan and Burghardt’s experiment is that there is much to be done, using simulation methods, that does not buy into the question of internal architectures, but looks at one phenotypic characteristic at a time and assesses its effect on the evolution of communication. For example, one could simulate a population of agents who were either communicators or mutes, and then allow that population to evolve under different cost and benefit regimes for communicative behaviour. We might expect that when both the sending and the receiving agent benefit from communicative behaviour, then communicators will come to dominate the population. But what about when only the receiver benefits, or when the sender’s benefit is relatively small? What happens when communicators will only signal to other communicators? This sort of simulation, taking up where the mathematical arguments of biologists such as Hamilton (1964) and Grafen (1990) leave off, would give us a sound basis for further investigations.

The best philosophical strategy for such future work is to adopt the intentional stance with respect to simulated organisms. Despite the fact that MacLennan and Burghardt at one point go too far, in ascribing high-level intentional phenomena such as language dialects to the simorgs (see section 4.2), we agree with Dennett (1987, p. 265) about intentional accounts: “. . . in a nutshell, *they work*. Not always, but gratifyingly often.” We are also confident that mechanistic explanations can peacefully co-exist with intentional ones; in the very simple simulations we are initially proposing, no doubt mechanistic accounts will predominate, with the balance gradually shifting as real-world complexity is incrementally approached.

MacLennan and Burghardt are at pains to *avoid* intentional talk when they define communication, and MacLennan (1991) criticises denotational (i.e. intentional) theories of meaning. Nevertheless they rely on an analogy featuring rational, intentional agents — the story of the hunters — to make sense of their simulation, and they use denotation matrices to index the meaning of symbols: an intentional technique, in the sense of “aboutness”, if ever there was a *intencanoess*” xt612000.3(e1)-999.349

- Burghardt, G. M. (1970). Defining 'communication'. In Johnston, Jr., J. W., Moulton, D. G., & Turk, A. (Eds.), *Communication by Chemical Signals*. Appleton-Century-Crofts, New York.
- de Bourcier, P., & Wheeler, M. (1994). Signalling and territorial aggression: An investigation by means of synthetic behavioural ecology. In Cliff, D.,

