

Evolving Communication without Dedicated Communication Channels

Matt Quinn

Centre for Computational Neuroscience and Robotics,
University of Sussex, Brighton, U.K.
`matthewq@cogs.susx.ac.uk`

Abstract. Artificial Life models have consistently implemented communication as an exchange of signals over dedicated and functionally isolated channels. I argue that such a feature prevents models from providing a satisfactory account of the origins of communication and present a model in which there are no dedicated channels. Agents controlled by neural networks and equipped with proximity sensors and wheels are presented with a co-ordinated movement task. It is observed that functional, but non-communicative, behaviours which evolve in the early stages of the simulation both make possible, and form the basis of, the communicative behaviour which subsequently evolves.

1 Introduction

The question of how communicative behaviour might have originated is an interesting one, and the transition from non-communicative to communicative behaviour has long been of interest to ethologists [2,4]. Artificial Life techniques, such as agent-based simulation models, are potentially useful tools for exploring questions and hypotheses related to this transition. In particular, they enable the simulation of co-evolving, interacting organisms at the level of changes in behaviour and perception. There are a number of models in the ALife literature which simulate the evolution of an organised communication system in an initially non-communicating population of agents (e.g., [11,6,1,5,3]). In all these models, communication is restricted to an exchange of signals over dedicated and functionally isolated communication channels. This feature, I wish to argue, severely reduces the explanatory value of a model of the evolutionary origins of communication in natural systems.

Dedicated channels are a reasonable feature of a model which assumes that individuals are already able to communicate. However, explaining the *origins* of communicative behaviour typically involves explaining how it could have evolved from originally non-communicative behaviours [2,4,7]. This kind of explanation is not possible with a model which restricts all potential communication to dedicated and functionally isolated channels. However, this problem is avoided if a model allows potentially communicative behaviour to be functional (and hence acquire selective value) in contexts other than communication. In order to illustrate this point, I present a model in which there are no dedicated communication channels. Agents are evolved to perform a non-trivial co-ordinated

movement task and are equipped only with proximity sensors and wheels. They are controlled by neural networks which map low-level sensory input directly onto motor output. It is observed that the functional but non-communicative behaviour which evolves in the early stages of the evolutionary simulation forms the basis of (and makes possible) the communicative behaviour which subsequently evolves to play an integral role in the successful completion of the task.

2 Orinary Explanations

At first sight, the mere existence of communication appears somewhat paradoxical when viewed from an evolutionary perspective. “Communication”, as Wilson writes, “is neither the signal by itself, nor the response, it is instead the relationship between the two” [12, p.176]. But how might such an obviously co-dependent relationship originate? What good would a signal be before there is a response for it to trigger? But equally, what possible use is a response to a signal unless that signal already exists? The apparent paradox is summed up succinctly by Maynard Smith: “It’s no good making a signal unless it is understood, and a signal will not be understood the first time it is made” [7, p.208]. Of course, this paradox is only apparent. There are at least two possible parsimonious explanations for the evolution of a communicative interaction. Firstly, organisms can often benefit from attending to many aspects of other organisms’ behaviour, even when that behaviour is not a signal. For example, there are obvious benefits to responding to movements indicative of, for example, imminent attack, rather than actually having to respond to the attack itself. Thus organisms may evolve to react to informative, but non-signalling behaviours; these behaviours may in turn evolve to capitalise on the response they now evoke, thereby becoming signals in their own right [2,4]. Secondly, most organisms are constantly responding to all manner of environmental stimuli. Hence a signal might evolve because it is able to trigger one such pre-existing response. For example, it has been argued that the courtship signal of the male water-mite *Neumania papillator*—namely, rapid leg movements causing water-borne vibrations—evolved to exploit an existing vibration response in the female which served to locate prey [8]. Note that both types of explanation resolve the apparent paradox by postulating an originally non-communicative role for either the signal or the response.

Bearing this in mind, consider the simulation models of the evolution of communication mentioned in the introduction [11,6,1,5,3]. Despite their many differences, these models share important common features. Firstly, in all of these models, agents are equipped with a set of behaviours that are designated as potential signals. These potential signals have no function or consequence in any context except signalling. Secondly, agents have one or more dedicated sensory channels with which to perceive these signals. These channels are sensitive to no other stimuli except signals, and signals register on no other sensory channels. What makes these features so problematic? It is not just the rather anomalous presence of specialised communication apparatus in a non-communicating population. More problematically, it is that these models clearly prevent all poten-

tially communicative behaviour from having any non-communicative function. Since neither the production nor the reception of a designated signal can confer any selective advantage outside of a signal-response relationship, these models can tell us nothing about why either the signal or the response exist before the existence of communication. Clearly then, if we wish to produce models that can provide such an explanation, we need to allow for the possibility of potential signals and responses initially evolving in non-communicative contexts. In the next section, I introduce a model which meets this requirement.

3 The Model

For this experiment, a co-operative co-ordination task previously implemented in [9] was adopted. Agents are simulated Khepera robots, equipped only with short-range active infra-red (IR) sensors and two motor-driven wheels (the simulator is described in [9]). The body plan of a Khepera robot is shown in figure 1. Each animat is controlled by an evolved neural network (described below). This network takes input directly from the sensors and uses this input to specify motor-speed outputs. Clearly there are no dedicated communication channels incorporated in this model. Furthermore, neither signals nor responses are initially incorporated into the model (indeed nor is any other sort of meaningful behaviour).

Agents are evolved in a single population, and are evaluated in pairs. Their task is as follows: The pair are placed in an obstacle-free environment in one of a number of possible initial configurations such that each is within sensor range of the other. They are then given 10 seconds in which to move at least 25 cm (approximately 10 agent radii) from their initial position whilst staying within sensor range and not colliding with one another. Whilst easy to describe, this task presents a non-trivial distributed control problem. Covering the required distance is unproblematic (agents have a top speed of 8cm/s). However, a number of other difficulties must be overcome. The IR sensors provide an animat with a noisy and non-linear indication of its distance from an object and have a range of just under 5cm. Should agents move beyond a distance of 5cm apart, they will have no indication of each other's location. Simply moving away and 'hoping' to be followed would thus be a bad strategy. Furthermore, since agents are cylindrical, their IR sensors do not reveal the direction in which another agent is facing. This means that orientation matching strategies, which form an integral component of flocking algorithms [10], cannot be utilised in this task. From [9] it appears that successful completion of this task appears to require some form of role differentiation between agents (e.g., a 'leader' and a 'follower'), and thus some interactive mechanism by which roles can be allocated. Communication is clearly one way in which this could be achieved.

The nature of the task facing the agents means that an individual's ability to perform well is significantly affected by its partner's behaviour. In addition, and particularly in the initial stages of evolution, variation in starting positions also has a large impact on the success of a pair of agents. For this reason, agents





Fig. 1. Left: Khepera mini-robot, showing the 8 IR sensors, and two motor driven wheels. **Right:** Parameters defining the agents' starting positions.



were given a large number of evaluations with many different partners and a fixed selection of starting positions. Starting positions were defined by three variables, d , the agents' distance apart, and α and β , the orientation of each agent relative to the other, as shown in figure 1. A sample set of starting configurations was chosen such that, $d \in \{1.5, 2.5, 3.5\text{cm}\}$ and $\alpha, \beta \in \{0, \frac{2\pi}{5}, \frac{4\pi}{5}, \frac{6\pi}{5}, \frac{8\pi}{5}\}$. Of these, 30 rotationally symmetrical duplicates were removed, leaving a set of 45. Each agent in the population was evaluated twice from each starting configuration, with each of these 90 trials undertaken with a different partner. The fitness score attributed to an agent was simply the mean of all the scores it achieved in these trials. In an individual trial, each agent of the evaluated pair received an equal score. There was therefore no conflict of interest between partners in an evaluation; each individual pairing was a co-operative venture. The score given to both agents at the end of a single trial is given by the following function:

$$\left| P. \left(\sum_{t=1}^T \left[g_t \cdot (1 + \tanh(s_t/10)) \right] \right) \right|$$

Here s_t is the amount by which the distance between the agents exceed the 5cm sensor range at time step t (if this is not exceeded $s_t = 0$). P is a collision penalty scalar, such that $P = \max(1 - c/(1 + c_{max}), 0)$, where c is the number of collisions that have occurred, and c_{max} is the maximum number of collisions allowed, here $c_{max} = 10$. Finally the function g_t is a measure of any improvement the pair have made on their previous best distance from the starting point. This is a function of d_t , the euclidean distance between the pair's location (i.e., its centre-point) at time t and its location at time $t = 0$, and D_{t-1} is the largest value that d_t had attained prior to time t . The value of g_t is zero if the pair are not improving on their previous best distance or if they have already exceeded the required 25cm, otherwise $g_t = d_t - D_{t-1}$. Note that scores are normalised, so the maximum possible score is 1.

Agents are controlled by artificial neural networks. These networks comprised 8 sensor nodes, 4 motor nodes and some number of artificial neurons, connected together by directional excitatory and inhibitory weighted links. The thresholds, weights, decay parameters and the size and connectivity of the network were genetically determined (as detailed in [9]). At any time-step, a neuron's output, O_t , is determined by the value of its 'membrane potential', m_t . If m_t exceeds the neuron's threshold then $O_t = 1$ (the neuron fires) otherwise $O_t = 0$. Here m_t is a function of a neuron's weighted, summed input(s), and the value of m_{t-1}

scaled by a temporal decay constant, such that:

$$m_t = \begin{cases} (1 - \gamma_A)m_{t-1} + \sum_{n=0}^N w_n i_n & \text{if } O_{t-1} = 0 \\ (1 - \gamma_B)m_{t-1} + \sum_{n=0}^N w_n i_n & \text{if } O_{t-1} = 1 \end{cases}$$

where the decay constants γ_A and γ_B are real numbers in the range $[0:1]$ and w_n designates the weight of the connection from the n^{th} input, i_n . Each sensor node outputs a real value in the range $[0.0:1.0]$, which is simple linear scaling of its associated IR sensor. Motor outputs consist of a ‘forward’ and a ‘reverse’ node for each wheel. If the summed weighted inputs to an output node are positive its output will be 1, otherwise 0. The output for each wheel is attained by subtracting its reverse node output from its forward node output.

4 Analysis

A simple generational evolutionary algorithm was employed to evolve populations of 180 (initially random) genotypes. A total of 30 runs were carried out, with each population being allowed to evolve for 2000 generations. Of these runs 27 produced successful solutions to the task (where success was defined as an evaluation score consistently in excess of 0.975). The solutions found by the successful runs had a number of similarities. In each case a pair were successful primarily because one agent adopted a ‘leader’ role whilst its partner adopted the role of ‘follower’. However, these roles were not pre-assigned. All successful runs ultimately produced homogeneous solutions; thus, neither agent was intrinsically biased toward adopting either role. Investigation revealed that it was the interactions between agents which served to establish which role each would adopt. This section focusses in some detail on a single run. It starts with a description of the evolved behaviour of the agents, with particular reference to those aspects which appear to function as signal and response and to co-ordinate the allocation of ‘follower’ and ‘leader’ roles. Next, analysis of the evolved neural network controller is presented to confirm that the signal and response do indeed perform the causal roles ascribed to them. Having established that communicative behaviour has indeed evolved its origins are then addressed. Analysis of the early stages of the evolutionary run is presented in order to show that the evolution of communication in this simulation affords a satisfactory explanation of the kind set out in section 2.

4.1 Evolved Behaviour

The successful achievement of the task can be very briefly described as follows. Initially, each agent rotates counter-clockwise until it is almost facing its partner. I shall refer to this process as ‘alignment’. Once agents have become aligned, and after a (generally) brief interaction, one agent reverses away from the other, and is almost immediately followed by its partner which moves forwards (see figure 2(iv)). For the remainder of the task the pair maintain this configuration,

with the leader moving backwards and the follower moving forwards, together achieving the required distance. How are the roles of leader and follower allocated? Observation of the team from each of the starting positions shows that whenever there is a difference in alignment times, this difference plays an important part in role allocation. Figure 2 shows an example of two agents which become aligned at different times. The first agent to become aligned moves very close to its still-rotating partner and then waits, oscillating back and forth. Subsequently the second agent becomes aligned and reverses away, closely followed by its partner.





Fig. 2. An example interaction: (i) Both agents rotate anti-clockwise; (ii) Agent B becomes aligned first, and moves toward A; (iii) Agent B then remains close to A, moving backward and forward staying between 0.25–2.0cm (approx.) from A. (iv) Once agent A becomes aligned it reverses and is followed by B.

It seems then, that the actions of the first aligned agent serve as a signal to the second. If an agent perceives the signal whilst still rotating, it will adopt the leader role. However, if it becomes aligned without having perceived the signal, it will itself perform the signalling action and subsequently take the follower role. Such a strategy would clearly serve to coordinate role allocation and aid in the successful completion of the task¹.

Analysis of the evolved neural network was undertaken to ensure that the behaviours identified as signal and response did indeed perform the causal roles suggested above. Figure 3 shows the evolved neural network after all non-functional neurons and connections have been removed. Note that the network only utilises two sensors. Sensor 0 is the sensor immediately in front of the left wheel, and sensor 3 is the right-hand sensor of the front-most pair of sensors (see figure 1). Since agents rotate counterclockwise, sensor 0 will normally be the first to register the presence of another agent. What the analysis presented below will demonstrate is this: In cases where there is a difference in alignment times, if sensor 0 is saturated (i.e. fully activated) *prior* to the activation of sensor 3, the result is that an agent will reverse, this constitutes perception of, and response to, the signal. However if sensor 3 is activated without the prior saturation of

¹ From a minority of starting positions there is insufficient difference in alignment times for this strategy to be effective; the procedure by which roles are then allocated is more complicated. However, analysis of the general case is sufficient for the purposes of this paper.

sensor 0 then an agent will move forwards, thereby producing the signal. This clearly does not constitute a full explanation of the network’s functionality, but is sufficient for the purposes of the argument being presented.



Neuron	T	γ_A	γ_B
0	24.36	0.00	0.00
1	0.14	1.00	0.00
2	10.22	1.00	0.84
3	-34.07	1.00	1.00

Fig. 3. The evolved neural network after pruning. Solid lines are excitory links and dashed are inhibitory. The table above gives the threshold, T , and decay constants, γ_A and γ_B , of each neuron. All values shown to 2 decimal places.

The first thing to note about this network is that in the absence of any sensory input, it will cause an agent will rotate counterclockwise. This constitutes an agents ‘base behaviour’, occurring because neuron 3 fires constantly (it has a negative threshold and no inputs) and excites the LB and RF motor nodes. Let us consider how the base behaviour is modulated by sensory input. As an agent rotates toward its partner, sensor 0 will be the first sensor to be activated. This sensor, it was suggested, receives the signal. Note that it strongly inhibits neuron 1, thus preventing any forward movement. It may also activate neuron 0. Given its non-firing (γ_A) decay rate and its recurrent inhibitory connection, neuron 0 will fire if sensor 0 is saturated (or near saturation) for 3 consecutive timesteps. If this occurs, neuron 0 will in turn activate neuron 2 and these, combining with always active neuron 3, will cause the agent to reverse. Thus, the reversing response of the ‘leader’ occurs because the extreme proximity of its partner saturates sensor 0. If however, an agent’s partner is not close enough to trigger the response, the agent will continue rotating and sensor 3 will subsequently become active. Only minimal activation of sensor 3 is required for neuron 1 to fire, and when this occurs the agent will move forwards. Since the agent is not completely facing its partner, forward movement causes sensor 3’s activation to decrease until rotation recommences, turning the agent back toward its partner, reactivating sensor 3, thus causing further forward movement. Repetition of this cycle moves the agent closer to its partner, producing the initial arcing motion shown in figure 2(ii). Collision is averted because, at extreme proximity, sensor 0 will become sufficiently active to cause the agent to reverse briefly, producing the ‘oscillating’ movement shown in figure 2(iii).

4.2 Origins

From the previous section it is evident that agents have evolved behavioural sequences which function as signal and response. This section aims to show the

non-communicative origins of this behaviour through a description of the first 400 generations of the evolutionary simulation.

Successful agents in initial generations of the run were, unsurprisingly, those which were able to produce any form of consistent movement. Such behaviour received some score, since it displaced the pair's centre-point, but it also resulted in agents rapidly losing sensor contact with one another. By around generation 20, and through until around generation 50, agent behaviour consisted essentially in switching on both motors and proceeding, ballistically, in a straight line. The initial spread of this behaviour is reflected in figure 4b in the rapid increase, over the first 20 generations, of both the mean distance travelled by a pair during a trial, and the mean distance between them at the end of a trial.




Fig. 4. Fitness and performance statistics over the first 500 generations

In the period between generations 50 and 100 there is an increase in population fitness which is strongly correlated with a drop in the collision rate (figures 4a and 4c). This reflects the progressive development of basic collision avoidance strategies. Collision avoidance generally took the form of rotating, halting both motors, or halting one and reversing the other. It increased in effectiveness over this period as more sensors came to be used to alter the underlying ballistic behaviour. One consequence of these developments is particularly relevant here. This occurred at positions where both agents' underlying ballistic behaviour led them directly toward each other. In such cases the various halting and turning responses to close proximity typically led to a form of 'deadlock' situation, in which each agent remained close to the other, moving towards and then backing away from its partner, whilst turning to the left and right.

At generation 110, the situation from most starting positions is unchanged. However, a new behaviour has appeared which allows the deadlock situation to be broken. This takes place as follows. The deadlock situation commences and initially proceeds as described above. However, after some time, one agent backs away from its partner. The partner, its collision avoidance behaviour effectively deactivated, recommences its forward progress. The pair move away jerkily, with one agent moving forward and the other moving in reverse. Such pairings rarely maintain formation for the full extent of the required distance. In addition, the reversing agent and the following agent are phenotypically different. Only one type of agent was capable of backing away from a deadlock situation, and then

only when in combination with the other type of agent which itself was not capable of this reversing behaviour. Combinations of two agents of the same type remained deadlocked. Nonetheless, this behaviour—reversal in response to sustained proximity—is the origin of the response behaviour described in the previous section. Although modified in form over the remainder of the run, it is to be observed in all subsequent generations.

From this point, up until around generation 370, the polymorphism of two interdependent types persists. The population comprises those which do reverse in response to proximity ('leader' types) and those that do not ('follower' types). This period is marked by a progressive specialisation of each type to its role. The ability of mixed-type pairs to maintain the follower-leader formation improves markedly. Another significant development occurs. Agents begin to circle in the absence of sensory input, and in later generations, to rotate. This behaviour serves to increase the probability of encountering what was previously a 'deadlock' situation, but which is now potentially an opportunity for an improved score. The gradual increase in population fitness over this period is somewhat deceptive. Individuals engage in both same-type and mixed-type interactions; the increase in mixed-type performance is offset by a decrease in same-type scores over this period. The increasing tendency for pairs to either remain deadlock (same-type) or move away in formation (mixed-type) is reflected in the steady drop in the final distance between agents at the end of each trial.

In summary, agents initially evolved ballistic behaviour and then basic obstacle avoidance. In combination these led to 'deadlock' situations. Agents which subsequently came to reverse in response to sustained proximity capitalised on both of these behaviours. Situations involving sustained proximity only arose after the evolution of obstacle avoidance, and reversal was only adaptive because the original ballistic behaviour caused the non-reversing agent to 'follow'. Once the response was in existence, it came to play an important role in shaping future behaviour. Agents began to evolve strategies which increased the probability that they would be able to trigger the response, and strategies which increased the probability that they would be in a position to respond. It was another 1000 generations before the behaviours described in the previous section were fully formed. However, it should be clear from the above that origins of those behaviours can be satisfactorily explained within the context of this model.

5 Conclusion

Open any textbook which gives an evolutionary account of animal behaviour and you will find hypothetical reconstructions of the processes surrounding the origins and early evolution of communicative behaviour which revolve around non-communicative behaviours acquiring selective value in a communicative context (for example, [2,4,7]). Artificial Life models are potentially very well suited to exploring these and related hypotheses and to critically evaluating the assumptions on which they are based. However, it is difficult see how such models

can even begin to do this if they are implemented in a way that prevents non-communicative behaviour from acquiring a communicative function.

The model described in this paper has not set out to test any particular hypothesis. It is intended simply as a proof of concept. It demonstrates firstly that it is possible to evolve communication in a model without dedicated channels, and secondly, that an explanation of how communication evolves in such a model is far more relevant to the evolution of communication in natural systems than those afforded by previous models.

Acknowledgements. Many thanks to Jason Noble and Seth Bullock for their comments.

References

1. D. Ackley and M. Littman. Altruism in the evolution of cooperation. In R. Brooks and P. Maes, editors, *Proc. Artificial Life IV*. MIT Press, 1994.
2. R. Hinde. *Animal Behaviour*. McGraw-Hill, London, 1966.
3. H. Kawamura, M. Yamamoto, K. Suzuki, and A. Ochuchi. Ants war with evolutive pheromone style communication. In D. Floreano, J-D. Nicoud, and F. Mondada, editors, *Proc. 5th European Conf. on Artificial Life*. Springer Verlag, 1999.
4. J. Krebs and N. Davies. *An Introduction to Behavioural Ecology*. Blackwell, 1981.
5. M. Levin. The evolution of understanding: A genetic algorithm model of the evolution of communication. *BioSystems*, 36:167–178, 1995.
6. B. MacLennan. Synthetic Ecology: An approach to the study of communication. In C. Langton et. al., editor, *Proc. Artificial Life II*. Addison-Wesley, 1991.
7. J. Maynard Smith. *The Theory of Evolution*. C.U.P., 3rd edition, 1997.
8. H. Proctor. Courtship in the water-mites *neumania papillator*: males capitalize on female adaptations for predation. *Animal Behaviour*, 42:589–598, 1991.
9. M. Quinn. A comparison of approaches to the evolution of homogeneous multi-robot teams. In *Proc. Int. Congress on Evolutionary Computing*, 2001. (In press).
10. C Reynolds. Flocks, herds, and schools: A distributed behavioral model. *Computer Graphics*, 21(4), 1987.
11. G. Werner and D. Dyer. The evolution of communication in artificial organisms. In C. Langton et. al., editor, *Proc. Artificial Life II*. Addison-Wesley, 1991.
12. E.O. Wilson. *Sociobiology: The New Synthesis*. Belknap Press, Harvard, 1975.