

The origins of communication revisited

Jordi Arranz, Jason Noble and Eric Silverman

Institute for Complex Systems Simulation
University of Southampton
jordi.arranz@soton.ac.uk

Abstract

Quinn (2001) sought to demonstrate that communication between simulated agents could be evolved without pre-defined communication channels. Quinn’s work was exciting because it showed the potential for ALife models to look at the real origin of communication; however, the work has never been replicated. In order to test the generality of Quinn’s result we use a similar task but a completely different agent architecture. We find that qualitatively similar behaviours emerge, but it is not clear whether they are genuinely communicative. We extend Quinn’s work by adding perceptual noise and internal state to the agents in order to promote ritualization of the nascent signal. Results were inconclusive; philosophical implications are discussed.

Introduction

Artificial life researchers have been modelling the evolution of communication for some time now (for early examples see MacLennan, 1992; Werner and Dyer, 1992; Noble and Cliff, 1996). Communication is of interest in our field for a range of overlapping reasons, most notably because it is associated with two of the major transitions in evolution (Maynard Smith and Szathmáry, 1995): the jump from solitary to social living; and the later development of language and culture in our own species. ALife’s agent-based simulations are a natural match for this research area as they can provide emergent explanations of communication and related co-evolutionary phenomena that are not possible using more traditional modelling techniques.

However, prior to the publication of a seminal paper by Quinn (2001), computational models of the origins of communication and language were missing an important opportunity. Influenced by game theory, by the long shadow of Shannon and Weaver (1949), and by what Lakoff and Johnson (1980) called “the conduit metaphor” for communication, modellers tended to assume that a signalling channel already existed between the relevant agents, and that the thing to be explained was how and why that signalling channel would come to be used for honest, coherent, and reliable communication. MacLennan’s (1992) early work, for example, imagined agents with eight possible world states, each

matched with one of eight preferred responses, and a convenient library of eight ready-made symbols that had to be mapped, over evolutionary time, in a way that would allow pairs of agents to communicate and thus perform optimally.

These kinds of models ignored the apparently vicious circle involved in the evolution of natural communication systems: for a signal to have any meaning, for it to be worth producing, there has to be a community of responders. But why would the appropriate response behaviour already exist if the signal itself has not evolved yet?

This paradox had been noted, and resolved, many years earlier by the ethologists (Tinbergen, 1964). The two key concepts in the ethological picture of the evolution of communication are “intention movements” — non-signals which provide the raw materials for signal evolution — and the subsequent “ritualization” of the nascent signal. Intention movements have not been selected for *per se*; they are simply a physically necessary step in performing some action, e.g., an animal that intends to bite an opponent *must* bare its teeth before doing so. Intention movements thus provide information about future behaviour, and it is not difficult to see how such movements, coupled with the complementary ability to recognize them, might provide the seeds for the evolution of a communication system. Ritualization is what happens when an initially irrelevant movement such as teeth-baring starts to be of informational value to other animals. The ethologists, assuming that the reliable transmission of information would always carry a selective advantage, thought that the original cue would then be exaggerated or stylized in the interests of reducing ambiguity.

Inspired in part by the ethological perspective, Quinn (2001) sought to demonstrate in a simulation that communication between agents could be evolved without pre-defined communication channels; in other words, he hoped to produce a genuine account of the *origin* of communication. Quinn’s point was that by supplying a signalling channel and a library of signals, most of the previous models were assuming the existence of exactly what it was they should be trying to explain. He began with pairs of agents that were linked only by basic sensory-motor interaction, i.e., if one

agent moved this could be detected by the visual system of the other agent. The agents were then faced with an explicitly cooperative task: moving their joint centre of mass as far as possible within a time limit. A genetic algorithm (GA) was used to select for agents that, when paired with another member of the population, managed to coordinate their behaviour and score highly on the task. Quinn interpreted his results as showing that communication had evolved in the form of a dance-like negotiation process between the agents that was followed by matched movement away from their starting positions. Note that no explicit role allocation had been forced on the agents: each one was equally likely to end up as the leader or the follower in the movement phase.

Quinn’s work was exciting because it showed the potential for ALife models to look at the real origin of communication, rather than just the conditions under which it could be maintained in a system where it was already possible. The model is appealing in that it provides a great example of the kind of emergent explanation that ALife can provide, and a potential bridging account between two levels of description (i.e., the level of raw sensory-motor interaction and the level of symbols and reference). It is also a valuable contribution to the biological literature on communication because it lends support to the ethological theory of intention movements and ritualization. Finally, Quinn (2001) is a very popular paper, having been cited 113 times as of April 2011, according to Google Scholar.

However, Quinn’s work has never been replicated. We feel that precisely because Quinn’s approach is so promising, it is important to establish its generality before going further: one goal of the current paper is to check whether Quinn’s central result is robust. Quinn was working in the area of evolutionary robotics and used a fairly detailed model of a real robot; he also employed a continuous-time recurrent neural network (CTRNN) as the evolvable control architecture. What if his result was a freak occurrence, and turned out to be contingent on some detail of the robot’s sensory system or cognitive architecture? The general finding should be robust across these specific details if it is going to be of any value, and therefore we have attempted to replicate Quinn’s work using a different model of agent perception and movement, as well as a different evolvable control architecture.

We also want to ask: did Quinn pick the right task? He showed the emergence of (at least) a coordination protocol between pairs of agents, but did he definitively show the evolution of communication? This in turn raises questions about how to define communication and how to distinguish it from “mere” coordinated behaviour; we will address these issues below. Scheutz and Schermerhorn (2008) make the point that in many simple ALife scenarios, there may not in fact be any selective pressure for a communicative solution, and we feel this may regrettably apply in the Quinn case.




Figure 1: The layout of the ray-cast sensors of our agents. Note that this is not an exact replication of Quinn’s simulated Khepera robots. The diagram is not to scale: robot diameter is 55 mm and maximum sensor range is 50 mm.

The model

Our goal in the first instance is to find out how general Quinn’s result was, and thus we have set up a similar task but used a completely different agent architecture. Quinn’s agents were fairly realistic simulations of a *Khepera* — these are small, low-cost cylindrical robots, 55 mm in diameter and 30 mm in height, with two independent motors driving two wheels, and a set of eight infra-red (IR) proximity sensors giving the robot the ability to perceive nearby objects. We constructed our own 2D simulator that was less detailed than Quinn’s. Our agents are of the same size and shape as a Khepera robot but the sensors are of a different kind, number, and position: see figure 1 for details. Most of the changes we have made to Quinn’s design are arbitrary, and that is exactly the point. We need to keep certain basic features the same so that the coordinated movement task is both recognizable and feasible, but beyond that our simulation will work best as a measure of the generality of Quinn’s result if it is as different as possible.

The agents have been simplified in several ways, but the cylindrical shape has been kept in order to make them rotationally invariant and thus prevent any simple short-cuts that would allow one agent to detect the orientation of its partner. The drive wheels of the Khepera, and details such as inertia and friction, are no longer simulated. Movement and rotation are simply transforms in the two-dimensional simulated environment; agents are moved and rotated around their centre-point. The eight IR sensors have been replaced

by ten ray-cast sensors. They operate by throwing a ray of a certain length and infinitesimal width along the vector the sensor is pointing to. If the ray collides with the other agent (there is nothing else in the environment) within its 50 mm range, the sensor reports the collision distance. If the ray gets to the end of its range without colliding with anything (either because the other agent is not present in that direction, or is more than 50 mm away) the sensor reports 51 mm. Since the amount of space covered by these sensors is significantly smaller than the overlapping fan-shaped response areas of the IR sensors, two additional sensors have been added, bringing the total to ten per agent.

Instead of using a CTRNN as a controller, as Quinn did, our model is based on a simple production-rule system. This is much like a classifier system but with the real-time learning capability removed. Every rule or classifier is composed of a set of ten sensor threshold values, logical operators linking each of these, a comparator condition describing how the sensor values should be compared with the thresholds (less than, greater than, or equal to) and an associated behaviour. Classifiers are fired when the sensory input of the agent at a given time-step matches the classifier condition. When no classifier can be matched a default behaviour is chosen. Every classifier also has a “weight” to avoid clashes when more than one classifier matches the sensory input. In such cases the highest-weighted classifier is fired. Note that the weight of a classifier is not altered by experience: it is a purely random value which can be affected by mutation as can the rest of the classifier.

Agents with internal state are introduced later on in the paper (the initial agents do not have internal state) and they are effectively finite-state machines. Classifiers are specific to a particular internal state of the agent, and when a classifier is fired the state of the agent changes to the output state of the classifier. If no classifier can be matched, there is a default output state, and thus any time the default behaviour is fired, the agent switches to this default state.

In order to make this a replication, the task the agents face is exactly the same as in Quinn’s work: a pair of agents must move their joint centre of mass as far as possible while staying within each other’s sensor range and without colliding. We used much the same type of GA as Quinn did to evolve the population of agents, but some of the parameters employed, as well as the way fitness is computed, are different. As in the original model, there is no predefined role allocation. Agents are drawn randomly from the population and evaluated in pairs. Each pair is given a certain amount of time to solve the task. Evaluation is performed in discrete time steps; at every time step new sensor values are computed for both agents; and finally the agent behaves according to its sensory input through the activation of the highest-weighted matching classifier. Each agent in the pair gets the same score depending on their joint performance. A selection process keeps the best 40% of agents and deletes

the rest in every generation, with new agents being created through recombination and mutation of the successful individuals of the previous generation.

The fitness of every pair of agents is computed as an average over two different terms. The first term measures whether or not the agents are in each other’s sensory range and is itself averaged across all simulation time steps. This term is important in shaping effective solutions, as the agents are effectively very short-sighted and moving out of sensor range is usually a disaster for the over-arching goal of moving the joint centre of mass in a consistent direction. The score of an agent on a given time step is computed as an exponential decay function on the distance to the other agent. If an agent is in sensor range the fitness obtained is 1.0, otherwise fitness decreases exponentially with the distance. The maximum distance is computed as the maximum linear distance an agent could achieve given its linear velocity and the overall simulation time.

At the end of the simulation the second fitness term is computed: it measures the distance that the agents have travelled. If *either* agent has travelled at least 250 mm then this component of fitness is 1.0. If the agents have travelled a shorter distance from their starting positions, this fitness component will be the quotient of the distance travelled by whichever agent has travelled the furthest, over the target distance. Note that an agent could travel approximately 500 mm — double the target distance — during the time available if it moved away in a straight line, which means that the fitness function allows the agents a reasonable amount of time for potential communication before movement begins in earnest.

Even though the overall goal is moving the joint centre of mass, we do not measure this directly. Optimal performance is achieved by staying in sensor range *and* moving as far as possible. The final fitness score is the average of the two terms described above, and thus the maximum score is 1.0. Fitness scores of 0.5 are relatively easily achieved by either not moving at all (thus staying in sensor range and scoring highly on the first component) or moving off in random directions at full speed (scoring highly on the second component).

Finally, at the end of a generation the final fitness of each agent is equal to the average of its scores across many different evaluations with different partners and in different initial positions. Note that all initial positions have the agents starting inside each other’s 50 mm sensor range.

Replication results

We ran our simulation 30 times, with each run lasting 2000 generations. Quantitatively our mean and maximum fitness values were similar to Quinn’s despite the differences in the agent architecture, the GA, and the fitness function. Some of the agents scored very high and even perfect fitness levels although these could not be maintained in the long run as

	<i>Front</i>				<i>Right</i>				<i>Rear</i>				<i>Left</i>							
	S0		S1		S2		S3		S4		S5		S6		S7		S8		S9	Behaviour
Rule 1	49	∨	9	∧	35	∧	37	∧	29	∨	13	∧	33	∨	29	∧	18	∧	19	Forward
Rule 2	14	∨	16	∧	47	∧	37	∧	47	∧	20	∧	38	∨	10	∧	23	∨	26	Backward
Rule 3	23	∧	39	∨	49	∨	14	∧	40	∨	16	∨	35	∨	8	∧	27	∧	20	Rotate CCW
Rule 4	49	∨	9	∧	2	∧	37	∧	29	∧	19	∧	4	∨	37	∧	26	∨	19	Forward
Rule 5	3	∧	9	∧	15	∨	43	∧	14	∧	8	∧	6	∨	44	∧	38	∨	20	Backward
Default																				Rotate CCW

Table 1: The production rule set of a high-performing strategy found during the replication runs (this rule set leads to perfect performance on the task). For each rule, the table lists the threshold value in mm for each of the ten sensors, and the logical operator (either “and” or “or”) used to link them. In every case the rules used the “less than” comparator, i.e., the value would be true if the other agent was detected at the given distance or closer. Note the mix of “forward”, “backward”, and “rotate counter-clockwise” behaviours that combine to produce coordinated movement. An agent that could detect nothing within sensor range would fall through to the default behaviour of counter-clockwise rotation.

mutation pressure prevented the population as a whole from adopting an optimal strategy. Table 1 shows one of the best rule sets evolved.

Qualitatively, we have analyzed in detail the kinds of strategies that evolved in the most successful runs. Although many different strategies evolved that could accomplish the coordination task, we found that the most common and the most successful one we observed fits reasonably well with the main strategy described by Quinn. Figure 2 illustrates the sequence of behaviours. Both agents start rotating counter-clockwise (A) until the first agent (shown in brown) reaches its favoured alignment relative to the second agent (shown in white) and starts moving one step forward and one step backwards in order to “signal” its readiness and direction to the second agent (B). In the meantime, the second agent keeps rotating counter-clockwise until it matches the first agent’s alignment (C). When both agents are aligned and pointing in opposite directions, the first agent starts moving backwards while the second agent starts moving forward, and thus they move together until the end of the time frame (D). Many variations on this strategy exist, with varying degrees of speed and reliability in achieving alignment. Some of these strategies include several intermediate steps in order to restart the synchronization if something goes wrong, as well as different types of orbiting behaviours after the two agents have been aligned. Quinn also notes that the strategy he picked to illustrate the behaviour of the agents is just one of the simplest cases among many variants observed.

The change in the number of collisions over evolutionary time also matches Quinn’s results. The collision rate is extremely high in the early generations but rapidly decreases as fitness increases. Sudden decreases in the collision rate usually match fitness jumps even though our implementation does not include an explicit penalty for collisions. We can also confirm that, as Quinn stated, the evolution of successful behaviours is extremely sensitive to the initial condi-

tions used (the starting distance between the two agents and their relative orientations) as well as to how the agents are evaluated. In essence every agent has to be evaluated with every possible angle and distance: random runs in which every agent is evaluated with different randomly chosen starting distances and relative orientation angles are completely unsuccessful.

There were many differences introduced between Quinn’s setup and our own, notably the use of a different sensory system and control architecture. Nevertheless we managed to replicate Quinn’s findings: very similar behaviours evolved. We therefore suggest that the emergence of coordinated (and possibly communicative) behaviour to solve this type of task is likely to be a general and framework-independent finding.

Re-examination of the Quinn paradigm

In the previous section we reported the successful replication of a dance-like negotiation phase between the pairs of agents. This is a pleasing result as it goes some way towards showing that Quinn’s findings are general. However, we did not observe any unequivocal “ritualization” process by which the signal became more exaggerated over time. This led us to wonder whether our agents were really communicating at all.

So what do we mean by communication anyway? Should we expect a sharp dividing line between coordinated behaviour and “true” communication? Some ideas from the philosopher Millikan (1984) will be useful here. She argues that although there is no *sharp* line between those two categories, there is certainly a distinction worth making. Millikan lays out four classes of representational phenomena, in order of increasing sophistication: tacit suppositions, intentional icons, inner representations, and mental sentences. The first two are all we will need given the simplicity of our agents. (Millikan’s typology was initially directed at the issue of what might count as an internal representation within a single organism but it is relevant to our purposes as she




Figure 2: Illustration of the evolved sequence of behaviours in a typical case. A: agents rotate until one reaches a favoured orientation. B: the first agent to achieve this starts moving backwards and forwards. C: the second agent orbits the first until it is aligned in the opposite direction. D: the two agents move away together, with the second agent moving in reverse.

sees communication as simply the exchange of representations between organisms.)

Tacit suppositions occur when the design of an organism meshes so neatly with a feature of the environment that it is tempting to say the design “represents” that feature. For example, if a biological clock produces a cycle close to 24 hours then we may be tempted to say that the clock mechanism somehow represents the length of the day. Millikan refers to such adaptations as tacit suppositions because they presuppose certain facts about the environment in order that their evolved function is fulfilled.

For a system to qualify as minimally representational, it must involve more than tacit suppositions. Firstly, there must be something identifiable as the representation itself: an “icon”. Furthermore, the icon must have a “producer” and a “consumer”. It must be the function of the producer to generate the icon in accordance with a mapping rule that relates one or more dimensions of possible variance in the icon to variance in the environment. It must be the evolved function of the consumer to use or be guided by the icon in some way. If all of these conditions are met, Millikan

suggests that the system involves an “intentional icon”. For example, the waggle dance of the honeybee is a paradigm case of an intentional icon: the dance itself is the icon, the dancing bee is the producer, and a mapping rule relates the angle and duration of the dance to the direction and distance to a food source. The watching bees are the consumers of the icon, because it is the adaptive function of the dance to guide them to the food source. The important point is that there is a difference between tacitly supposing that the world — including your interaction partners — regularly works in a certain way, and evolving a distinct behaviour or trait that has been selected for on both sides (production and reception) precisely because it conveys information from one agent to the other.

Consider the difference between two scenarios. In the first one, you are at home, and it is my job to pick you up in a car. I drive to your house: you are not outside, so I drive around the block repeatedly and check again for your presence each time I go by. Assuming I do that reasonably reliably, you can tacitly suppose the existence of my strategy, and go out into the street whenever you see my car going by. I will then

see you and stop to pick you up on the next cycle. Both of us have strategies that rely on the other one acting a certain way but neither strategy has been exaggerated into a signal. We are coordinated but not communicative. The second scenario is exactly the same, except that through some adaptive process we have arrived at a communicative solution: I honk the horn three times in quick succession, and you come outside in response.

These two scenarios demonstrate the difficulty of showing that Quinn’s (or our) observed behaviours are anything more than coordinated. Each agent is tacitly supposing that the other will rotate, align, move forwards and/or backwards, etc. The dance-like movement is, on the surface, reminiscent of the bee dance, and we suspect this resemblance has made many readers of Quinn’s original paper confidently interpret the behaviour as communication. However, it is important to note that there is no mapping rule and no clear referential signalling going on.

What would it take to make Quinn’s negotiation dance a signal? In answering this question, Millikan would agree closely with the ethologists. Non-signalling behaviours must provide the seed for signalling behaviours — how could it be otherwise? So the thing to look for in classifying something as “real communication” is a history of selection for exaggeration on both sides, both in the production of the signal and the sensitivity or scale of the response. In Quinn’s paradigm we do not really see this: as far as we can tell from historical analyses of our runs, the agents hit on their coordination strategy and it remains essentially unchanged.

Quinn’s dance in its current form appears to be a borderline case: it surely qualifies as an intention movement, and is quite possibly ripe for exaggeration into a signal. In the next section we try to push things towards communication by adding both perceptual noise and internal state to the agents. Noise may make a difference in that we can imagine the “dance signal” being exaggerated or strengthened to make sure it cannot be misunderstood in a noisy environment. State is a slightly different story: our stateless agents are necessarily reactive. It is not clear whether Quinn’s CTRNN agents had any internal state; they might have, due to the possibility of recurrent connections. If we add state bits and find that this improves performance, that means that the task was “state-hungry”, which in turn suggests a potential interpretation in terms of intentional icons, i.e., that the agents could be communicating about their current internal state.

Results of the extended model

We extended our replication of Quinn’s model to try to assess whether or not the evolved behaviours really qualify as communicative. In order to do so we have added two new features: perceptual noise and internal state. The addition of Gaussian noise to the sensory inputs adds ambiguity to the perceptual world of the agents and would seem likely to

make the task more difficult. Thus it might be a driver for more explicitly communicative strategies. The second extension is the addition of 1 and then 2 bits of internal state to the agents. The acquisition of internal state enhances the cognitive capabilities of the agents, giving them more behavioural options than a purely reactive agent. This should make it easier for the agents to sequence their coordinated behaviours over time, but for our purposes it may also give them something to communicate *about*, i.e., their current internal state values.

We added 17 new sets of 30 runs each, employing different noise values (0%, 1%, 2%, 4%, 8% and 16%) and adding either 0, 1 or 2 bits of internal state to the agents. In the end we have a total of 18 run sets (including the original noiseless-stateless run) exploring every combination of noise and number of state bits. In order to reflect the increased range of behavioural possibilities that come with having internal state, we have also increased the number of classifiers from 5 (in stateless runs) to 10 (for 1-bit state runs) and 15 (for 2-bit state runs). The remaining parameters of the simulation including mutation rate, population size, initial conditions and length of runs remain the same.

The results are presented in figure 3; the general pattern is in line with our expectations. We can see that the addition of noise decreases the performance of the agents in solving the task. On the other hand, the addition of state seems to make the task easier: the 2-bit condition is only slightly superior to the 1-bit condition, but both are significant improvements on performance in the stateless case.

When looking at the different runs individually, we find that state-equipped agents evolved more robust strategies than stateless ones. In fact, some of the 2-bit state solutions reach consistently optimal performance across the lower noise levels. In such cases, the mean fitness of the population reaches a sustained score of 1.0 with only occasional perturbations due to the randomness added by mutation. Since we have not observed such robust performance in any of the of the stateless runs we take this as evidence that the task chosen by Quinn, despite its simplicity, is “state-hungry”.

In order to qualitatively assess whether or not the agents were evolving genuinely communicative solutions, we looked for the equivalent of Tinbergen’s “intention movements” in the early stages of each evolutionary time line, and looked also for their ritualization or exaggeration into proper signals. We have found *some* suggestive cases of exaggeration in state-equipped runs, in particular for the forward-backwards movement that Quinn originally highlighted as a suspected signal. The movement sometimes becomes exaggerated just before the population starts to score perfect fitness scores. The exaggeration consists of a two-steps-forward-two-steps-backwards routine instead of the former one-step-forward-one-step-backwards. Despite this interesting result, we found no indications of a general trend. Fur-




Figure 3: Results from the extended model. Mean fitness reached at the end of the evolutionary run is shown for various combinations of perceptual noise level and the number of state bits. Each plotted data point is an average across 30 replications. Standard errors across these replications are not shown for reasons of clarity, but the mean size of the standard error was 0.013. The general result is that performance is reduced as noise levels increase, and that at least one bit of state leads to better performance on the task.

thermore, the runs that evolved such exaggerated “signals” were not among the runs with the highest overall average fitness (although, on the other hand, this kind of exaggeration never appeared in stateless runs). It may be that the task picked by Quinn is not “communication hungry”, i.e., it does not require explicit information transmission between the agents in order to achieve optimal performance levels (see Scheutz and Schermerhorn, 2008).

Conclusions

We have achieved one of our goals, in demonstrating the generality of Quinn’s (2001) finding that sensory-motor interaction with no pre-defined communication channels can lead to coordinated behaviour. The result does not seem to be dependent on specific details of Quinn’s setup such as the CTRNN control architecture.

We also asked some critical questions as to whether the dance-like coordination behaviour should be seen as communicative. We extended Quinn’s model to include increased levels of perceptual noise, and internal state for the agents. This was done with the intention of pushing the agents into developing exaggerated signalling and response behaviours over evolutionary time that would more clearly fit the definition of communication. Unsurprisingly,

we found that higher levels of noise make the task more difficult. We also found that adding one or two bits of internal state improved performance, indicating that Quinn’s task is somewhat “state hungry”. Unfortunately we were not able to get consistent evidence of signal exaggeration and ritualization. We have to conclude that the dance-like coordination behaviour exhibited by the agents is at best a borderline case of true signalling.

The difficulty is that Quinn’s chosen task simply appears not to provide selective pressure for communication in Millikan’s sense of producing intentional icons. Scheutz and Schermerhorn (2008) have noted that this is true of many of the simple scenarios employed by ALife researchers. If we look at the world inhabited by our agents, it becomes clear that there is effectively not much to talk about: they always begin their interaction within sensor range of each other, the other agent is the only feature in the world and thus the only thing that can be detected by their sensors, and the cooperative goal of joint movement is always consistent. Once a coordinated solution has been evolved, the agents are already performing near-optimally and there is no evolutionary pressure towards any exaggeration of the signal. We suspect that a promising direction for future work in this area is to use tasks in which referential communication about distant ob-

jects is essential for optimal performance (see e.g., Williams et al., 2008).

Why does all this matter anyway? Coordination or communication — what’s the difference? We believe it matters because there are two very different messages one can take from Quinn’s original finding. On the one hand you may see Quinn’s result as showing how the *appearance* of communication can be explained away as being just the result of mechanical feedback loops in a physical system. Some “enactivist” thinkers in ALife appear to endorse this position. The hope is to eventually demonstrate that human-level intelligence is really made up of a toolkit of sensory-motor tricks and hacks; Beer’s (2003) dynamical systems approach is a good example.

On the other hand, Quinn’s result can be seen as an attempt to bridge two levels of description. Quinn published his paper out of frustration with previous ALife work on signalling that constantly presupposed the very thing it was trying to explain, but that does not mean that he hoped to render talk of signals and channels irrelevant. If a model like Quinn’s could successfully show that communication can indeed emerge from sensory-motor interactions, we could take that not as undermining the concept of communication but as explaining how one level of description (L2: that of signals, symbols, and representations) can emerge from another (L1: the mechanics of sensory-motor feedback).

It has been argued (de Pinedo and Noble, 2008) that in explaining the behaviour of evolved agents, both agent- and sub-agent-level explanations will be necessary — and models like Quinn’s seem a useful step in that direction. Having established that L1 can give rise to L2, we thus establish that every subsequent simulation which incorporates L2-type communication does not need to provide direct evidence of the origins of that communication — we are safe in assuming that said communication would evolve in some fashion or other. Models like Quinn’s can thus provide *bridging explanations*: they verify the relationship between L1 and L2 and then allow those interested in L2 alone to get on with simulating phenomena at that level, confident that L2’s origins are understood.

The question as to which of these two views of communication and reference will ultimately prevail is of course still open. However, we are convinced that ALife simulations such as Quinn’s provide a uniquely valuable testing ground for working out the consequences of either approach.

Acknowledgements

JA was supported by the UK’s Engineering and Physical Sciences Research Council via Southampton’s Institute for Complex Systems Simulation and doctoral training centre. JN was partly funded by research project FFI2010-19455 from the Spanish Ministry of Science. ES was partly supported by the EPSRC-funded *Care Life Cycle* project.

References

- Beer, R. D. (2003). The dynamics of active categorical perception in an evolved model agent. *Adaptive Behavior*, 11(4):209–244.
- Bullock, S., Noble, J., Watson, R., and Bedau, M. A., editors (2008). *ALife XI: Proceedings of the Eleventh International Conference on Artificial Life*. MIT Press.
- de Pinedo, M. and Noble, J. (2008). Beyond persons: Extending the personal / subpersonal distinction to non-rational animals and artificial agents. *Biology and Philosophy*, 23(1):87–100.
- Lakoff, G. and Johnson, M. (1980). *Metaphors We Live By*. University of Chicago Press.
- Langton, C. G., Taylor, C., Farmer, J. D., and Rasmussen, S., editors (1992). *Artificial Life II*. Addison-Wesley, Redwood City, CA.
- MacLennan, B. (1992). Synthetic ethology: An approach to the study of communication. In Langton et al. (1992), pages 631–658.
- Maynard Smith, J. and Szathmáry, E. (1995). *The Major Transitions in Evolution*. W. H. Freeman, Oxford.
- Millikan, R. G. (1984). *Language, Thought, and Other Biological Categories*. MIT Press / Bradford Books, Cambridge, MA.
- Noble, J. and Cliff, D. (1996). On simulating the evolution of communication. In Maes, P., Mataric, M., Meyer, J.-A., Pollack, J., and Wilson, S. W., editors, *From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*, pages 608–617. MIT Press, Cambridge, MA.
- Quinn, M. (2001). Evolving communication without dedicated communication channels. In Kelemen, J. and Sosik, P., editors, *Advances in Artificial Life: Sixth European Conference on Artificial Life (ECAL’01)*, Lecture Notes in Artificial Intelligence, pages 357–366. Springer, Berlin.
- Scheutz, M. and Schermerhorn, P. (2008). The limited utility of communication in simple organisms. In Bullock et al. (2008), pages 521–528.
- Shannon, C. E. and Weaver, W. (1949). *The Mathematical Theory of Communication*. University of Illinois Press, Urbana.
- Tinbergen, N. (1964). The evolution of signalling devices. In Etkin, W., editor, *Social Behavior and Organization Among Vertebrates*, pages 206–230. University of Chicago Press, Chicago.
- Werner, G. M. and Dyer, M. G. (1992). Evolution of communication in artificial organisms. In Langton et al. (1992), pages 659–687.
- Williams, P. L., Beer, R. D., and Gasser, M. (2008). Evolving referential communication in embodied dynamical agents. In Bullock et al. (2008), pages 702–709.