Decisions and noise: the scope of evolutionary synthesis and dynamical analysis

By Ezequiel A. Di Paolo and Inman Harvey

Department of Informatics, School of Science and Technology
University of Sussex, Brighton BN1 9QH, U.K.
Email: {ezequiel,inmanh}@sussex.ac.uk, *Fax: +44-1273-671320*

The agents evolved by Randall Beer for active object discrimination perform an undeniably interesting task. It may seem modest from the point of view of an external human observer, but for a robot with limited sensory capability trying to work out two-dimensional shapes from a flux of one-dimensional data, it is indeed challenging. Already in this fact we find a serious answer to critics who think that simple experiments in mobile robotics like this one have little relevance for understanding cognition as a natural phenomenon. Whether a task is cognitively interesting cannot be judged in a vacuum, or only by human standards, but depends on the dynamical, bodily and environmental conditions with which an agent must cope. If by his initial judgement the designer cannot foresee a trivial way to perform the task given the resources provided to the agent, then the synthesis of successful behaviour is an event from which something can certainly be learned. Such is the motivation that is intuitive for many people working within autonomous and evolutionary robotics, but which often seems to escape those who think that these disciplines are solely focused on synthesizing human level intelligence from the bottom up.

Of course, the question remains whether what is learned by understanding the mechanisms that give rise to an agent's behaviour is useful, suggestive and of potential generality, or whether it is instead applicable only to one particular case or the product of an undesired design constraint. Arguably, Beer's thorough analysis provides us with one example for each option, though interestingly both examples contain something to learn from.

Let's take the second case first. Beer's evolved agents, at least those analysed in the paper, rely heavily on the maximum apparent diameter of an object to perform the discrimination. Beer demonstrates this by studying catching and avoidance performance for different profiles and dropping a variety of shapes with the same diameter as the diamonds and obtaining the same response from the agents. Disappointing? Perhaps. Indeed, if we were expecting some sophisticated visual shape discrimination this may not be what we had in mind. Evolutionary robotics is full of examples (often unreported) where evolution manages to bypass the complexities of a problem by taking advantage of what at first look like innocuous design assumptions – in this case the fact that the shapes always have different maximum diameters and that these are fixed.

But unlike situations where the evolved solutions rely on a design artefact or a bug in the simulation code, here we simply get what we asked for. Discrimination is successful and utilizes a readily available environmental variable. The clues provided by this variable must still be exploited appropriately using active strategies that may well be described as visual scanning. How this is done is not trivial. The situation is not different from the many tricks we find in animal intelligence where environmental regularities are exploited opportunistically – it also reinforces a message that unfortunately still needs to be repeated today particularly in a cognitive science not yet recovered from decades of cognitivism: intelligent behaviours need not arise from intelligent mechanisms. 'Dumb' underlying mechanisms only disappoint if this message has not sunk in. When W. Ross Ashby presented his Homeostat at the ninth Macy conference on cybernetics in 1952 the reception was sceptical, as if they were being confronted with a sleight of hand because learning and adaptation seemed implausibly complex to be supported only by random undirected mechanisms. The confusion of mechanistic and behavioural languages, it seems, is an old sin of the field.

Still, it would be interesting to see agents capable of discriminating *shape* as suggested (but not entailed) by the description of the experiment. The flexibility of the evolutionary method makes this an easy task in principle. If discrimination based on diameter relies on a constant difference in maximum diameters between the two classes of object, then one should make those maximum diameters equal. Additionally, to be certain that discrimination is based on the shape of the object and not on anything else, one should make the size of the objects random. Perform the evolutionary search again and, if successful, there is a guarantee that the robot will discriminate objects based on shape only.

Fortunately, the changes that must be done to the original scheme are trivial and shape discriminators are easily evolved using a similar set-up as Beer's (with minimal parametrical changes). Figure 1 shows the performance of a successful agent evolved in 1000 generations capable of discriminating circles and diamonds with a range of sizes (diameter within [13.5,22.5]) as a function of object initial position and size (brighter shade means fitter, white means more than 90% success). Even better results are obtained for the converse task of approaching diamonds and avoiding circles (figure 2). No analysis has been made yet on the 14-neuron CTRNN controller of these agents, but from the figures it is clear that they are not relying on absolute differences in diameter, but need to use some other geometrical properties of the objects. To bring the performance of these agents even closer to real shape discrimination we should also introduce inter-trial variability in object orientation (significant here only for diamonds). We shall come back to the role of randomness and noise in evolutionary design.

A more suggestive, and possibly more general, discovery is revealed by Beer's investigation of changing shapes at different times during a trial and the subsequent effect on the 'decision' to catch or to avoid. The results suggest that 'decision making' is not a discrete event occurring at a clearly demarcated point in time before the action takes place. It is rather a temporally extended process entangled with the rest of the behaviour of the agent. We may criticize the use of terms like 'decision' in this context, yet the fact remains that this example makes it easy to conceive that something similar could possibly be going on in more complex situations where 'decision' is the right word. How do we take a decision to walk on one or the other side of the road? Possibly before we start walking, but conceivably (and now much more easily conceivably) in a way that cannot be separated from walking itself. The usefulness of this minimalistic approach becomes apparent in this case, not because it points to the specific mechanisms, nor because it provides testable predictions, but because it exercises the researcher's mind into playing with possibilities that he may not have previously considered.

Yet, *if* we were interested in decision making as a discrete event, we find again that the method is flexible enough to allow for a simple extension to study this particular case. Imagine that the agent has two light bulbs that must be turned on or off depending on whether it is going to catch or avoid a falling object respectively. One light bulb has the role of indicating 'a decision to catch', the other 'a decision not to catch'. Once either bulb is turned on, it may not be turned off. We may include as a fitness criterion the requirement that the correct bulb must be turned on as soon as possible within the trial. The behaviour of successful agents would thus show a clearly demarcated time by which a 'decision' must have been made before the behaviour is fully executed.

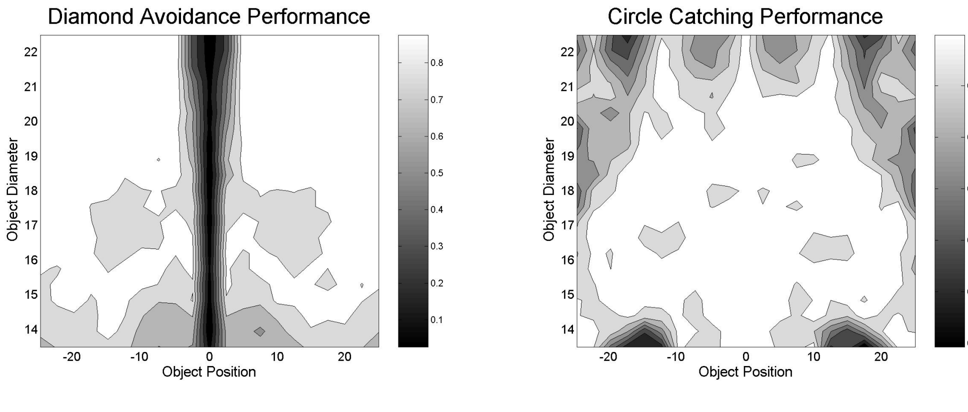

(a) (b)

Figure 1. Performance for avoiding diamonds (left) and catching circles (right) for the best agent evolved using random object size as a function of object maximum diameter and position. Data taken for 400 evenly distributed points each representing the average of 50 trials. Bright means high fitness. Symmetry with respect to position is not perfect because of sensor and motor noise, which is however not sufficient to disambiguate diamonds falling right on top of the robot. Robot diameter = 8, 8 sensor rays evenly distributed in 60 degrees, laterally symmetric CTRNN controller, 8 sensor neurons, 2 motor neurons, 4 fully recurrent interneurons.

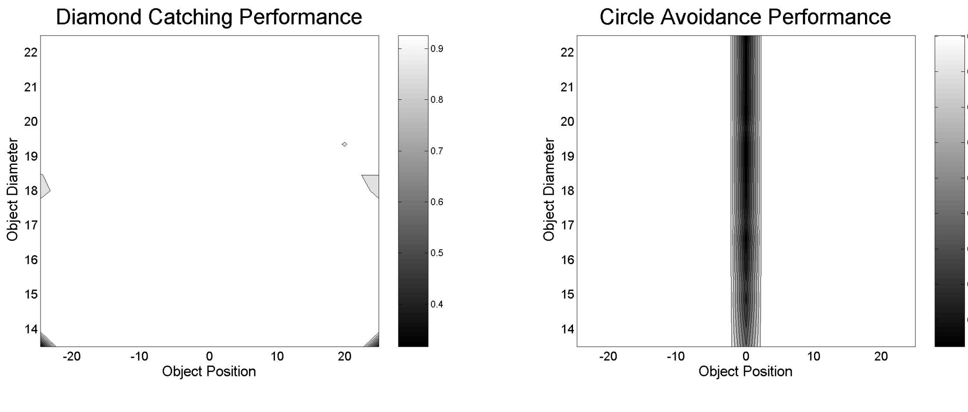

(a) (b)

Figure 2. Same as figure 1 but for an agent evolved to catch diamonds and avoid circles.

Something that is not evident from the paper, but apparent from the above comments, is that there are many degrees of freedom within which the designer can apply a variety of subtle constraints that may alter the final result sometimes in significant ways. The practice of working with the evolutionary method is often more sophisticated than it seems, and if used without care it will unwittingly exploit subtleties and possible ambiguities in such everyday concepts as 'making a decision'. Evolutionary robotics is a harsh taskmaster in this respect, and practitioners must learn to be rigorous in defining such concepts operationally. This, unfortunately, is not an ideal situation if evolutionary synthesis is to be used reliably as a tool in cognitive science and steps should be taken to better formalize the design process.

One such step has to do with the use of variability between evaluations and noise during a same evaluation. These factors are particularly relevant for cognition. The only variability that is introduced between trials in Beer's model is in the initial position of the falling object. Inter-trial variability of object size results in further robustness as shown above, producing true shape discriminators (except perhaps for the invariant diamond orientation). Other forms of variability, such as random initial conditions in the network activation and other robot variables, uncertainty in sensor readings and motor output (also used for the data in figure 1), inaccuracies in sensor position, etc., could also be added with the expected result of mechanisms that are robust to such uncertainties.

It is not necessary for Beer to explore all these different possibilities to make his main points. However, we think that for the case of cognition, even minimal cognition, noise and uncertainty are fundamental factors that can only be excluded at the risk of obtaining very clean but brittle solutions to the cognitive task. Moreover, when assuming that a test is for 'circle/diamond' discrimination these terms are *only* strictly appropriate if the test cases vary in all possible features *apart from* shape. Otherwise the evolutionary method may and probably will pick up on another accidentally correlated feature. Hence, noise is not an optional extra, but an essential requirement.

Noise and inter-trial variability are often used for the practical purpose of enhancing the transfer of neural controllers from simulation to real robots, a specific aim of evolutionary robotics but not one of this paper. However, we could argue that uncertainty and noise are both realities that all cognitive systems have to deal with, and may also provide *positive* mechanisms for producing robust and adaptive behaviour.

Consider for instance a further twist in the light bulb story above. Let's suppose we change the fitness requirements so that instead of turning on the correct bulb *after* an object has been identified the agent should 'decide' by itself whether to catch either diamonds or circles. The agent must use the light bulbs to 'express its future intention' *before* an object is presented and then stick to its decision. A successful agent must therefore perform either one out of two different behaviours with equal probability, and the decision must be made while the sensory input shows no object (i.e., the input is symmetric with respect to the outcome decision). How can this task be achieved if there are no sources of symmetry-breaking mechanisms, like neural noise, or variability in initial conditions?

Noise is often avoided because it may complicate further analysis. We wonder if this is necessarily true in this case. Much of the analysis presented in the paper could easily be done in the presence of noise, perhaps with minimal modifications.

Indeed, Beer provides a paradigmatic example of how a dynamical explanation should be constructed. Contrary to common assumptions, understanding the neural dynamics is the last step to be taken since the complexity involved, even for these rather small networks, can be daunting. To undertake a blind analysis often leads nowhere. In contrast, it is much easier to take advantage of the flexibility of the simulation to perform a series of psychophysical tests that lead to partial explanations and hypotheses about the agent's behaviour. These may later be supported or falsified by further crucial experiments or neural analysis. A dynamical explanation is therefore built by partial steps that support one another like typical explanations in science. The presence of noise would not change any of these steps.

The resulting explanation can still be overwhelming. The objector to a dynamical systems strategy towards understanding cognitive systems may see this as proving her case. If such paraphernalia of psychophysical, parametrical, transient and attractor studies is necessary to explain this relatively simple behaviour, then what are the odds of this approach succeeding for 'higher' forms of cognition?

We should consider three aspects of this criticism. First of all, it must be formulated in a fair context. An embodied, dynamical systems perspective on cognition, to the degree that it may exclude a representational, computationalist alternative (and this is by no means generally accepted), has been proposed because it offers a better entry point for unlocking the investigation of many cognitive phenomena at different levels – cases in which computationalism either fails or presents more problems than it solves. So the above question about the expected generality of the dynamical systems approach must always be a comparative one. If we point to a case where dynamicism will supposedly fail, we must equally show that in this case computationalism will uncontroversially succeed.

The second aspect worth considering is that, as shown in this paper, a partial dynamical explanation is still useful. It narrows the field of possibilities and concentrates the effort of the investigation. No one is ruling out that explanations of a functional kind may also complement a dynamical systems story, but these will not be just any functional explanation, only those that work within the constraints set out by the dynamical findings.

Finally, it is precisely because such complexity belongs to cognitive systems themselves (rather than exclusively to the dynamical method) that the 'frictionless brains' idea of minimalistic agents and tasks is potentially so useful. Because complete dynamical explanations may be possible only in cases like these, we must attempt them so as to train the mind in the exercise of explaining fully integrated cognitive agents, and not just sub-systems in a reductionistic manner with no subsequent attempt at integration. In this sense, Beer's minimalistic approach, when carried out with this objective in mind, is far from being another set of toy problems but a fundamental addition to the methods of cognitive science.