

A Role for Action Selection in Consciousness: An Investigation of a Second-Order Darwinian Mind

Robert H. Wortham and Joanna J. Bryson
Dept of Computer Science, University of Bath
Claverton Down, Bath, BA2 7AY, UK
Email: {r.h.wortham, j.j.bryson}@bath.ac.uk

Abstract—We investigate a small footprint cognitive architecture comprised of two reactive planner instances. The first interacts with the world via sensor and behaviour interfaces. The second monitors the first, and dynamically adjusts its plan in accordance with some predefined objective function. We show that this configuration produces a Darwinian mind, yet aware of its own operation and performance, and able to maintain performance as the environment changes. We identify this architecture as a second-order Darwinian mind, and discuss the philosophical implications for the study of consciousness. We use the Instinct Robot World agent based modelling environment, which in turn uses the Instinct Planner for cognition.

BIOLOGICALLY INSPIRED COGNITIVE ARCHITECTURES

From the 1950's through to the 1980's the study of embodied AI assumed a cognitive symbolic planning model for robotic systems — SMPA (Sense Model Plan Act) — the most well known example of this being the Shakey robot project [1]. In this model the world is first sensed and a model of the world is constructed within the AI. Based on this model and the objectives of the AI, a plan is constructed to achieve the goals of the robot. Only then does the robot act. Although this idea seemed logical and initially attractive, it was found to be quite inadequate for complex, real world environments.

In the 1990's Rodney Brooks and others [2] introduced the then radical idea that it was possible to have intelligence without representation [3]. Brooks developed his subsumption architecture as a pattern for the design of intelligent embodied systems that have no internal representation of their environment, and minimal internal state. These autonomous agents could traverse difficult terrain on insect-like legs, appear to interact socially with humans through shared attention and gaze tracking, and in many ways appeared to possess behaviours similar to that observed in animals. However, the systems produced by Brooks and his colleagues could only respond immediately to stimuli from the world. They had no means of focusing attention on a specific goal or of executing complex sequences of actions to achieve more complex behaviours. Biologically inspired approaches are still favoured by many academics, although a wide gap exists between existing implementations and the capabilities of the human mind [4]. Today, the argument persists concerning whether symbolic, sub-symbolic or hybrid approaches are best suited for the creation of powerful cognitive systems [5]. Here we concern ourselves more specifically with action selection as a core component of any useful cognitive architecture.

From Ethology to Robots

Following in-depth studies of animals such as gulls in their natural environment, ideas of how animals perform action selection were originally formulated by Nico Tinbergen and other early ethologists [6], [7]. Reactions are based on pre-determined drives and competences, but depend also on the internal state of the organism [8]. Bryson [9] harnessed these ideas to achieve a major step forwards with the POSH (Parallel Ordered Slipstack Hierarchy) reactive planner and the BOD (Behaviour Oriented Design) methodology, both of which are strongly biologically inspired. A POSH plan consists of a *Drive Collection (DC)* containing one or more *Drives*. Each *Drive (D)* has a priority and a releaser. When the *Drive* is released as a result of sensory input, a hierarchical plan of *Competences*, *Action Patterns* and *Actions* follows. POSH plans are authored, or designed, by humans alongside the design of senses and behaviour modules. An iterative approach is defined within BOD for the design of intelligent artefacts — these are known as agents, or if they are physically embodied, robots.

Kinds of Minds

Daniel Dennett [10] elegantly outlines a high level ontology for the kind of minds that exist in the natural world. At the most basic level, the Darwinian mind produces 'hardwired' behaviours, or phenotypes, based on the genetic coding of the organism. The Skinnerian mind is plastic, and capable of 'ABC' learning — Associationism, Behaviourism, Connectionism. The Popperian mind runs simulations to predict the effect of planned actions, anticipating experience. It therefore permits hypotheses "to die in our head" rather than requiring them to be executed in the world before learning can take place. Finally the Gregorian mind (after the psychologist Richard Gregory) is able to import tools from the cultural environment, for example language and writing. Using these tools enables the Gregorian mind, for example the human mind, to be self-reflective.

However, perhaps the simple Darwinian mind might also be arranged to monitor itself, and in some small and limited sense to be aware of its own performance and act to correct it. Bryson suggests that consciousness might assist in action selection [11], and here we investigate whether action selection achieved through reactive planning might parallel one of the

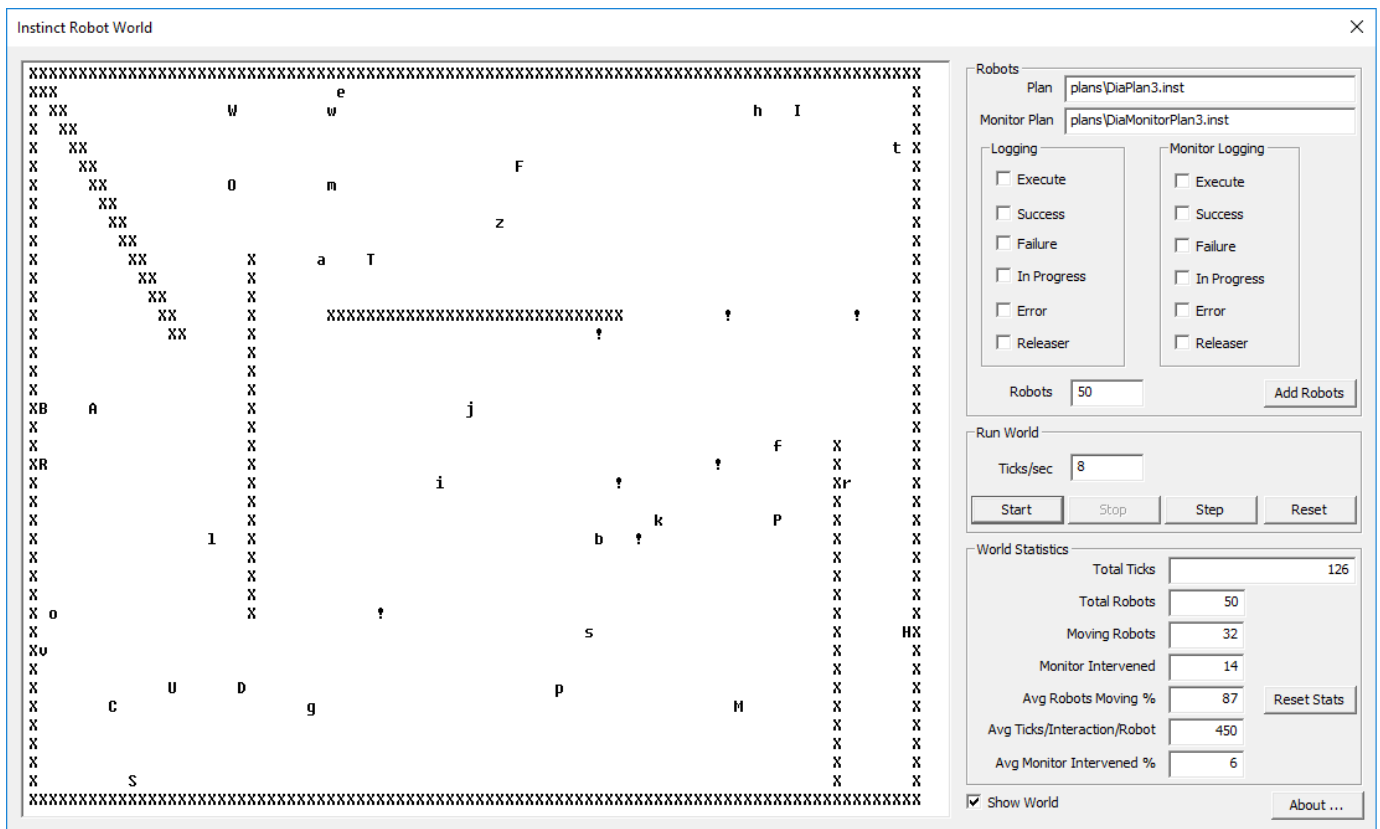


Fig. 1: Screen shot of the Instinct Robot World in operation. Each robot is represented as a single character within the display. Robots are labelled with letters and numbers to distinguish them. When a robot's monitor plan becomes active the robot representation changes to the shriek character (!). The top right section of the screen is used to control the robots and the plans they use. The bottom right section displays statistics about the world as it runs.

commonly accepted characteristics of consciousness; that is to be self-reflective and regulating [12].

Instinct and the Robot World

The *Instinct Planner* [13] is a biologically inspired reactive planner specifically designed for low power processors and embedded real-time AI environments. Written in C++, it runs efficiently on both ARDUINO and MICROSOFT VC++ environments and has been deployed within the R5 low cost maker robot to study AI Transparency [14].

Its unique features are its tiny memory footprint and efficient operation, meaning that it can operate on a low powered micro-controller environment such as ARDUINO. Alternatively, as in this experiment, many planners can run within one application on a laptop PC.

The *Instinct Robot World* is a new agent based modelling tool, shown in Figure 1. This is an open source project and all code and configuration files are available online ¹. Each virtual 'robot' within the Robot World uses an Instinct Planner to provide action selection. Strictly, since these virtual robots are not physically embodied, we should refer to them as agents. However, we have chosen to use 'robot' throughout, as intuitively these cognitive entities appear to be virtually embodied within the Robot World, and this choice of language seems more natural. In the final section of this paper we discuss

¹<http://www.robwortham.com/instinct-planner/>

future work where we may realise physical embodiment of this architecture.

The Robot World allows many robots to be instantiated, each with the same reactive plan, or with a variety of plans. The robots each have senses to sense the 'walls' of the environment, and other robots. The reactive plan invokes simple behaviours to move the robot, adjust its speed and direction, or interact with robots that it encounters within the world as it moves. Most importantly for this investigation, each robot also has a second Instinct Planner instance. This planner monitors the first, and is able to modify its parameters based on a predefined plan.

The Instinct Robot World provides statistical monitoring to report on the overall activity of the robots within the world. These include the average percentage of robots that are moving at any one time, the average number of time units (ticks) between robot interactions, and the average amount of time that the monitor intervenes to modify the robot plan. We use the Instinct Robot World to investigate the idea of *Reflective Reactive Planning* — one reactive planner driving behaviour based on sensory input and predefined drives and competences, and another reactive planner monitoring performance and intervening to modify the predefined plan of the first, in accordance with some higher level objective. This simple combination of two Darwinian minds, one monitoring the other, might also be considered to be a second-order Darwinian mind.

Reflective Reactive Planning - A 2nd Order Darwinian Mind

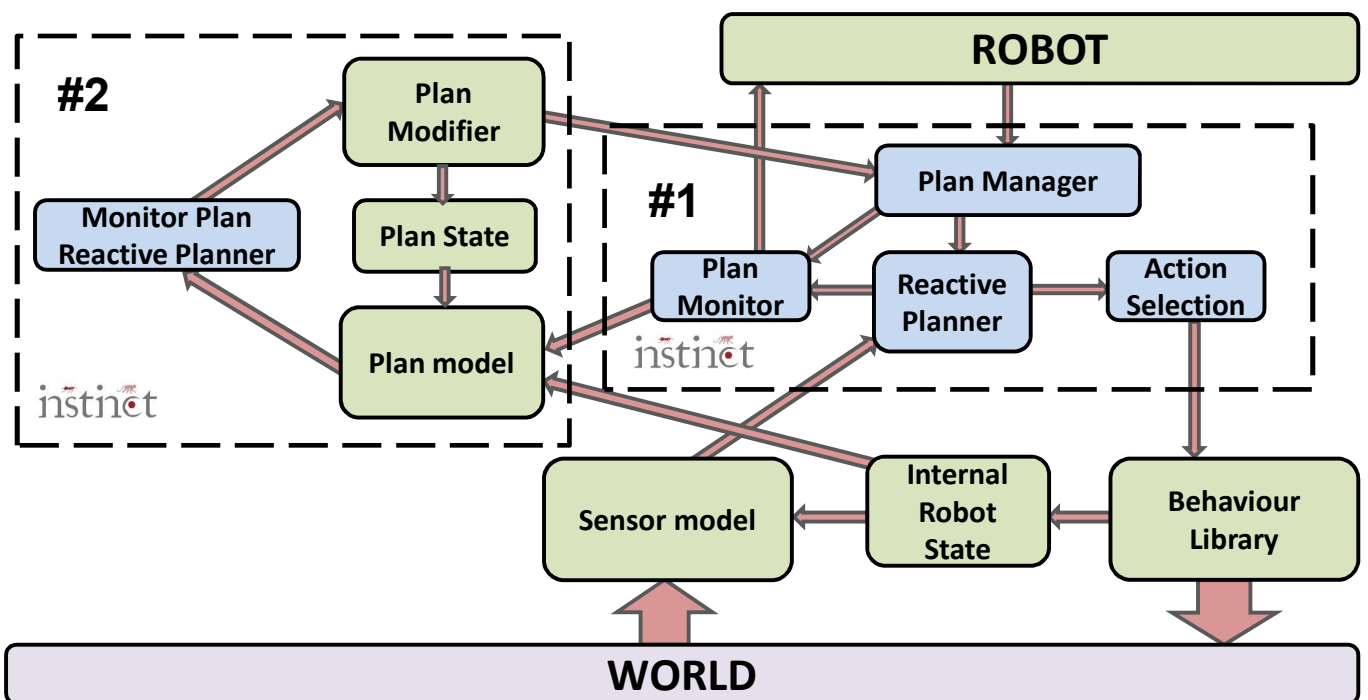


Fig. 2: Architecture of the second-order Darwinian mind. The robot is controlled by the Instinct Reactive Planner, as it interacts with the Sensor model and Behaviour Library. In turn, a second instance of Instinct monitors the first, together with the Internal robot state, and dynamically modifies parameters within the robot's planner. The overall effect is a robot that not only reacts to its environment according to a predefined set of goals, but is also to modify that interaction according to some performance measure calculated within the Plan model.

CONJECTURES

We expect that second-order Darwinian minds will outperform first order minds when the environment changes, because the monitor planner is concerned with achieving higher order objectives, and modifies the operation of the first planner to improve its performance. We also hypothesise that this architecture will remain stable over extended periods of time, because by restricting ourselves to the reactive planning paradigm we have reduced the number of degrees of freedom within which the architecture must operate, and previous work shows that first-order minds produce reliable control architectures [14]. Finally, we expect that such a second-order system should be relatively simple to design, being modular, well structured and conceptually straightforward.

METHODS

Figure 2 shows the Reflective Reactive Planning architecture implemented within the Instinct Robot World, and controlling the behaviour of each robot within that world. The robot plan has the following simple objectives, each implemented as an Instinct Drive.

- Move around in the environment so as to explore it.
- Avoid objects i.e. the walls marked as 'X' in Figure 1.
- Interact when another robot is 'encountered' i.e. when another robot is sensed as having the same coordinates within the grid of the Robot World. This interaction causes the robot to stop for 200 clock cycles or 'ticks'.

While the robot is in the 'Interacting' state it is shown as a shriek character (!) within the Robot World display. Once the robot has interacted its priority for interaction decreases, but ramps up over time. This may be likened to most natural drives, for example mating, feeding and the need for social interaction.

The Monitor Plan is designed to keep the robot exploring when it is overly diverted from social interactions. It achieves this by monitoring the time between interactions. If, over three interactions, the average time between interactions reduces below 1000 ticks, then the Monitor Planner reduces the priority of the interaction Drive. After 1000 ticks the priority is reset to its original level. We might use alternative intentional language here to say that the Monitor Planner 'notices' that the robot is being diverted by too many social interactions. It then reduces the priority of those interactions, so that the robot is diverted less frequently. After some time the Monitor Planner ceases to intervene until it next notices this situation re-occurring.

The Robot World is populated with varying numbers of robots (2, 3, 5, 10, 20, 50, 100, 200, 500, 1000), and for each number the experiment is run twice, once with a monitor plan, and once without. For each run, the environment is allowed to run for some time, typically about 10 minutes, until the reported statistics have settled and are seen to be no longer changing over time.

OUTCOMES

The results are most elegantly and succinctly presented as simple graphs. Firstly, the average number of robots moving at any one time within the world is shown in Figure 3. In both




Fig. 3: This graph shows the average percentage number of robots that are moving at any one time within the world, for a given total number of robots in the world. It can be seen that the addition of the monitor plan maintains more robots moving as the number of robots increases. Note the log scale for robots in world.

cases, as the number of robots within the world increases, the amount of time that the robot spends moving reduces. However the Monitor Planner acts to reduce the extent of this reduction from 60% to less than 20% over the full range of two to a thousand robots within the world. Similarly, in Figure 4 we see that as more robots are introduced into the world, the average time between interactions naturally reduces. However, the action of the Monitor Planner progressively limits this reduction, so that with 1000 robots the time between interactions is almost trebled, from 310 to 885 ticks per interaction. Interestingly, in both these graphs we see smooth curves both with and without the action of the monitor plan.




Fig. 4: This graph shows the average time between robot interactions, both with and without the monitor plan. The addition of the monitor plan reduces the variance in interaction time as robot numbers vary. Again, note the log scale.

The final graph, Figure 5 also shows a smooth, sigmoid like increase in activation of the Monitor Planner as the number of robots increases, plotted on a logarithmic scale.




Fig. 5: This graph shows the average percentage number of robots whose monitor plan is activated at any one time, for a given number of total robots in the world. Note the log scale.

The Instinct Robot World was found to be a stable, reliable platform for our experiments, and the results it achieved were repeatable. The application is single threaded, and so uses only one core of the CPU on the laptop PC on which it was run. Nevertheless, it was possible to simulate 1000 robots with both reactive planners active operating in the world at the rate of 70 clock cycles (ticks) per second.

DISCUSSION

From the results we can see that by using a second Instinct instance to monitor the first, we can achieve real-time learning within a tiny-footprint yet nevertheless symbolic cognitive architecture. In addition, since this learning modifies parameters from a human designed plan, the learning can be well understood and is transparent in nature. This contrasts strongly with machine learning approaches such as neural networks that typically learn offline, are opaque, and require a much larger memory workspace. Despite the stochastic nature of the environment, the performance graphs show smooth curves over a wide range of robot populations.

This relatively simple experiment also provides further fuel for the fire concerning the philosophical discussion of the nature of consciousness. Critics may say that when we use the intentional stance [15] to describe the behaviour of the Monitor Planner as 'noticing' something, we are merely using metaphor. They might argue that there is in fact no sentience doing any noticing, and in fact the only 'noticing' that is happening here is us noticing the behaviour of this human designed mechanism, which itself is operating quite without any sentience and certainly without being conscious [16]. But that is to miss the point. We are not claiming that this architecture is conscious in the human or even significant sense of the word, merely that our architecture is inspired by one aspect of how biological consciousness appears to operate. However, having shown that this architecture can indeed provide adaptive control, and drawing on the knowledge that gene expression produces behaviours which can be modelled using reactive planning, we might also consider whether consciousness in animals and humans may indeed

arise from complex hierarchical mechanisms. These mechanisms are biologically pre-determined by genetics, and yet in combination yield flexible, adaptive systems able to respond to changing environments and optimise for objective functions unrelated to the immediate competences of preprogrammed behavioural responses. This is not to argue for some kind of emergence [17], spooky or otherwise, but more simply to add weight to the idea that the ‘I’ in consciousness is nothing more than an internal introspective narrative, and such a narrative may be generated by using hierarchical mechanisms that notice one another’s internal states, decision processes and progress towards pre-defined (phenotypic) objectives.

We could certainly envisage a much grander architecture, assembled at the level of reactive planners, using maybe hundreds or thousands of planners each concerned with certain objectives. Many of these planners may be homeostatic in nature, whilst others would be concerned with the achievement of higher level objectives. We must remember that planners merely coordinate action selection, and say nothing about how sensor models may be formed, nor how complex behaviours themselves may be implemented. However, all dynamic architectures need some kind of decision centric ‘glue’ to bind them together, and reactive planning seems to be a useful candidate here, as evidenced by practical experiment and biological underpinning.

Machine transparency is a core element of our research. We have shown elsewhere [14] that reactive planners, particularly the Instinct Planner, are able to facilitate transparency. This is due to the human design of their plans, and the ability to gather meaningful symbolic information about internal system state and decision processes in real-time as the planner operates. This ability to inspect the operation of the architecture may assist designers in achieving larger scale cognitive implementations. Equally importantly, transparency is an important consideration for users and operators of intelligent systems, particularly robots, and this is highlighted in the EPSRC Principles of Robotics [18].

The human brain does not run by virtue of some elegant algorithm. It is a hack, built by the unseeing forces of evolution, without foresight or consideration for modularity, transparency or any other good design practice. If we are to build intelligent systems, the brain is not a good physical model from which we should proceed. Rather, we should look at the behaviours of intelligent organisms, model the way in which these organisms react, and then scale up these models to build useful, manageable intelligent systems.

Whilst our Reflective Reactive Planner is a very simple architecture, it does share many of the characteristics cited for architectures that are worthy of evaluation, such as efficiency and scalability, reactivity and persistence, improvability, and autonomy and extended operation [19]. We hope that our work with reactive planners might strengthen the case for their consideration in situations where decision centric ‘glue’ is required.

CONCLUSIONS AND FURTHER WORK

We have shown that a second-order Darwinian mind may be constructed from two instances of the Instinct reactive

planner. This architecture, which we call Reflective Reactive Planning, successfully controls the behaviour of a virtual robot within a simulated world, according to pre-defined goals and higher level objectives. We have shown how this architecture may provide both practical cognitive implementations, and inform philosophical discussion on the nature and purpose of consciousness.

The Instinct Robot World is an entirely open source platform, available online. We welcome those interested in agent based modelling, cognitive architectures generally, and reactive planning specifically, to investigate these technologies and offer suggestions for new applications and further work. One possibility might be to apply this architecture to the Small Loop Problem [20], a specific challenge for biologically inspired cognitive architectures.

We continue to develop robot applications for the Instinct Planner, together with the Instinct Robot World. We are investigating the use of a small robot swarm to build a physically embodied version of this experiment. To this end, we are currently working with the University of Manchester’s Mona robot².

REFERENCES

- [1] N. J. Nilsson, “Shakey the Robot,” SRI International, Technical Note 323, Tech. Rep., 1984.
- [2] C. Breazeal and B. Scassellati, “Robots that imitate humans,” *Trends in Cognitive Sciences*, vol. 6, no. 11, pp. 481–487, 2002.
- [3] R. A. Brooks, “Intelligence Without Representation,” *Artificial Intelligence*, vol. 47, no. 1, pp. 139–159, 1991.
- [4] A. V. Samsonovich, “Extending cognitive architectures,” *Advances in Intelligent Systems and Computing*, vol. 196 AISC, pp. 41–49, 2013.
- [5] A. Lieto, A. Chella, and M. Frixione, “Conceptual Spaces for Cognitive Architectures : A lingua franca for different levels of representation,” *Biologically Inspired Cognitive Architectures*, no. November, pp. 1–9, 2016. [Online]. Available: <http://dx.doi.org/10.1016/j.bica.2016.10.005>
- [6] N. Tinbergen, *The Study of Instinct*. Oxford, UK: Oxford University Press, 1951. [Online]. Available: <https://books.google.co.uk/books?id=WqZNkgEACAAJ>
- [7] N. Tinbergen and H. Falkus, *Signals for Survival*. Oxford: Clarendon Press, 1970. [Online]. Available: <http://books.google.co.uk/books?id=5LHwAAAAMAAJ>
- [8] J. J. Bryson, “The study of sequential and hierarchical organisation of behaviour via artificial mechanisms of action selection,” 2000, M.Phil. Thesis, University of Edinburgh.
- [9] —, “Intelligence by design: Principles of modularity and coordination for engineering complex adaptive agents,” Ph.D. dissertation, MIT, Department of EECS, Cambridge, MA, June 2001, AI Technical Report 2001-003.
- [10] D. C. Dennett, *Kinds of minds: Towards an understanding of consciousness*. Weidenfeld and Nicolson, 1996.
- [11] J. J. Bryson, “A Role for Consciousness in Action Selection,” in *Proceedings of the AISB 2011 Symposium: Machine Consciousness*, R. Chrisley, R. Clowes, and S. Torrance, Eds. York: SSAISB, 2011, pp. 15—20.
- [12] J. W. Sherman, B. Gawronski, and Y. Trope, *Dual-Process Theories of the Social Mind*. Guilford Publications, 2014. [Online]. Available: <https://books.google.co.uk/books?id=prtaAwAAQBAJ>
- [13] R. H. Wortham, S. E. Gaudl, and J. J. Bryson, “Instinct : A Biologically Inspired Reactive Planner for Embedded Environments,” in *Proceedings of ICAPS 2016 PlanRob Workshop*, London, UK, 2016. [Online]. Available: <http://icaps16.icaps-conference.org/proceedings/planrob16.pdf>
- [14] R. H. Wortham, A. Theodorou, and J. J. Bryson, “Robot Transparency : Improving Understanding of Intelligent Behaviour for Designers and Users,” in *Proceedings of TAROS 2017*. Guildford, UK: {accepted for publication}, 2017.
- [15] D. C. Dennett, *The intentional stance*. MIT press, 1989.

²<http://www.monarobot.uk/>

- [16] P. O. A. HAIKONEN, "Consciousness and Sentient Robots," *International Journal of Machine Consciousness*, vol. 05, no. 01, pp. 11–26, 2013. [Online]. Available: <http://www.worldscientific.com/doi/abs/10.1142/S1793843013400027>
- [17] J. H. Holland, *Emergence: From Chaos to Order*, ser. Popular science / Oxford University Press. Oxford University Press, 2000. [Online]. Available: <https://books.google.co.uk/books?id=VjKtpujRGuAC>
- [18] M. Boden, J. Bryson, D. Caldwell, K. Dautenhahn, L. Edwards, S. Kember, P. Newman, V. Parry, G. Pegman, T. Rodden, T. Sorell, M. Wallis, B. Whitby, and A. Winfield, "Principles of robotics," The United Kingdom's Engineering and Physical Sciences Research Council (EPSRC), April 2011, web publication.
- [19] P. Langley, J. E. Laird, and S. Rogers, "Cognitive architectures: Research issues and challenges," *Cognitive Systems Research*, vol. 10, no. 2, pp. 141–160, jun 2009. [Online]. Available: <http://linkinghub.elsevier.com/retrieve/pii/S1389041708000557>
- [20] O. L. Georgeon and J. B. Marshall, "The small loop problem: A challenge for artificial emergent cognition," *Advances in Intelligent Systems and Computing*, vol. 196 AISC, pp. 137–144, 2013.