Hummel, J. E. and Holyoak, K. J., 1997, Distributed representations of structure: A theory of analogical access and mapping. *Psychological Review*, **104**: 427–466.

Kanerva, P., 1998, Encoding structure in Boolean space. In *Proceedings of the International Conference on Artificial Neural Networks*, pp. 387–392.

Love, B. C., 1999, Utilising time: Asynchronous binding. In M. S. Kearns, S. A. Solla and D. A. Cohn (eds) *Advances in Neural Information Processing*, Vol. 11, pp. 38–44, (Cambridge, MA: MIT Press).

Plate, T., 1995, Holographic reduced representations. *IEEE Transactions on Neural Networks*, **6**: 623–641.

Rachkovskij, D. A. and Kussul, E. M., 2001, Bindng and normalisation of binary sparse distributed representations by context-dependent thinning. *Neural Computation*, **13**: 411–452.

Trehub, A., 1991, *The cognitive brain* (Cambridge, MA: MIT Press).

Antony Browne
School of Computing, Information Systems and Mathematics
London Guildhall University
100 Minories, London EC3N 1JY, UK
email: abrowne@lgu.ac.uk

Evolutionary Robotics: The Biology, Intelligence and Technology of Self-organizing Machines. By Stefano Nolfi and Dario Floreano (MIT Press, Cambridge, MA, 2000), 284 pp., US $50.00/£34.95, ISBN 0-262-14060-5 .

Evolutionary robotics (ER) has been an active field of research for more than a decade, although it has deeper historical roots in cybernetics and pre-symbolic artificial intelligence (AI). Alan Turing, W. Ross Ashby, Gordon Pask and W. Grey Walter all toyed with related ideas and, more recently, Valentino Braitenberg (1984) has directly hinted at and inspired much of the initial work in this field. ER's methodology relies on the automatic synthesis of robot controllers by means of search algorithms akin to natural evolution. It has proved to be quite successful in producing autonomous robots working in real environments of a diverse variety, and is today one of the hottest areas in AI, taught and practised only in a handful of research centres around the globe.

Stefano Nolfi and Dario Floreano, two of the main researchers in the field, have written the first book dedicated exclusively to the practice of ER, providing useful and integrated information, not just about research topics and methods, but also about many of the underlying currents of thought that motivate ER—currents that strongly resonate with dynamical, embodied and situated views of cognition. Contrary to common misconceptions, ER is not just about evolving insect-like behaviour and hoping it will eventually scale up—it is rather a disciplined way of understanding adaptive systems from a minimalist point of view; a method for constantly keeping one's assumptions in check and being frequently surprised by how simple, biologically inspired mechanisms can interact to give rise to interesting and sometimes quite complex behaviour without necessarily resembling our preconceived ideas about them. Analysis of evolved controllers, as many examples in the book explicitly show, can lead to insights about the origins of intelligent behaviour, which often lie not just within the robot's controller, but also in its body and the active exploitation of environmental regularities. Such implications can extend beyond the simple behaviour evolved for the robot and serve as intuition pumps for a variety of problems in AI and cognitive science.

The book describes a series of detailed case studies, drawn from the authors' own research and from that of others. It is accessible to advanced undergraduate and postgraduate students as well as being aimed at researchers in AI, cognitive science, neuroscience and animal behaviour with interests in ER as a modelling technique. The studies are organized around themes such as reactive navigation, non-reactive controllers, learning and evolution, co-evolved pursuit evasion and complex robot morphologies including walking robots and evolvable hardware. Useful technical details and boxed marginal comments on interesting effects, controversies and general implications are distributed along the text in such a way that method, technique, result and insight are difficult to separate. This is just as well, as it reflects the real experience of someone who engages for the first time in this mode of research (and also the experience of the expert researcher). A significant amount of work in ER, as in other disciplines, is not necessarily reflected in the finished, published product, but goes on in the process of obtaining the desired results. It is a constant learning curve and the skilful practitioner can sometimes be described as a cross between a curious engineer and rigorous artist. Nolfi and Floreano manage to reflect aspects of this process in the structure of their book. The slightly undesirable side-effect, despite the very good introductory chapters, is that the reader must be aware that motivations, technical details and broader implications are not cleanly separated, but will be found next to each other and throughout the text.

What is the rationale behind ER? It is the use of an automated design method to synthesize structures that conform to either broad or specific constraints. For the most part these constraints tend to be purely behavioural—a certain robot performance is sought—and the degree to which the desired performance is achieved serves to define an objective fitness function that is used in the process of selecting for increasingly better controllers. The structure of a robot's controller (which could include features of the robot's body as well as activity-dependent functions of time such as developmental rules or rules for plastic change) is encoded in a genotype (a string of binary or real-valued numbers). After a developmental stage, which could be a straightforward mapping or a complex epigenetic encoding—a whole chapter is dedicated to this topic—robots are tested and selected according to their fitness scores. New controllers are generated by mutating and recombining the fitter genotypes and the process is repeated for a number of iterations.

Although typically behavioural, fitness constraints can sometimes be more complex; they can, for instance, restrict solutions only to certain types of internal dynamics, body morphologies, developmental histories, plastic rules, as well as statistics over many trials (such as in the case of learning robots where one seeks improvement over a series of performances), or they can even describe more relaxed 'higher-level' requirements. This means that there is a significant and multi-dimensional range of designer's presence in the evolutionary process, from the very involved (e.g. fixing a neural network controller to be symmetric) to the far removed (e.g. two identical robots must co-ordinate their actions in a role-allocation task, no matter how they achieve this (Quinn 2001)).

Designing a working robot by hand for interesting, non-trivial and robust behaviour is an increasingly hopeless task. This alone would justify some sort of automatic synthetic process like ER on practical grounds. But there is more than this. Evolutionary algorithms can find ways of solving problems that cut across conceptual boundaries and can teach us a lot about the complex causal networks that give rise to co-ordinated behaviour. Such results have the effect of awakening us to our own preconceptions; and this is always a desired effect, but it also brings some new problems.

If one approaches ER with a scientific aim, for instance to model certain behaviours in natural creatures, or to test hypotheses about the nature of the mechanisms underlying certain kinds of behaviour, one is immediately confronted with the difficulty of analysis. This difficulty correlates precisely with the fact that the evolutionary algorithms work with a minimum of conceptual assumptions. Very rapidly the obtained results are complex enough that one must resort to neuro-ethological and psychophysical experiments on evolved robots to begin to understand what is going on; direct analysis of the controller is too difficult because cleanly divided, modular solutions are not found very often unless one constrains evolution to search for them. Studying evolved controllers is still easier than trying to understand the mechanisms of fully fledged animal behaviour in the wild but, nevertheless, not a trivial task.

The other side of this problem, if one comes to the field with engineering purposes in mind, is that because evolution can in principle take advantage of any regularity during the fitness evaluations, robots may end up performing the right behaviour but for the 'wrong' reasons, that is only because they rely on contingent regularities of their test environment, such as the fact that they start to move from a biased initial orientation, or a certain fluorescent lamp was on during all the evaluations, or sensor, motors and controllers are updated in a specific order, and so on. When the robot is tested under different conditions it fails to perform. This is a typical problem when simulated environments are used to speed up the evolutionary process and then the evolved controller is tested on a real robot and in a real environment. Fortunately, the book describes techniques that have been developed to avoid such problems, such as Nick Jakobi's minimal simulations (Jakobi 1997), where the key idea is to make contingent details as unreliable as possible from one trial to the next with the clever use of noise so that the best evolutionary strategy is to ignore them.

The examples in the book are not limited to 'classics'. ER is a very active field and the book also gives an account of cutting edge research, such as the evolution of plastic controllers by Urzelai and Floreano, where local rules of activity-dependent plastic change are encoded genetically instead of directly encoding the weight values for a neural network controller. Learning and performance in successful robots happen in an integrated manner and cannot be fully separated. Successful performance is also obtained when the same plastic controller is transplanted into a different robot's body.

This and other pieces of work are part of a recent trend towards using synthetic techniques not just to arrive at more interesting robot behaviour, but to understand what the global effects are of biologically inspired micro-mechanisms when these are implemented in controllers for well understood tasks. Today, the current direction of much work in ER is not so much 'upward' in the scale of complexity, but 'inward' in the detailed understanding of the relations between the structure of behaviour and the spatio-temporal structure of dynamical controllers as well as their overall relation to evolvability and fitness function design. Examples of this trend include visual shape discrimination using neural controllers with diffusible 'gaseous' neuromodulators (Husbands *et al.* 1998), visually guided navigation using spiking neurons (Floreano and Mattiussi 2001) and radical visual adaptation using homeostatic regulation of local plasticity (Di Paolo 2000). Increasingly, as these examples show, the purpose of work in ER is less centred on trying to obtain more 'cognitively' complex performance—a goal that has not been abandoned—and more on understanding other dimensions of adaptation and the role of different kinds of underlying mechanisms. The design and study of novel integrated systems of this sort may well be one way for evolutionary robotics to contribute useful information back to biology, especially neuroscience, in the proximate future.

As these recent trends show, ER is a rapidly moving field, novel ideas are developing very fast and the potential of evolutionary methods is far from being fully mapped. However, by covering both the basics as well as more complex topics in a solid, down-to-earth manner, this book is bound to remain the definitive introduction for years to come.

References

Braitenberg, V., 1984, *Vehicles: Experiments in Synthetic Psychology* (Cambridge MA: MIT Press).

Di Paolo, E. A., 2000, Homeostatic adaptation to inversion of the visual field and other sensorimotor disruptions. In J. A. Meyer, A. Berthoz, D. Floreano, H. Roitblat and S. Wilson (eds) *From Animals to Animats 6: Proceedings of the 6th International Conference on the Simulation of Adaptive Behavior*, Paris, France (Cambridge MA: MIT Press).

Floreano, D., and Mattiussi, C. 2001, Evolution of spiking neural controllers for autonomous vision-based robots. In T. Gomi (ed.) *Evolutionary Robotics IV* (Berlin: Springer).

Husbands, P., Smith, T., Jakobi, N., and O'Shea, M., 1998, Better living through chemistry: evolving GasNets for robot control. *Connection Science*, **10**: 185–210.

Jakobi, N., 1997, Evolutionary robotics and the radical envelope-of-noise hypothesis. *Adaptive Behavior*, **6**: 325–368.

Quinn, M., 2001, Evolving communication without dedicated communication channels. In J. Kelemen and P. Sosík (eds) *Advances in Artificial Life: Proceedings of the Sixth European Conference on Artificial Life* (Berlin: Springer), pp. 357–366.

Ezequiel A. Di Paolo
School of Cognitive and Computing Sciences,
University of Sussex,
Brighton BN1 9QH, UK
email: ezequiel@cogs.susx.ac.uk