



Chapter 25

Evolutionary robotics

Dylan Evans and Walter de Back

Introduction

Robotics has been heavily influenced by cognitive psychology and evolutionary biology, but the child prodigy of these two disciplines, evolutionary psychology, has not yet had such a large impact on the field. Here we describe how evolutionary psychology may come to have a greater impact on robotics in the near future. First, however, we will briefly outline the role that the two parent disciplines of evolutionary psychology have played in the development of robotics, and describe the few lessons that have been drawn so far by roboticists from evolutionary psychology itself.

Cognitive psychology

The cognitive revolution swept through psychology in the 1960s, displacing the behaviourist paradigm that had held sway since the 1920s. Its origins, however, lie in the 1950s. If one day had to be singled out as the birthday of cognitive science, it is surely 11 September 1956. It was on that day that three seminal papers were presented at a historic meeting at the Massachusetts Institute of Technology (MIT). Allen Newell and Herbert Simon spoke about a 'logic theory machine', inaugurating the modern discipline of Artificial Intelligence (Newell and Simon 1956). Noam Chomsky described 'three models for the description of language' in a paper that has been described as marking the birth of modern linguistics (Dennett 1995, p. 384; Chomsky 1956). Finally George Miller presented a paper about short-term memory that is now recognized as one of the foundational papers of cognitive psychology (Miller 1956).

It is somewhat harder to date the origins of robotics, but the pioneering work of Grey Walter and other figures in the cybernetics movement in the late 1940s has a good claim to be considered as laying the foundations for the discipline as we know it today. The mechanical tortoises built by Walter are noticeably different from the robots after the rise of cognitive science a decade later (Walter 1951). For one thing, Walter made no attempt to build a 'mind' out of various functional components such as memory, perceptual processing, and so on. Rather, his emphasis was on the behaviour of the robot as a whole, and on seeking the simplest internal structure that would enable complex behaviour to emerge when the robot was placed in the right kind of environment. This approach may have been largely due to the fact that the integrated circuit had not yet been invented, which meant that Walter had to work entirely with conventional analogue electronic components, large numbers of which would have been required to build components with a distinct psychological function.

The invention of integrated circuits, in which large numbers of components were packed onto a small silicon chip, gave roboticists the means, and the simultaneous rise of cognitive science, the motive, to pay more attention to the internal computational structure of their machines. An early example of the new generation of 'cognitive robots' was Shakey. Developed at the Artificial Intelligence Center at Stanford Research Institute between 1966 and 1972, Shakey had a complex



1 cognitive architecture in which distinct functions such as perception, planning, and natural
2 language processing were implemented by separate programs, which reflected the emphasis of
3 cognitive psychology on the functional decomposition of mental processes.

4 **Evolutionary biology**

5 **Biorobotics and evolutionary robotics**

6 By the late 1980s, the cognitive approach to robotics was coming under increasing criticism from
7 Rodney Brooks and his colleagues at MIT. Brooks argued that reasoning and planning via symbol
8 systems, as in Shakey, was often unnecessary, since similar behaviours could be achieved more
9 simply and economically by building robots out of distinct modules, each directly linking sensors
10 to actuators. Unlike cognitive robots, in these 'behaviour-based' robots there was no central
11 processing unit through which all information would be routed. Instead, complex behaviour
12 emerged from the interactions between largely autonomous modules organized into a hierarchy
13 where higher levels could inhibit lower ones.

14 Brooks and his colleagues were inspired directly by ideas from evolutionary biology. He drew a
15 parallel between his approach to designing robots incrementally, and the stepwise manner by
16 which natural selection builds complexity in nature:

17 The advantage of this approach is that it gives an incremental path from very simple systems to com-
18 plex autonomous intelligent systems. At each step of the way, it is only necessary to build one small
19 piece, and interface it to an existing, working, complete intelligence.

(Brooks 1991, p. 403)

21 Others have criticized the analogy between the construction of behaviour-based robots and the
22 evolution of organisms. Paul Griffiths (1999, p. 51) is surely right when he states that it is 'implau-
23 sible that our brains evolved by adding separate mechanisms subserving new functions'. The
24 mind may be massively domain-specific as many evolutionary psychologists contend, but if it is,
25 the various mechanisms surely evolved in parallel, just like the organs of the body.

26 Brooks also likes to compare the development of artificial intelligence with the history of life on
27 earth, suggesting that it would not be possible to build robots with human-level intelligence with-
28 out first endowing them with less complex capacities such as locomotion. From this perspective,
29 the construction of Deep Blue, a chess-playing computer which won a six-game match against
30 world champion Garry Kasparov in 1997, was an ambiguous achievement. Since many people find
31 chess cognitively-challenging, it was perhaps natural for the pioneers of artificial intelligence to
32 regard chess-playing as a legitimate goal for computers, but Brooks viewed such projects as dead
33 ends. While IBM was building Deep Blue, other researchers were struggling to build robots that
34 could walk over anything except a perfectly flat surface without falling over. Brooks argued that the
35 flexibility of human intelligence would only be replicated in robots that had first acquired a bed-
36 rock of capacities, such as locomotion, that had evolved much earlier than abstract reasoning.

37 The robots built by Brooks were inspired by general principles of evolutionary biology, and
38 sometimes by generic features of large classes of organisms such as hexapod locomotion, but were
39 not generally intended to replicate particular species. During the 1990s, however, researchers
40 began to use robots to test detailed hypotheses about the biological mechanisms of individual
41 species. Barbara Webb, for example, developed a spiking neuron network to model the process
42 by which crickets processed acoustic signals, and mounted the circuit on a mobile robot to test
43 her hypotheses about how female crickets locate males by detecting their songs (Webb 1995).
44 Other roboticists have used a similar approach to test hypotheses about whisking in rats and
45 navigation in ants, giving rise to a dynamic subfield known as 'biorobotics' (Webb 2001).



1 A related area of research, known as ‘evolutionary robotics’ (Nolfi and Floreano 2000), uses
2 genetic algorithms and other types of evolutionary computation to develop artificial neural net-
3 works (and sometimes robot morphology). Most experiments in evolutionary robotics follow the
4 same basic format. First, a population of artificial chromosomes that code for the control system
5 and/or the morphology of a robot are randomly generated. Next, these chromosomes are decoded
6 and the resulting robot (which may be physical or simulated) is set free to act in a given environ-
7 ment while its performance on various tasks is automatically evaluated according to a predefined
8 criterion known as the fitness function. The fittest robots are then allowed to reproduce by gen-
9 erating copies of their genotypes, with the addition of changes introduced by some genetic opera-
10 tors such as mutation and recombination. This process is then repeated for several generations,
11 allowing selection to enhance the average fitness of the population in the same way as it has done
12 so effectively in the history of life as we know it.

13 Most current research in biorobotics and evolutionary robotics tends to fall into one of the
14 following categories:

15 **Navigation**

16 A capacity for navigation is one of the most basic requirements for an autonomous mobile robot,
17 so it is unsurprising that this remains a focus of evolutionary robotics research. Typical tasks that
18 robots are set include obstacle avoidance and navigating toward a target area. Since most animal
19 navigation is visually guided, this field of research includes a great deal of work on visual sensing,
20 including such areas as object recognition and discrimination, though many robots still use infra-
21 red sensors or sonar instead of, or in addition to, a camera.

22 **Competitive co-evolution**

23 Most experiments in the evolution of robot navigation involve testing each member
24 of each generation, one-by-one, in a fixed physical environment. Experiments in co-evolution, by
25 contrast, involve a continually changing dynamic environment consisting of other agents.
26 Biologists have long speculated that such an environment may enhance the adaptive power
27 of natural selection. For example, Dawkins and Krebs (1979) argued that competing populations
28 may reciprocally drive one another to increasing levels of behavioural complexity by
29 producing an evolutionary arms race. Researchers in artificial life began to develop virtual
30 simulations of such arms races in the early 1990s (Miller and Cliff 1994), and more recently,
31 researchers in evolutionary robotics have developed fully-embodied physical models of
32 predator–prey dynamics.

33 **Cooperative collective robots**

34 Experiments in collective robotics need not be entirely competitive. A thriving strand of research
35 studies the evolution of cooperation among teams of mobile robots. Such experiments often
36 involve modelling the behaviour of social insects (e.g. ants, bees) which display sophisticated col-
37 lective intelligence. This swarm intelligence (Bonabeau et al. 1999) emerges from the networks of
38 interactions among agents that are individually quite simple. Research in this area includes work
39 on social signalling; Luc Steels (1998), in a pioneering series of experiments, used robots to
40 explore various hypotheses about the origins of language.

41 **Evolvable hardware**

42 Most research in evolutionary robotics takes the robot body as a fixed parameter, focusing exclu-
43 sively on evolving robot control systems. However, a few projects have applied the evolutionary



1 process to robot bodies. The most striking experiment so far in this field involved the
2 evolution of simple locomotive systems composed of bars and actuators, and then used rapid-
3 prototyping technology to produce the multi-linked structures, so that the only human input
4 needed was the final attachment of snap-on motors (Lipson and Pollack 2000). Another approach
5 in evolvable hardware is evolving robot controllers in reconfigurable electronic circuits such
6 as field-programmable gate arrays (FPGAs). These are argued to be analogue dynamical
7 continuous-time systems and thereby avoid the constraints of discrete digital design
8 (Thompson 1997).

9 The explicit discussion of evolvable hardware draws attention to an aspect of evolutionary
10 robotics that may not be apparent from the discussion so far: the fact that experiments can be
11 done by means of computer simulations as well as with real, physical robots. The relative advan-
12 tages and disadvantages of each approach are the subject of some debate in evolutionary robotics.
13 Some argue that physical robots are better models since they incorporate real problems in coping
14 with physical forces like sensory noise, energy consumption, damage, and inertia. However,
15 experiments involving physical robots are very time-consuming. Simulated robots avoid this
16 problem, but also circumvent the physical problems that may well be essential to the structure of
17 natural minds. Another important problem with physical robots is the inability to evolve bodily
18 structures (evolvable hardware as described above is a rare exception), which is essential for the
19 modelling of co-evolution of body and brain. This is so far only possible in simulation. Hybrid
20 approaches perform the most time-consuming periods of evolution in simulation, and subse-
21 quently evolve the controllers further in physical robots.

22 **Evolutionary psychology**

23 The previous two sections have provided a brief summary of the huge impact that the two parent
24 disciplines of evolutionary psychology—cognitive psychology and evolutionary biology—have
25 each had on the development of robotics. Given this legacy, it is somewhat surprising that evolu-
26 tionary psychology has itself had little influence so far on robot research. In this section we discuss
27 some ways in which evolutionary psychology may be applied to the development of robots in the
28 near future.

29 **Synthetic evolutionary psychology**

30 The methods used by evolutionary psychologists so far may be described as analytic, in the sense
31 that they start with a real system (the minds of modern humans and/or of various ancestral spe-
32 cies) and attempt to collect data about this system that might permit us to infer the internal
33 structure of this system. There is nothing inherently wrong with analytic methods—indeed, they
34 are the backbone of most modern science—but researchers are increasingly aware of their draw-
35 backs. One important drawback is that the difficulty of analysing a system grows exponentially
36 with the complexity of the system. Since minds, especially the minds of advanced primates such
37 as our own and those of our recent ancestors, are notoriously complex systems, it follows that all
38 analytic methods are very hard to apply to the study of the human mind in a way that all research-
39 ers agree upon. One has only to look at the voluminous polemical literature that has arisen in
40 response to the version of the Wason selection task originally described by Cosmides and Tooby,
41 to realize how difficult it is to derive uncontroversial conclusions from analytic methods in evo-
42 lutionary psychology (Cosmides and Tooby 1992).

43 Valentino Braitenberg (1984) has proposed, in respect of such difficulties, that when it comes
44 to complex systems it is often easier to discover their internal structure by synthetic methods. In
45 other words, if we wish to discover how some system works, it is often easier to do so by building



1 successively more complex models, rather than by attempting to infer the mechanism from mere
2 observation:

3 It is pleasurable and easy to create little machines that do certain tricks. It is also quite easy to observe
4 the full repertoire and behaviour of these machines even if it goes beyond what we had originally
5 planned, as it often does. But it is much more difficult to start from the outside and to try to guess
6 internal structure just from the observation of behaviour.

7 (Braitenberg 1984, p. 20)

8 Braitenberg refers to this generalization as the law of uphill analysis and downhill synthesis. But
9 this way of putting things is misleading to the extent that it suggests that the researcher is faced
10 with a choice between analytic and synthetic methods. In reality, analytic and synthetic methods
11 are not alternatives, but complementary aspects of a dialectic that involves moving back and forth
12 between the analysis of empirical data and the construction of simple models of underlying
13 mechanisms.

14 Synthetic methods involve building models of the system under investigation and then observ-
15 ing their behaviour. The more closely the behaviour of the model corresponds to the behaviour
16 of the target system, the more confident we can be that the internal structure of the model cor-
17 responds to the internal structure of the target. Because we have built the model ourselves, its
18 internal structure is transparent, and need not be inferred by analysis.

19 With a few notable exceptions (such as Geoffrey Miller, Gerd Gigerenzer, and Douglas Kenrick),
20 most evolutionary psychologists have not attempted to translate their hypotheses about human
21 mental structure into working machines. This might seem to exclude evolutionary psychology
22 from cognitive science, a core feature of which is often taken to be the design-based approach.
23 However, this conclusion is too quick, for emphasis on design does not imply that all cognitive
24 scientists must take an active part in building artificial minds. It simply requires that cognitive
25 scientists propose models of the mind that are computational enough to permit computer pro-
26 grams to be readily designed on the basis of the models. Many of the domain-specific mechanisms
27 proposed by evolutionary psychologists take such a form; they are not specified in terms of any
28 programming language, but they are often spelled out in a form that would be relatively easy to
29 convert into a computer program. The models proposed by evolutionary psychologists thus count
30 as fully cognitive, and evolutionary psychology is firmly within the fold of cognitive science.

31 Even so, it seems a shame that evolutionary psychologists have not taken more interest in trans-
32 lating their models into real machines. The tools of artificial intelligence and computational
33 modelling might well offer them ways of testing their hypotheses and thus enable them to answer
34 the common charge of telling 'just-so stories'. Critics of evolutionary psychology frequently dis-
35 miss it on the grounds that it promulgates untestable theories. Evolutionary psychologists
36 acknowledge that there are methodological difficulties posed by investigating the history of the
37 mind, but point out that most of these difficulties are not particular to their discipline. Most of
38 them are common problems faced by all those who wish to investigate evolutionary hypotheses,
39 so to be consistent the critics should also dismiss the whole of evolutionary biology. Such general
40 arguments, however, would be strengthened if evolutionary psychologists could also point to
41 experimental ways of testing their hypotheses. Robotics could supply evolutionary psychology
42 with just such experimental techniques.

43 Computer simulations of evolutionary processes are already common. In one early example,
44 Thomas Ray (1992) designed a virtual world called Tierra and populated it with a simple digital
45 organism. This was a simple self-replicating program which occasionally made mistakes in the
46 copying process. These 'mutations' led to an increasingly diverse population of digital organisms.
47 Competition for limited memory space on the computer's hard disk ensured that there was





1 differential survival. The conditions for natural selection were therefore all in place, and Ray was
2 able to observe numerous cases of digital evolution complete with virtual viruses, parasite resist-
3 ance, and other surprisingly 'natural' features.

4 Tierra is only a virtual world, and thus subject to the criticisms of roboticists like Rodney Brooks,
5 who argue that it is all too easy in such simulations to make some crucial but unnoticed simplifica-
6 tion that renders the simulation invalid. In order to avoid this potential danger, Brooks recommends
7 that cognitive scientists work with actual physical robots. By using the techniques of evolutionary
8 robotics, evolutionary psychologists could observe artificial evolution in the real world, giving rise to
9 a genuine lineage of robots evolving by natural selection. Perhaps these methods can provide a way
10 for evolutionary psychologists to test their hypotheses about mental evolution.

11 As an example of such an approach, den Dulk et al. (2003) studied the evolution of dual-route
12 dynamics for affective processing. In particular, they used the standard methods of evolutionary
13 robotics to examine the evolutionary justification given by LeDoux (1998) for his dual-route
14 model of fear-processing. LeDoux has found evidence that, in many mammals, fear is processed
15 simultaneously by two neural pathways, one subcortical and the other largely cortical. The sub-
16 cortical route is faster but generates many false positives, while the cortical route is slower but
17 more accurate. LeDoux argues that this dual-route mechanism evolved by natural selection
18 because it allowed animals to get the best of both worlds by escaping quickly when necessary, but
19 not wasting too much effort on false alarms. By allowing agents to evolve in a simple environment
20 consisting of predators and food, den Dulk et al. (2003) found that agents did indeed evolve a
21 dual-route mechanism similar to that proposed by LeDoux, but only when certain conditions
22 were met: the food and the predator had to be relatively hard to distinguish, and information
23 must take significantly longer to propagate via the cortical route than via the subcortical route.

24 A programme of research in synthetic evolutionary psychology might use much of the same
25 tools as those used by mainstream research in evolutionary robotics, but the issues for investiga-
26 tion would be substantially different. As we have already seen, evolutionary robotics is generally
27 concerned with navigation, co-evolution, cooperation, and evolvable hardware. Evolutionary
28 psychologists would instead be more interested in using robots to explore psychological questions
29 such as the history and structure of the various mental modules that comprise the human mind.

30 A representative sample of the kinds of problems that are thought by evolutionary psycholo-
31 gists to have led to the evolution of mental modules include: 1) finding food and discriminating
32 between nutritious substances and toxins; 2) detecting and avoiding predators; 3) eliminating
33 infectious agents from the body; 4) establishing and maintaining strategic alliances; 5) communi-
34 cating with conspecifics; 6) finding and keeping mates; and 7) rearing offspring.

35 Experiments in evolutionary robotics have already begun to explore the first two problems in
36 this list, and have touched on the problem of communication, but have left the other problems
37 virtually unexplored. This is partly due to technological limitations, but such considerations
38 apply only to conducting experiments with real physical robots. Computer simulations of evol-
39 ving robot populations could explore the other areas without much difficulty. There is great scope,
40 then, for a broad research programme to explore all of these problems in a systematic way, per-
41 haps by a graded approach that tackles each problem in the order in which they were faced by our
42 ancestors. Such an incremental approach might also throw light on the way in which prior adap-
43 tations may be co-opted by natural selection as the basis for solutions to later problems.

44 **Open-ended evolution**

45 A well-known and important difference between artificial and natural evolution is that the former
46 is goal-directed, while the latter is open-ended. Artificial evolution is a technique originally





1 devised to optimize certain parameters. Natural evolution, in contrast, does not converge to a
2 single solution and then stop: it is open-ended and continuous. There are no optimal solutions in
3 nature because the problems, arising from the biological context which includes many co-evolving
4 creatures, are constantly changing. Evolutionary psychologists are interested in the mechanisms
5 that result from natural, and therefore open-ended, evolution. A synthetic approach to evolution-
6 ary psychology should attempt to replicate this process. Over the years, various technical proce-
7 dures have been proposed that partly overcome the major differences between artificial and
8 natural evolution, such as massive co-evolution, variable genotype length, and complex geno-
9 type/phenotype mapping schemes.

10 The goal-directed nature of artificial evolution is mainly due to the use of fitness functions,
11 which are mathematical formulae used to calculate the relative fitness of an individual. Selection
12 mechanisms operate on the basis of the relative differences between fitness values: better indi-
13 viduals are selected for reproduction. After many iterations, evolution thus selects the individuals
14 that optimize the components of the fitness function. In evolutionary robotics, the task of robot
15 is thus implicitly coded in the fitness function (although the behavioural components are not
16 described). Floreano and Urzelai (2000) propose the fitness space as a framework in which fitness
17 functions can be positioned consisting of three dimensions: 1) *Functional/behavioural*: is fitness
18 dependent on the internal functioning of the controller or the behavioural manifestation?
19 2) *Explicit/implicit*: are there many or very few variables included in the fitness function?
20 3) *External/internal*: are the variables in the fitness function only available to the fitness evaluator
21 or to the evolving robot as well? According to this framework, the key to approximate open-
22 ended evolution as found in nature is to make a fitness function as behavioural, implicit, and
23 internal as possible.

24 **Alternative evolutionary schemes**

25 In evolutionary theory, fitness is defined in terms of relative reproductive success. This is incom-
26 patible with standard artificial evolution on pain of tautology: selection for reproduction cannot
27 be based on relative fitness if this is itself based on reproductive success, because it implies that
28 reproduction already took place. A more biologically plausible evolutionary scheme for experi-
29 ments in synthetic evolutionary psychology would therefore be one in which selection is not
30 based on explicit fitness values. Instead, selection should emerge from the interaction of the
31 agents with their environment. Fitness values should be used solely for the purpose of cross-
32 comparison and other analyses of the experimental results. This will have impact on the overall
33 set-up of an experiment, which will render it more biologically plausible, without making it nec-
34 essarily more detailed or complex.

35 An experiment of this kind should have populations of robots roaming the environment at the
36 same time, in contrast to standard evolutionary robotics where individual controllers are often
37 tested one at a time on the same physical robot. Having simultaneous individuals introduces
38 problems of survival and reproduction if resources are limited. The adaptations to cope with such
39 problems are precisely the ones in which we are interested in evolutionary psychology.
40 Furthermore, having simultaneous populations of potential mates implies continuous reproduc-
41 tion. This results in a gradual generational change, in contrast to many evolutionary robotics
42 experiments where generations are discrete in the sense that they are separated by a selection
43 operation. Gradual generational change and simultaneous populations of individuals means hav-
44 ing different (grand)parents and offspring in the same environment, which together with the
45 problem of survival (and growth), introduces the problem of parental care. This would also allow
46 adaptations for social transmission and cultural evolution to evolve. Even a few steps closer to a
47 biologically plausible context should allow sexual selection to influence morphology, having the



1 environment include predation, or even a full-fledged co-evolving food chain. The possibilities of
2 this kind of experiment are of course limited by the constraints imposed by physical design
3 (e.g. appropriate sensors and effectors), adaptability of robot controllers (e.g. sufficiently com-
4 plex neural networks), mechanical issues (e.g. energetic autonomy), and many other technical
5 limitations.

6 **Experimental design in synthetic evolutionary psychology**

7 Incorporating the alternative evolutionary scheme just described would have a substantial impact
8 on experimental design. Because fitness functions are omitted in the experimental set-up, there
9 are no means to formulate a specific task, which is the first step in designing a standard experi-
10 ment in evolutionary robotics. Research in synthetic evolutionary psychology should proceed
11 differently. Although specifying an entire experimental framework of synthetic evolutionary psy-
12 chology is beyond the scope of this chapter, and involves long and thorough experimentation,
13 a few important aspects are noted here.

14 A hypothesis about the emergence or structure of a particular mental module or behavioural
15 trait is generally based on the following aspects: (changing) environmental conditions, evolution-
16 ary forces, and already existing mental structures. That is, an adaptation emerges from new or
17 changed situations in the individual's environment; it benefits the organism's ability to cope with
18 the evolutionary forces; and it is based on existing mental structures or neural substrates that
19 facilitate the emergence of new abilities. Verifying such a hypothesis involves synthesizing these
20 aspects in the experiment and analysing evolving individuals for emerging adaptations in coher-
21 ence with changing environments.

22 Changing environmental conditions have played a central role in the evolution of all species.
23 Either it gradually forces adaptation in a species by slow environmental change such as diminish-
24 ing food resources, or it radically forces selection by historical natural accident like a meteor
25 impact or epidemic. Researchers disagree about the relative importance of these two kinds of
26 environmental pressures, but everyone agrees that both have played a part in the evolution of the
27 human mind. It is relatively easy to synthesize both kinds of environmental changes in conditions
28 in synthetic evolutionary psychological experiments. This can be used to test the robustness of
29 certain adaptations, to cause the emergence of new adaptations, and to provide insight into the
30 trade-off between the importance of robust processes and actual sequences in explaining the his-
31 tory of the mind. Experimenters could slowly increase or decrease the presence of predators, food
32 resources, and mates. Historical accidents can be simulated through sudden changes in availabil-
33 ity of resources or removal of individuals.

34 The existing mental structures that facilitate the emergence of a certain adaptation can be syn-
35 thesized in two ways that resemble the different goals in evolutionary psychology. If one is par-
36 ticularly interested in exploring the historical trajectory of the mind, the mental structures that
37 facilitate certain adaptations are themselves the results of earlier evolutionary experiments. That
38 is, a graded or incremental approach is taken to evolve increasingly complex robot controllers.
39 This is achieved by increasing the difficulty to survive and reproduce by adding more and more
40 environmental complexity. On the other hand, if one is more interested in the structure of the
41 mind, a somewhat more pragmatic approach can be adopted. In these experiments, the existence
42 of particular structures can be assumed. These structures can, for example, be synthesized by
43 means of neural controllers that are trained by using supervised feedback, without concern for the
44 evolutionary plausibility of their existence. Although this undermines the historical account of
45 adaptation, it can prove to be a fruitful and time-saving approach to investigate the emergence of
46 higher-level adaptations.

1 Once these three conditions have been specified and synthesized, the individual creatures must
2 be designed with the appropriate sensors, motors, and neural controllers. Experimenters subse-
3 quently look for evidence for adaptations in changing behaviour, examine internal functioning of
4 the controller and robustness to changing conditions, analyse fitness changes, make comparisons
5 with naturally evolved creatures, or perform lesions and look for dissociations. Plausible stories
6 about the historical evolution of situated, embodied minds can then be supported or eliminated
7 by hard evidence.

8 Synthetic evolutionary psychological research has the potential to yield new, or confirm exist-
9 ing, hypotheses about the evolution of specific adaptations. Furthermore, it can provide insight
10 about the importance of the different variables in the evolution of natural minds. It can provide
11 us with increasingly complex machines with needs comparable with those of living creatures. And
12 perhaps these machines will exhibit cognitive abilities, evolved out of the most primal forces in
13 nature, comparable with ours.

14 **The application of evolutionary psychology** 15 **to human–robot interaction**

16 Some might argue that the research programme we have outlined in the previous section is out of
17 place in a book about ‘applied evolutionary psychology’, on the grounds that this programme is
18 more about the application of robotics to evolutionary psychology than vice versa. We would
19 disagree with such a claim, since our suggestions about open-ended evolution and alternative
20 evolutionary schemes, which are directly inspired by evolutionary theory, would clearly modify
21 the existing experimental protocols in evolutionary robotics. Furthermore, some roboticists are
22 already carrying out experiments along the lines we have described, and have been directly
23 inspired by ideas from evolutionary psychology in doing so. For example, Alan Winfield and
24 Frances Griffiths are building an artificial society of robots in order to model the mechanisms by
25 which a group of organisms might make the transition from the merely social to the cultural
26 (Winfield and Griffiths 2010). The project is inspired by a memetic theory of culture (Blackmore
27 1999), itself inspired by Richard Dawkins’ speculations about non-genetic replicators at the end
28 of *The Selfish Gene* (Dawkins 1976).

29 Nevertheless, if we want a less ambiguous example of how robotics might benefit from ideas
30 drawn from evolutionary psychology, we should turn to the study of human–robot interaction.
31 When robots are designed specifically to have rich psychological interactions with humans, as in
32 pedagogical robots and artificial companions (Wilks 2010), rather than merely physical interac-
33 tions such as surgical robots and military robots, the designers ignore evolutionary psychology at
34 their peril, or must laboriously rediscover for themselves what they could have learned more
35 quickly from evolutionary psychologists. A case in point is the so-called ‘uncanny valley hypoth-
36 esis’ first advanced by the Japanese roboticist Masahiro Mori (1970). He proposed that the rela-
37 tionship between a robot’s degree of resemblance to a human and a human observer’s emotional
38 response was highly non-linear. Robots lacking humanoid features would, he speculated, produce
39 a neutral response in observers, but as robots become progressively more humanlike, we would
40 experience a more positive emotional reaction to them. However, at the point just before it
41 became impossible to distinguish between robots and humans, the robot would evoke in us a
42 sense of revulsion and horror. In reference to the shape of the graph he produced to illustrate his
43 hypothesis (see Figure 25.1), Mori dubbed this zone ‘the valley of the uncanny’.

44 Whether Mori’s use of the word ‘uncanny’ was intended to evoke Freud’s earlier essay on this
45 particular feeling is unclear, but in that essay Freud (1955 [1919]) specifically refers to ‘ingen-
46 uously constructed dolls and automata’ as prime examples of things that seem spooky. Citing a






Fig. 25.1 The valley of the uncanny.

1 1906 essay by Ernst Jentsch, Freud notes that ‘in telling a story one of the most successful devices
 2 for easily creating uncanny effects is to leave the reader in uncertainty whether a particular figure
 3 in the story is a human being or an automaton’. (However, Freud’s explanation for this phenom-
 4 enon, which traces it to the fear of castration, is bizarre and unconvincing.)

5 Various evolutionary explanations for the uncanny valley have been proposed, all of which are
 6 more convincing than Freud’s, and which are not mutually inconsistent. One explanation focuses
 7 on the psychological mechanisms for mate selection. According to this view, automatic, stimulus-
 8 driven appraisals of uncanny stimuli elicit aversion by activating an evolved cognitive mechanism
 9 for rejecting potential mates with low fertility, poor hormonal health, or ineffective immune sys-
 10 tems, based on visible features of the face and body that are predictive of those traits (Green et al.
 11 2008). Another explanation focuses on pathogen avoidance: uncanny stimuli may activate a cog-
 12 nitive mechanism that originally evolved to motivate avoidance of potential sources of pathogens
 13 by eliciting a disgust response. Thus, the visual anomalies of android robots have the same effect
 14 as those of corpses and visibly diseased individuals: the elicitation of alarm and revulsion (Green
 15 et al. 2008).

16 Those who build robots intended for rich psychological interactions with humans could also
 17 benefit from evolutionary psychology when writing the software for such machines. For example,
 18 evolutionary psychology has focused attention, in a way that other branches of psychology have
 19 not, on the cognitive mechanisms that underlie reciprocity and fairness. Those who wish to build
 20 robot companions might glean much from this literature, perhaps yielding some counter-
 21 intuitive guidelines for robot design. For example, it may be that a robot companion who does
 22 not punish its owner (in some appropriate way) when the owner treats it unfairly is less engaging
 23 than one who does (Evans 2010).

24 Conclusion

25 During its relatively brief history, evolutionary psychology has used many different methods,
 26 from experimental manipulation of human behaviour in the laboratory to observation of indig-
 27 enous peoples and analysis of archaeological data. All these methods may be called analytic, in the
 28 sense that they collect data about already-existing systems and then analyse them. We have argued
 29 that evolutionary psychologists could benefit from extending their methodological repertoire to
 30 include synthetic methods, which involve constructing artificial systems. Such artificial systems
 31 can provide useful models of evolved minds and evolutionary histories that might provide

1 evolutionary psychologists with additional data to test hypotheses about mental structure and
2 evolutionary trajectories. One kind of synthetic method in which evolutionary psychologists have
3 so far shown little interest, is evolutionary robotics. We have argued that, by ignoring this field,
4 evolutionary psychologists are missing out on a valuable research tool, and we have sketched out
5 a research programme involving the use of robots to test evolutionary psychological hypotheses.

6 We have referred to some initial work that uses robots to test hypotheses about human evolu-
7 tion. At the moment, however, this work is too rare and patchy to warrant description as a full-
8 fledged research programme. Our manifesto for a new field of synthetic evolutionary psychology
9 remains, for the moment at least, more of a vision than a reality. Time will tell whether this vision
10 is embraced by enough researchers to make it a permanent and fruitful addition to the evolu-
11 tary psychologist's repertoire of methodologies.

12 Acknowledgements

13 We are grateful to Geoffrey Miller, Verena Hafner, Lars Zwanepol, Tijn van der Zant and Rens
14 Kortmann for their comments. Dylan Evans' work was supported by a Platform Grant from the
15 Engineering and Physical Sciences Research Council, EPSRC GR/M97503.

16 References

- 17 Blackmore, S. (1999). *The meme machine*. Oxford University Press, Oxford.
- 18 Bonabeau, E., Dorigo, M., and Theraulaz, G. (1999). *Swarm intelligence: from natural to artificial systems*.
19 Oxford University Press, New York.
- 20 Braitenberg, V. (1984). *Vehicles: experiments in synthetic psychology*. MIT Press, Cambridge, MA.
- 21 Brooks, R.A. (1991). Intelligence without representation. In: J. Haugeland (ed.), *Mind design II: philosophy,*
22 *psychology, artificial intelligence*, pp. 395–420. MIT Press, Cambridge, MA.
- 23 Chomsky, N. (1956). Three models for the description of language. *IRE Transactions on Information*
24 *Theory*, 2, 113–24.
- 25 Cosmides, L. and Tooby, J. (1992). Cognitive adaptations for social exchange. In: J. Barkow, L. Cosmides,
26 and J. Tooby (eds), *The adapted mind: evolutionary psychology and the generation of culture*,
27 pp. 163–228. Oxford University Press, Oxford.
- 28 Dawkins, R. (1976). *The selfish gene*. Oxford University Press, Oxford.
- 29 Dawkins, R. and Krebs, J.R. (1979). Arms races between and within species. *Proceedings of the Royal Society*
30 *B*, 205, 489–511.
- 31 den Dulk, P., Heerebout, B.T., and Phaf, R.H. (2003). A computational study into the evolution of
32 dual-route dynamics for affective processing. *Journal of Cognitive Neuroscience*, 15, 194–208.
- 33 Dennett, D. (1995). *Darwin's dangerous idea: evolution and the meanings of life*. Allen Lane, London.
- 34 Evans, D. (2010). Wanting the impossible: the dilemma at the heart of intimate human-robot relationships.
35 In: Y. Wilks (ed.), *Close engagements with artificial companions. Key social, psychological, ethical and*
36 *design issues*, pp. 75–87. John Benjamins, Amsterdam.
- 37 Floreano, D. and Urzelai, J. (2000). Evolutionary robots with on-line self-organization and behavioral
38 fitness. *Neural Networks*, 13, 431–43.
- 39 Freud, S. (1955 [1919]). The 'Uncanny'. In: J. Strachey (ed.), *The standard edition of the complete*
40 *psychological works of Sigmund Freud*, Vol. XVII, pp. 217–56. Hogarth Press, London.
- 41 Green, R.D., MacDorman, K.F., Chin-Chang, H., and Vasudevan, S. (2008). Sensitivity to the proportions
42 of faces that vary in human likeness. *Computers in Human Behavior*, 24, 2456–74.
- 43 Griffiths, P. (1999). Author's response. *Metascience*, 8, 49–62.
- 44 LeDoux, J. (1998). *The emotional brain: the mysterious underpinnings of emotional life*. Weidenfeld and
45 Nicholson, London.



- 1 Lipson, H. and Pollack, J. (2000). Automatic design and manufacture of robotic lifeforms. *Nature*,
2 **406**, 974–8.
- 3 Miller, G.A. (1956). The magical number seven, plus or minus two: some limits on our capacity for
4 processing information. *Psychological Review*, **63**, 81–97.
- 5 Miller, G.F. and Cliff, D. (1994). Protean behaviour in dynamic games: arguments for the co-evolution of
6 pursuit-evasion tactics in simulated robots. In: D. Cliff, P. Husbands, J. A. Meyer, and S. Wilson (eds),
7 *From animals to animats 3: Proceedings of the Third International Conference on Simulation of Adaptive*
8 *Behaviour*, pp. 411–20. MIT Press, Cambridge, MA.
- 9 Mori, M. (1970). Bukimi no tani. [The uncanny valley] (translated by K.F. MacDorman and T. Minato).
10 *Energy*, **7**, 33–5.
- 11 Newell, A. and Simon, H.A. (1956). The logic theory machine- a complex information processing system.
12 *IRE Transactions on Information Theory*, **2**, 61–79.
- 13 Nolfi, S. and Floreano, D. (2000). *Evolutionary robotics: the biology, intelligence and technology of*
14 *self-organizing machines*. MIT Press, Cambridge, MA.
- 15 Ray, T.S. (1992). An approach to the synthesis of life. In: M.A. Boden (ed.), *The philosophy of artificial life*,
16 pp. 111–45. Oxford University Press, Oxford.
- 17 Steels, L. (1998). The origins of syntax in visually grounded robotic agents. *Artificial Intelligence*,
18 **103**, 133–56.
- 19 Thompson, A. (1997). Artificial evolution in the physical world. In: T. Gomi. (ed.), *Evolutionary robotics:*
20 *from intelligent robots to artificial life*, pp. 101–25. AAI Books, Ottawa, ON.
- 21 Walter, W.G. (1951). A machine that learns. *Scientific American*, **185**, 60–3.
- 22 Webb, B. (1995). Using robots to model animals: a cricket test. *Robotics and Autonomous Systems*,
23 **16**, 117–34.
- 24 Webb, B. (2001). Can robots make good models of biological behaviour? *Behavioral and Brain Sciences*, **24**,
25 1033–50.
- 26 Wilks, Y. (ed.) (2010). Close engagements with artificial companions. *Key social, psychological, ethical and*
27 *design issues*. John Benjamins, Amsterdam.
- 28 Winfield, A.F.T. and Griffiths, F. (2010). Towards the emergence of artificial culture in collective robot
29 systems. In: P. Levi and S. Kernbach (eds), *Symbiotic multi-robot organisms*, pp. 431–9. Springer, Berlin.



