

MIRO: A Robot “Mammal” with a Biomimetic Brain-Based Control System

Ben Mitchinson and Tony J. Prescott^(✉)

Department of Psychology and Sheffield Robotics, University of Sheffield,
Western Bank, Sheffield, S10 2TN, UK
{b.mitchinson,t.j.prescott}@sheffield.ac.uk

Abstract. We describe the design of a novel commercial biomimetic brain-based robot, MIRO, developed as a prototype robot companion. The MIRO robot is animal-like in several aspects of its appearance, however, it is also biomimetic in a more significant way, in that its control architecture mimics some of the key principles underlying the design of the mammalian brain as revealed by neuroscience. Specifically, MIRO builds on decades of previous work in developing robots with brain-based control systems using a layered control architecture alongside centralized mechanisms for integration and action selection. MIRO’s control system operates across three core processors, P1-P3, that mimic aspects of spinal cord, brainstem, and forebrain functionality respectively. Whilst designed as a versatile prototype for next generation companion robots, MIRO also provides developers and researchers with a new platform for investigating the potential advantages of brain-based control.

1 Introduction

Many robots have been developed that are animal-like in appearance; a much smaller number have been designed to implement biological principles in their control systems [1, 2]. Of these, even fewer have given rise to commercial platforms that demonstrate the potential for brain-based, or *neuromimetic*, control in real-world systems. Building on more than two decades of research on robots designed to emulate animal behavior and neural control [3–7]—that has developed key competences such as sensorimotor interaction, orienting, decision-making, navigation, and tracking—we teamed with an industrial designer, experts in control electronics, and a manufacturer, to create an affordable animal-like robot companion. The resulting platform, MIRO, was originally designed to be assembled in stages, where each stage constitutes a fully-operational robot that demonstrates functionality similar to that seen in animals. These stages loosely recapitulate brain development, as well as, to some extent, brain/phylogenetic evolution. with the finalized robot emulating some of the core functionality of a generalized mammal. Previous publications have reported on the potential of the MIRO robot as a biomimetic social companion [8] and as a platform for education and entertainment [9]. In this article we (i) describe the principles of brain-based control that have inspired MIRO, (ii) outline the morphology and hardware design, and (iii) detail the three key levels of the MIRO control architecture and the functionality to which they give rise.

We end our article by briefly discussing some of the trade-offs we have made in developing a functioning brain-based robot as a commercial product.

2 Principles of Mammal-Like Brain-Based Control

Living, behaving systems display patterns of behavior that are integrated over space and time such that the animal controls its effector systems in a coordinated way, generating sequences of actions that maintain homeostatic equilibrium, satisfy drives, or meet goals. How animals achieve behavioral integration is, in general, an unsolved problem in anything other than some of the simplest invertebrates. This has also been called the problem of architecture, and it is equally as problematic for robots as it is for animals [10]. Today's robots are notoriously "brittle" in that their behavior—which may appear integrated and coordinated with respect to a well-defined task—can rapidly break down and become disintegrated when task parameters go outside those anticipated by the robot's programmers. Animals can also go into states of indecision and disintegration when challenged by difficult situations [11], but generally show a robustness and capacity to quickly adapt that is the envy of roboticists [12].

We believe that neuroscience and neuroethology have important lessons for robotics concerning the problem of architecture. Specifically, theoretical and computational analyses of animal nervous systems point to the presence of "hybrid" control architectures that combine elements of reactive control with integrative mechanisms that operate both in space, coordinating different parts of the body, and in time, organizing behavior over multiple time-scales (for discussion, see, [1, 13–15]).

One key principle, whose history dates at least to the 19th century neurologist John Hughlings Jackson [16], is that of layered architecture. A layered control system is one in which there are multiple levels of control at which the sensing apparatus is interfaced with the motor system [17]. It is distinguished from hierarchical control by the constraint that the architecture should exhibit *dissociations*, such that the lower levels still operate, and exhibit some sort of behavioral competence, in the absence (through damage or removal) of the higher layers but *not* vice versa. A substantial body of the neuroscience literature can be interpreted as demonstrating layered control systems in the vertebrate brain; layering has also been an important theme in the design of artificial control systems, for instance, for autonomous robots [18]. The notion of a layered architecture has been mapped out in some detail in the context of specific types of behavior. For example, in [15], we described how the vertebrate defense system—the control system that protects the body from physical harm—can be viewed as being instantiated in multiple layers from the spinal cord (reflexes), through the hindbrain (potentiated reflexes), midbrain (coordinated responses to species-specific stimuli), forebrain (coordinated responses to conditioned stimuli), and cortex (modification of responses according to context). In this system the higher layers generally operate by modulating (suppressing, potentiating, or modifying) responses generated by the lower layers.

Whilst the brain shows clear evidence of layered control there are other important governing principles in its organization. Indeed, a system that worked by the principles of layered control alone would be too rigid to exhibit the intelligent, flexible behavior

that mammals are clearly capable of. One proposal, stemming from the research of the neurologist Wilder Penfield, is of a centralized, or *centrencephalic*, organizing principle whereby a group of central, sub-cortical brain structures serves to coordinate and integrate the activity of both higher- and lower-level neural systems [19]. Candidate structures include the midbrain reticular formation—which may be important in integrating behavior within the brainstem, and in regulating behavior during early development—and the basal ganglia, a group of mid- and forebrain structures that we have argued play a critical role in action selection. We have previously developed several embodied models of these brain systems (see Fig. 1) and have demonstrated their sufficiency to generate appropriate behavioral sequences for mobile robots engaged in activities such as simulated foraging [4, 20].




Fig. 1. Neurobotic models of control architectures. Left: [4] embedded a model of the vertebrate basal ganglia in a table-top robot and showed its ability to control action selection and behavioural sequencing for a simulated foraging task. Right: Shrewbot [5] is one of series of whiskered robots developed to explore the effectiveness of brain-based control architectures in generating life-like behaviour.

Our research on biomimetic robot control architectures is predicated on the notion that the principles of both centrencephalic organization and layered control are at work in mammalian brains and can be co-opted to generate coordinated and robust behavior for robots. Over recent years we have developed a number of neurobotic models to further test this proposition [5, 21], of which the MIRO robot is the first commercial instantiation.

A further question with regard to the problem of control architecture concerns the fundamental units of selection. The neuroethology literature suggests a decomposition of control into behavioral sub-systems that then compete to control the animal (see [15, 22], an approach that has been enthusiastically adopted by researchers in behavior-based robotics (see, e.g. [23])). An alternative hypothesis emerges from the literature on spatial attention, particularly that on visual attention in primates including humans [24]. This approach suggests that actions, such as eye movements and reaches towards targets, are generated by first computing a ‘salience map’ that integrates information about the relevance (salience) to the animal of particular locations in space into a single topographic

representation. Some maximization algorithm is then used to select the most salient position in space towards which action is then directed. Of course, the approaches of behavioral competition and salience map competition are, again, not mutually exclusive and it is possible to imagine various hierarchical schemes, whereby, for instance, a behavior is selected first and then a point in space to which the behavior will be directed. In the mammalian brain, sensorimotor loops involving the cortex, superior colliculus, basal ganglia, and midbrain areas such as the periaqueductal gray, interact to control how the animal orients towards or away from different targets and what actions and behaviors are then selected with respect to these targets [15]. Other structures provide contextual information based on past experience—the hippocampal system, for instance, contributes to the animal’s sense of time and place—thereby promoting better decisions in the here-and-now [25].

In the following we briefly describe the physical instantiation of MIRO as a robot platform and then return to the question of how MIRO has been designed to support a brain-based control architecture.

3 The MIRO Platform

The MIRO platform (see Fig. 2) is built around a core of a differential drive base and a three degree-of-freedom (DOF) neck (lift, pitch, yaw). Additional DOFs include two for each ear (curl, rotate), two for the tail (droop, wag), and one for the eyelids (open/close). Whilst these latter DOFs target only communication, the movements of the neck and body that serve locomotion and active sensing play a significant role in communication as well. The platform is also equipped for sound production and with two arrays of colored lights, one on each side, both elements serving communication and/or emotional expression.

All DOFs in MIRO are equipped with proprioceptive sensors (potentiometers for absolute positions and optical shaft encoders for wheel speed). Four light level sensors are placed at the corners of the base, two task-specific ‘cliff sensors’ point down from its front face, and four capacitive sensors are arrayed along the inside of the body shell providing sensing of direct human contact. In the head, stereo microphones (in the base of the ears) and stereo cameras (in the eyes) are complemented by a sonar ranger in the nose and an additional four capacitive sensors over the top and back of the head (behind the ears). Accelerometers are present in both head and body.

MIRO has a three-level processing stack (see below). Peripheral components are reached on an I2C bus from the ‘spinal processor’ (ARM Cortex M0), which communicates via SPI with the ‘brainstem processor’ (ARM Cortex M0/M4 dual core), which in turn communicates via USB with the ‘forebrain processor’ (ARM Cortex A8). All peripherals and some aspects of processing are accessible from off-board through WiFi connectivity (with MIRO optionally configured as a ROS—Robot Operating System—node), and the forebrain processor can be reprogrammed if lower-level access is required (lower processors can be re-programmed if desired, though with more onerous requirements to respect the specifics of the platform).



Fig. 2. The MIRO prototype companion robot. Some example MIRO behavior can be seen at <https://youtu.be/x4tya6Oj5sU>

4 Control Architecture of the MIRO Robot

As explained above, a fundamental feature of the MIRO control architecture is its layered form as further illustrated in Fig. 3 (which is not exhaustive but includes the key architectural elements). Processing loops are present at many different levels, generating actuator control signals based on sensory signals and current state. In addition, higher systems are able to modulate the operation of loops lower down (a few examples are shown in the figure) thus implementing a form of subsumption [18]. Each layer builds upon the function of those below, so that the architecture is best understood from the bottom-up. We place the loops into three groups, each loosely associated with a broad region of the mammalian central nervous system, as follows.

Spinal Cord

The first layer, which we denote “spinal cord”, provides two types of processing. The first is signal conditioning—non-state-related transformations that can be applied unconditionally to incoming signals. This includes robot-specific operations such as removing register roll-overs from the shaft encoder signals, but also operations with biological correlates. An example of the latter is automatic acquisition of the zero-point of accelerometer signals (accounting for variability of manufacture) which is functionally comparable to sensory habituation in spinal cord neurons. The resulting “cleaned” or “normalized” signals provide the input to the second type of processing in this layer—reflex loops. A bilateral “cliff reflex” inhibits forward motion of each wheel at the lowest level if the corresponding cliff sensor does not detect a floor surface. A parallel “freeze reflex” watches for signals that might indicate the presence of another agent (a tilting acceleration-due-to-gravity vector, or touch on any of the touch sensors) and inhibits all motions, sounds, and lighting effects when

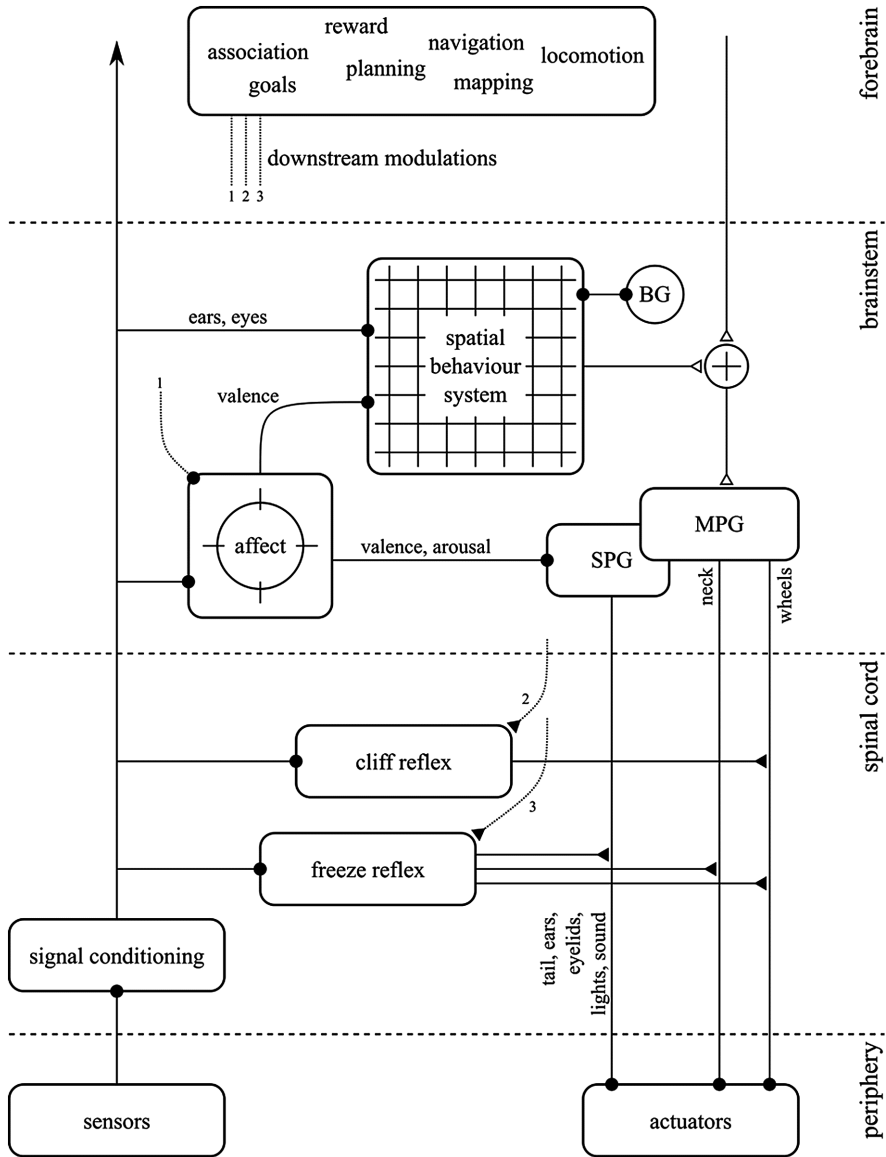


Fig. 3. Control architecture of MIRO loosely mapped onto brain regions (spinal cord, brainstem, forebrain). Signal pathways are excitatory (open triangles), inhibitory (closed triangles), or complex (closed circles). See text for description of components.

triggered. If left alone, MIRO will slowly recover and begin to move and vocalize once more.

All of the reflexes can be inhibited by higher systems, allowing them to be “switched off” if a higher-level understanding of MIRO’s context demands it. Overall, this layer can be characterized as implementing “reactive control”.

Brainstem

We group some of the most central elements of MIRO’s biological control system into the second layer, denoted “brainstem”. This layer is concerned with simple action selection, the computation and maintenance of affective state, simple spatial behaviors and the generation of motor patterns to drive the actuators.

Affect is represented using a circumplex model derived from affective neuroscience [26], that comprises a two-dimensional state representing valence (unpleasantness, pleasantness) and arousal. Fixed transforms map events arising in MIRO’s sensorium into changes in affective state: for example, stroking MIRO drives valence upwards, whilst striking him on the head drives valence down. Baseline arousal is computed from a number of sources including the real-time clock. That is, MIRO has a circadian rhythm, being more active during daylight hours. General sound and light levels also affect baseline arousal, whilst discrete events cause acute changes of affective state (very loud sound events raise arousal and decrease valence, for example).

MIRO expresses affect in a number of ways. Most directly, a set of “social” pattern generators (SPG) drive the light displays, as well as movement of the ears, tail, and eyelids, so as to indicate affective state [8]. Meanwhile, MIRO’s vocalization model, a complete generative mechano-acoustic model of the mammalian vocal system [27], is modulated by affect, so that MIRO’s voice can range from morose to manic, angry to relaxed. More indirectly, MIRO’s movements are modulated also by affect: low/high arousal slows/speeds movement, and very low arousal leads to a less upright posture of the neck.

The other major system in the brainstem layer is a spatial behavior system modeled on the management of spatial attention and behavior in superior colliculus and related nuclei in mammals [28]. This system comprises a topographic salience map of the space around MIRO’s head which is driven by aspects of both visual and aural inputs. One filter generates positive salience from changes in brightness in camera images, so that movement is typically a key generator of salience. Another, alongside, uses a Jeffress model [29] to localize the source of loud sound events so that a representation of their intensity can be added to the salience map at the appropriate location. Other aspects of these sensory streams, as well as signals from other sensory modalities, can be configured to contribute to this global salience map in a straightforward way [30].

Simple hard-coded filters generate behavioral plans from this map: “where” is computed as the maximum of the map; “what” is computed by combining MIRO’s current affective state with the nature of the stimulus (for example its size, location, or temporal nature). The system generates behavioral plans including “orient” (turn to visually “foveate” the stimulus), “avert” (turn away from the stimulus), “approach” and “flee” (related behaviors with locomotion components), and assigns a priority (a scalar value) to each plan. A model of the basal ganglia (BG) [4, 22] is then used to select, with persistence and pre-emption, one of these plans for execution by the motor plant at any one time. Overall, this system corresponds closely to similar, hard-wired, behavior

systems that have been identified in several animal species, including rodents [31] and amphibians [14].

This loop is closed through a motor pattern generator (MPG) that takes as input behavioral plans and generates time series signals for the actuators. Any behavioral plan is encoded as an open-loop trajectory for a point in the frame of reference of one of the robot's kinematic links. In all current plans, the point chosen corresponds to a "generalized sensory fovea" [30, 31] just in front of the nose; thus, MIRO is "led by the nose" as a behavioral plan executes. The MPG comprises a kinematic self-model that is computed by moving the guided point and then identifying the remaining parameters (undriven) of the model through a principle of "least necessary movement", starting with the most distal DOFs. This computation is performed using a non-iterated coordinate descent procedure. The lack of iteration limits the quality of the approximate solution, but is very cheap to compute, biologically plausible, and performs reasonably well. In previous work, we have used an adaptive filter model as a pre-processing stage to this MPG, greatly improving accuracy, and suggested that this may be a role played by mammalian cerebellum [31].

A second, distinct, kinematic self-model is used to estimate MIRO's configuration for the interpretation of sensory signals. The model combines motor efferent signals with sensory afferent (proprioceptive) signals through a complementary filter [32] to derive a timely estimate of MIRO's instantaneous configuration. This configuration is available as an input to the analysis of data with a spatial component; for example, it determines the optical axis of the cameras when a video frame was captured.

The brainstem layer contains several other sub-systems that have biological correlates. For instance, sleep dynamics are implemented as a relaxation oscillator, with wakefulness and exhaustion the two oscillator states. Thus, MIRO spends around five in every twenty minutes "asleep", expressed by closed eyes and a lowered head. Motor reafferent noise is present in MIRO's sensory streams in several forms—particular sources include obstruction of the cameras by blinking of the eyelids, corruption of video frames through self-motion (blurring), and the presence of audio noise whilst motors are active. All of these forms of noise are eliminated from the incoming data streams by gating, based on efferent and afferent cues of their presence. Thus, for example, MIRO will not attempt to detect motion when it is, itself, in motion. Selective suppression of sensory streams during some forms of motion is also a feature of biological vision [33].

Forebrain

MIRO's forebrain control systems are under present development. Figure 2 gives an indication of the character of components that are anticipated for this layer. The nature of the control architecture allows that these "higher" systems can be built on top of the existing layers, taking advantage of already-implemented functionality. For example, a higher system intended to perform task-specific orienting would not need to replicate the orienting system that is already present. Rather, a suitable modulation can be applied to the existing spatial salience filters, or additional filters added, and the orienting behavioral plan can be "primed" [30]. The result is a tendency to perform orienting towards the primed region of signal space (or physical space). In MIRO, all lower systems are amenable to modulation; highlighted in the diagram are implemented

modulation routes allowing affect to be driven by influences from the forebrain layer or reflexes to be inhibited completely allowing the recovery of direct control. Current research is directed at implementing a spatial cognition module modeled on the mammalian hippocampus that will support inhibition-of-return during exploratory behavior and will allow the robot to learn about, and navigate to, important sites such as a home ‘bed’.

The centrality of the basal ganglia model to any extension of the motor repertoire is notable. Since there is only one motor plant, only one motor pattern should be selected at any one time (simultaneous activation of multiple motor plans through the same output space constituting a motor error). Therefore, some selection mechanism is required so that only one plan is disinhibited at any one time. There is substantial support for the hypothesis that the vertebrate basal ganglia is such a centralized selection mechanism, that may implement a form of optimal decision-making between competing actions, that operates across the different layers of the neuraxis, and that has contributed to the flexibility and scalability of the vertebrate brain architecture [4, 10, 22, 31].

Processing Stack of MIRO Robot

The layers of MIRO’s biomimetic control architecture are mirrored in their implementation distributed across three on-board processors as shown in Fig. 4. A fourth level of processing, denoted “P4”, is available by inclusion of off-board systems into the control stack. The rationale for this arrangement has a pedagogic aspect (ease of understanding) but the key benefits are functional.

One important feature is that the control latency of loops through the lowest reprogrammable processor, P1, can be as low as a few milliseconds. This contrasts very favorably with the control latency through an off-board processor, P4, which—even under favorable conditions—can be hundreds of milliseconds. The inherent unreliability of wireless communications means that off-board latency can, on occasion, be longer still. Thus, safety critical aspects of the control policy, such as that implemented by the cliff reflex, will display superior performance if implemented in P1 versus, say, P4. There is, unsurprisingly, a continuum of latencies from P1 (~10 ms) through P2 (~30 ms), P3 (~50-200 ms), and P4 (100 ms or more).

Conversely, computational power (as well as energy consumption) increases as we move upwards through the processing stack (see Figure). This means that there is also a continuum of “competence”, or control sophistication. P1 can respond fast, but lacks the power to make sophisticated decisions. P2 is able to perform spatial processing, and respond quickly to the spatial nature of events, but lacks the power to perform pattern discriminations or image segmentation, say. P3 is more capable still, but as an on-board processor on a battery-powered mobile robot still has tight computational constraints. The characteristic of increasing latency and control sophistication as we move up through the different levels of layered architecture is shared by vertebrate brains [15]. Latencies of escape reflexes implemented in spinal cord can be ~10 ms, but reflex responses are relatively unsophisticated and involve minimal signal processing; meanwhile, midbrain responses to visual events begin after 50 ms [34], whilst classification of objects by human visual cortex begins to emerge at around 100 ms [35] but allows a much more sophisticated response.

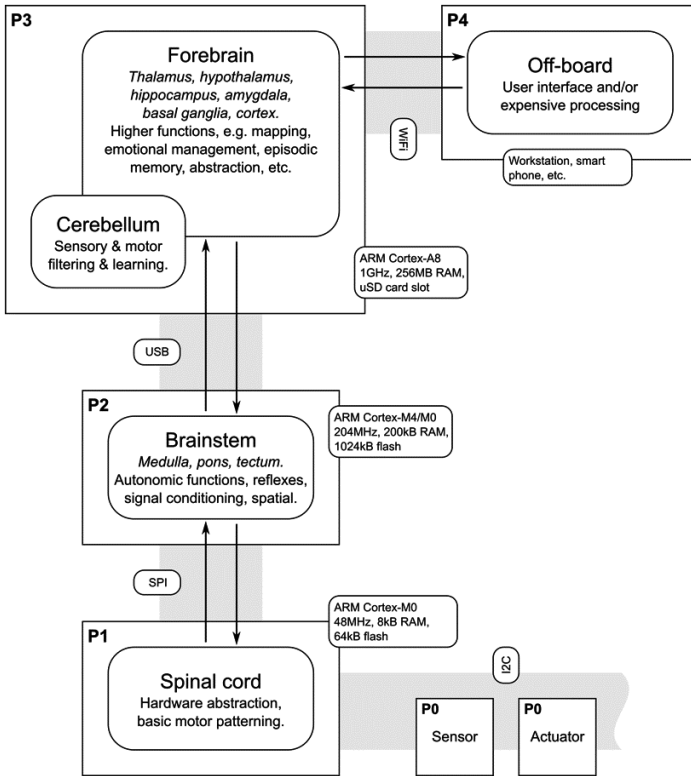


Fig. 4. The biomimetic MIRO control architecture is implemented across three on-board processors, P1-3, each loosely associated with a broad region of the mammalian brain. The design physically displays dissociation, since P2 (and, to a more limited extent, P1) is able to control the robot independently of higher processing layers. “P4” denotes off-board processing, and “P0” non-reprogrammable peripheral-specific processors.

This distribution of substrates from “fast and simple” through to “slow and sophisticated” may be a potentially useful design element for many robots. One aspect of “simple”, that can easily be forgotten, is “less likely to fail”. Especially during development, sophisticated systems such as P3 are highly prone to transient failures; having lower-level systems that protect the robot from possible damage will be beneficial. At the same time, higher processors can be put to sleep when they are not required, saving power and leaving lower processors to watch for events that may turn out to be behaviorally-relevant. The downside to this tiered processing stack is design complexity. Cost, however, may not be a serious concern, since the simpler processors are rather cheap parts.

A concern specific to robotics is increasing accessibility as we move up the stack. Running new control code on P4 can be as simple as pressing a key or clicking a mouse; on P3, at least a network file transfer will be required, and perhaps also a cross-compile step; changing the control code in P2 requires reprogramming the majority of the sectors

of the on-board FLASH, an operation that takes a few seconds, and requires the harnessing of P3 as a mediator; reprogramming P1 requires that the robot be powered down and undergo a minor wiring change first, before being updated, reconfigured and powered back up. Changing the code in any P0 (non-reprogrammable) processor requires installing a new part. Thus, the development cycle tends to favour placing code that is changing often (typically sophisticated) higher up, and code that is more stable (typically simple) lower down. Whilst brain evolution is fundamentally different from this style of robot design, there is a similar tendency in nature towards conservation of structure and function towards the lower end of the neuraxis (spinal cord/brainstem), and increase in flexibility and adaptability at the upper end (cortex) [15].

5 Conclusion

We began with the goal of creating an affordable animal-like robot in which we could embed a biomimetic control architecture that we had previously developed on expensive bespoke robotic platforms. Important constraints in the design process, that were later relaxed but still strongly influenced the outcome, were the need to have a platform that could operate in an integrated way at multiple stages of construction, and that no single component should cost more than \$10. We have found that a brain-based design is actually well-suited to these challenges of incremental construction and use of cheap, off-the-shelf parts. During the course of evolution, the mammalian brain has adapted and scaled to many different body types and ecological niches; the MIRO robot shows that future living machines, built of non-biological components such as plastic and silicon, can also make use of layered control architectures inspired by, and abstracted from, those we find in animals.

Acknowledgments. The development of the MIRO robot was funded by EagleMoss Publishing and Consequential Robotics with contributions from Sebastian Conran, Tom Pearce, Victor Chen, Dave Keating, Jim Wyatt, Maggie Calmels, and Emily Collins. Our research on layered architectures was also supported by the FP7 WYSIWYD project (ICT-612139), and the EPSRC BELLA project (EP/I032533/1).

References

1. Prescott, T.J., et al.: Embodied Models and Neurorobotics. In: Arbib, M.A., Bonaiuto, J.J. (eds.) *From Neuron to Cognition via Computational Neuroscience*. MIT Press, Cambridge (in press)
2. Floreano, D., Auke, J., Ijspeert, J., Schaal, S.: Robotics and neuroscience. *Curr. Biol.* **24**(18), R910–R920 (2014)
3. Prescott, T.J., Ibbotson, C.: A robot trace-maker: modeling the fossil evidence of early invertebrate behavior. *Artif Life* **3**, 289–306 (1997)
4. Prescott, T.J., et al.: A robot model of the basal ganglia: behaviour and intrinsic processing. *Neural Netw.* **19**(1), 31–61 (2006)
5. Pearson, M.J., et al.: Biomimetic vibrissal sensing for robots. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* **366**(1581), 3085–3096 (2011)

6. Lepora, N.F., et al.: Optimal decision-making in mammals: insights from a robot study of rodent texture discrimination. *J. R. Soc. Interface* **9**(72), 1517–1528 (2012)
7. Mitchinson, B., et al.: Biomimetic tactile target acquisition, tracking and capture. *Robot. Auton. Syst.* **62**(3), 366–375 (2014)
8. Collins, E.C., Prescott, T.J., Mitchinson, B.: Saying it with light: a pilot study of affective communication using the MIRO robot. In: Wilson, S.P., Verschure, P.F.M.J., Mura, A., Prescott, T.J. (eds.) *Living Machines 2015*. LNCS, vol. 9222, pp. 243–255. Springer, Heidelberg (2015)
9. Collins, E.C., et al.: MIRO: a versatile biomimetic edutainment robot. In: *12th Conference on Advances in Computer Entertainment*, Iskandar, Malaysia (2015)
10. Prescott, T.J.: Forced moves or good tricks in design space? landmarks in the evolution of neural mechanisms for action selection. *Adapt. Behav.* **15**(1), 9–31 (2007)
11. Hinde, R.A.: *Animal Behaviour: a Synthesis of Ethology and Comparative Psychology*. McGraw-Hill, London (1966)
12. McFarland, D., Bossert, T.: *Intelligent Behaviour in Animals and Robots*. MIT Press, Cambridge (1993)
13. Verschure, P.F.M.J., Krose, B., Pfeifer, R.: Distributed adaptive control: the self-organization of structured behavior. *Robot. Auton. Syst.* **9**, 181–196 (1992)
14. Arbib, M.A., Liaw, J.S.: Sensorimotor transformations in the worlds of frogs and robots. *Artif. Intell.* **72**(1–2), 53–79 (1995)
15. Prescott, T.J., Redgrave, P., Gurney, K.N.: Layered control architectures in robots and vertebrates. *Adapt. Behav.* **7**(1), 99–127 (1999)
16. Jackson, J.H.: Evolution and dissolution of the nervous system. In: Taylor, J. (ed.) *Selected Writings of John Hughlings Jackson*. Staples Press, London (1884/1958)
17. Prescott, T.J.: Layered control architectures. In: Pashler, H. (ed.) *Encyclopedia of Mind*, pp. 464–467. Sage, London (2013)
18. Brooks, R.A.: A robust layered control system for a mobile robot. *IEEE J. Robot. Autom.* **RA-2**, 14–23 (1986)
19. Penfield, W.: Centrencephalic integrating system. *Brain* **81**, 231–234 (1958)
20. Humphries, M.D., Gurney, K., Prescott, T.J.: Is there a brainstem substrate for action selection? *Philos. Trans. R. Soc. Lond. B Biol. Sci.* **362**(1485), 1627–1639 (2007)
21. Prescott, T.J., et al.: Whisking with robots: From rat vibrissae to biomimetic technology for active touch. *IEEE Robot. Autom. Mag.* **16**(3), 42–50 (2009)
22. Redgrave, P., Prescott, T., Gurney, K.N.: The basal ganglia: a vertebrate solution to the selection problem? *Neuroscience* **89**, 1009–1023 (1999)
23. Brooks, R.A.: New approaches to robotics. *Science* **253**, 1227–1232 (1991)
24. Gandhi, N.J., Katnani, H.A.: Motor functions of the superior colliculus. *Annu. Rev. Neurosci.* **34**, 205–231 (2011)
25. Fox, C., et al.: Technical integration of hippocampus, basal ganglia and physical models for spatial navigation. *Front. Neuroinformatics*, **3**(6) (2009)
26. Posner, J., Russell, J.A., Peterson, B.S.: The circumplex model of affect: an integrative approach to affective neuroscience, cognitive development, and psychopathology. *Dev. Psychopathol.* **17**(3), 715–734 (2005)
27. Hofe, R., Moore, R.K.: Towards an investigation of speech energetics using ‘AnTon’: an animatronic model of a human tongue and vocal tract. *Connection Sci.* **20**(4), 319–336 (2008)
28. Mitchinson, B., Prescott, T.J.: Whisker movements reveal spatial attention: a unified computational model of active sensing control in the rat. *PLoS Comput. Biol.* **9**(9), e1003236 (2013)

29. Jeffress, L.A.: A place theory of sound localization. *J Comp Physiol Psychol.* **41**(1), 35–39 (1948)
30. Mitchinson, B.: Attention and orienting. In: Lepora, P.T.J.N., Verschure, P.F.M.J. (eds.) *Living Machines: A Handbook of Research in Biomimetic and Biohybrid Systems* OUP: Oxford (in press)
31. Prescott, T.J., et al.: The robot vibrissal system: understanding mammalian sensorimotor coordination through biomimetics. In: Krieger, P., Groh, A. (eds.) *Sensorimotor Integration in the Whisker System*, pp. 213–240. Springer, New York (2015)
32. Higgins, W.T.: A comparison of complementary and Kalman filtering. *IEEE Trans. Aerosp. Electron. Syst.* **AES-11**(3), 321–325 (1975)
33. Burr, D.C., Morrone, M.C., Ross, J.: Selective suppression of the magnocellular visual pathway during saccadic eye movements. *Nature* **371**(6497), 511–513 (1994)
34. Redgrave, P., Prescott, T.J., Gurney, K.: Is the short-latency dopamine response too short to signal reward error?. *Trends Neurosci.* **22**(4), 146–151 (1999)
35. Thorpe, S.J.: The speed of categorization in the human visual system. *Neuron* **62**(2), 168–170 (2009)