



Deposited via The University of Sheffield.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/114167/>

Version: Accepted Version

Proceedings Paper:

Prescott, T.J. and Redgrave, P. (1996) Layered control architectures in natural and artificial systems. In: Workshop on Learning in Robots and Animals. AISB'96, 31 Mar - 02 Apr 1996, Brighton, UK. AISB, Brighton.

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Layered control architectures in natural and artificial systems*

Tony J. Prescott and Peter Redgrave
Department of Psychology,
University of Sheffield,
Psychology Building,
Western Bank,
Sheffield, S10 2TP, UK.

Tel: +44 114 282 6547/6562
Fax: +44 114 276 6515

Email: T.J.Prescott, P.Redgrave @sheffield.ac.uk.

Abstract

We review recent research in robotics and neuroscience with the aim of highlighting some points of agreement and convergence. Specifically, we compare Brooks' [9] subsumption architecture for robot control with a part of the neuroscience literature that can be interpreted as demonstrating hierarchical control systems in animal brains. We focus first on work that follows the tradition of Hughlings Jackson [23] who, in neuroscience and neuropsychology, is particularly associated with the notion of layered competence. From this perspective we further argue that recent work on the defense system of the rat can be interpreted by analogy to Brooks' subsumption architecture. An important focus is the role of multiple learning systems in the brain, and of hierarchical learning mechanisms in the rat defense system.

1 Introduction

The field of adaptive behaviour seeks a convergence of ideas from the different disciplines that study biological and artificial autonomous systems. Demonstrating convergence allows the interchange of concepts and ideas and enriches our understanding of both the natural and the synthetic [2, 32]. This paper reviews research in robotics and neuroscience with the aim of highlighting some points of agreement. Specifically, we aim to show a convergence between Brooks' [9] subsumption architecture for robot control and a part of the neuroscience literature can be interpreted as demonstrating hierarchical control systems in animal brains. We focus first on work that follows the tradition of Hughlings Jackson [23] who, in neuroscience and neuropsychology, is particularly associated with the notion of layered competence. From this perspective we then further argue that recent work on the defense system of the rat can be interpreted by analogy to Brooks' subsumption architecture. A various points through the paper we also consider the role of learning in modular control systems in animals and robots.

2 Brooks' subsumption architecture

Rod Brooks' [9-12] work in engineering robot 'creatures' needs little introduction to researchers in adaptive behaviour. In the mid-eighties Brooks introduced a new methodology—based on an analogy with natural evolution—for building 'self-sustaining' mobile robots that operate in real-time and in un-customised human environments. This work has had enormous influence in robotics, and together with other contemporary work that proposed a move towards more distributed and situated systems (e.g. [7, 33]), has inspired a new research paradigm in Artificial Intelligence (see e.g. [28, 32]). One of the key contribution of Brooks' work is his proposal for layered, distributed control architecture for mobile robots, termed the *subsumption architecture* (SA). The SA was introduced in [9] and has been subject to only limited modification since [10]. The same basic architecture has been employed

*Presented at the AISB96 Workshop on Learning in Robots and Animals, Brighton UK, April 1996.

in the design and control of a large number of mobile robots (see [11]) and has been widely copied. The key aspects of SA, of relevance here, are as follows:

Distributed, layered control

Control in a Brooks' robot is distributed across a hierarchy of 'layers', each composed of multiple modules often mounted on different processors. Layers operate in parallel and asynchronously. Within a layer there is no central control module.

Behavioural decomposition with a hierarchy of 'levels of competence'

The different layers of the control system are viewed as implementing specific *behaviours* [10] , the task of controlling the robot is thus decomposed behaviourally rather than functionally. Within a layer there may be a functional decomposition, for instance, into sensor and actuator components, however, different layers will use different decompositions based on specialised sub-sets of the available sensori-motor apparatus. Behaviours at ascending levels in the hierarchy implement increasingly sophisticated solutions to increasingly abstract and complex tasks—with each added layer the robot therefore attains a higher level of overall competence. Damage or failure at a higher level reduces the robot to functioning at the level of the next highest layer.

Incremental construction

A key constraint on the design process is that, as each additional level of competence is incorporated, the total system should be “a strict augmentation of the previous one” [10] .

Conflict resolution and communication between levels by subsumption mechanisms

Higher levels of the control system ‘subsume’ the roles of lower ones, principally by suppressing their outputs and (optionally) substituting their own. Each lower level system continues to function as higher levels are added, ‘unaware’ that the levels above may be interfering with its data paths.

Little sensor fusion and no central models

Brooks' asserts that there is rarely the need to fuse data from multiple sensors (so as to determine the most accurate estimate possible of world state). Rather, he argues that conflicts between different sources of information can be resolved by arbitration at the motor command level. One consequence of this view is the principle that the robot should have no need for central world models into which all available sense data is compiled. Rather than exploiting shared representations, behaviours at different levels are separated by “abstraction barriers” [11] , unable to share state or influence each others internal workings by anything more than simple subsumption mechanisms.

Learning

Many of Brooks' early robots were not designed to learn autonomously. However, Brooks' [12] has noted the value of learning for such tasks as sensor/actuator calibration, modification and tuning of behaviours, handling interactions between behaviours, acquiring task-relevant world knowledge (e.g. building spatial maps), and the acquisition of new behavioural modules. Some of these learning capacities have been investigated on Brooks' robots (see [12]) but it is recognised that the investigation of learning in subsumption architectures is still at an early stage. One of the aims of this paper is to identify some of the characteristics of learning in natural systems that may suggest ways of incorporating learning into layered robot control systems.

Despite the huge impact of Brooks' research in robotics and Artificial Intelligence, its influence on the study of natural autonomous systems has been more limited. Work in this wider area, that acknowledges the influence of Brooks' approach, includes studies of human perception and motor control [4, 25] , human development, and research in computational neuroethology [3, 15, 18] .

That the main impact of Brooks' research has been in AI is perhaps unsurprising. Brooks' explicitly states that, although the SA draws on an evolutionary metaphor, it is not a biological model, he also warns of the dangers of treating biological intelligence as a lodestar for AI [12] . However, Brooks also insists that his interest is in general intelligence [12] , that he sees animal intelligence as an important “existence proof of the possibility of intelligent entities” [11] , and that we should expect to gain insights for robot design by looking at the nervous systems of animals [12] . The search for further links between Brooks' robot architectures and our understanding of animal intelligence therefore fits naturally with the situated robotics approach.

3 The Jacksonian perspective in neuroscience

In 1884, in a famous lecture on the “evolution and dissolution of the nervous system” the British neurologist John Hughlings Jackson [23] proposed a hierarchical view of the nervous system, in which the brain is seen as implementing multiple layers of sensori-motor competence. Jackson's hierarchical view of the brain was based not

on the usual morphological divisions (spinal cord, medulla oblongata, and so on) but on functional grounds, or has he put it “as to the degree of indirectness with which each [division] represents the body, or part of it”. He divided the nervous system into lower, middle, and higher centres, and proposed that this hierarchy represented a progression from the ‘most organised’ (most fixed) to the ‘least organised’ (most modifiable), from the ‘most automatic’ to the ‘least automatic’, and from the most ‘perfectly reflex’ to the least ‘perfectly reflex’. This progression sees an increase in competence in a manner that we might now understand as a behavioural decomposition—higher centres are concerned with same sort of sensori-motor coordinations as those below, though in a more indirect and complex fashion:

“That the middle motor centres represent over again what all the lowest motor centres have represented, will be disputed by few. I go further, and say that the highest motor centres (frontal lobes) represent over again, in more complex combinations, what the middle motor centres represent. In recapitulation, there is increasing complexity, or greater intricacy of representation, so that ultimately the highest motor centres represent, or, in other words, coordinate, movements of all parts of the body in the most special and complex combinations.” ([23] p. 53)

Jackson viewed the evolution of the nervous system as an incremental process in which lower levels are retained intact but are suppressed by higher systems. Within his hierarchy of centres Jackson further considered there to be functionally distinct ‘layers’. He argued for a dissociation of higher layers from those below such that a breakdown at a higher layer—a “dissolution” in Jackson’s terminology—caused a reversion to the next highest layer of control.

There are further important parallels between Jackson’s writing and contemporary approaches in robotics, for instance, he was an early advocate of the notion of distributed representation. Overall, his writings show a conviction that ‘higher’ thought is grounded in perception and action—a perspective which, while radical for his era¹, is clearly in sympathy with recent proponents of ‘situated action’ (e.g. [11, 14]):

“A man physically regarded is a sensori-motor mechanism. I particularly wish to insist that the highest centres—physical basis of mind or consciousness—have this kind of constitution, that they represent innumerable different impressions and movements of all parts of the body [...] It may be rejoined that the highest centres are “for mind”. I assert that they are “for body”, too. If the doctrine of evolution be true, all nervous centres must be of sensori-motor constitution.” ([23] p. 63)

A number of researchers have explicitly embraced a ‘Jacksonian perspective’ on the hierarchical nature of the nervous system, with an experimental program founded on these principles [8, 13, 20, 38, 44, 45] ². Before proceeding to a detailed account of a specific system (the defense system of the rat) a range of research in support of this position will be briefly reviewed.

The Jacksonian view predicts a dissociation between higher and lower level components of a system where the lower level competence is left intact by damage at a higher level. Examples of dissociations between cortical and sub-cortical systems include that between the ‘two visual systems’ in the superior colliculus and the visual cortex [39] , between ‘proximal cue’ (mid-brain) and ‘distal cue’ (hippocampus) spatial navigation systems [34] , and between taste systems in the neocortex and mid-brain [8] . Teitelbaum et al. [45] provide a review of evidence on dissociation between higher cortical functions and intermediate level ‘sub-modules’ that implement coordinated sets of ‘allied’ reflexes. Like ourselves, they argue that an understanding of these systems could be useful in robot design.

A further implication of a Jacksonian view is that during ontogeny (development) the brain matures through the addition of higher centres [45] . This developmental sequence has been observed post-natally, in the maturation of rats and rhesus monkeys (see below). A related phenomenon is the ‘warm-up’ sequence observed in the recovery of locomotory behaviours of brain-damaged adult rats [19] and of infant rats placed in an open field [17] where different types of movement are recruited in a characteristic order. This sequential addition of competences can be interpreted as the successive involvement of higher level brain centres [17] .

¹The prevailing view was that the brain-stem and spinal cord controlled motor functions whilst the cerebral cortex was reserved for higher cognitive functions [22, 23] .

²The influence of Hughlings Jackson is also acknowledged by Arbib who has given Jacksonian analysis of a model of optic flow processing [1] .

Multiple learning systems

Razran [36] (see also [21]) reviewed a wide range of literature in comparative psychology and argued for four “superlevels” of learning appearing at different evolutionary stages. At the lowest level of his ‘learning hierarchy’ were the ‘reactive’ learning mechanisms capable of habituation and sensitization that exist in all animals. A second level of ‘associative’ learning systems, present in vertebrates and some invertebrates, allowed most forms of classical (stimulus-reward) and instrumental (stimulus-response) conditioning, while a third level of ‘integrative’ learning, present only in higher vertebrates, allowed the ability to acquire responses to compound stimuli and to learn the relationships between stimuli. A final level of ‘symbolic’ learning was restricted to humans and was primarily concerned with language acquisition. Razran proposed that in phylogenetically later animals multiple levels coexist with the higher levels dominating lower ones.

The distinction between ‘associative’ and ‘integrative’ learning systems, appearing under a number of different labels and guises, has gained considerable attention in research in psychology and neuroscience. Patterns of learning impairment in human amnesiacs and lesion studies with animals (see [21, 41] for review) indicate that ‘integrative’ learning abilities rely on specific cortical structures in the brain, in particular, the hippocampus. In contrast the simpler associative forms of learning are supported by neural systems that evolved much earlier. Dissociations between multiple learning systems have recently been demonstrated by McDonald and White [30] who investigated parallel learning systems in the hippocampus, amygdala, and basal ganglia. Lesions to any two of these three learning systems in the rat left intact the learning capability of the third.

The development sequence in the acquisition of higher brain centres, noted above, is also clearly demonstrated in the maturation of learning systems. For instance, in rats, reflexive responses to stimuli (visual, auditory, or gustatory) have been shown to mature several days before the same stimuli are able to mediate learned behavioural reactions [38]. Unlearned reflexive responses can be generated by the brain-stem components whereas associative learning requires higher level brain centres, this evidence is therefore consistent with a Jacksonian progression. Evidence that this developmental sequence extends to cortical systems comes from spatial learning, where it has been demonstrated that the rat’s ability to use proximal cues (probably involving the striatal learning system) is prior to spatial learning with respect to multiple distal cues (involving the hippocampal learning system) [38].

As is demonstrated further below, learning in higher level systems often involves the coordination, potentiation or suppression of outputs at the next lowest layer, thus learning in higher level systems is concerned with modulating coordinated responses of increasing sophistication and complexity.

5 Subsumption in the rat defense system

Brooks’ SA provides a candidate architecture for distributed and hierarchical robot control inspired by evolutionary principles; Jacksonian neuroscience shows empirical evidence for layers of competence in animal brains. It is our view that, if the parallels between the two disciplines can be clearly drawn, neuroscience should be able to suggest improvements for layered robot control architectures. In addition, SA could help to further our understanding of the architecture of biological nervous systems. In this section of the paper we are therefore concerned to look for a such a parallel, specifically, we review empirical findings concerning the defense system of the rat and argue that this shows remarkable architectural similarities to Brooks’ SA. Although the lowest levels of the defense system are reflexive, it will be shown that learning plays an increasingly important role at higher levels.

A substantial body of research into the neural defense mechanisms of the rat can be interpreted as showing a hierarchically organised set of functional modules—higher level components provide increasingly sophisticated solutions to the problems of reducing and avoiding harm. Each level in this architecture has its own sensory inputs and connections to motor/autonomic outputs, and when activated by appropriate stimuli, each level is capable of delivering a set of appropriate adaptive reactions. The output of all levels appears to be relayed to a restricted set of motor, autonomic, and hormonal effector units, which delivers a relatively stereotyped “defense reaction” [49]. The basic architecture is illustrated in the following figure:




Figure: Hierarchical organisation of neural mechanisms of defense: higher level components provide increasingly sophisticated solutions to problems of reducing and avoiding harm.

The following reviews the structures at each level in the hierarchy beginning with the lowest.

Level 0—low level avoidance of actual harm

The bottom level of the hierarchy is characterised by simple reactions to intense, harmful or potentially harmful, contact stimuli. Reactions of this type are seen in both invertebrates and vertebrates. In mammals such mechanisms are typically simple withdrawal reflexes motivated by noxious stimuli. Touching something hot initiates a di-synaptic withdrawal reflex mediated at the level of the spinal cord [40]. Similarly, a puff of air directed at the eye elicits an eye-blink [50]. Both reactions are examples of reflexes which involve relatively few muscle groups and protect the organism from actual harm. In evolutionary terms they are mediated by the oldest parts of the nervous system, namely, the brainstem and spinal cord.

Level 1—intermediate level systems organise co-ordinated responses to harm

Intermediate level systems, represented principally in the midbrain, add sophistication to both sensory and motor aspects of defensive behaviour.

First, systems have developed to detect stimuli which, themselves are not harmful, but give warning of harm; vision, of course, is the best example. Experiments, reviewed by the Blanchards [6] show that midbrain visual systems are critical for the ‘distance-defense’ sub-system observed in rats. It is likely that a set of filters in the rat (midbrain) superior colliculus monitors the distance of a potential threat. Each filter would be tuned to threats at a particular distance range and separately wired to motor systems delivering appropriate responses: “freeze” is the characteristic response at long-range, “flight” at intermediate distances, and “attack” at short range.

Second, compared with simple withdrawal, escape and avoidance often involves the vigorous use of many muscles. This requires the development of: (a) systems that can co-ordinate the complicated movements involved in “flight” and “fight”; and (b) autonomic and endocrine systems that can sustain the increased energy requirements of

vigorous activity. The midbrain periaqueductal grey (PAG) is generally viewed as the principle intermediate level structure responsible for co-ordinating the wide range of somatomotor, autonomic, and endocrine reactions involved in defensive behaviour [5]. For example, it appears that the outputs of nociceptors (pain receptors) located deep within the body selectively activate regions of the PAG associated with behavioural suppression and quiescence, while nociceptors from the body surface activate parts of the PAG producing active avoidance and escape [24].

It is particularly interesting to note that these intermediate level systems evolved in creatures like fish, frogs and lizards but anatomical and physiological evidence suggests that many have been retained in mammals including humans. In addition to the mid-brain components noted above, the (fore-brain) basal ganglia are also implicated in the control of intermediate level defense reactions. It is known that the basal ganglia are among the oldest regions of the forebrain and are present, with the same basic connections and cell types, in all jawed vertebrates including fish [31].

Level 2—high level systems learn that arbitrary stimuli can be predictive of harm.

At higher levels of the defense hierarchy arbitrary events can be associated with harm and used as triggers for the full suite of defense reactions. Thus a neutral light (conditioned stimulus) can be associated with a harmful footshock (unconditioned stimulus). After several pairings, defensive reactions, previously elicited by footshock, are evoked when light is presented alone. This is termed a *conditioned emotional response* [26]. The limbic system, and the amygdala in particular, appear to play the key role in the acquisition of these associations. To perform this task the amygdala has the following critical features:

- It receives sensory information, from both low-level and high-level sensory structures, encoding (conditioned) stimuli that can be simple or complex and of any modality (visual, auditory, tactile, or olfactory). On a further channel (probably via the Parabrachial nucleus) it receives primitive signals associated with pain and distress (unconditioned stimuli).
- It has connections to output mechanisms at lower levels of the hierarchy (hypothalamus, midbrain, pons and medulla) through which the various behavioural, autonomic and endocrine reactions to defend the animal are expressed.
- Sensory stimuli coding situational variables (e.g. presence of a light) that arrive in the amygdala at the same time as signals signifying pain (footshock) rapidly acquire the ability to activate defense-related output channels (by simple classical conditioning).

This analysis suggests the amygdala operates by putting an ‘emotional stamp’ on any sensory input which is present at the same time or occurs just before something nasty. In other words, if that stimulus occurs again in the future, it is given access to the defence circuitry.

The ability to acquire such responses in the case of relatively primitive, uni-modal stimuli, appears to be dissociated from the highest cortical centres. For example, LeDoux (1992) has found that relatively primitive subcortical auditory structures provide the input for the conditioned emotional response to a simple tone. Similarly, Davis (1992) established that subcortical (thalamic) visual input was required for a light induced emotionally-modulated defense reflex. In both cases lesions of the higher level visual and auditory cortex failed to abolish the respective responses. More complex sensory events require more sophisticated sensory analyses performed by sensory neocortex and sensory association areas, these structures consequently have direct output projections to the amygdala. The hippocampal learning system also projects to the amygdala (via the subiculum), providing access to additional contextual knowledge that may allow the animal to distinguish between threatening and non-threatening situations.

Level 4—cortical systems may unlearn associations between arbitrary stimuli and harm

Such are the vagaries of life that an acquired association between a stimulus and harm can later cease to be appropriate. LeDoux [27] has recently suggested that a yet higher level of competence is required to un-learn such associations between arbitrary stimuli and harm. This proposal arose from experiments on the conditioned emotional response that began by pairing a tone with an electric shock. After relatively few trials defense reactions to the tone were well established, at which point the shock was turned off so that the tone no longer predicted the aversive stimulus. Very gradually, over several weeks, the conditioned emotional reactions to the tone subsided and eventually disappeared. Once extinction had occurred restricted lesions were placed in the frontal cortex. On the first trial following these lesions the conditioned emotional response to the tone (alone) re-emerged at full strength.

This remarkable result suggests that acquired associations between conditioned and unconditioned stimuli at the level of the amygdala remained intact throughout the extinction period. In other words, the extinction of the emotional response did not reflect a gradual weakening of the link between the tone input and defense output within the amygdala. Instead, it appears that the frontal cortex gradually learned that the tone was no longer dangerous and established an appropriate inhibitory connection to the amygdala. This connection suppressed the output (in response to the tone) from the amygdala to efferent defense areas. When the frontal cortex was lesioned, the modulating inhibitory input to the amygdala was lost releasing the emotional link between the tone and defense reactions. These findings may have important implications for devising rational strategies for dealing with clinical problems such as phobias and post-traumatic stress disorders. However, for the present purposes it suffices to say that the frontal cortex may represent an even higher level module in the hierarchy of systems determining defense behaviour.

6 Is the rat defense system a subsumption architecture?

The above paints a fascinating picture of the defense system in the rat brain. The similarities to a subsumption architecture are clear and can be summarised as follows:

Distributed control

The defense system is composed of multiple sub-systems operating in parallel. Each level makes an important contribution to the functionality of the whole system.

Behavioural decomposition with a hierarchy of competences

At ascending levels the system implements defense reactions of increasing complexity and sophistication. Lower level defense reactions do not depend on higher level structures and will operate when triggered unless actively inhibited by higher levels. There is a dissociation between levels such that damage to higher level structures leaves lower level behaviours unimpaired.

Conflict resolution and communication between levels by subsumption mechanisms

Anatomical studies [37, 43, 46] indicate that outputs from higher level modules such as the amygdala and cerebral cortex converge on lower somatomotor, autonomic and endocrine output units. Thus higher level systems act by modulating the outputs generated at lower layers. In the defense hierarchy we see several examples of subsumption-like mechanisms that implement this modulation.

First, higher-level systems can *coordinate* multiple lower level mechanisms. This is observed in the coordinated control, by mid-brain defense sub-systems, of sets of lower-level reflex mechanisms.

Second, higher-level systems can provide *substitute* inputs to lower-level output mechanisms in the absence of primary noxious sensory information impinging on the lower layer. This is evident in the conditioned emotional response where low-level defense reactions can be triggered by arbitrary and potentially complex configurations of non-noxious sensory stimuli. A further example, is the task of orienting to the position of a remembered, and previously dangerous, target. In such situations, high level systems clearly must impose control without low-level sensory input. Thus, the low-level gaze positioning system (the superior colliculus) operates regardless of whether its input denotes a real or a remembered object [48].

Third, higher level systems can *inhibit* those below. This is essential if you are to avoid assaulting the dentist while he/she is drilling your teeth! A more formal example, described above, is the continuous inhibitory signal³ from the frontal cortex to the amygdala that allows unlearning of the conditioned emotional response.

Finally, lower level reactions can be modulated by simple mechanisms such as *gain control*. An example of this is the potentiated startle response [16] in which the amplitude of a low level reflex behaviour is dramatically enhanced when the animal is in a 'fearful' state initiated by a higher level module. This *potentiation* of the startle response is caused by an output from the amygdala that taps into the lower level reflex arc.

If high-level systems were able to get inside and really 'fiddle around' with the internal workings of lower ones (i.e. by using more than simple subsumption-like mechanisms to modulate the lower-level activity) then this would weaken the position of SA as an analogy for neural architecture. This appears not to be the case with the rat defense system, although it may be true elsewhere in the rat brain.

³It is interesting to note that this exact form of subsumption control (continuous, active inhibition) was a refinement that was added to the original SA (Brooks 89) to replace inhibitory mechanisms based on time-outs.

Little sensor fusion and no central models

At lower and intermediate levels, sensory inputs are largely unfused and subject to basic processing, however, this is generally not the case in higher level brain systems [42]. For instance, recent evidence suggests that solutions to the spatial correspondence between visual and auditory maps are computed at cortical levels and used to enhance the representation, at lower levels, of near threshold, weak visual and auditory stimuli. Thus, if two sources correspond in place, signal-enhancement can be observed in intermediate level somatomotor output mechanisms [47]. If the branch twitches and the leaves rustle *at the same location*, then a defense reaction will be much more likely.

The dissociation between layers of the defense system demonstrates that lower-level mechanisms are not dependent on any unified representations that might be established at higher levels. In general, each level has its own sensor channels and sensor processing mechanisms, thus Brooks' notion of "abstraction barriers" is at least partially respected.

Learning

The principle learning system in the defense system is the amygdala. Its key role is to allow arbitrary stimuli access to the largely 'hard-wired' coordinated sets of defense reactions. The primary channels for inputs to this system are from mid-brain sensory areas, however, additional sensory-processing and learning systems in the cortex and hippocampus, sit above the amygdala and provide it with more sophisticated sensory and contextual inputs. One notable characteristic of this system is that extinction appears not to involve the eradication of a learned response but active inhibition of that response by a higher level system. The knowledge acquired through learning is perhaps too valuable for it to be easily erased.

In summary, the above correspondences would appear sufficient to justify our view that the architecture of the defense system is closely analogous to Brooks' SA.

7 A hypothetical robotic model

We conclude by describing a SA for a hypothetical autonomous robot designed to mirror the structure of the defense system of the rat.

At level 0, we have a brain-stem/spinal-cord sub-system that produces a rapid reflex reaction to painful, contact stimuli. For the robot an analogous sensori-motor reflex might be a collision detector circuit linking a bump sensor to an emergency stop response.

At level 1 we have a mid-brain sub-system that responds to simple stimuli emanating from nearby objects or sources, that are not themselves harmful, but are commonly associated with danger. These sensor systems are linked to motor systems that implement a set of coordinated reflexive responses. For the robot an analogous system might be one that used sonar to detect nearby obstacles, formed a thresholded vector sum of repulsive forces, then moved in the resultant direction.

The example robot sub-systems chosen at these first two levels are low-level local navigation competences implemented in many of Brooks' robots (see e.g. [11]).

At level 2 we have a forebrain sub-system that allows arbitrary events to be associated with stimuli that are harmful or predictive of harm. These acquired associations provide access to lower level efferent mechanisms producing defense reactions. For the robot an analogous system might use reinforcement learning to learn to detect specific objects or areas to be avoided. For instance, the robot might learn to detect and avoid situations, for instance awkward obstacle configurations, where movement is hazardous—a learning system that acquires a competence of this type is described in [35].

At level 3 cortical structures allow relatively complex configurations of stimuli to be recognised. The outputs of this sub-system mediate lower level responses. For the robot an analogous system might learn to recognise a specific obstacle configuration as a goal which should be approached. In appropriate circumstances the system would then act by inhibiting the avoidance responses generated at lower levels (for a system capable of a similar level of competence see [29]).

This specific robot design has not been built but clearly could be constructed using existing techniques in situated robotics. Were it to be built, its control architecture could be viewed as an abstract, but embodied, model of the defense system of the rat.

Acknowledgement

This research was partially supported by Wellcome trust grant 038011/Z/93 to P. Redgrave.

8 References

1. Arbib, M.A., *The Metaphorical Brain*. 1989, New York: Wiley and Sons.
2. Arbib, M.A., *Cooperative computation in brains and computers*, in *Natural and Artificial Parallel Computation*, M.A. Arbib and J.A. Robertson, Editors. 1990, MIT Press: Cambridge, MA.
3. Arbib, M.A. and A. Cobas. *Schemas for prey-catching in frog and toad*. in *From Animals to Animats: Proceedings of the First International Conference on the Simulation of Adaptive Behaviour*. 1990. Paris.
4. Ballard, D.H., M.M. Hayhoe, and S.D. Whitehead. *Hand-eye coordination during sequential tasks*. in *Proceedings of the Royal Society*. 1992.
5. Bandler, R. and A. Depaulis, *Midbrain periaqueductal gray control of defensive behaviour in the cat and the rat*, in *The midbrain periaqueductal gray matter*, A. Depaulis and R. Bandler, Editors. 1991, Plenum Press: New York. p. 175-198.
6. Blanchard, R.J. and D.C. Blanchard, *An ethoexperimental analysis of defense, fear and anxiety*. In: *Anxiety*, (eds) McNaughton N and Andrews G, University of Otago Press., 1990: p. 124-133.
7. Braitenberg, V., *Vehicles: experiments in synthetic psychology*. 1986, Cambridge, MA: MIT Press.
8. Braun, J.J., P.S. Lasiter, and S.W. Kiefer, *The gustatory neocortex of the rat*. *Physiological Psychology*, 1982. **10**(1): p. 13-45.
9. Brooks, R.A., *A robust layered control system for a mobile robot*. *IEEE Journal on Robotics and Automation*, 1986. **RA-2**: p. 14-23.
10. Brooks, R.A., *A robot that walks: emergent behaviour from a carefully evolved network*. *Neural Computation*, 1989. **1**(2): p. 253-262.
11. Brooks, R.A., *Elephants don't play chess*, in *Designing Autonomous Agents: Theory and Practice from Biology to Engineering and Back*, P. Maes, Editor. 1990, MIT Press: Cambridge, MA.
12. Brooks, R.A., *Intelligence without reason*, in *The Artificial Life Route to Artificial Intelligence*, L. Steels and R. Brooks, Editors. 1995, Lawrence Erlbaum: Hillsdale, New Jersey.
13. Campbell, B.A., A. Sananes, and J.R. Gaddy, *Animal models of Jacksonian dissolution of memory in the aged*, in *Senile dementia of Alzheimer type*, J. Traber and W.H. Gispen, Editors. 1985, Springer-Verlag: Berlin. p. 283-291.
14. Chapman, D., *Vision, Instruction, and Action*. 1991, Cambridge, MA: MIT Press.
15. Cliff, D.T., *Computational neuroethology: a provisional manifesto*, . 1990, University of Sussex.
16. Davis, M., *The role of the amygdala in fear and anxiety*. *Annual Review of Neuroscience*, 1992. **1992**(15): p. 353-75.
17. Eilam, D. and I. Golani, *The ontogeny of exploratory behaviour in the house rat (Rattus rattus): the mobility gradient*. *Developmental Psychobiology*, 1988. **21**(7): p. 679-710.
18. Franceschini, N., *From insect vision to robot vision*. *Philosophical Transactions of the Royal Society of London. B*, 1992. **337**: p. 283-294.
19. Golani, I., D. Wolgin, and P. Teitelbaum, *A proposed natural geometry for recovery from akinesia in the lateral hypothalamic rat*. *Brain Research*, 1979. **164**(237-267).
20. Grill, H.J. and K.C. Berridge, *Taste reactivity as a measure of the neural control of palatability*, in *Progress in Psychobiology and Physiological Psychology*, J.N. Sprague and A.N. Epstein, Editors. 1985, Academic Press: New York.
21. Hintzman, D.L. *Twenty-five years of learning and memory: was the cognitive revolution a mistake*. in *Fourteenth International Symposium on Attention and Performance*. 1992. Michigan: Bradford Books.
22. Humphrey, D.R., *Representation of movements and muscles within the primate precentral motor cortex: historical and current perspectives*. *Federation Proceedings*, 1986. **45**: p. 2687-2699.
23. Jackson, J.H., *Evolution and dissolution of the nervous system*, in *Selected Writings of John Hughlings Jackson*, J. Taylor, Editor. 1884/1958, Staples Press: London.
24. Keay, K.A. and R. Bandler, *Deep and superficial noxious stimulation increases Fos-Like immunoreactivity in different regions of the midbrain periaqueductal grey of the rat*. *Neurosci. Lett.*, 1993. **154**: p. 23-26.
25. King, S.M., et al., *Use of a distracting task to obtain defensive head movements to looming visual stimuli by human adults in a laboratory setting*. *Perception*, 1992. **21**: p. 245-259.

26. Ledoux, J.E., *Emotion and the amygdala*, in *The Amygdala: Neurological Aspects of Emotion, Memory, and Mental Dysfunction*, J.P. Aggleton, Editor. 1992, Wiley-Liss: New York. p. 339-351.
27. LeDoux, J.E., *Emotion: clues from the brain*. Annual Review of Psychology, 1995. **46**: p. 206-35.
28. Maes, P. *Behaviour-Based Artificial Intelligence*. in *From Animals to Animats: Proceedings of the Second International Conference on Simulation of Adaptive Behaviour*. 1992. Honolulu, USA.
29. Mataric, M.J. *Navigating with a rat brain: A neurologically-inspired model for robot spatial representation*. in *From Animals to Animats: Proceedings of the First International Conference Simulation of Adaptive Behaviour*. 1990. Paris: MIT Press.
30. McDonald, R.J. and N.M. White, *A triple dissociation of memory systems: hippocampus, amygdala, and dorsal striatum*. Behavioural neuroscience, 1993. **107**(1): p. 3-22.
31. Medina, L. and A. Reiner, *Neurotransmitter organization and connectivity of the Basal Ganglia in vertebrates: implications for the evolution of the basal ganglia*. 1995.
32. Meyer, J. and A. Guillot. *Simulation of adaptive behaviour in animats: review and prospect*. in *From Animals to Animats: Proceedings of the First International Conference Simulation of Adaptive Behaviour*. 1990. Paris.
33. Minsky, M.L., *The Society of Mind*. 1986, New York: Simon & Schuster.
34. Morris, R.G.M., *Spatial localisation does not require the presence of local cues*. Learning and Motivation, 1981. **12**: p. 239-260.
35. Prescott, A.J. and J.E.W.M. Mayhew, *Adaptive local navigation*, in *Active Vision*, A. Blake and A. Yuille, Editors. 1992, MIT Press.
36. Razran, G., *Mind in Evolution*. 1971, Boston, MA: Houghton Mifflin.
37. Rizvi, T.A., et al., *Connections between the central nucleus of the amygdala and the midbrain periaqueductal gray - topography and reciprocity*. J Comp Neurol, 1991. **303**: p. 212-131.
38. Rudy, J.W., S. Stadler-Morris, and P. Albert, *Ontogeny of spatial navigation behaviours in the rat: dissociation of "proximal"- and "distal"- cue-based behaviours*. Behavioural Neuroscience, 1987. **101**(1): p. 62-73.
39. Schneider, G.E., *Two visual systems*. Science, 1969. **163**: p. 895-902.
40. Schouenborg, J. and H.R. Weng, *Sensorimotor transformation in a spinal motor system*. Exp Brain Res, 1994. **100**: p. 170-174.
41. Squire, L.R., *Memory and the Hippocampus: A synthesis from findings with rats, monkeys and humans*. Psychological Review, 1992. **99**(2): p. 195-231.
42. Stein, B.E. and M.A. Meredith, *The merging of the senses*. 1993, Cambridge, MA.: The MIT Press.
43. Takeuchi, Y., et al., *Amygdaloid pathway to the trigeminal motor nucleus via the pontine reticular formation in the rat*. Brain Res. Bull., 1988. **21**: p. 829-833.
44. Teitelbaum, P., V.C. Pellis, and S.M. Pellis. *Can allied reflexes promote the integration of a robot's behaviour*. in *From Animals to Animats: Proceedings of the First International Conference on the Simulation of Adaptive Behaviour*. 1990. Paris: MIT Press.
45. Teitelbaum, P., T. Schallert, and I.Q. Whishaw, *Sources of spontaneity in motivated behaviour*, in *Handbook of Behavioural Neurobiology*, P. Teitelbaum and E. Satinoff, Editors. 1983, Plenum Press: New York. p. 23-66.
46. Veening, J.G., I.W. Swanson, and P.E. Sawchenko, *The organization of projections from the central nucleus of the amygdala to brainstem sites involved in central autonomic regulation: a combined retrograde transport-immunohistochemical study*. Brain Res, 1984. **303**: p. 337-357.
47. Wallace, M.T. and B.E. Stein, *Cross-Modal synthesis in the midbrain depends on input from cortex*. J Neurophysiol, 1994. **71**: p. 429-432.
48. Wurtz, R.H. and O. Hikosaka, *Role of the basal ganglia in the initiation of saccadic eye movements.*, in *In: Oculomotor and skeletal-motor systems.*, Prog Brain Res vol 64 ed Freund. 1986. p. 175-189.
49. Yardley, C.P. and S.M. Hilton, *The hypothalamic and brainstem areas from which the cardiovascular and behavioural components of the defence reaction are elicited in the rat*. J Autonomic Nerv Syst, 1986. **15**: p. 227-244.
50. Yeo, C.H., *Cerebellum and classical conditioning.*, in *Cerebellum and Neuronal Plasticity*, M. Glickstein, C. Yeo, and J. Stein, Editors. 1987, Plenum Publishing Corporation: London. p. 321-338.