

Action and Adaptation: Lessons from Neurobiology and Challenges for Robot Cognitive Architectures

Rodrigo Ventura

Institute for Systems and Robotics
Instituto Superior Técnico, Lisbon, PORTUGAL
yoda@isr.ist.utl.pt

Abstract

In this paper we review recent findings from neurobiology concerning action and adaptation, from the level of motor control up to the level of decision making. At the level of motor control, proposals concerning multiple forward and inverse internal models in the cerebellum are described. Particular consideration is devoted to the adaptive mechanisms of these models. Then, the decision making processes taking place in the basal ganglia and the orbitofrontal cortex are discussed. Under the light of recent findings, these two structures correspond to two levels of decision making with distinct characteristics. The amygdala plays an important role in these decision making processes, for it provides affective evaluation to the action options being considered, as well as an adaptive role in associating novel stimuli with affective states. Then, we discuss possible contributions that these findings may provide for the development of cognitive architectures for robots. In particular, we focus on the following issues: (1) the integration of information coming from different levels, with distinct natures and time frames (the binding problem), and (2) the nature of the described internal models, in the sense that they model the relationship of an embedded and embodied being with its environment. We finally discuss future directions for the research on cognitive architectures, taking into account the discussed neurobiological adaptive mechanisms.

Keywords: cognitive architectures, robotics, motor control, decision making, emotions.

Introduction

Every time engineering has sought to solve a problem that biology has already solved, even if partially, it understandably looks to nature for inspiration. Since the beginning of the field of Artificial Intelligence (AI) that biological intelligence has served as a natural source of inspiration. This inspiration has taken many forms — rational thought in the form of mathematical logic, economics, psychology, and neuroscience — usually consisting in functional explanations of how the mind works. However, it is important to note that, as airplanes share little similarity with birds (although the latter have inspired Wright's brothers early experiments), the optimal design for an intelligent machine is

most probably radically different from the architecture of the human mind. Taking into account the radical differences between biological (neurons, synapses, neurotransmitters, and a massive parallel architecture) and computational systems (gates, digital circuits, and serial computation), it is natural to assume that the best methodologies to attain a given set of design goals with an artificial system ought to be radically different from the ones emerging from biological evolution. For instance, massively search algorithms have proved successful by beating the chess grandmaster Gary Kasparov in 1997 (Campbell, Jr., and Hsu 2002). The machine (Deep Blue) is able to search over 100 million positions per second, while a human player supposedly examines far less positions. The comparison between the serial processing performed by computers, and the parallel nature of the brain (von Neumann 2000) is often mentioned to justify such diverse approaches. Taking human intelligence as inspiration is thus useful to the extent that no better models of intelligence seem to currently exist, and so, taking inspiration in biology can be considered a good starting point for novel advances.

This paper focuses on one particular form of inspiration: the neurosciences, encompassing both cognitive science and neurobiology. Taking a neurobiological point of view, we consider system-level models of the brain. From this point of view, the brain comprises several, functionally more or less distinct, subsystems. In this paper, three of these subsystems are examined: the *cerebellum*, which is responsible for motor control, the *prefrontal cortex*, namely in what concerns decision making, and the *amygdala*, in the role of providing affective evaluation of stimuli. Taking recent findings concerning the functionality of each one of these subsystems as a starting point, we discuss a set of ideas suggested by these results for the design of cognitive architectures. It should be noted that this research is particularly focused on cognitive architectures for robots, as they pose particular design constraints. Namely, the constraints imposed by the fact that any robot is situated in a physical environment with which it interacts (possibly including humans and other robots), it receives raw sensory input from sensors, and it is embodied in a physical structure.

The paper is organized as follows: the next section reviews recent results from the neurosciences relevant to this work, then the following section discusses possible contri-

butions that these results may provide for the development of cognitive architectures, and finally some conclusions are presented in the last section.

Lessons from neurobiology

Motor control

The cerebellum is a region of the brain located in the inferior posterior portion of the head, being mainly responsible for motor control. The computation model of the cerebellum proposed by Wolpert *et al.* (Wolpert, Miallb, and Kawato 1998) comprises multiple paired forward-inverse models. Each one of these models contains a forward internal model (FIM) and an inverse internal model (IIM) (Kawato 1999). The forward model predicts sensory consequences of motor commands given as input. For instance, given an arm movement command, it predicts the trajectory followed by the arm, as perceived by the subject. The inverse model outputs a motor command, given a desired trajectory. The intention to grasp an object, for instance, triggers an appropriate motor command, produced by the corresponding inverse model. These models are part of a complex control architecture, illustrated in figure 1. Three control loops can be identified in this architecture: a (1) *slow feedback loop* passing through the environment and the sensory system, which introduces an unavoidable latency in the loop, a (2) *fast feedback loop* involving the forward model, shunting the sensory system, and thus mitigating the effects of its latency, and an (3) *open loop* control, which allows for the fastest response, since it is independent of any feedback signal. The slow feedback is essential to bootstrap the internal models, that eventually allow the open loop control to work after sufficient adaptation. First, the forward model learns from sensory data to predict the sensory consequences of motor commands, with the advantage of being able to output these predictions ahead of time, and thus compensating for the sensory feedback latency. Second, the controller output can be used to adapt the inverse model to output the appropriate motor commands, without any feedback. The robustness of this architecture is provided by the capability of the sensory feedback loop to take over whenever something goes wrong.

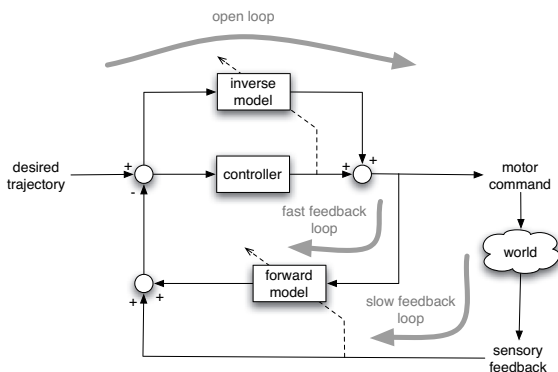


Figure 1: Sketch of the motor control mechanism of the cerebellum (adapted from (Kawato 1999)). Dashed lines denote an adaptive mechanism of the target module.

According to Wolpert *et al.* (Wolpert, Miallb, and Kawato 1998), a single pair of forward and inverse models is not sufficient to model all the diversity of motor tasks humans can perform. And thus he proposes the existence of multiple models, coordinated by a responsibility estimation mechanism. Depending on the context, this weights the preponderance of each model in the feedforward motor command, as well as its plasticity in each situation.

It should be noted that these models learn the dynamics of the whole system, comprising body and environment. They acknowledge the embeddedness and the embodiment of the organism in the world, since the dynamic behaviour of the body+world system depends on the physical properties of both as a whole.

Decision making

Two areas in the brain have been identified as directly implicated in the core process of decision making (Frank and Claus 2006): the basal ganglia and the orbitofrontal cortex. Decision making follows the consideration of one or more action options. These action options are represented and held in the premotor cortex, as representations of corresponding motor commands, during the decision making process.

The basal ganglia (BG) handles decision making on the basis of *go* and *no-go* responses (Schultz 2002; Frank and Claus 2006). It formulates predictions about rewards following the decision taken. When a mismatch occurs, a specific population of neurons¹ that releases dopamine is activated. Dopamine is a neurotransmitter that affects several brain areas. In particular, at the level of the BG, it promotes plasticity of the BG neurons responsible for signaling the *go* and *no-go* responses. Whenever the rewards follow the expected pattern, *i.e.*, either the expected occurrence of a reward, or no reward when none was expected, no plasticity occurs. However, when an unexpected reward takes place (as in the case of the acquisition of a conditioned response), there is an increase of dopamine release, leading to the facilitation of a *go* response. Conversely, when no reward takes place after one being expected, there is a decrease in dopamine release, thus facilitating a *no-go* response.

The acquisition of associations (conditioning) between cues and rewards is driven by the occurrence of initially unexpected rewards following the cue. These unexpected rewards elicit dopamine releases, which then facilitate *go* responses. After the acquisition of an association, the reward becomes expected, and the BG transfers the dopamine release from the reward (which becomes then predictable after the conditioned stimulus), to the occurrence of the cue itself. This mechanism resembles reinforcement based approaches in machine learning (Sutton and Barto 1998). On the one hand, the *go* and *no-go* responses can be matched against utility values, in the following sense: a rational agent, in decision theory, aims the maximization of the expected utility (von Neumann and Morgenstern 1944), and thus the practical consequence of the numerical values of utilities is that higher expected utilities facilitate the actions that lead to

¹Substantia nigra pars compacta.

them. On the other hand, the transfer of dopamine responses back in time (from rewards to conditioned stimuli) is comparable with the propagation of utility values to states that precede (anticipate) future rewards.

It is noteworthy to stress that the BG also promotes instrumental learning in the cortical areas (Frank and Claus 2006). As soon as this instrumental leaning is acquired, the BG does not have an active role (except of course in the case of any relevant unexpected event).

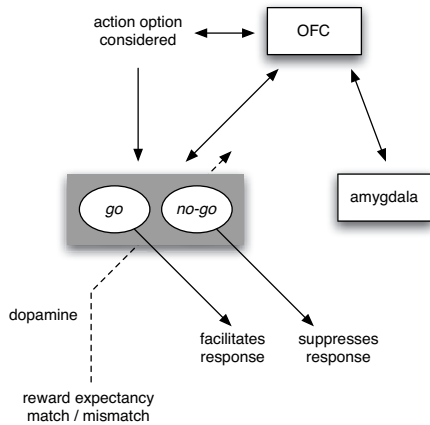


Figure 2: Illustration of the architecture involved in the decision making process (adapted from (Frank and Claus 2006)). The dashed line denotes plasticity effect (dopamine).

The other brain area implicated in decision making is the orbitofrontal cortex (OFC) (Rolls 2004; Frank and Claus 2006) (see figure 2). While the BG does not distinguish rewards of different natures, the OFC is responsible for holding in working memory gain-loss information associated with the considered action options. These associations include the nature of the action outcomes, as well as their temporal relationships (e.g., distinguishing between short-term and long-term rewards, and weighting them against each other). The information about the desirability of outcomes is given by connections from and to the amygdala, which performs affective evaluation of the predicted outcomes. These projections from the amygdala to the OFC are the basis of the contribution of emotions to decision making (Damasio 1994). Another important aspect of the OFC is that these associations are established very rapidly, thus allowing the brain to adapt quickly to new situations. It has been shown that the OFC is essential for good performance in reversal learning experiments, where the rewards are suddenly changed midway the experiment (Frank and Claus 2006).

The OFC contributes to the decision making process by activating either the *go* or the *no-go* areas of the BG (Frank and Claus 2006). The decision of performing an action is thus the result of the combination of two assessments: the one performed by the BG itself, whose adaptation is slowly driven by dopamine releases, and the one performed by the OFC (top-down), which holds in working memory recent information about gain-losses of action options.

We claim that these results suggest a division of tasks in

adaptation, from a machine learning point of view: on the one hand, we have a modality of learning based on the probability of future rewards (the slow adaptation performed by the BG since the estimation of probabilities requires repetition), and, on the other hand, we have a modality of learning based on past events (quick adaptation of the OFC). The latter is able to respond quickly to changes in the conditions under which rewards are obtained, while the former consolidates reward associations that are consistent along time. Decision making cannot be solely based on the recollection of past events, not only because the complex nature of the world would demand enormous storage and retrieval resources, but also because of the need to cope with novel situations.

Emotions

The amygdala is a brain structure primarily responsible for triggering emotional responses. It corresponds to a phylogenetically older part of the brain, thus suggesting that emotion processes predate cognition in the evolution of species (LeDoux 1996). Of the many emotions human are able to experience, fear has been specially the object of in depth research (LeDoux 2000), so that we now have a better understanding of its neural correlates than of any other emotion. The amygdala is able to respond innately to certain stimuli, as for instance the freezing behavior in rats after a loud tone. These responses provide basic survival. Besides, the amygdala is also able to associate novel stimuli with these basic responses. This adaptive capability is essential for a species to cope with complex environments. The classic example of these associations is the Pavlov classical conditioned experiment (Pavlov 1927): the pairing of the display of meat with a bell tone, so that the bell tone alone elicits salivation by the dog. It is well known that the amygdala is a central piece for the acquisition of these conditioned responses.

Sensory input originating in the sensory organs is relayed by the thalamus to both the cortex and the amygdala. In the cortex, it is subjected to cognitive processing, with the capacity of emotional elicitation following activation of the amygdala (see figure 3). These two paths are termed by LeDoux (LeDoux 1996) as the *high-* and the *low-road* to the amygdala, respectively, as they correspond not only to two distinct levels of complexity (refined in the high-road, and rough in the low-road), but also to very different response latencies (approx. 30 to 40 ms via the high-road, and 12 ms via the low-road). The amygdala not only receives sensory input from the thalamus, but also contextual information from the hippocampus (the Pavlov's dog salivates inside the lab, but not outside it). The eliciting of an emotional response has a myriad of consequences, ranging from regulatory functions of the body to cognition function (memory, attention, and decision making, to name a few).

The OFC relies on the amygdala (namely the basolateral nucleus; ABL) to obtain an affective evaluation of possible outcomes before a choice is made (Frank and Claus 2006). The ABL encodes associations between cues and affective value of outcomes, as in the case of the conditioned reflex responses. These associations are retrieved and held in work-




Figure 3: Illustration of the LeDoux's high- and low-road to the amygdala (adapted from (LeDoux 1996)).

ing memory by the OFC. The ABL thus anticipates emotional consequences of the choices being considered.

The contribution of the OFC is not limited to the phase prior to decision making, but it also extends to the processing of the outcomes of the choice made. This mechanism is essential, for instance, in the case of regret (Coricelli, Dolan, and Sirigu 2007), where humans take into account the chances of regretting a particular action option: the knowledge that the reward following a rejected option is better than the one obtained after the taken option. In this case (as in others), human behavior deviates from what is predicted by the expected utility maximization principle.

Bio-inspired cognitive architectures

Brief review of existing approaches

Biological inspiration has provided a fertile ground for the development of many cognitive architectures in the domain of AI. A broad range of proposals can be found in the literature (see (Vernon, Metta, and Sandini 2007) for a review). Prominent earlier ones include SOAR (Laird, Newell, and Rosenbloom 1987) and ACT-R (Anderson et al. 2004). The SOAR takes inspiration mainly in psychology, whereas ACT-R originates from neurophysiology studies. However, both of these architectures have production rules at the core of their implementations, and follow a symbolic approach (the physical symbol hypothesis (Newell and Simon 1976)). The symbolic approach to A.I. has been the object of severe criticism (Brooks 1991; Dreyfus 2007), and it is widely accepted that it is inadequate to robotics (except in very constrained and structured environments). Radically different approaches have been proposed, such as the subsumption architecture by Brooks (Brooks 1991). This has also shown limitations, namely in what concerns scaling to more complex levels of competence. Since then, many other cognitive architectures have been proposed. From the ones with an explicit neurobiological inspiration, we have chosen to take a closer look at the following: the Shanahan's Global Workspace (Shanahan 2006), the Christensen's SDAL (Christensen and Hooker 2000), the Weng's SASE (Weng 2004), and the Kirchmar's Darwin architecture (Krichmar and Edelman 2006).

The Global Workspace cognitive architecture proposed by Shanahan (Shanahan 2006) mimics in an explicit way the interplay between the BG and the amygdala. Its decision

making process comprises two levels: a first level, reactive level, where the control loop is closed through the environment, and a second level, where considered actions are internally rehearsed prior to the actual execution. The architecture is based on Baar's Global Workspace theory (Baars 2002), according to which cognition consists of several parallel modules, all competing for exclusive access to a global workspace. This competition results in a winner-take-all strategy for one module to take control over that global workspace, and the outcome of this control is then broadcasted back to all modules. Shanahan's implementation uses a neural network methodology. The architecture mimics the various brain structures involved (BG, amygdala, visual cortex) and their neural interconnections. Recent experimental results employ a Webots environment with several colored objects, where the (virtual) robot detects the objects' salience and rotates its head accordingly.

The Christensen's Self-Directed Anticipative Learning (SDAL) architecture (Christensen and Hooker 2000) follows a dynamical embodied approach, rooted in two main ideas: (1) specific self-directed processes to attain autonomous maintenance of the system state (e.g., survival), and (2) the use of prediction and anticipation for the construction of new skills. These ideas follow recent theories on anticipative-predictive neural learning, namely the role of dopamine in the process.

Taking a developmental approach, Weng proposed the Self-Aware Self-Effecting (SASE) architecture (Weng 2004). This architecture distinguishes the internal from the external environment of an agent. The agent interacts with the former using the internal sensors and effectors, and with the latter using the external sensors and effectors. The existence of an internal loop through the internal environment allows for the self awareness (and self effecting) proposed by the architecture (e.g., attention selection). This architecture has been implemented in robotic platforms. In these implementations, the sensorimotor system comprises three mappings in cascade: a sensory mapping, mimicking the hierarchical neural arrangement of the visual cortices, a cognitive mapping, including both innate and learnt behaviours, and a motor mapping that employs a subsumption paradigm to arbitrate the behaviours produced by the cognitive mapping. Robotic implementations of SASE include developmental approaches to the recognition from occluded views, to vision-guided navigation, to speech learning, to communicative learning, and to acquiring new capabilities (with or without a teacher).

Kirchmar *et al.* has been developing the idea of brain-based devices, and implementing these in robot platforms called Darwin (Krichmar and Edelman 2006; McKinstry, Edelman, and Krichmar 2006). These implementations follow closely models of neural dynamics and of the connectivity in the brain. The levels of competence they have attained with this approach include the recognition of visual targets, by mimicking several visual cortices, the navigation and categorization of textures using whiskers, based on the neuroanatomy of rats, and the development of spatial and episodic memory, based on models of the hippocampus and related regions.

Challenges

The findings reviewed here suggest several possible contributions of the neurosciences for the development of cognitive architectures.

First, we consider the binding problem, which consists in the problem of how the brain integrates, in a coherent way, information processed in different brain regions. It has been proposed that this is accomplished using event files (Hommel 2004), linking neural codes involved in both the perception of an object (the various perceptual features are themselves integrated into object files), and the action (if any) that is somehow related with that object (action files).

The brain presents a modular organization, but its modules do not function independently. However, cognitive architectures often feature different modules functioning more or less independently. A mechanism in cognitive architectures is thus needed to bind perceptions and actions causally connected in the interaction with the environment. Moreover, from a technical point of view, the implementation of such binding mechanism raises questions of how the continuous nature of motor control is coherently integrated with the discrete time, event based, process of considering courses of action, predicting outcomes, and making decisions. This coherence between these two time scales is attained by the brain, according to Kurby *et al.*, by a mechanism that segments perception into events (Kurby and Zacks 2008). According to them, this segmentation is based on a set of (local) predictors, that detect an event change whenever an significant error is detected between the predicted perception and the actual one.

Such predictors can be found at several brain regions (Schultz and Dickinson 2000): the forward models at the cerebellum predict future sensory consequences of motor commands (Wolpert, Miallb, and Kawato 1998), the reward expectancy at the BG (Schultz 2002), the predictive coding in the visual cortex (Rao and Ballard 1999), and so on. The ability of making predictions depends on models of some sort, involving the interaction of the organism within its environment. These models are therefore embedded and embodied. They are embedded because they model the organism together with the environment as a whole, and they are embodied because they depend on the existence of a body, that conditions the interaction with the world. This idea is not new, but we would like to stress that the prediction mechanisms at the higher cognitive levels should integrate the ability to predict the outcome of low level reactive processes. For instance, imagine a robot equipped with a reactive obstacle avoidance mechanism; a high level trajectory planning module should be able to predict that the robot would detour around an obstacle automatically, without direct intervention of the high level module.

We find it interesting to note that the discussed brain subsystems do not function independently. For instance, the cerebellum depends on prediction errors as evaluated by the visual cortices; the affective evaluation used by the OFC depends on the amygdala assessment of the action options; and the *go/no-go* response depends on the contribution of both the BG and the OFC. The economy of resources and the use of previously developed structures may provide a phylo-

genetic explanation for this organization. Some behaviour-based architectures (Arkin 1989) follow a different organization, as behaviours are implemented as independent modules, whose actions are then combined at the end.

Future directions

Two issues deserve further research from our point of view. The first one is the problem of perception. This includes, on the one hand, the problem of integrating the information from diverse sources and with various modalities: sensory information with recalled memory, cross-modal sensory information (e.g., visual with auditory), continuous time with discrete time. On the other hand, perception depends on predictive mechanisms that can be found in diverse brain subsystems. The design of such predictive mechanisms has to meet several challenges, for not only do they have to limit their scope to a subset of relevant aspects, but they also have to address the problem of how they are applied to new situations. Moreover, such predictive mechanisms can help segmenting time into events.

The second issue is the non-utilitarian nature of human decision making. Utilitarian approaches are founded on utility theory, according to which preferences can be represented by real numbers, and on decision theory, based on the maximization of the expected utility principle (von Neumann and Morgenstern 1944). These theories have provided fertile ground, for instance, to the research on sequential decision making (e.g., reinforcement learning (Sutton and Barto 1998)), sharing a solid mathematical foundation. However, there is evidence against the adequacy of utilitarian approaches to explain human decision making. One example already discussed in this paper is the case of regret (Coricelli, Dolan, and Sirigu 2007): experiments have shown that humans tend to prefer options with lower chance of future regret, than the one that mathematically maximizes expected utility. Another example is the compelling evidence that, in certain situations (namely of risk), humans are insensitive to the probability of outcomes, thus deviating from what could be predicted by utilitarian models (Loewenstein *et al.* 2001). These and other findings have contributed to the emerging field of neuroeconomics (Glimcher and Rustichini 2004), resulting in the convergence of economics, psychology, and neuroscience.

Conclusions

This paper takes a biologically inspired approach to the design of cognitive architectures. The biological inspiration is not new in the field of AI, and we have reviewed a selection of cognitive architectures that explicitly acknowledge this inspiration. However, we claim that progress in the neurosciences continues to provide interesting ideas for the AI research community. Moreover, we support the idea that the design of cognitive architectures for robots poses specific problems for which the biological inspiration has potential to be particularly useful. First, because robots are physically embodied, and thus relate with the world with the same constraints of situatedness as biological organisms, and second, because robot sensory apparatuses have to meet

the challenge of perception, encompassing the processing of raw sensory data up to the level of interpretation.

Acknowledgments

The author wishes to thank Patrícia Figueiredo for reviewing an early draft of this paper, as well as the comments from the anonymous reviewers. This work was partially supported by FCT (ISR/IST plurianual funding) through the POS_Conhecimento Program that includes FEDER funds.

References

- Anderson, J. R.; Bothell, D.; Byrne, M. D.; Douglass, S.; Lebiere, C.; and Qin, Y. 2004. An integrated theory of the mind. *Psychological Review* 111(4):1036–1060.
- Arkin, R. C. 1989. *Behavior-Based Robotics*. MIT Press.
- Baars, B. J. 2002. The consciousness access hypothesis: origins and recent evidence. *Trends in Cognitive Science* 6(1):47–52.
- Brooks, R. A. 1991. Intelligence without representation. *Artificial Intelligence* 47:139–159.
- Campbell, M.; Jr., A. J. H.; and Hsu, F. H. 2002. Deep blue. *Artificial Intelligence* 134:57–83.
- Christensen, W. D., and Hooker, C. A. 2000. An interactivist-constructivist approach to intelligence: self-directed anticipative learning. *Philosophical Psychology* 13(1):5–45.
- Coricelli, G.; Dolan, R. J.; and Sirigu, A. 2007. Brain, emotion and decision making: the paradigmatic example of regret. *Trends in Cognitive Sciences* 11(6):258–265.
- Damásio, A. R. 1994. *Descartes' Error: Emotion, Reason and the Human Brain*. Picador.
- Dreyfus, H. L. 2007. Why heideggerian AI failed and how fixing it would require making it more heideggerian. *Philosophical Psychology* 20(2):247–268.
- Frank, M. J., and Claus, E. D. 2006. Anatomy of a decision: Striato-orbitofrontal interactions in reinforcement learning, decision making, and reversal. *Psychological Review* 113(2):300–326.
- Glimcher, P. W., and Rustichini, A. 2004. Neuroeconomics: The consilience of brain and decision. *Science* 306(5695):447–452.
- Hommel, B. 2004. Event files: feature binding in and across perception and action. *Trends in Cognitive Sciences* 8(11):494–500.
- Kawato, M. 1999. Internal models for motor control and trajectory planning. *Current Opinion in Neurobiology* 9(6):718–727.
- Krichmar, J. L., and Edelman, G. M. 2006. Principles underlying the construction of brain-based devices. In Kovacs, T., and Marshall, J. A. R., eds., *Adaptation in Artificial and Biological Systems*, 37–42. Society for the Study of Artificial Intelligence and the Simulation of Behaviour.
- Kurby, C. A., and Zacks, J. M. 2008. Segmentation in the perception and memory of events. *Trends in Cognitive Sciences* 12(2):72–79.
- Laird, J. E.; Newell, A.; and Rosenbloom, P. S. 1987. SOAR: an architecture for general intelligence. *Artificial Intelligence* 33(1):1–64.
- LeDoux, J. 1996. *The Emotional Brain*. Simon & Schuster.
- LeDoux, J. 2000. Emotion circuits in the brain. *Annual Review of Neuroscience* 23:155–184.
- Loewenstein, G. F.; Weber, E. U.; Hsee, C. K.; and Welch, N. 2001. Risk as feelings. *Psychological Bulletin* 127(2):267–286.
- McKinstry, J. L.; Edelman, G. M.; and Krichmar, J. L. 2006. A cerebellar model for predictive motor control tested in a brain-based device. *Proceedings of the National Academy of Sciences of the United States of America* 103(9):3387–3392.
- Newell, A., and Simon, H. A. 1976. Computer science as empirical inquiry: Symbols and search. *Communications of the ACM* 19(3):113–126.
- Pavlov, I. P. 1927. Conditioned reflexes: An investigation of the physiological activity of the cerebral cortex. Classics in the History of Psychology web site. (Translated by G. V. Anrep).
- Rao, R. P. N., and Ballard, D. H. 1999. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience* 2(1):79–87.
- Rolls, E. T. 2004. Convergence of sensory systems in the orbitofrontal cortex in primates and brain design for emotion. *The Anatomical Record Part A* 281A(1):1212–1225.
- Schultz, W., and Dickinson, A. 2000. Neuronal coding of prediction errors. *Annual Review of Neuroscience* 23:473–500.
- Schultz, W. . 2002. Getting formal with dopamine and reward. *Neuron* 36(2):241–263.
- Shanahan, M. 2006. A cognitive architecture that combines internal simulation with a global workspace. *Consciousness and Cognition* 15:433–449.
- Sutton, R. S., and Barto, A. G. 1998. *Reinforcement Learning*. MIT Press.
- Vernon, D.; Metta, G.; and Sandini, G. 2007. A survey of artificial cognitive systems: Implications for the autonomous development of mental capabilities in computational agents. *IEEE Transactions on Evolutionary Computation* 11(2):151–180.
- von Neumann, J., and Morgenstern, O. 1944. *Theory of Games and Economic Behavior*. Princeton University Press.
- von Neumann, J. 2000. *The Computer and the Brain*. Yale Nota Bene, second edition edition.
- Weng, J. 2004. Developmental robotics: Theory and experiments. *International Journal of Humanoid Robotics* 1(2):199–236.
- Wolpert, D. M.; Miallb, R. C.; and Kawato, M. 1998. Internal models in the cerebellum. *Trends in Cognitive Sciences* 2(9):338–347.