

Opinion

Navigating the Affordance Landscape: Feedback Control as a Process Model of Behavior and Cognition

Giovanni Pezzulo^{1,*} and Paul Cisek²

We discuss how cybernetic principles of feedback control, used to explain sensorimotor behavior, can be extended to provide a foundation for understanding cognition. In particular, we describe behavior as parallel processes of competition and selection among potential action opportunities ('affordances') expressed at multiple levels of abstraction. Adaptive selection among currently available affordances is biased not only by predictions of their immediate outcomes and payoffs but also by predictions of what new affordances they will make available. This allows animals to purposively create new affordances that they can later exploit to achieve high-level goals, resulting in intentional action that links across multiple levels of control. Finally, we discuss how such a 'hierarchical affordance competition' process can be mapped to brain structure.

Human Cognition from the Perspective of Feedback Control

Cognitive science is defined as the study of the human mind, and its fundamental tenet is that 'thinking can be understood in terms of the representational structures in the mind and computational procedures that operate on those structures' [1]. This definition places cognition between the perceptual processes that provide its input and the motor processes that execute its output – sketching the shape of the serial sense–think–act model of behavior that has dominated psychological theories for more than 50 years. However, throughout that time, an alternative view has existed, proposing that the brain is a feedback **control system** (see [Glossary](#)) [2–5] whose primary goal is not to understand the world, but to guide interaction with the world. A feedback control system is one in which outputs are generated so as to control some variable whose value is measured via input. In the case of behavior, actions are performed to keep the animal in a desirable state (satiated, safe, etc.) and **perceptions** are used to evaluate that state [4,6]. In this opinion article, we discuss how that feedback control view of behavior can be extended to go beyond simple sensorimotor control, and how it can provide a conceptual foundation for understanding human cognition that better aligns with neurophysiological data than classic serial models.

Similar to other biological processes (e.g., thermoregulation), behavior is a feedback control process – we take actions so as to influence our state in the world [6,7]. Although overt behavior extends beyond the skin, it is nevertheless functionally organized like other biological feedback processes: it relies on predictable causal relationships between actions and outcomes (approach food → make food obtainable) and is self-regulating (eat food → satiate hunger/deplete food).

Trends

Traditional assumptions of cognitive psychology are increasingly questioned by neurophysiology, casting doubt on the classic framework of serial information processing.

For over 100 years there has existed an alternative framework, which describes behavior as a control system. Although originally applied to simple actions, it is increasingly being extended to address more sophisticated behavior, including intentional action.

The brain's ability to predict the consequences of actions enables it to link across levels of abstraction, and to bias immediate actions by the predicted long-term opportunities they make possible – hence supporting intentional action.

The organization of the brain, including the cerebral cortex, is increasingly viewed in terms of the species-typical activities that it evolved to support, as opposed to the hypothetical modules of cognitive psychology theory.

¹Institute of Cognitive Sciences and Technologies, National Research Council, Via San Martino della Battaglia 44, 00185 Rome, Italy

²Department of Neuroscience, University of Montréal, 2960 chemin de la tour, room 4117, Montréal Québec, H3T 1J4, Canada

*Correspondence: giovanni.pezzulo@istc.cnr.it (G. Pezzulo).

The basic framework of feedback control is widely recognized in studies of autonomic physiology [8], sensorimotor control [9–11], and natural animal behavior [12], but largely absent from theories of how human cognition operates. Here, we argue that this is an oversight and that even human cognition is best understood within the context of the feedback control theoretical principles that govern all biological systems. In this perspective, we take adaptive action control – and the problems faced by situated agents who pursue their goals in dynamic (yet structured) environments – as a central paradigm to understand human cognition.

Some biologically grounded models of embodied action and cognition, such as the ‘**affordance competition hypothesis**’ [13], ‘**active inference**’ [14], and others [15–17], incorporate control theoretical principles (Figure 1). However, these proposals have been mostly applied to simple scenarios and cognitive tasks that do not fully engage higher cognitive processes. Here, we discuss how these models can be extended beyond simple sensorimotor behavior to address the domain of **intentional action** and higher cognitive skills, while retaining important principles of feedback control at their core.

Hierarchical Affordance Competition

The affordance competition hypothesis [13, 18, 32] suggests that during interactive behavior, the brain simultaneously specifies the set of desirable actions currently available in its environment (‘**affordances**’), and decides what to do through a competition between representations of these actions, biased by the desirability of their predicted outcomes. Once a given action is selected, it is executed through continuous feedback control, using sensory information from the environment as well as internal predictions of expected feedback to fine-tune and update the ongoing action until completion. Furthermore, because alternative potential actions continue to be processed even during ongoing activity, the hypothesis proposes how animals can rapidly switch actions if the need or opportunity arises.

Importantly, in this view, potential actions are not distinct categorical entities, such as the button presses of a classical psychological experiment, but regions within a continuous landscape of actions – akin to a ‘**desirability density function**’ across the space of movement parameters (Figure 2A) [19]. That landscape is defined by the geometry of the external world and changed continuously by events in the environment and the animal's own actions. When choices emerge as distinct regions of desirable actions, adaptive behavior relies on the animal's ability to predict the future consequences of selecting one over another.

While affordance competition was initially described as a theory of how animals select between concrete and immediately available actions, it can be extended toward a more general theory of decisions made on multiple levels of abstraction [20]. Here, we discuss how such a ‘**hierarchical affordance competition**’ can address the domain of intentional action. Key to this proposal is the recognition that the brain's ability to predict the consequences of actions enables it to link across levels of abstraction, and to bias immediate actions by the predicted long-term opportunities they make possible.

According to this hypothesis, intentional action can be conceptualized as a (purposive) navigation in an ‘**affordance landscape**’: a temporally extended space of possible affordances, which changes over time due to events in the environment but also – importantly – due to the agent's own actions. The key for extending the simple competition among affordances toward intentional action is to recognize that brains are continuously engaged in generating predictions (e.g., about future opportunities) rather than just reacting to already available affordances [21–23]. Consider the example depicted in Figure 2B: a monkey is sitting on a tree branch, within reach of a small berry. This situation presents the monkey with the possible actions (among others) of reaching for the berry or walking outward on the tree branch. Because the monkey can predict

Glossary

Active perception: in ecological psychology (and beyond), perception is active in many senses. First, cognitive processing does not start with a passive stimulus-processing phase, but with an action that produces the next sensed stimulus (or an expectation to be met). In turn, the stimulus is used within the feedback loop to guide action toward its goal [72]. Second, objects are recognized through the actions we (can) make on them, see Affordance. Third, active perception refers to strategies of sensor or eye movement control that permit, for example, to implement the right stimuli for the current action context [73] or hypothesis testing [74].

Affordance: a potential action that is made possible to an agent by the environment around it [72]. While affordances are defined with respect to an agent's individual capabilities (e.g., a tree branch might have a ‘walkability’ affordance for a monkey, not necessarily for a sedentary man), they are objective in the sense that they do not depend on whether the agent perceives them, attends to them, or chooses to act on them, and can be recognized by anyone familiar with the agent's motor repertoire.

Control of perception: in analogy with homeostatic loops, behavior can be described as the ‘control of perception’ [4]: what is controlled is a perceptual state (e.g., when driving, the indicator of the speedometer is kept stable on ‘100 kph’), whereas actions (e.g., press brake versus gas pedal) are contextually selected to keep it in the desired range, counteracting ‘disturbances’ (e.g., hills).

Control system: a system that is able to keep one (or more) controlled variable(s) within a given range, to match a ‘reference signal’ (or ‘set point’) despite disturbances. Any mismatch between the reference signal and the perceptual (feedback) signal represents an error that the controller seeks to minimize by taking action – as in the case of a thermostat that opens or closes the furnace to keep room temperature (the controlled variable) within a prespecified range.

Feedback signal: a signal used to update a control system's estimate of the state of the controlled variable. It

that reaching for the berry has the desirable benefit of satisfying hunger, that action will be favored in the competition if nothing else (internal, e.g., satiety; or external, e.g., a predator) enters the picture. However, advanced brains can go beyond such simple choices. In particular, the monkey can predict that walking out on the tree branch will result in putting an apple within reach – a prediction from available affordances (a ‘walkable’ tree branch) to the expected affordances that those actions make available (a ‘reachable’ apple). Because reaching for an apple is highly desirable, this predicted affordance can now be linked to the current situation via a ‘top-down’ bias that favors the selection of walking over reaching for the berry.

Decision-Making within a Hierarchy of Control Loops

This idea of top-down biasing of decisions can be formalized by casting affordance competition as a hierarchy of control loops [4,7], with multiple competitions occurring in parallel at different hierarchical layers of the architecture and mutually influencing each other via top-down and bottom-up signals [20]. Here, prediction dynamics can elicit representations of future affordances and engender a competition at higher layers between action courses yielding different distal outcomes (berry versus apple), which in turn continuously biases in a top-down manner the competition occurring at lower layers between proximal actions (pick the berry versus walking); for example, by setting a ‘reachable apple’ subgoal as a desirable **set point** for the lower layer.

However, the competition at lower layers is part and parcel of the decision process and can feed back on the competition at higher layers. For example, despite the initial top-down bias, situational constraints might cause the lower layer competition to be won by a lower cost motor plan for picking the berry (e.g., if the monkey is fatigued or the tree branch is too wet). This creates a mismatch between two hierarchical levels: the affordance expected by the higher layer plan is not produced by the berry picking action selected by the lower layer. If the hierarchical architecture represents the outcomes of these plans (say) as density functions over animal/hand locations, their mismatch propagates bottom-up in the hierarchy as a feedback (or prediction error) signal and can eventually cause a revision of the apple reaching plan. This implies that ultimately a decision is not computed centrally but in a distributed manner [20].

More generally, the purposive aspect of the ‘affordance navigation’ process lies in the fact that an intentional agent is not limited to (reactively) pick up one of the currently available affordances, but can also (intentionally) create or destroy affordances, which can then be exploited to execute successive actions that ultimately achieve long-term goals. For example, for a climber, the right way to grasp a climbing hold is the one that is functional to reach the next hold in the sequence (and ultimately the top of the climbing route), thus the former movements serve to create affordances for the latter movements. This ‘affordance navigation’ emerges from a continuous climber–wall interaction, but it can be also – at least partially – planned before starting the climb, which requires the climber to predict (sequences of) affordances that are not yet available but can be created, in a similar way as the apple reaching plan in the previous example [24]. This process has to take into consideration situated and embodied aspects of the problem (e.g., body strength, limb length, fatigue, configuration of the climbing holds), epitomizing the kind of embodied decisions that animal continuously face in their daily activities [19,25].

Action Control within a Hierarchy of Control Loops

Intentional affordance navigation requires behavioral flexibility. For example, a boxer who wants to hit an opponent often needs to first move toward him to make the ‘hittability’ affordance available; but sometimes, when he is too close to the opponent, needs to move backward for the same purpose. This illustrates a hallmark of control theories [4]: what the control loop strives to keep stable is the controlled variable – here, the body scaled boxer–opponent distance, which in turn determines the ‘hittability’ affordance – while the action required to achieve this result is

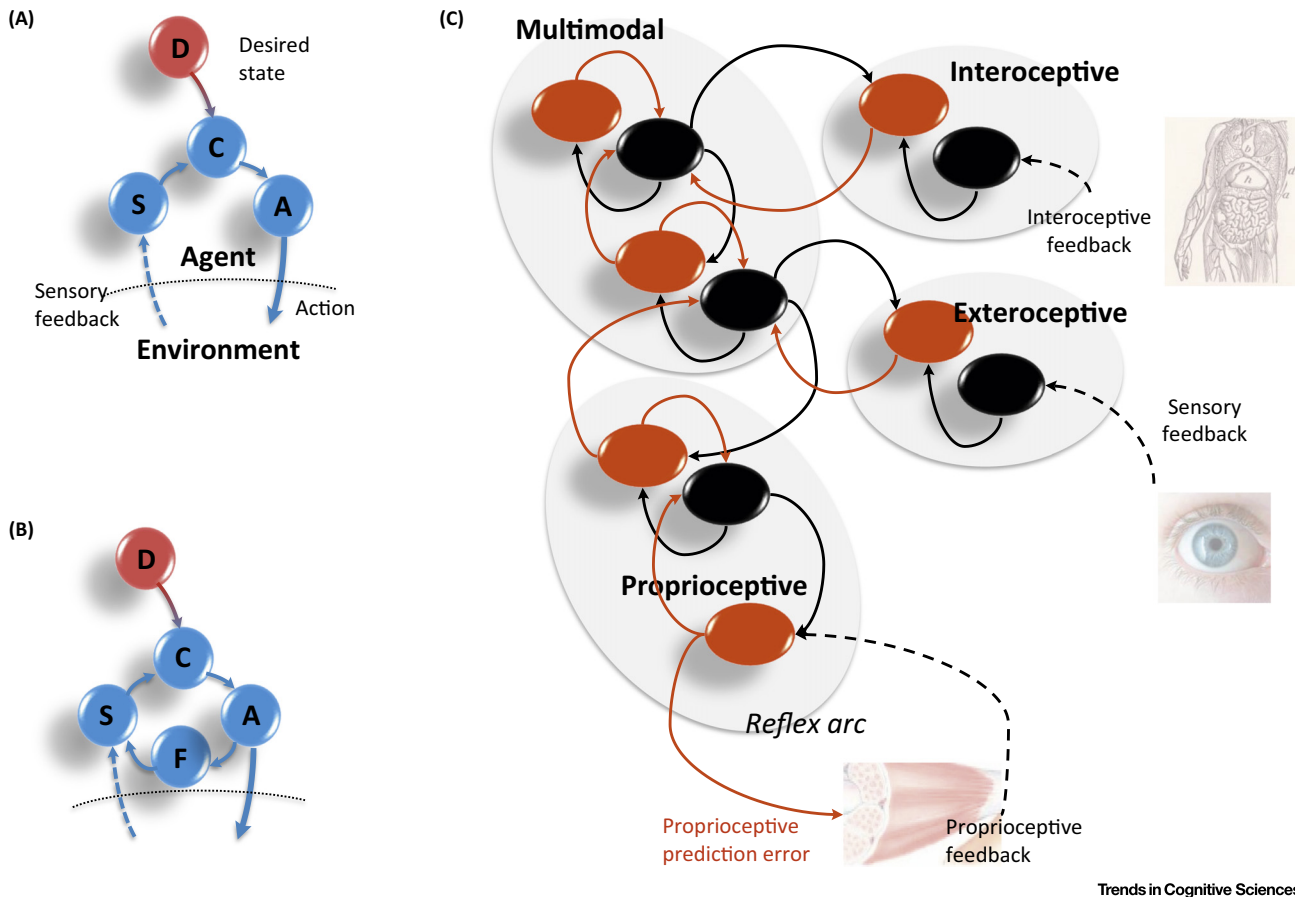
usually refers to sensory information arriving from the external world, but could also refer to internal feedback.

Forward model: a component of control systems that serves to predict the next state based on its current estimate and the selected action.

Intentional action: action that is taken to produce some intended effect, that is, achieve some desirable goal.

Reflex arc: a neural pathway that controls a reflex action. As Dewey noted [5], it should be conceived as a circular process in which a stimulus motivates an action as much as an action produces the next stimulus.

Set point: the desired or target value for a controlled variable of a system.



Trends in Cognitive Sciences

Figure 1. Examples of Feedback Control Schemes. (A) Schematic of a simple feedback controller. If the comparator c detects a discrepancy between desired state (or set point) d and current state s (e.g., between desired and actual car velocity), an action a is triggered (e.g., press brake) that causes changes in the world, which in turn result in sensory feedback (broken arrow) and influence the state s . This architecture can be replicated hierarchically, when the set point ('100 kph') is furnished by a hierarchically higher level that encodes more abstract goals (e.g., 'get home soon'). (B) A feedback system augmented with a forward model f that predicts the state that results from executing the action. In advanced control theoretical schemes, prediction is used, for example, to improve state estimation or to substitute missing (or delayed) feedback [11]. (C) In active inference, a hierarchy of prediction (black) and prediction error (red) units form a generative model for perception and action [75]. Goals encoded at high hierarchical levels (as prior preferences) generate a cascade of descending predictions (black edges) and ascending prediction errors (red edges) in various modalities: exteroceptive, proprioceptive, interoceptive. Descending predictions are compared with incoming sensations to generate prediction errors that are propagated backward. The architecture uses (precision-weighted) top-down and bottom-up dynamics to continuously suppress prediction errors until the external situation matches the goal – or the goal is revised. Action consists in minimizing (proprioceptive) prediction errors by engaging reflexes; but active inference also accommodates action plans (Box 2).

context-dependent (e.g., move forward or backward depending on the boxer–opponent distance), and is not simply a fixed response [26].

The hierarchical control architecture elucidated earlier has this flexibility. Higher hierarchical layers encoding more abstract goals ('land a punch') propagate top-down the expectations (or set points or reference signals) for the lower layers, such as for example 'maintain a given distance from the opponent', and in this way constrain the to-be-produced affordances ('hittability') – but, importantly, do not prescribe how the lower layers should produce them (e.g., by moving forward or backward). Conversely, the success or failure in engendering the required affordances produces a **feedback signal** or residual prediction error (e.g., a mismatch between the expected and actual affordance, which is low if the hittability affordance was successfully produced and high otherwise) that propagates bottom-up in the hierarchy and sometimes forces the boxer to revise his strategy.

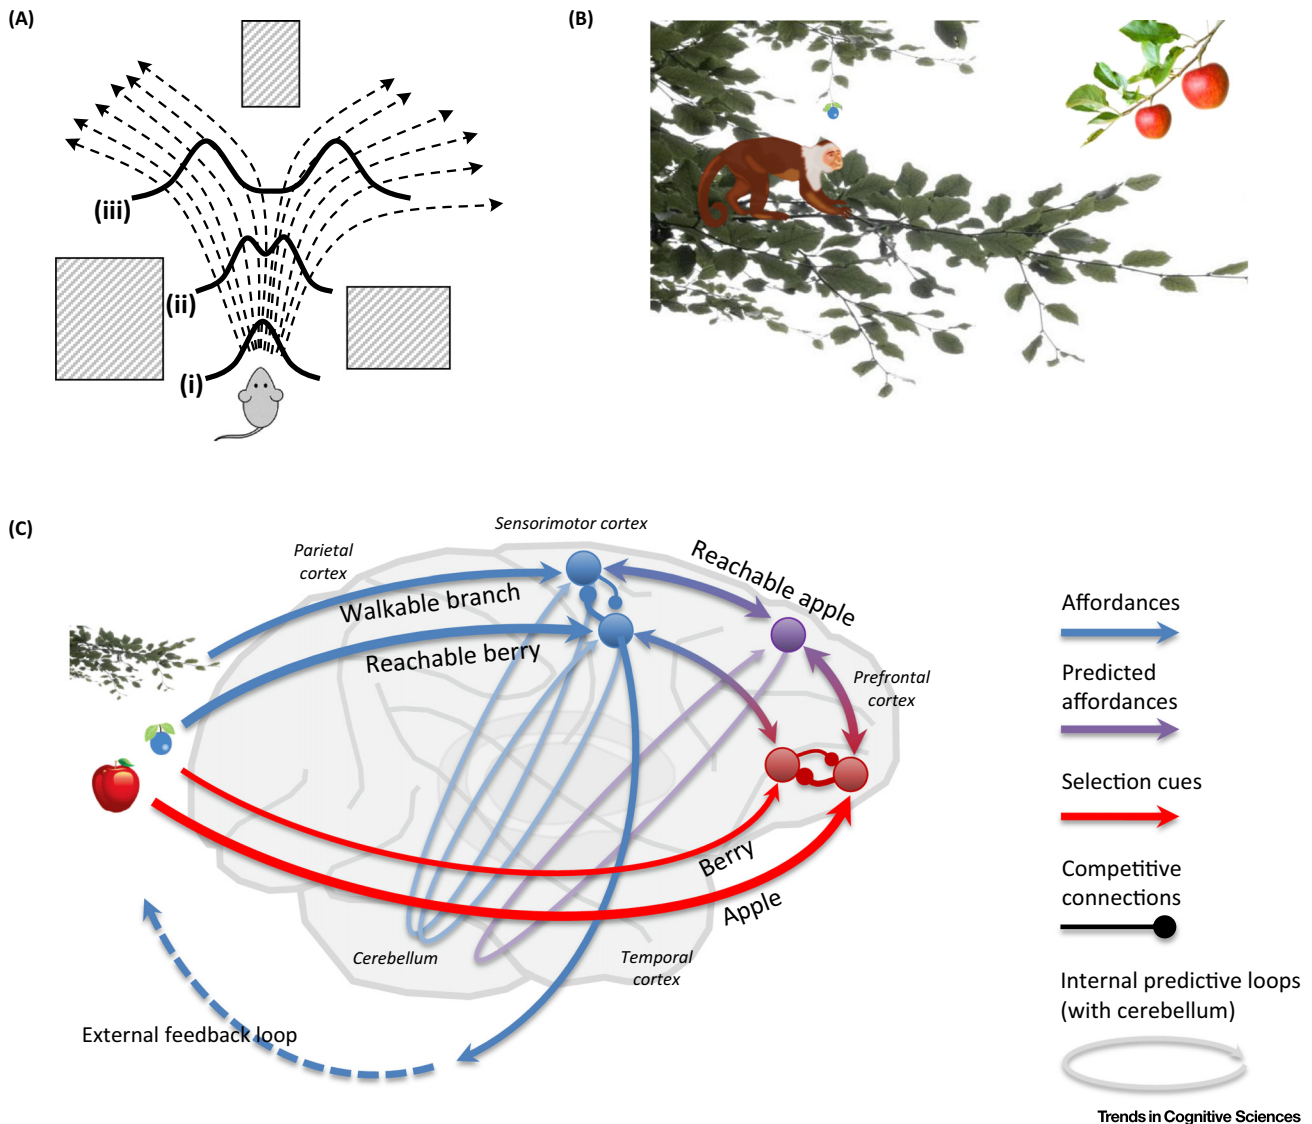


Figure 2. Intentional Navigation in an Affordance Landscape. (A) Schematic illustrating how, in ecological contexts, affordances and the choices between them emerge from the geometry of interactions between the agent and its environment. Broken lines indicate possible paths for the mouse to move between obstacles. Unbroken curves indicate distributions of potential directions at three moments in time. At (i), the distribution can be averaged into a single central direction – thus there is a single affordance and no competition. At (ii) the distribution begins to separate, but averaging is still possible. At (iii) the average is no longer viable and a decision must be made between the two possible choices that emerged: moving to the right or left. (B) A more complex choice situation that involves both proximal and distal outcomes: a monkey is faced with selecting between reaching for a berry versus walking outward on a tree branch toward an apple. According to hierarchical affordance competition, this situation maps to a multilevel ‘decision space’, where both proximal actions (pick the berry versus walking) and distal outcomes (berry versus apple) compete, but at different hierarchical levels of a multilevel feedback controller with prediction linking across levels. (C) The hierarchical affordance competition schematically mapped onto brain structures. Available affordances are specified along the parietal cortex (blue arrows) and compete in sensorimotor regions (blue circles). The competition is biased by goals specified in prefrontal cortex on the basis of visual processing in temporal cortex (red arrows) and predicted affordances (purple). Panel A is reproduced with permission from [19].

What is computed at higher levels is not a complete behavioral plan that is successively decomposed and executed downstream, but a nested cascade of expectations or reference signals that prescribe the next affordances to be produced (or directly exploited if already available), without necessarily specifying the lower level actions to be executed (Box 1). In other words, the higher levels bias the competition at lower levels but ultimately leave them significant autonomy in action selection. This implies that nested within a control process there is always a

Box 1. Functional Role of Top-Down and Bottom-Up Signals in Hierarchical Feedback Control Architectures

In hierarchical feedback control architectures, top-down and bottom-up signals can convey predictions and prediction errors, respectively (see [Figure 1](#) in main text). However, the functional interpretation of these signals is different if one focuses on control or decision-making tasks. In a control task, a top-down signal can be interpreted as a reference value (or set point) provided by a higher level plan (e.g., for reaching an apple) that constrains the affordance to be produced by a lower layer action (e.g., being close to the apple, or apple reachability). Conversely, a bottom-up signal reports the mismatch between this expected reference value and the affordances currently available. Feedback control ensures that this mismatch or residual prediction error is minimized by selecting an appropriate lower level action (i.e., walk the tree branch toward the apple). Instead, in a decision-making task, both top-down and bottom-up signals can be conceptualized as biases for the competitions occurring at different hierarchical levels. For example, the expectation of a long-term benefit associated with eating the apple can exert a top-down bias for the competition between walking a tree branch versus picking up a berry – operationally, this can be done by setting ‘walkability’ as an (expected) affordance to be produced by the lower layer competition, thus diminishing the value of the alternative action choices at that level. In a similar manner, the competition occurring at a given (lower) layer can bias the competition occurring at another (higher) layer, as it can generate a mismatch with the expected affordance (e.g., when a low cost action for picking up a berry is selected) that propagates bottom-up as a prediction error that cannot be easily minimized unless the monkey abandons the plan to reach for the apple. That both top-down and bottom-up signals have a biasing role in decision-making reflects the fact that ultimately feedback control needs to minimize mismatches (or prediction errors) at all levels. Still, the fact that some levels can have relatively more importance than others can be captured in these systems using mechanisms of gain (or precision-based) control, which essentially weigh prediction errors by their relative importance or reliability [76]. This also entails that when biases (or even priors) become strong enough, they can prevent other parts of the system to effectively contribute to the choice. This is a potential mechanism through which lower layer actions that are afforded sufficient gain or precision might become routinized and activated automatically by a given situation rather than by top-down signals [77].

continuous competition at the level below, among alternative ways to specify the demands of higher levels. Ultimately, the selected action can be executed through continuous feedback control.

The hierarchical feedback control view of behavior and cognition suggests a novel way to look at brain function and its neural organization. We have discussed how affordances can exist at multiple levels, including concrete actions that are immediately available (reaching for an object within range), desirable situations that could be made available through specific actions (being within reach of something), and high-level goals that satisfy current needs (eating a fruit). If the brain can represent its environment at these multiple levels, then it can link between them through prediction of consequences (from current to future, predicted affordances) and top-down biasing of choices (from plans to achieve distal goals to proximal actions). In the next section, we briefly discuss how this hierarchical control structure is reflected in the organization of the mammalian forebrain.

Brain Mechanisms Supporting Hierarchical Affordance Competition

It has been suggested that the expansion of the frontal cortex, especially prominent among primates, made possible the extension of planning to more abstract and more temporally extended activity [27]. Our hypothesis revisits this proposal of brain organization, although within the context of a system of nested feedback control loops.

The architecture for ‘hierarchical affordance competition’ has a multilevel structure, whereby low level mechanisms for competition among available affordances (e.g., pick the berry or not) are regulated by a higher level mechanism of competition among predicted states (e.g., the decision to walk the tree branch to create the reachability affordance for an apple) and expected outcomes (eat berry or apple) [28]. We propose that this hierarchical ‘decision and action space’ is reflected in brain physiology ([Figure 2C](#)). The lowest level mechanisms consist of parallel streams within reciprocally interconnected frontal and parietal cortical areas, which implement the sensorimotor control of different species-specific actions (in the case of primates, terrestrial and arboreal locomotion, reach-to-grasp actions, feeding behavior, defensive and

aggressive behavior, gaze control, etc.) [29–31]. Each of these low level systems resolves competition between the affordances to which it is sensitive (e.g., walkable branches, reachable objects, grasp types, gaze targets) within specialized maps of action space [28,32,33]. Further arbitration between these action types is resolved by closed loops between each cortical region and the dorsolateral striatum, pallidum, and thalamus [34–36]. Each of these systems also forms a specific cortico-cerebellar loop that predicts the sensory consequences of specific actions. The functional topology of specific cortical regions with dedicated cortico-striatal and cortico-cerebellar loops is recapitulated anteriorly in the frontal lobe: each frontal cortical region possesses a cortico-striatal and a cortico-cerebellar loop, and each is reciprocally connected both with its immediate posterior and anterior neighbors [37]. A growing body of neurophysiological studies has shown that the information processed in frontal cortex is increasingly abstract and domain-independent as one progresses from posterior to anterior frontal cortex [38–41].

Importantly, hierarchical control loops follow a temporal principle of organization, reflecting the fact that decisions should be made, and actions controlled, at a hierarchy of timescales: from shorter ones that only affect immediate interactions (e.g., pick a berry versus walk), to longer ones that have longer-term consequences (e.g., approach apple and reach for it). Actions at different levels of abstraction [42] are reflected in multiscale brain dynamics and internal models [43], and predictive dynamics link across them (Box 2). In principle, a hierarchy of cortico-subcortical loops is well placed to support multiscale brain dynamics, but several aspects of this idea remain to be assessed empirically (e.g., if and how cortico-cerebellar loops can generate predictions at multiple timescales). This principle of organization has profound consequences on the way we conceptualize brain structure and function.

First, brain cortico-subcortical hierarchies do not correspond to a model in which higher levels specify whole behaviors (e.g., a whole defensive movement) and lower levels decompose it into subunits (e.g., the component finger, hand, and head movements). Instead, growing evidence suggests that cortical structures are organized around complex and ethologically relevant

Box 2. Active Inference as a Modern Version of Cybernetic Theory

Active inference is essentially a (Bayesian) predictive coding architecture extended with **reflexes** [14,37,78]. Predictive coding was first proposed as a model of visual perception, in which the hierarchical layers are coupled through top-down and bottom-up signals, encoding predictions and prediction errors, respectively, and weighted by their precision (inverse variance). Top-down and bottom-up dynamics serve to suppress prediction errors (or free energy [75]); sensory mismatches at the lowest layer propagate upward and help revise (higher) perceptual hypotheses. In contrast to predictive coding, active inference can also minimize prediction error by acting: by engaging reflexes that suppress residual (proprioceptive) errors. For example, if one expects to see a berry but does not see it, not only can he revise the perceptual hypothesis ('there is no berry') but he can also put the berry in front of him or search for the berry by moving the eyes, until there is no more prediction error.

Active inference implements planning in a way that resembles the idea that distal affordances (e.g., apple reachability) can influence the competition between proximal actions (picking the berry versus walking) [14]. It uses a hierarchical generative (forward) model to predict action consequences, and the ensuing 'value' of possible action sequences (plans) by considering – iteratively – whether the distal states they make accessible approximate the goal states (encoded as prior preferences). These plan values enable the selection of immediate actions: the greater the plan's value, the more likely it is to specify the next action.

As these examples illustrate, active inference can be considered a biologically grounded synthesis of cybernetic ideas (on homeostasis and control) and the Bayesian brain hypothesis. This might seem odd – because cybernetic theory often dispenses with an 'inner model', while according to the Bayesian brain hypothesis, the brain is a statistical machine that learns world models. However, in active inference the necessity of models stems from control principles (e.g., the 'good regulator theorem' that the best regulator requires a model [79]). Furthermore, there is an essential contribution of the body and environment in structuring the content of generative models, because it needs to embody the structure of sensorimotor interactions. Although the representational aspects of active inference seem odd to 'radical' embodied theories, it is possible that within this scheme one can understand how representational abilities emerge that are relevant for interactive behavior [49,80].

movements – ethological action maps [29,31,44] – which group somatotopically the animal's behavioral repertoire (e.g., hand, mouth, and eye movements that collectively realize a feeding movement in one part of the motor cortex, and those that collectively realize a climbing movement in another part). These maps may be reiterated at multiple hierarchical layers. For example, the more abstract anterior cortical regions (with their associated subcortical loops) may be organized by ethologically relevant goals such as exploiting available goods, exploring the environment, and avoiding threats [45]. This perspective leads one to speculate that the functional roles of specific regions of prefrontal cortex may not be defined by specific computational tasks (e.g., a general working memory), but by the needs of specific behavioral strategies (e.g., maintaining a link between the action of walking and the predicted consequence of making the apple reachable).

Second, this view is incompatible with the widespread distinction between eminently cognitive or executive areas (e.g., prefrontal cortex), where the decision happens, and movement control areas (e.g., motor cortex) as the 'slave systems' that execute these decisions. Here, instead, different brain areas process in parallel various aspects of a decision. Not only (higher) decisions about action plans can bias immediate action selection (where action is not just the next movement), the result of a competition between affordances at any (low) point of the hierarchy can influence the choice at (higher) hierarchical levels, by creating a residual prediction error that cannot be minimized without changing a long-term plan. This architecture is configured for embodied decisions – the hallmark of adaptive behavior – in which situated aspects of the choice (e.g., affordances, motor costs scaled by the current fatigue level of the animal, physical distance from potential targets) must be part and parcel of the decision. Situational aspects of decisions may be continuously processed in 'lower' cortico-subcortical loops (involving sensorimotor brain areas), which are engaged in the decision process through reciprocal top-down and bottom-up exchanges with 'higher' cortico-subcortical loops.

Third, in this perspective all behavior is controlled toward specific goals (while leaving place for routinized behavior [37]) – contrary to the idea of specialized computational mechanisms for 'executive' or 'cognitive control' that can supersede default mechanisms of stimulus–response [46]. All processes engendered by lower or higher cortico-subcortical loops are controlled, and the main difference among them is a temporal one – that is, at higher levels, courses of actions are selected that last (and are controlled/monitored) for prolonged periods. The reason why executive functions are usually associated to prefrontal cortex might not be that they require separate computations, as commonly assumed, but that prefrontal cortex-based control loops last longer. While 'executive' (e.g., monitoring, inhibition) aspects of action control might be ubiquitous across the hierarchy, it might be that they become more apparent (or can be measured more effectively) only on the longer timescales of anterior prefrontal cortex-based control loops.

Fourth, this view suggests that higher cognitive processes do not need a separate neural substrate but might largely reuse the neuronal resources and computations of situated control, yet in an 'internally generated' or detached mode. At least since Piaget [47] it has been argued that cognitive operations can be based on an 'internalized', covert (or off-line) reuse of the brain mechanisms supporting overt sensorimotor loops [48–51]. In turn, these might be largely based on internally generated (not stimulus-bound) brain dynamics that use the same internal generative models and feedback processes, but when part of the feedback is self-generated and mediated by covert internal modeling dynamics (e.g., through cerebellar loops [52–54]) rather than overt sensorimotor engagement. Consider, for example, motor simulation processes in action prediction [55] or – in a different context – hippocampal 'replays' of (time-compressed) spatial trajectories, which are implied in navigational planning and memory consolidation [56–58]. These are all examples in which resources used in situated interaction can be temporarily disengaged from the demands of

action control and reused for cognitive operations [48,49]. In this perspective, higher cognitive abilities might depend on a balance between covert and overt processes using largely the same resources. Internally generated brain dynamics, which operate faster than the action–perception loop, might permit cognitive operations such as planning, deliberation, and means–ends reasoning to be ahead of time and provide them the necessary autonomy (detachment) from the current situated context of the animal. At the same time, these covert operations can be put into practice (to support situated action) by re-engaging brain dynamics operating at the appropriate timescale for action–perception loops [37,56,59].

Finally, ‘thinking’ might be conceptualized as a controlled process of prediction and imagination (of action possibilities and their outcomes), which engages covertly the same resources as overt interaction [7,60], rather than stemming from specialized computational procedures independent of perception and action systems [1,61]. As an example of a ‘controlled imagination’ process, an interior designer can compare, in the mind, different possible arrangements of the furniture in a room by considering their shape, color, and size, anticipate if the clients will be satisfied or not, and keep changing (controlling) the imagined arrangements until the desired goal is achieved (the desired design of the room). Although the designer’s thought processes are temporarily detached from the overt sensorimotor loop, they might use the same mechanisms of feedback control and **forward modeling**. An empirical prediction of this ‘embodied intelligence’ view is that even in seemingly abstract thinking processes one might find the signature of situated and embodied action and, in some cases, residual aspects of overt movements.

Concluding Remarks

The architectures for hierarchical feedback control of our evolutionary ancestors were arguably well configured to solve problems of situated choice and adaptive control. While psychological thought has long assumed that complex human behavior demands a different brain architecture, we argue that neurophysiological and neuroanatomical data motivate us to abandon that assumption. Rather, the basic design principles of our ancestors’ brains are largely conserved [62]; and higher cognition abilities such as planning, cognitive control, and thinking, traditionally considered to require specialized neural and computational resources, appear to be elaborations of the same basic control loops that underlie sensorimotor behavior. For example, forms of planning may imply a covert replay of actual experience [7,38,48,57,63]. Thus, we propose that concepts of feedback control, which underlie all biological systems, also provide a viable conceptual foundation to understand general human cognition – from the control of movement to the control of internal processes, such as planning, thinking, or attention [7,53,64]. Although such proposals have been made for more than a century, often resurfacing as in recent embodied approaches to cognition [65–67], what they often lack are detailed process models that link feedback control and specific brain computations, especially for higher cognition.

In this opinion article, we have contributed toward filling this gap by proposing how principles of hierarchical feedback control might apply beyond sensorimotor behavior, exploiting prediction dynamics to address the realm of intentional action and ‘navigation in an affordance landscape’. Furthermore, we briefly discussed how these nested control hierarchies might be implemented in specific neural structures. We proposed that forward simulations can link proximal actions and distal goals by predicting future affordances that are used as (sub)goal states. In other words, living organisms can ‘recognize’ affordances that are not there, but can be created, as in the example of the ‘reachable’ apple, or of a climber planning the best sequence of holds to climb a wall. Clearly, living organisms dwelling in realistic scenarios confront a combinatorial explosion and cannot simulate or search through all future possibilities. Computational approaches to this problem usually adopt some sort of (learned) internal model and/or intermediate state values to guide the search; and some of them begin to be able to address large problem spaces, for example, Monte Carlo tree search [68,69]. Similarly, some forms of ‘cognitive search’ may use

Outstanding Questions

Neuroscience has inherited a taxonomy of brain functions from cognitive psychology (distinguishing, e.g., attention, memory, decision-making, action control). Should we instead decompose brain function according to the various mechanisms of feedback control (e.g., reference point, comparator, feedback signal, internal model)?

If cognitive skills are sophistications of control mechanisms originally developed for situated action (without major evolutionary changes in brain organization), then what are the behavioral tasks one should study in the laboratory to gain insights into the mechanisms our brain uses to solve ecologically valid problems?

We discussed ethological action maps in motor and premotor areas. Is the same functional organization recapitulated at higher hierarchical layers? For example, is the prefrontal cortex organized in terms of maps of action-to-state (e.g., move to be near something, open door to make it passable)?

Affordance competition is organized around ‘action specification’ and ‘action selection’ circuits. Is this general architecture recapitulated several times in the brain, toward more abstract categories of ‘action’ as one proceeds forward from motor to frontal areas?

What are the hierarchies most useful for hierarchical control? In theories of visual processing, hierarchies are often ones between parts and wholes. Does this apply to control hierarchies? For example, a higher level control loop may be more extended in time, but not necessarily be any less concrete than a lower level control loop.

Does the idea of ‘navigating the affordance landscape’ apply to social behavior? Can the actions of others be conceived as social affordances? Can social interaction and communication be conceptualized in terms of feedback control – for example, we control the behavior of others to achieve our (individualistic or joint) goals?

model-based methods. Whether this is true and what specific forms of internal models and model-based search biological organisms adopt are important research objectives, which require a tight collaboration between empirical and computational methods. For example, it has been proposed that sophisticated forms of prospective cognition may require specific adaptations in primates associated with the appearance of granular prefrontal cortex [45], and computational studies may help to elucidate the underlying computational principles [41,70].

The emerging view is that adaptive action control – and the problems faced by situated agents who pursue their goals in dynamic (yet structured) environments – should be considered as a central paradigm to understand human cognition. Methodologically, this suggests designing experiments that reflect conditions that are as ecologically valid as possible, as opposed to conditions designed to face subjects with problems that do not capture the fundamental challenges to which the brain has adapted [71]. Furthermore, using principles of feedback control as metaphors (or better still, as implemented computational systems) for experimental design might help to address the numerous research questions that remain open (see Outstanding Questions) and arguably lead to a much-improved view of human cognition.

Acknowledgments

The authors wish to thank Karl Friston for comments on this material, and the editor and two anonymous reviewers for constructive suggestions. G.P. is supported by the European Community's Seventh Framework Programme, grant FP7-270108 (Goal-Leaders) and by the Human Frontier Science Program, grant RGY0088/2014. P.C. is supported by a Chercheur National Award from the Fonds de Recherche du Québec – Santé and operating grants from the Canadian Institutes of Health Research and the Natural Sciences and Engineering Research Council of Canada.

References

- Thagard, P. (2014) Cognitive science. In *The Stanford Encyclopedia of Philosophy (Fall 2014 edition)* (Zalta, E.N., ed.), Standord University
- Ashby, W.R. (1952) In *Design for a Brain* John Wiley & Sons
- Wiener, N. (1948) *Cybernetics: Or Control and Communication in the Animal and the Machine*, MIT Press
- Powers, W.T. (1973) *Behavior: The Control of Perception*, Aldine
- Dewey, J. (1896) The reflex arc concept in psychology. *Psychol. Rev.* 3, 357–370
- Cisek, P. (1999) Beyond the computer metaphor: behavior as interaction. *J. Conscious. Stud.* 6, 125–142
- Pezzulo, G. and Castelfranchi, C. (2009) Thinking as the control of imagination: a conceptual framework for goal-directed systems. *Psychol. Res.* 73, 559–577
- Cannon, W.B. (1939) *The Wisdom of the Body*, Norton
- Scott, S.H. (2012) The computational and neural basis of voluntary motor control and planning. *Trends Cogn. Sci.* 16, 541–549
- Yin, H.H. (2014) Action, time and the basal ganglia. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 369, 20120473
- Shadmehr, R. et al. (2010) Error correction, sensory prediction, and adaptation in motor control. *Annu. Rev. Neurosci.* 33, 89–108
- Hinde, R.A. (1966) *Animal Behaviour: A Synthesis of Ethology and Comparative Psychology*, McGraw-Hill Book Company
- Cisek, P. (2007) Cortical mechanisms of action selection: the affordance competition hypothesis. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 362, 1585–1599
- Friston, K. et al. (2012) Active inference and agency: optimal control without cost functions. *Biol. Cybernet.* 106, 523–541
- Verschure, P. et al. (2014) The why, what, where, when and how of goal-directed choice: neuronal and computational principles. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 369, 20130483
- Tani, J. (2014) Self-organization and compositionality in cognitive brains: a neurorobotics study. *Proc. IEEE* 102, 586–605
- Botvinick, M. and Toussaint, M. (2012) Planning as inference. *Trends Cogn. Sci.* 16, 485–488
- Cisek, P. and Kalaska, J.F. (2005) Neural correlates of reaching decisions in dorsal premotor cortex: specification of multiple direction choices and final selection of action. *Neuron* 45, 801–814
- Cisek, P. and Pastor-Bernier, A. (2014) On the challenges and mechanisms of embodied decisions. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 369, 1655
- Cisek, P. (2012) Making decisions through a distributed consensus. *Curr. Opin. Neurobiol.* 22, 927–936
- Friston, K. (2005) A theory of cortical responses. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 360, 815–836
- Bar, M. (2007) The proactive brain: using analogies and associations to generate predictions. *Trends Cogn. Sci.* 11, 280–289
- Pezzulo, G. (2008) Coordinating with the future: the anticipatory nature of representation. *Minds Mach.* 18, 179–225
- Pezzulo, G. et al. (2010) When affordances climb into your mind: advantages of motor simulation in a memory task performed by novice and expert rock climbers. *Brain Cogn.* 73, 68–73
- Lepora, N.F. and Pezzulo, G. (2015) Embodied choice: how action influences perceptual decision making. *PLoS Comput. Biol.* 11, e1004110
- Araújo, D. et al. (2006) The ecological dynamics of decision making in sport. *Psychol. Sport Exercise* 7, 653–676
- Fuster, J.M. (2002) Frontal lobe and cognitive development. *J. Neurocytol.* 31, 373–385
- Shadlen, M.N. et al. (2008) Neurobiology of decision making: an intentional framework. In *Better than Conscious? Decision Making, the Human Mind, and Implications for Institutions* (Engel, C. and Singer, W., eds), pp. 71–101, MIT Press
- Graziano, M.S. (2016) Ethological action maps: a paradigm shift for the motor cortex. *Trends Cogn. Sci.* 20, 121–132
- Stepniewska, I. et al. (2011) Optical imaging in galagos reveals parietal-frontal circuits underlying motor behavior. *Proc. Natl. Acad. Sci. U.S.A.* 108, E725–E732
- Brown, A.R. and Teskey, G.C. (2014) Motor cortex is functionally organized as a set of spatially distinct representations for complex movements. *J. Neurosci.* 34, 13574–13585

32. Cisek, P. and Kalaska, J.F. (2010) Neural mechanisms for interacting with a world full of action choices. *Annu. Rev. Neurosci.* 33, 269–298
33. Fluet, M.-C. *et al.* (2010) Context-specific grasp movement representation in macaque ventral premotor cortex. *J. Neurosci.* 30, 15175–15184
34. Redgrave, P. *et al.* (1999) The basal ganglia: a vertebrate solution to the selection problem? *Neuroscience* 89, 1009–1023
35. Gruber, A.J. and McDonald, R.J. (2012) Context, emotion, and the strategic pursuit of goals: interactions among multiple brain systems controlling motivated behavior. *Front. Behav. Neurosci.* 6, 50
36. Mink, J.W. (1996) The basal ganglia: focused selection and inhibition of competing motor programs. *Prog. Neurobiol.* 50, 381–425
37. Pezzulo, G. *et al.* (2015) Active inference, homeostatic regulation and adaptive behavioural control. *Prog. Neurobiol.* 134, 17–35
38. Mendoza, G. and Merchant, H. (2014) Motor system evolution and the emergence of high cognitive functions. *Prog. Neurobiol.* 122, 73–93
39. Badre, D. and D'Esposito, M. (2009) Is the rostro-caudal axis of the frontal lobe hierarchical? *Nat. Rev. Neurosci.* 10, 659–669
40. Badre, D. *et al.* (2010) Frontal cortex and the discovery of abstract action rules. *Neuron* 66, 315–326
41. Koehlin, E. (2014) An evolutionary computational theory of prefrontal executive function in decision-making. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 369, 1655
42. Bonini, L. *et al.* (2011) Grasping neurons of monkey parietal and premotor cortices encode action goals at distinct levels of abstraction during complex action sequences. *J. Neurosci.* 31, 5876–5886
43. Kiebel, S.J. *et al.* (2008) A hierarchy of time-scales and the brain. *PLoS Comput. Biol.* 4, e1000209
44. Harrison, T.C. *et al.* (2012) Distinct cortical circuit mechanisms for complex forelimb movement and motor map topography. *Neuron* 74, 397–409
45. Passingham, R.E. and Wise, S.P. (2012) *The Neurobiology of the Prefrontal Cortex: Anatomy, Evolution, and the Origin of Insight*. ((1st edn)), Oxford University Press
46. Mansell, W. (2011) Control of perception should be operationalized as a fundamental property of the nervous system. *Top. Cogn. Sci.* 3, 257–261
47. Piaget, J. (1954) *The Construction of Reality in the Child*, Ballentine Books
48. Buzsáki, G. *et al.* (2014) Emergence of cognition from action. *Cold Spring Harb. Symp. Quant. Biol.* 79, 41–50
49. Pezzulo, G. (2011) Grounding procedural and declarative knowledge in sensorimotor anticipation. *Mind Lang.* 26, 78–114
50. Grush, R. (2004) The emulation theory of representation: motor control, imagery, and perception. *Behav. Brain Sci.* 27, 377–396
51. Hesslow, G. (2012) The current status of the simulation theory of cognition. *Brain Res.* 1428, 71–79
52. Koziol, L.F. *et al.* (2014) Consensus paper: the cerebellum's role in movement and cognition. *Cerebellum* 13, 151–177
53. Ito, M. (1993) Movement and thought: identical control mechanisms by the cerebellum. *Trends Neurosci.* 16, 448–450
54. Schmahmann, J.D. (2004) Disorders of the cerebellum: ataxia, dysmetria of thought, and the cerebellar cognitive affective syndrome. *J. Neuropsychiatry* 16, 367–378
55. Jeannerod, M. (2001) Neural simulation of action: a unifying mechanism for motor cognition. *Neuroimage* 14, S103–S109
56. Pezzulo, G. *et al.* (2014) Internally generated sequences in learning and executing goal-directed behavior. *Trends Cogn. Sci.* 18, 647–657
57. Pfeiffer, B.E. and Foster, D.J. (2013) Hippocampal place-cell sequences depict future paths to remembered goals. *Nature* 497, 74–79
58. van der Meer, M. *et al.* (2012) Information processing in decision-making systems. *Neuroscientist* 18, 342–359
59. Pezzulo, G. (2012) An active inference view of cognitive control. *Front. Psychol.* 3, 478
60. Barkley, R.A. (2001) The executive functions and self-regulation: an evolutionary neuropsychological perspective. *Neuropsychol. Rev.* 11, 1–29
61. Miller, G.A. (2003) The cognitive revolution: a historical perspective. *Trends Cogn. Sci.* 7, 141–144
62. Krubitzer, L. (2007) The magnificent compromise: cortical field evolution in mammals. *Neuron* vol. 56, 201–208
63. Fuster, J.M. (1997) *The Prefrontal Cortex: Anatomy Physiology, and Neuropsychology of the Frontal Lobe*, Lippincott-Raven
64. Graziano, M.S.A. and Webb, T.W. (2015) The attention schema theory: a mechanistic account of subjective awareness. *Front. Psychol.* 6, 500
65. Barsalou, L.W. (1999) Perceptual symbol systems. *Behav. Brain Sci.* 22, 577–600
66. Pezzulo, G. *et al.* (2011) The mechanics of embodiment: a dialogue on embodiment and computational modeling. *Front. Cogn.* 2, 1–21
67. Glenberg, A.M. *et al.* (2013) From the revolution to embodiment: 25 years of cognitive psychology. *Perspect. Psychol. Sci.* 8, 573–585
68. Sutton, R.S. and Barto, A.G. (1998) *Reinforcement Learning: An Introduction*, MIT Press
69. Gelly, S. *et al.* (2012) The grand challenge of computer Go: Monte Carlo tree search and extensions. *Commun. ACM* 55, 106–113
70. Stoianov, I. *et al.* (2016) Prefrontal goal codes emerge as latent states in probabilistic value learning. *J. Cogn. Neurosci.* 28, 140–157
71. Pearson, J.M. *et al.* (2014) Decision making: the neuroethological turn. *Neuron* 82, 950–965
72. Gibson, J.J. (1979) *The Ecological Approach to Visual Perception*, Lawrence Erlbaum Associates
73. Ballard, D. (1991) Animate vision. *Artificial Intelligence* 48, 1–27
74. Friston, K. *et al.* (2012) Perceptions as hypotheses: saccades as experiments. *Front. Psychol.* 3, 151
75. Friston, K. (2010) The free-energy principle: a unified brain theory? *Nat. Rev. Neurosci.* 11, 127–138
76. Feldman, H. and Friston, K.J. (2010) Attention, uncertainty, and free-energy. *Front. Hum. Neurosci.* 4, 215
77. Friston, K. *et al.* (2015) Active inference and epistemic value. *Cogn. Neurosci.* 6, 187–214
78. Seth, A.K. (2015) The cybernetic Bayesian brain from interoceptive inference to sensorimotor contingencies. In *Open MIND* (Metzinger, T.K. and Windt, J.M., eds), pp. 1–24, MIND Group
79. Conant, R.C. and Ashby, W.R. (1970) Every good regulator of a system must be a model of that system. *Int. J. Systems Sci.* 1, 89–97
80. Clark, A. (2015) *Surfing Uncertainty: Prediction, Action, and the Embodied Mind*, Oxford University Press