

Article

Model-Based Approaches to Active Perception and Control

Giovanni Pezzulo *, Francesco Donnarumma, Pierpaolo Iodice, Domenico Maisto and Ivilin Stoianov

Institute of Cognitive Sciences and Technologies, National Research Council, Via S. Martino della Battaglia 44, 00185 Rome, Italy; francesco.donnarumma@istc.cnr.it (F.D.); pierpaolo.iodice@istc.cnr.it (P.I.); domenico.maisto@na.icar.cnr.it (D.M.); ivilin.stoianov@istc.cnr.it (I.S.)

* Correspondence: giovanni.pezzulo@istc.cnr.it; Tel.: +39-06-44595206; Fax: +39-06-44595243

Received: 3 May 2017; Accepted: 7 June 2017; Published: 9 June 2017

Abstract: There is an on-going debate in cognitive (neuro) science and philosophy between classical cognitive theory and embodied, embedded, extended, and enactive (“4-Es”) views of cognition—a family of theories that emphasize the role of the body in cognition and the importance of brain-body-environment interaction over and above internal representation. This debate touches foundational issues, such as whether the brain internally represents the external environment, and “infers” or “computes” something. Here we focus on two (4-Es-based) criticisms to traditional cognitive theories—to the notions of passive perception and of serial information processing—and discuss alternative ways to address them, by appealing to frameworks that use, or do not use, notions of internal modelling and inference. Our analysis illustrates that: an explicitly inferential framework can capture some key aspects of embodied and enactive theories of cognition; some claims of computational and dynamical theories can be reconciled rather than seen as alternative explanations of cognitive phenomena; and some aspects of cognitive processing (e.g., detached cognitive operations, such as planning and imagination) that are sometimes puzzling to explain from enactive and non-representational perspectives can, instead, be captured nicely from the perspective that internal generative models and predictive processing mediate adaptive control loops.

Keywords: embodied cognition; active perception; active inference; model-based systems

1. Introduction

Since the inception of the “cognitive revolution”, the brain has often been conceptualized as an information-processing device that implements a series of serial transformations from (input) stimuli to (internal) representations and sometimes (output) actions—or “a machine for converting stimuli into reactions” [1]. This perspective prompts the idea that cognition consists in computing the response to stimuli (also known as computationalism) and its most important aspects—the truly cognitive processes—are placed “in between” perception and action systems [2]. Methodologically, this perspective motivates a research program that studies, for example, how external stimuli are encoded in the brain, what computational procedures operate over the resulting internal representations, and how these computations implement various cognitive functions, such as perceptual categorization, economic choice, and reasoning.

Many researchers have, however, sidestepped this classical computationalism to embrace various forms of embodied, embedded, extended, and/or enactive cognition (“4-Es” theories). An extensive review of this diverse cluster of theories (sometimes called “4-Es”) is beyond the scope of this article; see [3–8]. Here, it suffices to say that these theories challenge in various ways some

central constructs of traditional cognitive science—including often the notions of computation and of internal representation—and propose (for example) that cognition is shaped by our bodies, extends beyond the brain, and encompasses brain-body-environment dynamics. While the importance of these challenges is increasingly recognized, there is still considerable debate on their effects on cognitive theory; for instance, whether “4-Es” theories are alternative to traditional cognitive theory or if, instead, the latter can (and should) be amended to accommodate aspects of the former; whether central notions of traditional theories such as computation and internal representation are still desirable or need to be re-conceptualized or abandoned; and what a “process model” of embodied or enactive cognition would look like and in what sense it would be different from (or have more explanatory power than) traditional cognitive models.

To address these questions, in the following section we focus on two key criticisms that are levelled at traditional cognitive theories by proponents of 4-E approaches to cognition. The first criticism regards passive perception—or the idea that perception consists in a largely passive and bottom-up transduction of external stimuli into neuronal representations. The alternative proposal would be that perceptual processing should be conceptualized in terms of an active (or interactive) framework in which sensory and motor processes form a closed loop, in agreement with the tenets of pragmatists [9–11]—perhaps in such a way to render internal representation superfluous [12,13]. The second criticism is a sort of extension of the first criticism, but goes beyond perception and touches the notion of serial (perception-representation-action) information processing and the ensuing conception of intentional action as a staged process [14]. The alternative proposal would be that intentional action is better described in terms of a control process than as a serial transduction from perceptual states to internal representations, and then actions. This second criticism exemplifies an action- or control-oriented view of cognition, according to which the primary role of the brain is guiding interaction with the environment rather than, for example, representing or understanding the world, *per se* [15,16].

Next, we will discuss alternative ways that have been advanced to address these criticisms, focusing on various proposals that are more “deflationary” (e.g., those that propose to abandon the notions of internal modelling and/or inference *tout court*) or more “conciliatory” (e.g., those that propose that traditional cognitive constructs such as internal models can be amended to better address active perception and control-oriented views of cognition). Our analysis will reveal that: (1) The two above criticisms to traditional cognitive theory are valid and relevant; (2) However, some aspects of these criticisms are often conflated and need to be teased apart; for example, the notion of active perception does not automatically entail a non-inferential or an ecological perspective [12]; (3) There are ways to incorporate the two criticisms within a family of models that use the notions of internal model and inference [17–22]; (4) The alternative formalisms (e.g., with or without internal models) have different features, powers, and limitations. For example, model-based solutions seem more suited to address the problem of detached cognition—or how living organisms can temporarily detach from the here-and-now, to implement (for example) future-oriented forms of cognition [21]; (5) The alternative formalisms have different theoretical implications, too; most notably, concerning the notion of an internal representation. Understanding the characteristics of alternative formalisms (e.g., with or without internal models) may help finessing embodied and enactive views of cognition, as well as assessing their relative theoretical and empirical merits.

2. Two Criticisms of “4-Es” Theories to Traditional Cognitive Theory

2.1. A Critique of Passive Perception

One domain where it is easier to exemplify the differences between contrasting theoretical perspectives is perceptual processing. Traditional computational theories assume that perception—including social perception—consists in the transduction from external (environmental) states to internal (neuronal) states, which successively act as internal representations of the external events and can be internally manipulated with computational procedures. Most traditional theories of perception are passive and input-dominated, in the sense

that they give prominence to the (bottom-up or feed-forward) flow of information from the sensory periphery and assume that perception, e.g., object perception, is achieved by progressively combining object features of increasing complexity at increasingly higher levels of the visual hierarchy, see [23] for a review.

A more recent perspective is that, to perceive a scene, the brain predicts it (rather than just integrating bottom-up sensory signals) and uses prediction errors (i.e., differences between top-down predictions and bottom-up sensations) to refine the initial perceptual hypotheses. More formally, in this “Bayesian brain” (or Helmholtzian) perspective, the brain instantiates an internal generative model of the causes of its sensory inputs, i.e., an internal model that describes how sensed stimuli are produced by (hidden) causes [24,25]. Such a model can encode, for example, the probability of some stimuli (e.g., I see something red and/or circular) given some hidden causes (i.e., the presence of an apple in front of me). Using a hierarchical Bayesian inference scheme—predictive coding [26]—the generative model would permit one the “hallucination” of an apple (e.g., predicting what I should see if an apple was in front of me). Furthermore, the same model can be inverted, to infer the probability of the “apple” hypothesis (or of alternative hypotheses) given my current sensory stimuli, e.g., the sight of a red colour (plus some prior information)—and the sensory evidence used for the inference can be weighted more or less, depending on its precision (or inverse uncertainty). Perception consists exactly in the “inversion” of such models, i.e., the inference of the hidden causes (the apple) given the stimuli (seeing something red) [27]. It is common to assume that, in this (simplified) generative scheme, internal hypotheses/hidden states and generative models correspond to an agent’s (neural) representations of the external world. This is implicitly assumed even in recent computational implementations of these ideas, using connectionist [28,29] or Bayesian [25] networks.

This “Bayesian brain” (or Helmholtzian) perspective resembles classical theories of perception in that it assumes brain internal models, but enriches them by emphasizing the integration of top-down and bottom-up flows of information and the generative nature of perceptual processing (and of cognitive processing more generally). Importantly, the generative/inferential scheme described above is pretty compatible with one of the most prominent embodied theories of perception, *perceptual symbol systems* (PSS) theory [30], which emphasizes the importance of situated simulations—and the re-enactment of aspects of previous experience in perceptual symbols—in the guidance of perceptual processing, prediction (as well as action), exactly with the same logic of the generative or “hallucinatory” process described under the Bayesian brain hypothesis (but note that while PSS assumes that perceptual symbols are modal or multi-modal but not amodal constructs, computational studies using generative models often leave this point unspecified).

Yet, some other embodied and enactive theories often claim (contra to traditional cognitive theories) that “perceiving” and “understanding” the environment (or other persons)—and more broadly, cognitive processing—are based on *interactive* dynamics, rather than using inferential mechanisms or internal representations. This line of thought dates back at least to the ecological approach to perception, which starts from the idea that living organisms do not internally represent the external world, but are configured in such a way to exploit an informational coupling with it, and exploit the informationally-rich ecological environment (not internal representations) to take action [12]. In a similar vein, enactive views of cognition highlight that understanding something is only achieved through interactive engagement with the entity and is, thus, action-based and not passive; similarly, social understanding is participatory rather than the (first-person) exercise of estimating or mirroring other’s mental states in one’s mind [4].

To better understand these theoretical arguments, it is useful to start from more mechanistic accounts of enactive (or interactive) perception: the theories of *sensorimotor contingencies* (SMCs) [13] and of *closed-loop perception* [31]. SMCs are contingencies between actions and ensuing sensory states (e.g., a sensation of softness given a grasping action), contingent to a given situation (e.g., the presence of a sponge). According to SMC theory, by exploiting learned SMCs, an agent is attuned with the external environment—in the sense that its motor and sensory patterns are coupled over time and become mutually interdependent while the agent grasps the sponge [9,12,32]. Perception

is, thus, the result of this progressive attunement process that unfolds over time by exploiting the agent's mastery of SMCs: it is by successfully exploiting SMCs (e.g., the grasp-softness contingencies) over time that the agent perceives an object such as a sponge, while it interacts with it. Yet another example is the perception of a *red* colour. In SMC theory, perceiving something as *red* does not depend on a static pattern of stimulation of the retina (e.g., light having a given wave-length), but the knowledge of SMCs, e.g., the ways an incoming stimulus would change when a red surface is inclined (thus changing light reflection), which is different from the way something blue or green would change under the same conditions [13]. From a convergent perspective, the perception of an object can be described as a closed-loop process that progressively "incorporates" the external object through multiple loops of *motor-sensory-motor contingencies* [31].

In sum, these theories (and others, see [33] for a review) emphasize the key features of an interactive view of perception and, most prominently, the mutual dependency between perception on action. We have mentioned above one core idea of the ecological approach to perception—that living organisms are informationally coupled to the environment and do not need to represent this information internally [12]. Yet, one possible criticism of this approach is that information in the environment can be too limited to be really useful for cognitive tasks such as recognizing an object or catching it. The two theories of *sensorimotor contingencies* and *closed-loop perception* clarify that the agent's actions create (at least part of) the task-relevant information and contribute not only to the success of the task at hand (e.g., catching a flying ball), but also keep perception stable and reliable during the task. They also imply that perception and understanding are interactive processes that require (inter)action rather than being just the presupposition or antecedent of an action (e.g., first recognize the sponge, then select a grasping action), as more often assumed by classical information-processing theories [14].

These theories have important consequences for neurophysiology, too. Perhaps the most important consequence of incorporating action components in perception is that such theories do not see perception as a property of (a fixed pattern of) stimuli, thus, providing a rationale for the dynamical and action-dependent character of sensory stimulations. For example, SMC theory explains why there is not a one-to-one relation between the pattern of sensory stimuli and perception (e.g., light with the same wave-length reflected on different surfaces is perceived as having different colours) and the theory of closed-loop perception explains how active sensing and epistemic behaviour, such as whisker movements in rodents steer dynamical neuronal patterns, which are key to perception, rather than impairing it [34].

The strengths of SMC and closed-loop theories of perception are increasingly well recognized. However, these theories entail (at least) two kinds of criticisms of traditional, passive views of perception, which are often conflated in the literature but need to be teased apart. The first criticism is that perception and understanding (but also, more generally, cognitive processing) are not passive processes, but have an action (or interaction) component. The second criticism is that perceptual (and cognitive) processing does not use internal models and/or inferential processes. This second criticism is related to "direct perception" theories and the idea that ecological information is self-sufficient to perform even complex tasks [12] and, thus, links more directly to anti-representationalism. These two criticisms can be kept separate; as we will see in the next section, there exist model-based solutions to the same problems of active perception highlighted here [35–37]. However, before discussing this point, in the remainder of this section we discuss a second criticism of traditional cognitive theories: a critique of serial information processing.

2.2. A Critique of Serial Information Processing

A second domain that allows us to compare different theoretical perspectives is intentional action, broadly construed (i.e., including relatively simpler actions such as grasping an object and relatively more complex actions, such as planning and then doing a daily trip; and considering both deliberation and action performance). The dominant scheme for intentional action in traditional cognitive theory is a serial transformation from sensory inputs to internal representations (possibly amodal representations, which can be internally manipulated using combinatorial rules, to derive

and select an action plan); and successively the overt execution of each element of the plan in a sequence [14]. In turn, action performance can be fractionated into relatively simple behavioural routines, which do not require attention, and more demanding executive processes [38].

Although the above description is necessarily simplified, it captures some of the essential elements that seem problematic from embodied or enactive perspectives. These include the fact that perception (or estimation), decision-making, and action planning (and/or execution) are implemented in serial and separate stages and use largely distinct neuronal processes (e.g., the neuronal resources implied in decision-making are not the same as those implied in perceptual or action processes). The presence of serial stages ‘breaks’ the action-perception loop that, as we have discussed above, is essential in ecological and enactive theories of perception and action, but also more generally in some theories of higher cognition [39–41]. Another criticism of the serial stages view derives from evolutionary arguments, and the recognition that our cognitive architecture derives from more primitive mechanisms permitting animals to face (often dangerous) situated choices and for this it could follow a different design: one in which all information is continuously integrated to specify and select multiple actions in parallel until one can be reliably selected—and in which decision and action (planning) are intrinsically linked. Neurobiological implementations of this idea, such as the “affordance competition” hypothesis [16,42,43] and the “intentional” framework of information processing [44] have received considerable empirical support. These ideas can be stretched even further, by considering that the serial stage idea is intrinsically flawed due to the backward influence from action to decision processes, leading to an “embodied choice” framework [45].

Finally, and importantly for our analysis, the “serial stage” idea suggests (although it does not imply) that the most relevant part of cognitive processing is the central (representation and decision) part, which is far removed from perceptual and action components—which has been sometimes called the “meat” of the cognitive (perception-representation-action) sandwich [2]. An alternative proposal that is more akin to the 4-Es camp starts from the idea that the brain is a control system whose main goal is guiding interaction with the environment rather than, for example, representing or understanding the world, *per se* [16,41,46]. This “control view” of brain and cognition has its historical roots in cybernetics [47–49], which emphasized the importance of studying control dynamics and feedback mechanisms in living organisms. It takes seriously the evolutionary arguments that our cognitive abilities were originally developed to make rapid, adaptive choices in situated contexts as part of our interaction with objects and other animals, not to solve lab tasks [50–52]; and even the most sophisticated (higher) cognitive abilities may be better seen as elaborations of the basic cognitive architecture of our early evolutionary ancestors. Therefore, this view immediately prompts a pragmatic (or action-centred) perspective on brain and cognition [15,53], which shifts the focus of investigation from “what happens in the brain in between the reception of a stimulus and the computation of a response?” to “how can the brain guide adaptive (inter)action?” and “can we trace back our sophisticated cognitive abilities to (action-perception) control loops or their elaborations?”.

Over the years, the “control view” of the brain and cognition has resurfaced many times and enjoyed some success in specific areas of psychology and neuroscience, such as movement neuroscience [54] and (active) perception [31]. The control view, which emphasizes control over and above representation and/or prediction, and which emphasizes closed perception-action loops and a tight coupling between agents and environments, seems particularly appealing from a non-representationalist, enactive perspective. Indeed, some of the most popular arguments for non-representationalist cognition are based on examples from control theory, such as the Watt governor, which is able to trigger complex patterns of behaviour without making use of internal models or internal representations [55]. A number of simulation studies have shown that interesting patterns of behaviour emerge by coupling a relatively simple agent controller (implemented sometimes as a feed-forward neural network whose weights are learned over time or evolved genetically, or as simple dynamical systems) with an environment or other simple agent controllers; sometimes, despite its simplicity, the agents can solve tasks for which traditional

cognitive theories would have supposed the necessity of categorization modules [56–59]. It has been variously proposed that we should take more seriously the possibility that even the far more complex patterns of behaviour that we observe in advanced animals, such as humans, may ultimately stem from the same class of dynamical solutions to control problems rather than from inferential processes operating on internal representations. At the same time, the control schemes that are nowadays most used in computational neuroscience, such as optimal control [60,61] and active inference [62–66], include two notions that are at least “suspect” from the 4-Es perspective. The first is the notion of an internal model—following the “good regulator” theorem, that “every good regulator of a system must be a model of that system” [67]. The second is the notion of (Bayesian) inference—following the demonstration that control problems can be cast equivalently as inference problems [61,68–72].

Thus, here we are in a similar situation for both criticisms: there exist alternative architectural solutions to both (active perception and control) problems identified by 4-Es theories, which are based on different theoretical assumptions—including more prominently the usage of internal models and inferential mechanisms. In the next two sections, we review more extensively model-based solutions to active perception and control problems (Section 4), and then compare solutions that use and do not use the notion of internal models (Section 5).

3. Model-Based Solutions to Active Perception and Control Problems

3.1. Active Perception from a Model-Based Perspective

While SMC and closed-loop theories are not usually associated with the notion of internal modelling, one can easily formalize SMCs in terms of internal models that encode (probabilistic) relations between series of actions and sensations over time and which permit one, for example, to predict the sensory consequence of an action pattern [73–75] (see also [66,76]). In action control, internal models have been long used for prediction of action consequences—appealing to the notion of a *forward model* [77]—but these ideas can be extended to cover the notions of SMCs and active perception.

Let us consider again the idea that grasping (and perceiving) an apple uses learned SMCs. Interactive success using an apple-grasping SMC, or a series of interconnected SMCs, indicates that the action presuppositions were true, e.g., there was indeed an apple to grasp [17,19,20,41,78]. Hence, the success of an interactive pattern (e.g., a grasping routine) can have epistemic or perceptual functions, as assumed in SMC theory [13]. When cast within a theory of internal models, one can imagine that an agent maintains a set of internal models, which encode different SMCs (or sets of SMCs) that are specialized (or parameterized) to interact with different objects, say grasping an apple versus a cup. In standard (passive) views of perception, one would recognize the object (apple) first and then trigger an apple-grasping routine. In active views of perception, instead, executing a sensorimotor routine is part and parcel of perceptual processing, as the success of the interactive process contributes to “perceiving” the object. In a model-based perspective, it would be quite natural to explicitly associate “beliefs” (intended in the technical sense of probability theory, e.g., probability distributions) to the different possibilities—e.g., about the presence of an apple or a cup—and update them depending on the interactive success of its internal model, e.g., by considering how much sensory prediction error the competing models generate over time. These beliefs can be used in many ways, for action selection (e.g., selecting the model that generates the least error over time), learning (e.g., to set an adaptive learning rate for the models) but also as explicit measures of an agent’s knowledge. What is interesting is that the explicit (belief) estimate would also have an associated confidence (inverse uncertainty) value, which effectively measures how supported the belief is, and which may have important psychological counterparts, e.g., a “feeling of knowing” whether an object is present or not, and whether or not one is executing the right action [79].

One may also enrich the above “model selection” idea with an explicit hypothesis-testing scheme, by considering an internal model for grasping an apple as a “hypothesis” (e.g., that there is

an apple) that competes with other internal models, or different parametrizations of the same internal model, which encode alternative hypotheses (e.g., that there is a glass or a cup). In this perspective, an action or a sequence of actions such as a grasp performed with the whole hand (*power grasp*) would play the (active) role of an “experiment” that updates the beliefs about the alternative hypotheses; and the experiment can be constructed in such a way that it (for example) disambiguates the alternative hypotheses in the best way [33]. Hence, belief updating would not stem from passively collecting motor-sensory statistics, but from a more active, hypothesis testing process—which constitutes an action-based metaphor for saccadic control [35,36] and haptic exploration [37].

These examples illustrate that one can cast an interactive (rather than a passive) view of perception using the notion of internal (generative) models in a way that is analogous to SMC theories—in the sense that the models primarily encode the statistics of motor and sensory events, conditioned on the current context. This view is compatible with the Helmholtzian perspective in that it includes internal models and inferential processes (roughly, of surprise minimization). At the same time, this view introduces two novel elements that make perceptual processing interactive. First, the generative models that are used for perceptual processing encode statistical regularities (contingencies) between action and sensory streams, not just the statistics of sensory streams as is more often assumed in traditional perceptual models. Second, there is an explicitly active component in perceptual processing, in that the agent selects the next action (partly) for perceptual and epistemic reasons, e.g., to disambiguate amongst perceptual hypotheses, to keep the stimulus constant or de-noise it, etc. (see also [33]).

3.2. Beyond Active Perception: Active Inference and the Embodied Nature of Inference

The framework of active inference goes beyond the mere recognition of a role of action in perception, and proposes that action is part and parcel of inference, in that it contributes to reduce prediction error (that in this framework is achieved by minimizing a *free energy* term [18]) in the same way model updates do [64]. To understand why this is the case, let us consider an agent who believes that there is an apple in its hand, and faces a significant prediction error (because there is no apple in its hand). Generally speaking, the agent has two ways to reduce this prediction error: it can either revise its hypothesis about grasping an apple (perception) or change the world and grasp an apple (action). In other words, both hypothesis revision and action make the world more similar to our predictions (hence, decreasing prediction errors)—although they operate in two opposite “directions of fit”: by updating the model to fit the world, or by changing the world to fit the model. Seen in this way, Active Inference is simply the extension of a predictive coding architecture with motor reflexes [18,64]. Casting perception and action in terms of the same prediction error (or free energy) minimization scheme may seem *prima facie* counterintuitive, but it makes the inferential architecture more integrated and the inferential process more “embodied”—in the sense that inference (and the model itself) spans across brain and body/action dynamics rather than being purely “internal”.

One key question then becomes how the agent “decides” (for example) to revise the apple-in-my-hand hypothesis or to grasp an apple. This problem is resolved in terms of a hierarchical (Bayesian) scheme, which weights the “strength” (or more formally, the precision) of priors at higher hierarchical levels, which play the role of goals (e.g., I want an apple) and of prediction errors coming from lower hierarchical levels: when the former dominates the latter, the architecture triggers a cascade of predictions (including perceptual, proprioceptive and interoceptive predictions about the apple-in-my-hand) that, in turn, guide perceptual processing and (through the minimization of proprioceptive prediction error) enslave action until the apple is really in the agent’s hand, or a change of mind occurs. This latter concept nicely extends to planning sequences of actions, by considering predictions about entire behavioural policies (e.g., reaching one of the different places where I can secure an apple or obtain cues about where to find apples) as opposed to considering only the current or the immediate next grasping action [35,80–86]. A related

body of work emphasizes proactive aspects of brain dynamics as well as interoceptive and bodily processes, such as the mobilization of resources in anticipation of future needs [65,87–89].

These simple examples illustrate that Active Inference realizes a synthesis between the ideas that “the brain is for prediction” (aka predictive processing) and that “the brain is for action” (aka control view); see [64–66] for more details. Given that it can simultaneously address domains of perception, action, and interoception [65,66,90], as well as of individual and social cognition [91–97] within a unitary theoretical framework, Active Inference has recently gained considerable prominence in computational and systems neuroscience [18], as well as in philosophy—although in the former field it is more commonly referred to as a “Free Energy Principle (FEP)” framework [18], while in the latter field it is more commonly referred to as a “Predictive Processing (PP)” [62,98–100] and/or “prediction error minimization (PEM)” [99] framework (henceforth, we will use these terms in an interchangeable way). Interestingly, the PP framework includes elements of both computational theories of cognition (e.g., inferential processes and internal models) and embodied and enactive theories of cognition (e.g., the contribution of action to cognitive processing and the importance of self-organizing processes and autopoiesis [101,102]) and it has been advocated by proponents of both representational and internalist theories [99] and ecological perspectives [103]; see also [98,104]. This points to the possibility of a useful convergence between theoretical approaches that are seen as mutually exclusive, but (despite their differences) have many elements in common: see Section 5.

4. Comparing Alternative Conceptualizations of Active Perception and Control

Our discussion so far exemplifies the fact that it is possible to characterize two key notions of 4-Es theories—active perception and control—using different approaches, some of which use model-based and inferential processes, and some of which dispense from using them—the latter being considered more “deflationist” compared to traditional cognitive theory. Yet, the problem of assessing the relative merits of these and other alternative proposals remains open.

Comparing different approaches is difficult given that they are often formulated at different levels of detail, e.g., at the theoretical level or as computationally implemented models. To mitigate this problem, we focused on examples for which detailed computational models have been proposed in the literature (see the above discussion). However, the mere existence of implemented computational or formal models does not solve all of the problems. Another problem in comparing different approaches is the usage of different terminologies or formal approaches. Indeed, it is possible that formal solutions that are commonly considered to be alternative are in fact mathematically equivalent—as in the case of the equivalence between control and inference problems [61]. A similar problem seems to exist when comparing computational and dynamical systems perspectives on cognitive phenomena—two approaches that are often considered to be mutually exclusive, especially by proponents of dynamical systems perspectives who support anti-representationalism [55]. As noticed by Botvinick ([105], p. 81) “*The message is that one must choose: One may either use differential equations to explain phenomena, or one may appeal to representation.*” However, this problem might be more apparent than real, at least in some cases. Botvinick [105] continues as follows:

“This strikes me as a false dilemma. As an illustration of how representation and dynamics can peacefully coexist, one may consider recent computational accounts of perceptual decision-making. Here, we find models that can be understood as implementing statistical procedures, computing the likelihood ratio of opposing hypotheses (read: representations), or with equal immediacy as systems of differential equations”.

and refers to two specific examples of models that have these characteristics [106,107]. Ahissar and Kleinfeld [34] (p. 53) provide another interesting illustration of duality between homeostatic (or control-theoretic) and computational perspectives:

“The operation of neuronal closed loops at various levels can be considered from either homeostatic or computational points of view. All closed loops have set-points at which the values of their state

variables are stable. Thus, feedback loops provide a mechanism for maintaining neuronal variables within a particular range of values. This can be termed a homeostatic function. On the other hand, since the feedback loops compute changes in the state variables to counteract changes in the external world, the change in state variables constitutes a representation of change in the outside world. As an example, we consider Wiener's description of the sensorimotor control of a stick with one finger. The state variables are the angle of the stick and the position (angle and pivot location) of the finger. When the stick leaves a set-point as a result of a change in local air pressure, the sensorimotor system will converge to a new set-point in which the position of the finger is different. The end result, from the homeostatic point of view, is that equilibrium is re-established. From the computational point of view, the new set-point is an internal representation of the new conditions, e.g., the new local air pressure, in the external world. (We note that the representation of perturbation by state variables may be dimensionally under- or over-determined and possibly not unique.) This internal representation is 'computed' by the closed-loop mechanism".

A similar case can be made for the duality between inference and agent-environment synchrony, if one considers the illustration of how Active Inference principles can be used to model dynamical or autopoietic systems [101]. In this example, the active inference framework is used to illustrate the emergence of (simplified forms of) "life" and self-organization from a sort of "primordial soup" in which particles having Newtonian and electrochemical dynamics interact over time and can self-organize. Technically speaking, the active inference agent has an internal model, whose internal states are kept separated from external (environmental) states by a so-called Markov blanket (a statistical construct that captures conditional independencies between nodes). By repeated interactions with the external environment, the agent's internal states "infer" the dynamics of the external environment. However, the very same process can be described both in terms of statistical inference (and free energy minimization) and of (generalized) synchrony between two dynamical systems—agent and environment—that is made possible by their continuous coupling.

This short illustration of the difficulties of comparing different approaches—and the possible errors one can incur if one naively maps different formal languages to different theories—is meant to suggest caution in the analysis, but not that all theories are equal. Rather, we suggest that different families of approaches (e.g., with or without internal models) to the problems we have focused on—active perception and control—have some elements in common but are different in other respects. In the rest of this section, we will discuss some of the theoretical implications of using, or not using, model-based and inferential approaches to problems of active perception and control, for what concerns the notion of internal representation and the way we conceptualize brain architecture.

4.1. Model-Based Approaches to Active Perception and Control: Conceptual Implications

As we have seen, it is possible to address related problems (e.g., active perception and control) and even appeal to similar constructs (e.g., sensorimotor contingencies) using a range of different architectural solutions. For example, one can cast active perception within a family of solutions rooted in dynamical systems theory (e.g., [13,31]) or, alternatively, within a family of solutions rooted in model-based and inferential computations (e.g., [35,36]). Both approaches implement perception as an interactive process, in which action dynamics (e.g., the routines for grasping an apple) probe whether the "presuppositions for action" (e.g., the presence of an apple) hold or not—hence, sensory and motor processes form a closed loop and not successive stages, in agreement with the tenets of pragmatists [9–11].

However, the appeal to similar pragmatist principles hides the theoretical differences between the two approaches. Enactive theories of cognition including SMC theory [13] tend to assume that perceptual processing depends on an implicit mastery of the rules of how sensations change depending on actions; thus, appealing to the notion of internal representation is not necessary or is even misleading, as it would divert the attention from the most important (interactive) components that make SMCs useful. In other words, enacting an apple-related action-perception loop is

sufficient for perception and a successful grasping: no internal apple representation is needed for this. This would make redundant the usage of notions such as “beliefs” or “hidden states” that model-based systems associate with perceptual hypotheses such as the presence of an apple or a cup, and of the notion of “inference” that often refers to maximizing the likelihood of (or minimizing surprise about) perceptual hypotheses. More specifically, one can argue that these notions would not be particularly problematic if used as technical constructs—as constituents of an adaptive agent architecture—but would instead become problematic if one assigned them a theoretical dignity, e.g., if one equated “belief” or “hidden state” to internal representation (it is, however, worth reminding that theories of ecological perception [12] would not accept the notion of “hidden states”—even if intended in a minimalistic sense—because they are not required under the assumption that perception is “direct” and sensory stimuli are self-sufficient for it, making the mediation of internal or hidden states unnecessary).

This point leads us to the question on whether a model-based approach like PP invites (or implies) a representational interpretation—an issue that is currently debated in philosophy, with contrasting proposals that highlight the relations between PP and various (e.g., internalist, externalist, or non-representationalist) epistemological perspectives [62,98,99,103,104,108,109]. This diversity of opinions recapitulates within the field of PP theories some long-lasting debates about the nature and/or the existence of representations. Our contribution to this debate is to review various existing examples of Active Inference agents, and discuss in which senses they may lend themselves to representationalist or anti-representationalist interpretations—with the obvious caveat that these interpretations may diverge, depending on the specific definition of representation.

Some theories of representation emphasize some form of correspondence or (exploitable) structural similarity between a vehicle and what it represents [108,110,111]. In this vein, a test for representation would be assessing the (structural) similarity between an agent’s internal generative model and/or hidden states (aka the vehicles) and the “true” environmental dynamics—or “generative process” in Active Inference parlance—which is unknown to the agent. When representation is conceived in this way, it seems natural to assign a representational status to (hidden) states within the agent’s Markov blanket, and to notice that the similarity between generative model and generative process is a guarantee for a “good regulator” [67]. In keeping, most implemented Active Inference agents have internal generative models that are very similar to the external generative process, and sometimes almost the same, see e.g., [82,84,112]. However, this is often done for practical purposes; and it is not necessary to assume a too strong (or naive) idea of similarity according to which internal models are necessarily copies of (or mirror) the external generative process.

In fact, in Active Inference systems generative models and generative processes can diverge in various ways, and for various reasons. The most obvious reason is that internal models are subject to imperfect learning procedures, whose objective is ultimately affording accurate control and goal-achievement or, in other words, mediating adaptive action, or permitting the agent to reciprocate external stimuli with appropriate adaptive actions in order to keep its internal variables within acceptable ranges (and minimize its free energy). Intuitively, given that biological agents have limited resources, and their ultimate goal is interacting successfully within their ecological niche, the “content” of their models will be biased by utilitarian considerations, with resources assigned to coding relevant aspects only, as evident (for example) in the fact that different animals perceive broader or narrower colour spectra [113]. Several learning procedures also have the objective to compress information, that is, to integrate in the models only the minimal amount of information necessary to solve a specific task [114]. One can reframe all these ideas within a more formal model comparison procedure, which is part and parcel of free energy minimization, and consider that when a simpler model (e.g., one that includes fewer variables) affords good control, then it may be privileged compared to a more complex model, which may putatively represent the environment more faithfully [81,115]. This might imply that, for example, an agent’s model may fail to encode differentially two external states or contexts, if they afford the same policy; and in the

long run, even a less discriminative model such as one that assumes that “all cows are black (at night)” can be privileged. If one additionally considers that Active Inference requires the active suppression of expected sensory consequences of actions to trigger movement, and it often affords a sort of optimism bias [116], it becomes evident that neither the agent’s generative model has to be identical to the generative process, nor the agent’s current beliefs (or inferred states) have to be necessarily aligned to external states. This is because, in Active inference, control demands have priority over the rest of inference. To what extent the above examples are compatible with a representational view that highlights some form of correspondence or structural similarity between a vehicle and what it represents remains a matter for debate [108,110,111].

The topic becomes even more controversial if one considers a slightly different way to construct the generative models for Active Inference, which appeals more directly to the notions of SMCs [13] and motor-sensory-motor contingencies [31]. For example, an agent’s generative model can be composed of a simple dynamical system [117,118] (e.g., a pendulum) that guides the active sampling of information, in analogy to rodent whisking behaviour. In this example, the pendulum may jointly support the control of a simple whisker-like sensor and the prediction of a sensory event following its protraction (with a certain amplitude)—or a sensorimotor contingency between whisker protraction and the receipt of sensory stimuli. Such mechanism would be sufficient to solve tasks such as the tactile discrimination or localization of some objects (e.g., walls versus open arenas) or distance discrimination tasks [119,120]. A peculiarity of this model is that pendulum would not be considered a model of the external object (e.g., a wall), but a model of the way an agent interacts with the environment or samples its inputs. Given its similarity with SMC theory, one can consider that the generative model in this example mediates successful interactive behaviour and dynamical coupling with the external environment, rather establishing a correspondence with—or represent—it. Alternatively, one might argue that the generative model deserves a representational status, as some of its internal variables (e.g., the angle of the pendulum) are related to external variables (e.g., animal-object distance); or alternatively, because the pendulum-supported active sampling can be part of a wider inferential scheme (hypothesis testing [35,36]), in which repeated cycles of whisking behaviour support the accumulation of evidence in favour or specific hypotheses or beliefs (e.g., whether or not the animal is facing a wall), which constitute representations. To the extent that representation is defined in relation to a consistent mapping between variables inside and outside the agent’s Markov blanket, then considering the inferential system as a whole (not just the pendulum) would meet the definition.

Having said so, it is important to recognize that—as our discussion exemplifies—the mapping between internal generative models (as well as hidden states and beliefs) and external, generative processes can be sometimes complex, even in very simple computational models using PP (and plausibly, much more in biological agents). There are also cases in which some aspects of agent-environment interactions don’t need to be modelled, because they are directly embedded in the way the body works; the field of morphological computation [121] provides several examples of this form of off-loading. In this perspective, one can even re-read the “good regulator” theorem [67] and consider that a good controller needs to *be* (not necessarily to *include*) a model of a system—hence, bodily and morphological processes can be part and parcel of the model. In sum, there are multiple ways to implement Active Inference agents and their generative models. This fact should not be surprising, as the same framework has been used to model autonomous systems at various levels of complexity, from cells that self-organize and show the emergence of life from a primordial soup [101] or of morphogenetic processes [102], to more sophisticated agent models that engage in cognitive [122,123] or social tasks [92]—while also appealing within these cognitive models to different inferred variables, including spatial representations [80], action possibilities within an affordance landscape [16], and internal (belief) states of other agents [96]. It is then possible that a reason for dissatisfaction with the above definition of representation is that it does not account for this diversity, and the possibility that we have different epistemological attitudes towards these diverse systems.

One can also sidestep entirely the question of whether or not an agent's generative model or internal states are similar to the external generative process, and ask in which cases they can be productively assumed to play a representational function—here, in the well-known (yet not universally accepted) sense of mediating interaction and cognition off-line, or “in the absence of” the entity or object that they putatively represent [8,11,17]. Using this (quite conservative) criterion of off-line usage and decouplability (or detachment), then different model-based or PP systems (and associated notions such as “belief”, “hidden state” or “generative model”) lend themselves to representational or non-representational interpretations, depending on how they are used within the system. If one considers again the aforementioned case of the emergence of life from a primordial soup [101], the agent's model is a medium for self-organization and synchrony between two coupled dynamical systems, and despite the presence of internal (hidden) states and inferential processes, the architecture does not invite a representational interpretation; see [98,101,103] for discussions and [102] for a related example. There are other cases in which beliefs and hidden states label components of internal models, which are transiently updated during the action-perception loop for accurate control or learning, but are not used or accessed outside it. The “beliefs” that are maintained within the model-based architecture might correspond, for example, to specific parametrizations of the system (e.g., joint angles of fingers) that need to be optimized during grasping (e.g., to produce the necessary hand preshape). In this example, it would seem too strong to assign such beliefs a truly representational status—at least if one assumes that a tenet of representational content is that it can be accessed and manipulated off-line [17,41].

However, other examples of model-based systems lend themselves to a representational interpretation, which would be precluded in some (e.g., enactivist) theoretical perspectives. Consider the same apple-grasping architecture described above, in which “beliefs” about hand preshape become systematically monitored and used outside the current online action-perception loop. These beliefs could be used in parallel for grasping the apple and for updating an internal virtual (physical) simulator, which might encode the position of objects, permitting an agent to remember them or to plan/imagine grasping actions when the objects are temporarily out of its view [124,125]. Another popular example in the “motor cognition” framework is the fact that some sub-processes implied in model-based motor control, and namely the forward modelling loop that accompanies overt actions, may be at times detached and reused off-line, in a neural simulation of action that supports action planning and imagination of movement [126]. The representational aspect of this process would not consist in the inference of current (latent/hidden) states, but in the process of anticipating action consequences—for example, in the anticipated softness of grasping a sponge or the anticipated sweetness of eating an apple. This view is compatible with the idea that representation is eminently anticipatory and consists in a set of predictions including action consequences and dispositions [19,78,127]. According to this analysis, the difference between non-representational and representational processes would not depend on the mere presence of constructs such as belief or hidden state or prediction, or even on their “content” (e.g., whether they encode states that “mirror” the external environment), but how they are used: for on-line control only or also for “detached” and off-line processes. A similar distinction can be found in some theories of “action-oriented” or “embodied” representation [17,19,20,41,128] as well as in conceptualizations of the differences between states that can, or cannot, be accessed consciously [129].

Another process of model-based systems that is usually associated with internal representation (and meta-cognitive abilities) concerns confidence estimation and the monitoring of internal variables such as beliefs and hidden states [79]. There are aspects of confidence estimation that are automatically available in probabilistic model-based systems—for example, a measure of the precision (or inverse variance) of current beliefs—but some architectures also monitor other variables, such as for example the quality of current evidence or the volatility of the environment [70,130,131]. These additional parameters (or meta-parameters) have multiple uses in learning control, such as adapting learning rates or the stochasticity of policy selection [112], but may also have psychological counterparts, such as a subjective “feeling” of confidence [79].

What would these (putatively representational) confidence ratings add to processes of active perception and control? It is worth remembering that in SMC or closed-loop perception theories, enacting the right sensorimotor program is sufficient to attune with (and, thus, perceive) the object. However, there is a potentially problematic aspect of this process: how can an organism know when it has found an object (say, an apple), and decide to disengage from it to search another object (say, a knife to cut the apple into pieces)? One possible answer is that the agent does not really need to “know” anything, but only to steer the appropriate (knife-search) routine at the right time. This is certainly possible, but not trivial except, for instance, in cases where action chains are routinized and external cues are sufficient to trigger the “next” action in a sequence. In their theory of closed-loop perception, [31] recognized this “disengagement” problem and proposed a possible solution, based on an additional component: a “confidence estimator”, which essentially measures when the closed-loop process converges and thus the agent can change task. In turn, they propose that the confidence estimator may be based on an internal model of the task—thus essentially describing a solution that resembles model-based control, but requires two components: one using internal models (for confidence ratings) and one not using internal models (for control). Despite this difference, the confidence-based system would essentially keep track of the precision (or inverse uncertainty) of the belief “there is an apple”, thus playing the same role as the more standard methods of confidence estimation considered above.

In sum, we have argued that inferential systems that use internal models and hidden states do not automatically invite a representational interpretation; this is most notably the case when the internal model only mediates agent-environment coupling and there is no separate access to the internal states or dynamics for other (e.g., offline) operations. In cases like this, it is sufficient to appeal to an “actionable” (generative) model that supports successful action-perception loops, affording agent-environment coupling. Under certain circumstances, however, a representational interpretation seems more appealing, and particularly when hidden states are used for off-line processing (e.g., remembering, imagining) and accessed in other ways (e.g., for confidence judgements), which are usually considered representational processes; and when the system shows similar dynamics when operating in the two dynamical regimes, on-line (coupled to action-perception loops) and offline (decoupled from them).

The possibility for generative models to operate in a dual mode—when coupled with external dynamics and when decoupled from them—presents a challenge for the interactive and anti-representational arguments of enactivists. This is because, if cognition and meaning were constitutively interactive phenomena, then they would be lost in the detached mode, when coupling is broken. The alternative hypothesis would be that generative models acquire their “meaning” through situated interaction, but retain it even when operating in a detached mode, to support forms of “simulated interaction”, such as action planning or understanding [126].

A useful biological illustration of a dual mode of operation of brain mechanisms is the phenomenon of “internally generated sequences” in the hippocampus, and beyond [132]. In short, dynamical patterns of neuronal activations that code for behavioural trajectories (i.e., sequence of place cells) are observed in the rodent hippocampus both when animals are actively engaged in overt spatial navigation, and when they are disengaged from the sensorimotor loop, e.g., when they sleep or groom after consuming a food—the latter depending on an internally-generated, spontaneous mode of neuronal processing that generally does not require external sensory inputs. Internally-generated sequences that mimic closely (albeit within different dynamical modes) neuronal activations observed during overt navigation have been proposed to be neuronal instantiations of internal models, which play multiple roles including memory consolidation and planning—thus illustrating a possible way the brain might reuse brain dynamics/internal models in a “dual mode”, across overt and covert cognitive processes [132–137]. An intriguing neurobiological possibility is that the internal models that produce internally generated sequences are formed by exploiting pre-existing internal neuronal dynamics that are initially “meaningless”, but acquire their “meaning” (e.g., code for a specific behavioural trajectory of the animal) through situated interaction, when the internal (spontaneous) and external dynamics become coupled [138].

From a theoretical perspective, this mechanism might be reiterated hierarchically—thus forming internal models whose different hierarchical levels capture interactive patterns at different timescales [139].

4.2. *Who Fears Internal Models?*

The discussion above should have contributed to demystify the notions of internal model and inference, showing, first, that they provide the basic mechanisms to construct interactive accounts of perception and action; second, that they lend themselves to non-representational or representational interpretations, depending on how they are used; and third (if one is interested in theories of representation that appeal to the notions of decouplability or detachment), that there is a useful way to think at coupled or decoupled (or detached) cognitive operations in terms of the same internal models operating in a dual mode.

In doing so, we have briefly addressed a common misunderstanding about internal models and associated internal (hidden) states: the idea that they need to be isomorphic to (or a “mirror” of) external reality—an assumption that would directly conflict with pragmatist ideas that motivate interactive views of perception and action. While in generic model-based systems there are no particular constraints on the form and content of internal models, in Active Inference (or similar approaches) internal models and inferential processes are shaped by the agent’s goals in significant ways, rather than being a mere “mirror” or replication of the external world or its dynamics. First of all, the most important role of internal models in an active inference (or similar) agent is affording accurate control of the environment, not mapping external states into internal states—the latter function is only required as long as it is functional to the former. Internal models can be organized around interactive patterns (of sensorimotor contingencies) or around ways the environment can be acted upon, and not around sensory regularities only, as in the example of the pendulum used as a model for active sampling in analogy to rodent whisking behaviour. The importance of these models becomes evident if one considers that internal models develop while the agent learns to interact with the external environment and to exercise its mastery and control over it. Encoding the statistics of external stimuli is not sufficient for this; what would be more useful, for example, is modelling the way external inputs are sampled, categorizing sensory-motor events in ways that afford goal achievement, or recognizing similarities in task space rather other than (only) in stimulus space [140–142]. The importance of drive- and goal-related processes to internal modelling and learning becomes even more evident in that the agent’s models develop in close cooperation with the process of fulfilling internal allostatic processes, and then progressively afford the realization of increasingly more abstract goal states [65]. In sum, internal models continuously depend (for both their acquisition and expression) on control and goal-related processes and need to support situated interaction over and above representation. Internal inferential processes are biased in a similar, goal-directed way: the active inference scheme assumes that in order to act, an agent must “believe” the expected results of its actions in the first place, and has an “optimism bias” concerning its action success [122]. Hence, the agent’s belief state needs to reflect its goals rather than just the external reality. All of these arguments lend support for a nuanced view of internal models and inferential processes, which support an agent’s adaptive behaviour—and sometimes also “in the absence of” external referents they putatively represent—in a way that is compatible with the tenets of control-oriented, pragmatist theories.

5. Conclusions

There is a recent trend in philosophy, cognitive science, and neuroscience to embrace embodied, enactive, and other related (4-Es) proposals that emphasize integrated brain-body-environmental dynamics over and above internal representations and inferential processes. The array of criticisms of 4-Es theories to classical cognitive science and computationalism is sometimes difficult to reconcile under a unitary perspective [3,6,51,143,144]. Here, we have focused on two criticisms—based around the ideas of passive perception and serial information processing—for which a variety of theoretical and computational proposals have been

advanced, often appealing to different perspectives grounded for example in dynamical systems or statistical learning theory. A common feature of these two criticisms is the appeal to the notion of action (or interaction) as constitutive of perception and cognition, which is in accordance to pragmatist principles that see a resurgence in current cognitive science and neuroscience [53].

Next, we have discussed different possible solutions to these criticisms, some of which appeal to notions that are more compatible with classical cognitive science, such as the notions of internal models and inference, and some of which do not. Our analysis suggests that the two above criticisms are often conflated and need to be teased apart; for example, the notion of active perception does not automatically entail a non-inferential or an ecological perspective [12]. We have also shown that there are ways to incorporate the two criticisms within a family of models that use the notions of internal models and inference [17–22]. We have then focused more specifically on model-based theories, discussing their differences from alternative solutions and their degree of compatibility with embodied cognition and enactive theories. We have proposed that it is possible to make model-based systems compatible with the tenets of control-oriented and pragmatist theories in such a way that they solve problems of active perception and control. While model-based systems can be constructed in ways that are more compatible with traditional notions of symbolic systems [145,146] or interactive accounts of perception and cognition [16,75], there is a trend—at least within proponents of PP—to increasingly recognize pragmatist, action- and control-oriented perspectives [53,62,147].

Yet, conceptually, the functioning of model-based systems is open to various interpretations, including representational and non-representational interpretations, and we have tried to dissect the cases in which the different interpretations are more or less viable. A particularly interesting case is the Active Inference framework—which can be productively considered a modern incarnation of cybernetic theory—which includes elements of both computational theories of cognition (e.g., inferential processes and internal models) and embodied and enactive theories of cognition (e.g., the contribution of action to cognitive processing and the importance of self-organizing processes and autopoiesis). By pointing to the fact that these apparently conflicting processes can both coexist in the same framework, we have suggested that some theoretical disagreements (between, e.g., computational and dynamical theories) might be only apparent—or at least possible to reconcile. At the same time, some disagreements between the different camps persist, most notably concerning the content and usage (especially offline usage) of generative models implied in Active Inference and similar inferential schemes. We have reviewed various implemented systems that—depending on their complexity, as well as on the theory of representation one assumes—would lend themselves more or less naturally to representational or anti-representational interpretations. For example, if one assumes that offline use and decouplability are strong tests for representational function (because, in their absence, there would also be alternative, non-representational interpretations), then not all existing examples of Active Inference systems using generative models would pass this test. However, we have also discussed how some models do—or in other words, the possibility for internal (generative) models to temporarily detach from the current sensorimotor loop may afford representational functions, in a way that is not easy to reconcile with non-representational enactive theories [17,19,20]. Thus, while the different formalisms discussed here (e.g., with or without internal models) have different features, powers, and limitations, model-based solutions seem more suited to address the problem of detached cognition—or how living organisms can temporarily detach from the here-and-now, to implement (for example) future-oriented forms of cognition [128,135].

In sum, our analysis illustrates that: an explicitly inferential framework can capture some key aspects of embodied and enactive theories of cognition; some claims of computational and dynamical theories can be reconciled rather than seen as alternative explanations of cognitive phenomena; some aspects of cognitive processing (e.g., detached cognitive operations such as planning and imagination) that are sometimes puzzling to explain from enactive and non-representational perspectives can instead be accommodated by internal generative models and predictive processing regimes that mediate adaptive control loops. It is worth noting that our

conclusion that PP can account for important theoretical points raised by 4-Es theories would not automatically entail that PP is compatible with the whole spectrum of 4-Es theories. This is because, despite we have treated 4-Es as a heterogeneous but coherent set of theories, some 4-Es theories have been regarded as partially discordant, or even mutually exclusive; this is the case, for example, of different embodied and enactive theories, which emphasize or de-emphasize representations (see Section 5), but also of embedded versus extended cognition—the latter assuming (contra to the former) that aspects of the extra-neural environment forms part of the mechanistic substrate that realizes cognitive phenomena [7]. Hence, the degree of compatibility of PP and specific 4-Es theories remains to be established case-by-case. It is also worth noting that the importance of action for cognition is exemplified in many more domains than discussed in this paper. For example, there are important demonstrations that action dynamics are required to stabilize perceptual learning [148] and sequence learning [149] and that action should be considered as part and parcel of decision processes—for example, decision-makers consider both rewards and action costs jointly, and action dynamics feed back on decisions [45,50], and one can offload decisions to one's own behaviour [83,150]. All of these examples, and other, suggest that action is part of cognitive processing and not just a consequence of it [53]. PP theories seem well suited to address all these domains, but their adequacy remains to be fully demonstrated.

The model-based approach to active perception and control exemplified here also has significant implications for brain architecture. In keeping with embodied and enactive theories, the model-based approach to active perception and control maintains that cognition does not boil down to the internal manipulation of symbols completely separated from perception and action systems. Furthermore, it incorporates the pragmatist idea that the brain is designed for embodied interactions in space and time, not for the passive contemplation of objects or choices outside of sensorimotor and situated contexts—hence, proposing a focus on control rather than representational processes [16,46]. At the same time, model-based and inferential approaches emphasize the neuronal instantiation of internal generative models and explicit predictive processes that mediate adaptive control loops—an idea that is becoming increasingly influential in theoretical neuroscience [18] and (sometimes under the label of “predictive processing”) in philosophy [100]. These assumptions are not easy to reconcile with some enactive theories, which emphasize coupling rather than internal modelling [4,13] or which focus on implicit processes of anticipatory synchronization rather than explicit prediction [151]. While the conceptual and empirical scrutiny of these and alternative proposals continues, we hope we have contributed to shed light on the most significant differences between competing approaches—those that are worth subjecting to (active) hypothesis testing.

Acknowledgments: We gratefully acknowledge the support of HFSP (Young Investigator Grant RGY0088/2014).

Author Contributions: Giovanni Pezzulo, Francesco Donnarumma, Pierpaolo Iodice, Domenico Maisto and Ivilin Stoianov conceived the article; Giovanni Pezzulo drafted the article; Giovanni Pezzulo, Francesco Donnarumma, Pierpaolo Iodice, Domenico Maisto and Ivilin Stoianov revised the article. All authors have read and approved the final manuscript.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. James, W. *The Principles of Psychology*; Dover Publications: New York, NY, USA, 1890.
2. Hurley, S. The shared circuits model (SCM): How control, mirroring, and simulation can enable imitation, deliberation, and mindreading. *Behav. Brain Sci.* **2008**, *31*, 1–22.
3. Clark, A. *Being There. Putting Brain, Body, and World Together*; The MIT Press: Cambridge, MA, USA, 1998.
4. Gallagher, S. *How the Body Shapes the Mind*; Clarendon Press: Oxford, UK, 2005.
5. Thompson, E.; Varela, F.J. Radical embodiment: Neural dynamics and consciousness. *Trends Cogn. Sci.* **2001**, *5*, 418–425.
6. Wilson, M. Six views of embodied cognition. *Psychon. Bull. Rev.* **2002**, *9*, 625–636.
7. Rupert, R.D. Challenges to the hypothesis of extended cognition. *J. Philos.* **2004**, *101*, 389–428.

8. Haugeland, J. Mind Embodied and Embedded. In *Mind and Cognition: 1993 International Symposium*; Academica Sinica: Taipei, Taiwan, 1993.
9. Dewey, J. The Reflex Arc Concept in Psychology. *Psychol. Rev.* **1896**, *3*, 357–370.
10. Peirce, C.S. *Philosophical Writings of Peirce*; Dover Publications: New York, NY, USA, 1897.
11. Piaget, J. *The Construction of Reality in the Child*; Routledge: Abingdon, UK, 1954.
12. Gibson, J.J. *The Ecological Approach to Visual Perception*; Houghton Mifflin Harcourt: Boston, MA, USA, 1979;
13. O'Regan, J.K.; Noe, A. A sensorimotor account of vision and visual consciousness. *Behav. Brain Sci.* **2001**, *24*, 883–917.
14. Newell, A.; Simon, H.A. *Human Problem Solving*; Prentice-Hall: Upper Saddle River, NJ, USA, 1972.
15. Engel, A.K.; Maye, A.; Kurthen, M.; König, P. Where's the action? The pragmatic turn in cognitive science. *Trends Cogn. Sci.* **2013**, *17*, 202–209.
16. Pezzulo, G.; Cisek, P. Navigating the Affordance Landscape: Feedback Control as a Process Model of Behavior and Cognition. *Trends Cogn. Sci.* **2016**, *20*, 414–424.
17. Clark, A.; Grush, R. Towards a Cognitive Robotics. *Adapt. Behav.* **1999**, *7*, 5–16.
18. Friston, K. The free-energy principle: A unified brain theory? *Nat. Rev. Neurosci.* **2010**, *11*, 127–138.
19. Grush, R. The emulation theory of representation: Motor control, imagery, and perception. *Behav. Brain Sci.* **2004**, *27*, 377–96.
20. Pezzulo, G. Grounding Procedural and Declarative Knowledge in Sensorimotor Anticipation. *Mind Lang.* **2011**, *26*, 78–114.
21. Pezzulo, G.; Castelfranchi, C. The Symbol Detachment Problem. *Cogn. Process.* **2007**, *8*, 115–131.
22. Toussaint, M. Probabilistic inference as a model of planned behavior. *Kuenstliche Intell.* **2009**, *23*, 23–29.
23. Churchland, P.S.; Ramachandran, V.S.; Sejnowski, T.J. A critique of pure vision. In *Large-Scale Neuronal Theor. Brain*; The MIT Press: Cambridge, MA, USA, 1994; pp. 23–60.
24. Doya, K.; Ishii, S.; Pouget, A.; Rao, R.P.N. (Eds.) *Bayesian Brain: Probabilistic Approaches to Neural Coding*, 1st ed.; The MIT Press: Cambridge, MA, USA, 2007.
25. Friston, K. A theory of cortical responses. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* **2005**, *360*, 815–836.
26. Rao, R.P.; Ballard, D.H. Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field effects. *Nat. Neurosci.* **1999**, *2*, 79–87.
27. Von Helmholtz, H. Concerning the perceptions in general. In *Treatise on Physiological Optics*; Southall, J.P.C., Ed.; Dover: New York, NY, USA, 1866; Volume 3.
28. Hinton, G.E. To recognize shapes, first learn to generate images. *Prog. Brain Res.* **2007**, *165*, 535–547.
29. Hinton, G.E. Learning multiple layers of representation. *Trends Cogn. Sci.* **2007**, *11*, 428–434.
30. Barsalou, L.W. Perceptual symbol systems. *Behav. Brain Sci.* **1999**, *22*, 577–600.
31. Ahissar, E.; Assa, E. Perception as a closed-loop convergence process. *eLife* **2016**, *5*, e12830.
32. Gibson, J.J. *The Senses Considered as Perceptual Systems*; Houghton Mifflin: Boston, MA, USA, 1966.
33. Bajcsy, R.; Aloimonos, Y.; Tsotsos, J.K. Revisiting Active Perception. *arXiv* **2016**, arXiv:1603.02729.
34. Ahissar, E.; Kleinfeld, D. Closed-loop Neuronal Computations: Focus on Vibrissa Somatosensation in Rat. *Cereb. Cortex* **2003**, *13*, 53–62.
35. Donnarumma, F.; Costantini, M.; Ambrosini, E.; Friston, K.; Pezzulo, G. Action perception as hypothesis testing. *Cortex* **2017**, *89*, 45–60.
36. Friston, K.; Adams, R.A.; Perrinet, L.; Breakspear, M. Perceptions as hypotheses: Saccades as experiments. *Front. Psychol.* **2012**, *3*, 151, doi:10.3389/fpsyg.2012.00151.
37. Lepora, N.F. Biomimetic Active Touch with Fingertips and Whiskers. *IEEE Trans. Haptics* **2016**, *9*, 170–183.
38. Norman, D.A.; Shallice, T. Attention to action: Willed and automatic control of behaviour. In *Consciousness and Self-Regulation: Advances in Research and Theory*; Davidson, R.J., Schwartz, G.E., Shapiro, D., Eds.; Springer: Berlin/Heidelberg, Germany, 1986; pp. 1–18.
39. Barkley, R.A. The executive functions and self-regulation: An evolutionary neuropsychological perspective. *Neuropsychol. Rev.* **2001**, *11*, 1–29.
40. Fuster, J.M. *The Prefrontal Cortex: Anatomy, Physiology, and Neuropsychology of the Frontal Lobe*; Lippincott-Raven: Philadelphia, PA, USA, 1997.
41. Pezzulo, G.; Castelfranchi, C. Thinking as the Control of Imagination: A Conceptual Framework for Goal-Directed Systems. *Psychol. Res.* **2009**, *73*, 559–577.
42. Cisek, P. Cortical mechanisms of action selection: The affordance competition hypothesis. *Philos. Trans. R. Soc. B* **2007**, *362*, 1585–1599.

43. Cisek, P.; Kalaska, J.F. Neural mechanisms for interacting with a world full of action choices. *Annu. Rev. Neurosci.* **2010**, *33*, 269–298.
44. Shadlen, M.N.; Kiani, R.; Hanks, T.D.; Churchland, A.K. Neurobiology of Decision Making: An Intentional Framework. In *Better than Conscious?: Decision Making, the Human Mind, and Implications for Institutions*; Engel, C., Singer, W., Eds.; The MIT Press: Cambridge, MA, USA, 2008.
45. Lepora, N.F.; Pezzulo, G. Embodied Choice: How action influences perceptual decision making. *PLoS Comput. Biol.* **2015**, *11*, e1004110.
46. Cisek, P. Beyond the computer metaphor: Behavior as interaction. *J. Conscious. Stud.* **1999**, *6*, 125–142.
47. Ashby, W.R. *Design for a Brain*; Wiley: Oxford, UK, 1952; Volume ix.
48. Powers, W.T. *Behavior: The Control of Perception*; Aldine: Chicago, IL, USA, 1973.
49. Wiener, N. *Cybernetics: Or Control and Communication in the Animal and the Machine*; The MIT Press: Cambridge, MA, USA, 1948.
50. Cisek, P.; Pastor-Bernier, A. On the challenges and mechanisms of embodied decisions. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* **2014**, *369*, 20130479.
51. Pezzulo, G.; Barsalou, L.W.; Cangelosi, A.; Fischer, M.H.; McRae, K.; Spivey, M. The Mechanics of Embodiment: A Dialogue on Embodiment and Computational Modeling. *Front. Cogn.* **2011**, *2*, 1–21.
52. Verschure, P.; Pennartz, C.M.A.; Pezzulo, G. The why, what, where, when and how of goal-directed choice: Neuronal and computational principles. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* **2014**, *369*, 20130483.
53. Engel, A.K.; Friston, K.J.; Kragic, D. *The Pragmatic Turn: Toward Action-Oriented Views in Cognitive Science*; The MIT Press: Cambridge, MA, USA, 2016.
54. Shadmehr, R.; Smith, M.A.; Krakauer, J.W. Error correction, sensory prediction, and adaptation in motor control. *Annu. Rev. Neurosci.* **2010**, *33*, 89–108.
55. Port, R.; van Gelder, T. *Mind as Motion: Explorations in the Dynamics of Cognition*; The MIT Press: Cambridge, MA, USA, 1995.
56. Beer, R.D. The dynamics of adaptive behavior: A research program. *Robot. Auton. Syst.* **1997**, *20*, 257–289.
57. Hope, T.; Stoianov, I.; Zorzi, M. Through neural stimulation to behavior manipulation: A novel method for analyzing dynamical cognitive models. *Cogn. Sci.* **2010**, *34*, 406–433.
58. Nolfi, S. Behavior and cognition as a complex adaptive system: Insights from robotic experiments. In *Handbook of the Philosophy of Science: Philosophy of Complex Systems*; Hooker, C., Gabbay, D.M., Thagard, P., Woods, J., Eds.; Elsevier: Amsterdam, The Netherlands, 2009; Volume 10.
59. Nolfi, S.; Floreano, D. *Evolutionary Robotics. The Biology, Intelligence, and Technology of Self-Organizing Machines*; The MIT Press: Cambridge, MA, USA, 2001.
60. Todorov, E.; Jordan, M.I. Optimal feedback control as a theory of motor coordination. *Nat. Neurosci.* **2002**, *5*, 1226–1235.
61. Todorov, E. Optimality principles in sensorimotor control. *Nat. Neurosci.* **2004**, *7*, 907–915.
62. Clark, A. *Surfing Uncertainty: Prediction, Action, and the Embodied Mind*; Oxford University Press: Oxford, UK, 2016; ISBN 978-0-19-021701-3.
63. Friston, K. What is optimal about motor control? *Neuron* **2011**, *72*, 488–498.
64. Friston, K.; Samothrakis, S.; Montague, R. Active inference and agency: Optimal control without cost functions. *Biol. Cybern.* **2012**, *106*, 523–541.
65. Pezzulo, G.; Rigoli, F.; Friston, K.J. Active Inference, homeostatic regulation and adaptive behavioural control. *Prog. Neurobiol.* **2015**, *134*, 17–35.
66. Seth, A.K. The Cybernetic Bayesian Brain: From Interoceptive Inference to Sensorimotor Contingencies. In *Open MIND*; Metzinger, T., Windt, J.M., Eds.; MIND Group: Frankfurt, Germany, 2014.
67. Conant, R.C.; Ashby, W.R. Every good regulator of a system must be a model of that system. *Int. J. Syst. Sci.* **1970**, *1*, 89–97.
68. Kappen, H.J.; Gómez, V.; Opper, M. Optimal control as a graphical model inference problem. *Mach. Learn.* **2012**, *87*, 159–182.
69. Penny, W.D.; Zeidman, P.; Burgess, N. Forward and Backward Inference in Spatial Cognition. *PLoS Comput. Biol.* **2013**, *9*, e1003383.
70. Pezzulo, G.; Rigoli, F.; Chersi, F. The Mixed Instrumental Controller: Using Value of Information to combine habitual choice and mental simulation. *Front. Cogn.* **2013**, *4*, 92, doi:10.3389/fpsyg.2013.00092.
71. Pezzulo, G.; Rigoli, F. The value of foresight: How prospection affects decision-making. *Front. Neurosci.* **2011**, *5*, 79, doi:10.3389/fnins.2011.00079.

72. Solway, A.; Botvinick, M.M. Goal-directed decision making as probabilistic inference: A computational framework and potential neural correlates. *Psychol. Rev.* **2012**, *119*, 120–154.
73. Butz, M.V. Toward a Unified Sub-symbolic Computational Theory of Cognition. *Front. Psychol.* **2016**, *7*, 925, doi:10.3389/fpsyg.2016.00925.
74. Hemion, N.J. Discovering Latent States for Model Learning: Applying Sensorimotor Contingencies Theory and Predictive Processing to Model Context. *arXiv* **2016**, arXiv:1608.00359v1.
75. Maye, A.; Engel, A.K. A computational model of sensorimotor contingencies for object perception and control of behavior. In Proceedings of the IEEE International Conference on Robotics and Automation (ICRA 2011), Shanghai, China, 9–13 May 2011.
76. Seth, A.K. A predictive processing theory of sensorimotor contingencies: Explaining the puzzle of perceptual presence and its absence in synesthesia. *Cogn. Neurosci.* **2014**, *5*, 97–118.
77. Wolpert, D.M. Computational approaches to motor control. *Trends Cogn. Sci.* **1997**, *1*, 209–216.
78. Bickhard, M.H. Representational content in humans and machines. *J. Exp. Theor. Artif. Intell.* **1993**, *5*, 285–333.
79. Meyniel, F.; Schlunegger, D.; Dehaene, S. The Sense of Confidence during Probabilistic Learning: A Normative Account. *PLoS Comput. Biol.* **2015**, *11*, e1004305.
80. Friston, K.; FitzGerald, T.; Rigoli, F.; Schwartenbeck, P.; Pezzulo, G. Active Inference: A Process Theory. *Neural Comput.* **2016**, *29*, 1–49.
81. Friston, K.; FitzGerald, T.; Rigoli, F.; Schwartenbeck, P.; O'Doherty, J.; Pezzulo, G. Active inference and learning. *Neurosci. Biobehav. Rev.* **2016**, *68*, 862–879.
82. Pezzulo, G.; Cartoni, E.; Rigoli, F.; Pio-Lopez, L.; Friston, K. Active Inference, epistemic value, and vicarious trial and error. *Learn. Mem.* **2016**, *23*, 322–338.
83. Pezzulo, G.; Ognibene, D. Proactive Action Preparation: Seeing Action Preparation as a Continuous and Proactive Process. *Motor Control* **2011**, *16*, 386–424.
84. Pio-Lopez, L.; Nizard, A.; Friston, K.; Pezzulo, G. Active inference and robot control: A case study. *J. R. Soc. Interface* **2016**, *13*, doi:10.1098/rsif.2016.0616.
85. Maisto, D.; Donnarumma, F.; Pezzulo, G. Nonparametric Problem-Space Clustering: Learning Efficient Codes for Cognitive Control Tasks. *Entropy* **2016**, *18*, 61, doi:10.3390/e18020061.
86. Donnarumma, F.; Maisto, D.; Pezzulo, G. Problem Solving as Probabilistic Inference with Subgoalings: Explaining Human Successes and Pitfalls in the Tower of Hanoi. *PLoS Comput. Biol.* **2016**, *12*, e1004864.
87. Barrett, L.F.; Quigley, K.S.; Hamilton, P. An active inference theory of allostasis and interoception in depression. *Phil. Trans. R. Soc. B* **2016**, *371*, 20160011.
88. Pezzulo, G. Why do you fear the Bogeyman? An embodied predictive coding model of perceptual inference. *Cogn. Affect. Behav. Neurosci.* **2013**, *14*, 902–911.
89. Seth, A.K.; Friston, K.J. Active interoceptive inference and the emotional brain. *Philos. Trans. R. Soc. B* **2016**, *371*, 20160007.
90. Adams, R.A.; Shipp, S.; Friston, K.J. Predictions not commands: Active inference in the motor system. *Brain Struct. Funct.* **2013**, *218*, 611–643.
91. Kilner, J.M.; Friston, K.J.; Frith, C.D. Predictive coding: An account of the Mirror Neuron system. *Cogn. Process.* **2007**, *8*, 159–166.
92. Friston, K.; Frith, C. A Duet for one. *Conscious. Cogn.* **2015**, *36*, 390–405.
93. Dindo, H.; Donnarumma, F.; Chersi, F.; Pezzulo, G. The intentional stance as structure learning: A computational perspective on mindreading. *Biol. Cybern.* **2015**, *109*, 453–467.
94. Donnarumma, F.; Dindo, H.; Pezzulo, G. Sensorimotor coarticulation in the execution and recognition of intentional actions. *Front. Psychol.* **2017**, *8*, 237.
95. Donnarumma, F.; Dindo, H.; Iodice, P.; Pezzulo, G. You cannot speak and listen at the same time: A probabilistic model of turn-taking. *Biol. Cybern.* **2017**, *111*, 165–183.
96. Friston, K.; Mattout, J.; Kilner, J. Action understanding and active inference. *Biol. Cybern.* **2011**, *104*, 137–160.
97. Pezzulo, G.; Iodice, P.; Donnarumma, F.; Dindo, H.; Knoblich, G. Avoiding accidents at the champagne reception: A study of joint lifting and balancing. *Psychol. Sci.* **2017**, doi:10.1177/0956797616683015.
98. Allen, M.; Friston, K.J. From cognitivism to autopoiesis: Towards a computational framework for the embodied mind. *Synthese* **2016**, *1–24*, doi:10.1007/s11229-016-1288-5.
99. Hohwy, J. *The Predictive Mind*; Oxford University Press: Oxford, UK, 2013.

100. Metzinger, T.; Wiese, W. (Eds.) *The Philosophy of Predictive Processing*; Open Mind: Frankfurt, Germany, 2017.
101. Friston, K. Life as we know it. *J. R. Soc. Interface* **2013**, *10*, 20130475, doi:10.1098/rsif.2013.0475.
102. Friston, K.; Levin, M.; Sengupta, B.; Pezzulo, G. Knowing one's place: A free-energy approach to pattern regulation. *J. R. Soc. Interface* **2015**, *12*, 20141383.
103. Bruineberg, J.; Kiverstein, J.; Rietveld, E. The anticipating brain is not a scientist: The free-energy principle from an ecological-enactive perspective. *Synthese* **2016**, *2016*, 1–28.
104. Gallagher, S.; Allen, M. Active inference, enactivism and the hermeneutics of social cognition. *Synthese* **2016**, *2016*, 1–22.
105. Botvinick, M. Commentary: Why I Am Not a Dynamicist. *Top. Cogn. Sci.* **2012**, *4*, 78–83.
106. Beck, J.M.; Pouget, A. Exact inferences in a neural implementation of a hidden Markov model. *Neural Comput.* **2007**, *19*, 1344–1361.
107. Bogacz, R.; Brown, E.; Moehlis, J.; Holmes, P.; Cohen, J.D. The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks. *Psychol. Rev.* **2006**, *113*, 700–765.
108. Kiefer, A.; Hohwy, J. Content and misrepresentation in hierarchical generative models. *Synthese* **2017**, *2017*, 1–29.
109. Orlandi, N. Bayesian Perception Is Ecological Perception. Available online: <http://mindsonline.philosophyofbrains.com/wp-content/uploads/2015/09/Orlandi-Minds-2015.pdf> (accessed on 8 June 2017).
110. Gładziejewski, P. Predictive coding and representationalism. *Synthese* **2016**, *193*, 559–582.
111. Cummins, R.C. *Meaning and Mental Representation*; The MIT Press: Cambridge, MA, USA, 1989.
112. Friston, K.; Rigoli, F.; Ognibene, D.; Mathys, C.; Fitzgerald, T.; Pezzulo, G. Active inference and epistemic value. *Cogn. Neurosci.* **2015**, *6*, 187–214.
113. Montague, P.R.; King-Casas, B. Efficient statistics, common currencies and the problem of reward-harvesting. *Trends Cogn. Sci.* **2007**, *11*, 514–519.
114. Rubin, J.; Ulanovsky, N.; Nelken, I.; Tishby, N. The Representation of Prediction Error in Auditory Cortex. *PLoS Comput. Biol.* **2016**, *12*, e1005058.
115. FitzGerald, T.H.; Dolan, R.J.; Friston, K.J. *Model Averaging, Optimal Inference, and Habit Formation*; Frontiers Media SA: Lausanne, Switzerland, 2014.
116. Friston, K.; Schwartenbeck, P.; FitzGerald, T.; Moutoussis, M.; Behrens, T.; Dolan, R.J. The anatomy of choice: Active inference and agency. *Front. Hum. Neurosci.* **2013**, *7*, 598, doi:10.3389/fnhum.2013.00598.
117. Friston, K.; Shiner, T.; FitzGerald, T.; Galea, J.M.; Adams, R.; Brown, H.; Dolan, R.J.; Moran, R.; Stephan, K.E.; Bestmann, S. Dopamine, Affordance and Active Inference. *PLoS Comput. Biol.* **2012**, *8*, e1002327.
118. Kanai, R.; Komura, Y.; Shipp, S.; Friston, K. Cerebral hierarchies: Predictive processing, precision and the pulvinar. *Philos. Trans. R. Soc. B* **2015**, *370*, 20140169.
119. Saraf-Sinik, I.; Assa, E.; Ahissar, E. Motion Makes Sense: An Adaptive Motor-Sensory Strategy Underlies the Perception of Object Location in Rats. *J. Neurosci.* **2015**, *35*, 8777–8789.
120. Voigts, J.; Herman, D.H.; Celikel, T. Tactile object localization by anticipatory whisker motion. *J. Neurophysiol.* **2015**, *113*, 620–632.
121. Pfeifer, R.; Bongard, J.C. *How the Body Shapes the Way We Think*; MIT Press: London, UK, 2006.
122. Friston, K.; Schwartenbeck, P.; FitzGerald, T.; Moutoussis, M.; Behrens, T.; Dolan, R.J. The anatomy of choice: Dopamine and decision-making. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* **2014**, *369*, 20130481.
123. Pezzulo, G. An Active Inference view of cognitive control. *Front. Theor. Philos. Psychol.* **2012**, *3*, 487, doi:10.3389/fpsyg.2012.00478.
124. Roy, D. Semiotic schemas: A framework for grounding language in action and perception. *Artif. Intell.* **2005**, *167*, 170–205.
125. Roy, D.; Hsiao, K.; Mavridis, N.; Gorniak, P. Ripley, Hand Me the Cup: Sensorimotor Representations for Grounding Word Meaning. Available online: <https://www.media.mit.edu/cogmac/publications/asru03.pdf>. (accessed on 9 June 2017).
126. Jeannerod, M. Neural simulation of action: A unifying mechanism for motor cognition. *NeuroImage* **2001**, *14*, S103–S109.
127. Pezzulo, G. Coordinating with the Future: The Anticipatory Nature of Representation. *Minds Mach.* **2008**, *18*, 179–225.
128. Pezzulo, G. *Tracing the Roots of Cognition in Predictive Processing*; Open MIND: Frankfurt, Germany, 2017.

129. Jeannerod, M. *Motor Cognition*; Oxford University Press: Oxford, UK, 2006.
130. Behrens, T.E.J.; Woolrich, M.W.; Walton, M.E.; Rushworth, M.F.S. Learning the value of information in an uncertain world. *Nat. Neurosci.* **2007**, *10*, 1214–1221.
131. Mathys, C.; Daunizeau, J.; Friston, K.J.; Stephan, K.E. A bayesian foundation for individual learning under uncertainty. *Front. Hum. Neurosci.* **2011**, *5*, 39, doi:10.3389/fnhum.2011.00039.
132. Pezzulo, G.; van der Meer, M.A.A.; Lansink, C.S.; Pennartz, C.M.A. Internally generated sequences in learning and executing goal-directed behavior. *Trends Cogn. Sci.* **2014**, *18*, 647–657.
133. Buzsáki, G.; Peyrache, A.; Kubie, J. Emergence of Cognition from Action. *Cold Spring Harb. Symp. Quant. Biol.* **2014**, *79*, 41–50.
134. Buzsáki, G.; Moser, E.I. Memory, navigation and theta rhythm in the hippocampal-entorhinal system. *Nat. Neurosci.* **2013**, *16*, 130–138.
135. Pezzulo, G.; Kemere, C.; van der Meer, M. Internally generated hippocampal sequences as a vantage point to probe future-oriented cognition. *Ann. N. Y. Acad. Sci.* **2017**, *1396*, 144–165.
136. Pfeiffer, B.E.; Foster, D.J. Hippocampal place-cell sequences depict future paths to remembered goals. *Nature* **2013**, *497*, 74–79.
137. Redish, A.D. Vicarious trial and error. *Nat. Rev. Neurosci.* **2016**, *17*, 147–159.
138. Buzsáki, G. *Rhythms of the Brain*; Oxford University Press: Oxford, UK, 2006; ISBN 978-0-19-530106-9.
139. Friston, K. Hierarchical Models in the Brain. *PLoS Comput. Biol.* **2008**, *4*, e1000211.
140. Barsalou, L.W. Ad hoc categories. *Mem. Cogn.* **1983**, *11*, 211–227.
141. Rigoli, F.; Pezzulo, G.; Dolan, R.; Friston, K. A Goal-Directed Bayesian Framework for Categorization. *Front. Psychol.* **2017**, *8*, 408, doi:10.3389/fpsyg.2017.00408.
142. Stoianov, I.; Genovesio, A.; Pezzulo, G. Prefrontal Goal Codes Emerge as Latent States in Probabilistic Value Learning. *J. Cogn. Neurosci.* **2015**, *28*, 140–157.
143. Anderson, M.L. Embodied Cognition: A Field Guide. *Artif. Intell.* **2003**, *149*, 91–130.
144. Pezzulo, G.; Barsalou, L.W.; Cangelosi, A.; Fischer, M.H.; McRae, K.; Spivey, M.J. Computational Grounded Cognition: A new alliance between grounded cognition and computational modeling. *Front. Psychol.* **2013**, *3*, 612, doi:10.3389/fpsyg.2012.00612.
145. Thaker, P.; Tenenbaum, J.B.; Gershman, S.J. Online learning of symbolic concepts. *J. Math. Psychol.* **2017**, *77*, 10–20.
146. Tenenbaum, J.B.; Kemp, C.; Griffiths, T.L.; Goodman, N.D. How to grow a mind: Statistics, structure, and abstraction. *Science* **2011**, *331*, 1279–1285.
147. Clark, A. Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behav. Brain Sci.* **2013**, *36*, 181–204.
148. Verschure, P.F.M.J.; Voegtlin, T.; Douglas, R.J. Environmentally mediated synergy between perception and behaviour in mobile robots. *Nature* **2003**, *425*, 620–624.
149. Wu, Z.; Yamaguchi, Y. Input-dependent learning rule for the memory of spatiotemporal sequences in hippocampal network with theta phase precession. *Biol. Cybern.* **2004**, *90*, 113–124.
150. Carvalho, J.T.; Nolfi, S. Cognitive Offloading Does Not Prevent but Rather Promotes Cognitive Development. *PLoS ONE* **2016**, *11*, e0160679.
151. Stepp, N.; Turvey, M.T. The Muddle of Anticipation. *Ecol. Psychol.* **2015**, *27*, 103–126.

