
Tracing the Roots of Cognition in Predictive Processing

Giovanni Pezzulo

Can PP (Predictive Processing) help us understand “the roots of cognition”, and how we may have acquired (during evolution and/or development) our sophisticated cognitive abilities from the relatively simpler adaptive control mechanisms of our early evolutionary ancestors? Here I make the case that some cognitive operations may be constructed as *detached actions* — where the detachment process rests on the construction of generative PP models, which permit one to *internalize* action-environment dynamics. I provide three examples. The first example focuses on the role of internally-generated sequences of (hippocampal) neuronal activity across goal-directed navigation and detached tasks such as planning. This example illustrates how neuronal sequences (putatively forming an internal model for spatial navigation) may have a “dual use”: they may support both overt navigation and covert cognitive operations while running, respectively, in stimulus-based and internally-generated modes. Furthermore, the latter (internally-generated) mode may be considered a form of internalization of the former (stimulus-based) mode. The second example focuses on actions to resolve epistemic uncertainty, and the formal similarity in PP between *epistemic actions* that are executed overtly (e.g., exploration) and those executed covertly (e.g., episodic retrieval). This example illustrates the possibility of defining *mental actions*, such as reducing one’s uncertainty before making a choice, as internalized information foraging acts that have the same intentionality as externally-directed actions. The third example focuses on the detachment of cognitive goals such as “eating in a fancy restaurant” from homeostatic drives such as “being satiated”. This example illustrates that by internalizing regulatory loops within hierarchical PP models, one can build cognitive goals that can in turn enjoy some form of detachment — for example, one can go to a restaurant or buy food even when one is not hungry. I discuss these examples in relation to alternative theories of how higher cognition originates from (or is independent of) action-perception loops, including various versions of action-oriented, embodied and enactivist views.

Keywords

Active inference | Embodied cognition | Internalization | Predictive processing | Reuse

Acknowledgements

I would like to thank Thomas Metzinger, Wanja Wiese, Lucy Mayne and the anonymous reviewers for useful suggestions and editorial help.

1 Introduction

Predictive processing (PP) — especially in its most comprehensive version, Karl Friston’s *free-energy principle* (Friston 2010) — has recently become an influential framework for understanding brain activity and cognition across many disciplines, including systems neuroscience, cognitive science, philosophy, psychology and psychiatry. The most important constructs of PP — predictions and prediction errors, generative models, and precision — are increasingly mentioned in all these disciplines. This is sometimes in descriptive ways (i.e. to describe behavioural or neuronal regularities without committing to the ontological validity of these constructs) but is more often done with an implicit or explicit assumption that brains implement these mechanisms neurally.

PP suggests that the brain is a “prediction machine”. However, PP is not a unitary theory but rather refers to a variety of approaches. These approaches need to specify, for example, which brain functions are predictive and which are not, what exactly the brain predicts (e.g., the unfolding of a visual scene or action-perception contingencies) and at which timescales, and which computational mechanisms

(e.g., forward models, predictive coding) implement prediction and what their neuronal underpinnings are. One domain where PP was initially applied was (visual) perception. In this domain, *predictive coding* (Rao and Ballard 1999) emphasized the importance of a hierarchical (Bayesian) scheme, in which higher levels convey predictions to lower levels, and lower levels convey prediction errors to higher levels, with this process being iterated until prediction error is minimized, thereby disambiguating the most supported perceptual hypothesis. In this article, I will focus on *active inference* (under the Free Energy principle, (Friston 2010)), which starts from a predictive coding scheme but extends it to cover the domain of action control. Active inference uses an approximate Bayesian inference scheme and assumes that action control consists in producing proprioceptive predictions and successively fulfilling them by acting, rather than specifying motor commands (as is more commonly assumed in computational neuroscience and optimal control theory). In turn, proprioceptive (and other) predictions stem from priors encoded at high hierarchical levels, which thus essentially play the role of goal representations rather than the perceptual hypotheses of predictive coding.

PP is widely recognized in the two aforementioned domains — perception (e.g., *predictive coding*) and action (e.g., *active inference*). However, the PP framework is also increasingly used to explain a wide variety of cognitive phenomena of varying complexity, which go beyond action-perception loops and target abilities that have been traditionally considered to be the province of “higher cognition” (including for example planning, mindreading, foresight and cognitive control) as well as other domains including interoception, awareness and consciousness (Clark 2015; Clark 2016; Donnarumma et al. 2017; Hohwy 2013; Friston et al. 2013; Friston and Frith 2015a; Friston and Frith 2015b; Friston et al. 2016a; Pezzulo and Rigoli 2011, Seth 2013; Stoianov et al. 2016). This is appealing as, in principle, one can use the language of PP across multiple domains of cognition and even across different disciplines. However, a gap remains between the domains of (relatively simpler) action-perception loops and (relatively more complex) higher cognitive abilities. The former have been characterized in formal and quantitative terms using PP, whereas explanations of the latter tend to appeal to the same PP concepts but often lack a comprehensive quantitative and computational characterization. Thus, it remains to be seen if PP really “scales up” to higher cognition domains.

A second, related question is how exactly we should construct a PP theory of higher cognition. The mere fact that one can apply the principles of PP to action-perception loops and higher cognitive abilities leaves open the question of whether and how these domains are interconnected. A first logical possibility is that both action-perception loops and higher cognitive abilities comply with PP principles at an abstract level but use distinct sets of (neuro-computational) mechanisms, with higher cognition therefore being independent from action-perception loops. This “modular” perspective (or “a theory of two brains”) is compatible with more traditional cognitive theories that segregate perception, action and cognition (and their neuronal underpinnings). For example, one may assume that children possess innate modules for language or “intuitive theories” of physics and psychology, and although these abilities may be described using PP principles, these are fully distinct from the PP mechanisms involved in action-perception loops¹. A second logical possibility is that higher cognitive abilities are elaborations of action-perception loops which have never become (fully) segregated from them, and hence higher cognition remains functionally dependent on action-perception loops both during development and (at least in some cases) afterwards. This second, more “gradualist” perspective is compatible with various (stronger or weaker) forms of embodied or enactive cognition. Within this view, the existence of “cognitive mediators” — or sets of mechanisms that are shared across

¹ Modularism need not necessarily postulate innatism. One possibility is that the (evolutionarily more primitive) mechanisms supporting action perception loops scaffolded (the evolutionarily more recent) higher cognitive abilities during evolution and/or development, with the latter later becoming segregated and functionally independent from the former. This “scaffolding” view implies PP mechanisms in the construction of abstract (and possibly amodal) cognitive domains, which at some point became independent from their scaffold and hence became functionally modular. For example, one may assume that language acquisition initially leveraged (e.g., speech) action-perception loops but this process resulted in amodal representations that fully support language processing. This view would be partially in agreement with some versions of embodied theories but not with others that reject amodal symbols (Barsalou 1999).

action-perception cycles and higher cognitive abilities — has often been postulated. One example I will discuss below is the idea that internal forward modeling is used on-line for action prediction and off-line for action simulation².

In principle, PP can be used to construct both modularist and gradualist theories. However, PP, and in particular active inference, has often been conceptualized in embodied or enactive terms that invite a gradualist view. There is however a problem that any gradualist PP theory has to face.

1.1 The Problem of “Detachment”

Traditional cognitive theories have been attacked for their inability to deal with the “symbol grounding” problem, that is, how abstract knowledge and internally manipulated symbols acquire their semantics, and how abstract cognitive operations link to the action-perception loops that realize them (Harnad 1990). Embodied and enactivist theories of cognition are better placed to solve the grounding problem because semantics can be directly grounded in the predictive mechanisms underlying action-perception loops. Ironically, however, these theories face the opposite problem: detachment. Because they assume functional and/or causal relations between action-perception loops and higher cognitive abilities, embodied and enactivist accounts need to explain 1) how the latter originated from the former during ontogenesis and/or phylogenesis; 2) whether and how the latter become functionally autonomous (or “detached”) from the former, as exemplified by the fact that one can imagine an action without executing it; and 3) what detachment implies at the mechanistic level, i.e., whether imagining (or observing) an action engages action-perception loops covertly, recruits other mechanisms, or engages a combination of the two. All these problems are widely debated in cognitive and computational neuroscience, psychology and philosophy. How can PP help to shed light on these questions?

In this article I will discuss how PP can help us understand how living organisms could develop higher cognitive abilities from the mechanisms supporting adaptive action control. Central to this proposal are two interconnected ideas: 1) generative models that support PP in action-perception loops can progressively *internalize* aspects of agent-environment interactions; and 2) these generative models can be used in a “dual mode”, with one stimulus-tied mode supporting overt action control, and another internally-generated or spontaneous mode supporting covert and detached forms of cognition.

The internalization process plausibly operates at an evolutionary timescale as it requires building sophisticated internal models, but it can sometimes also operate (or be completed) during development. Either way, the result of the internalization process is a form of cognition that retains embodied and even enactivist aspects within the usual inferential scheme that PP uses to explain action and perception — and it is in this sense that one can trace back the roots of cognition in PP. From this perspective, the distinction between overt and covert processes — using the same generative model — is sufficient to explain the differences between simpler forms of action-perception loops and some forms of higher cognition.

Similar ideas regarding the internalization and reuse of predictive mechanisms have been advanced multiple times in cognitive science, at least since Piaget’s proposals on the construction of intelligence from sensorimotor experience (Piaget and Cook 1952); but also in many more recent variants (Clark and Grush 1999; Cotterill 1998; Hesslow 2002; Grush 2004) and especially under the theoretical umbrella of action-oriented representation (AOR). In the rest of this article, I will firstly summarize some ideas on the reuse of predictive mechanisms across motor control and higher cognition from the perspective of AOR, pointing out some limitations of these proposals. Then, I will provide

² For brevity, I cannot focus on the many other versions of this idea. For example, a “radical” version of embodied and enactivist approaches is that there is no real distinction between action-perception loops and higher cognition because the latter reduces to the former, i.e., higher cognition is fully supported by action-perception loops that comply to PP principles, or is possibly off-loaded to the external environment when it is necessary to implement forms of “symbolic processing”.

three concrete examples of how generative models may support the internalization and reuse of PP dynamics across action-perception loops and detached cognitive abilities. The first example discusses rodent goal-directed navigation and highlights the importance of *internally generated sequences* of hippocampal neurons (place cells) that support both online navigation and the off-line “replay” of experience, along with the possibility of understanding this latter phenomenon as the internalization of a generative model. The second example discusses the formal similarities in PP between epistemic behaviour or information foraging in the external world (e.g., overt exploration) and the internal milieu (e.g., a mental action that lowers the uncertainty of a belief state), discussing how the latter may be an internalization of the former. The third example demonstrates how the internalization of homeostatic mechanisms (that, e.g., satisfy hunger) may lead to the formation of cognitive goals (e.g., buying food) at a higher level of a PP hierarchy, and how the latter may support detached cognition — for example, buying food even when one is not currently hungry. These three examples will clarify that the PP perspective does not simply recapitulate previous proposals but brings new and significant insights, extending the notion of internalization over and above the off-line engagement of internal forward models implied in action performance (a key tenet of AOR). Finally, I will briefly discuss the implications of this proposal, in particular in relation to embodied cognition and enactivist theories.

2 The Reuse of Predictive Dynamics for Higher Cognition in the Action Oriented Representation (AOR) Framework

There have been many attempts to discuss the origins of cognition as elaborations of a predictive control system, in particular appealing to the idea of *action-oriented representation* (AOR) or — often with a similar meaning — of “motor cognition” (Clark and Grush 1999; Cruse and Schilling 2015; Grush 2004; Jeannerod 2006; Pezzulo et al. 2011). AOR theories propose that the motor prediction and control architecture of our early ancestors was gradually improved to afford higher cognitive functions such as cognitive control, executive function, imagery, planning and declarative knowledge — and in parallel, joint action and communication in the social domain — but that these higher cognitive abilities retain important “signatures” of their situated origins, thus making even higher cognition embodied to some extent. A core mechanism for extending primitive architectures to more complex, higher cognitive domains is the reuse of motor predictions in an off-line mode, to support (for example) “what if” simulations in decision-making or the covert simulation of another’s actions to understand her intentions. The basic idea is that, while engaged in an action-perception loop, agents also run another loop in parallel — a predictive loop (using a “forward model”) to aid action control (e.g., to compensate for delays) which mimics an action-perception loop. However, under certain circumstances, such as when external inputs and external outputs are inhibited, the forward model can also operate in isolation from the action-perception loop. It is in such cases that agents perform covert (cognitive) operations such as action simulation or imagination. Unlike enactivist theories (Gallagher 2005; Varela et al. 1992), AOR theories emphasize the importance of internal models in supporting covert cognitive operations while the agent is disengaged from online interactions with the environment (including other agents). In summary, AOR theories constrain the space of cognitive operations to those that can effectively use forward models that were originally developed for online interaction. For this, according to AOR, higher cognition *retains essential features of* online interactions (i.e., forward models) although *it does not consist in* online interaction.

These and other proposals within AOR (or related frameworks) have highlighted the importance of prediction in the development of higher cognition from sensorimotor control. However, several aspects remained underspecified. It is unclear whether internalization exclusively regards forward models supporting action control, or whether it is a broader phenomenon. It is also unclear which aspects of action-perception loops can be internalized. Furthermore, AOR theories have been constructed on top of a process model of sensorimotor action that stems from optimal control theory (or its variants),

and it is unclear whether this is the right foundation for understanding cognitive operations. This question is pressing as optimal control theory does not easily include some aspects of active inference that are appealing from embodied or enactivist perspectives. These include *epistemic* aspects of behaviour (e.g., epistemic foraging for information (Pirolli and Card 1999) or hypothesis testing) which may be important for explaining a range of actions (including mental actions) that change an agent's informational or belief state as opposed to a state of the external world (Friston et al. 2015). Finally, in AOR, the relations between action (sensorimotor) control and the adaptive processes of homeostatic regulation (and associated sensorimotor loops) have rarely been investigated, but they may be important for linking actions and motivations and for constructing notions of cognitive goals that go beyond the execution of simple responses (Pezzulo et al. 2015; Pezzulo and Cisek 2016).

Can PP help understand the “roots of cognition”? Is the framework of active inference generalizable to higher cognitive abilities, and how can the relations of these cognitive abilities to action-perception loops be conceptualized? Is the notion of internal generative model useful for understanding how animals may “detach” from the here-and-now and engage in sophisticated forms of (retrospective or prospective) cognition? Below I will address these and other questions by discussing three examples of how generative models may support the internalization and reuse of PP dynamics across action-perception loops and detached cognitive abilities. Each of these examples is supported by a convergence of empirical studies and modeling studies using PP. They are: 1) goal-directed navigation and the role of hippocampal internally generated sequences (IGSs) within it; 2) information foraging or epistemic actions in the external world and in mental space; and 3) the detachment of goal states from homeostatic drives.

3 Internally Generated Sequences (IGSs) in Goal-Directed Navigation

The first example concerns the role of hippocampal dynamics in rodent goal-directed navigation. This example is relevant because the rodent hippocampus can process sequences of neuronal activity in two modes: a *stimulus-tied* mode while the animal is actually foraging in the environment, and an *internally-generated* (or spontaneous) mode in the partial or even total absence of external stimuli (e.g., while the animal sleeps).

The stimulus-tied mode of activity is evident when a rodent is engaged in a navigation task (e.g., when it actively explores its environment). During navigation, the animal's spatial position can be decoded by considering the so-called “place cells” in the hippocampus, which fire preferentially in specific portions of the environment (i.e., have localized place fields). Place cells are sequentially activated as the animal visits successive spatial positions corresponding to the cells' place fields. Therefore, at the population level, place cells form sequences that code for the animal's current spatio-temporal trajectory. At the behavioural timescale of rodent navigation, sequences occur in the presence of external cues or landmarks (O'Keefe and Dostrovsky 1971).

However, sequential neuronal activity can also arise in the hippocampus due to a distinct, internally-generated mode that is self-organized in the sense that it operates in the (complete or partial) absence of changing cues or feedback. These *internally generated sequences* (IGSs) of place cells correspond to temporally compressed representations of particular spatio-temporal trajectories that the animal has taken (recently or remotely), or might take (Diba and Buzsáki 2007; Foster and Wilson 2006; Pezzulo et al. 2014). Recent evidence suggests that IGSs play pivotal roles across a variety of cognitive tasks such as memory function (e.g., consolidation) and future-oriented cognition (e.g., route planning) (Pfeiffer and Foster 2013).

There are at least two important forms of IGSs. The first form of IGSs is the “replay” of spatial trajectories during sleep or when the animal is in the delay period of a memory task. Replay implies that the same sequence of neurons that coded for spatial locations during the actual rodent navigation (place cells) can be reactivated endogenously in the absence of triggering stimuli, in a time-compressed way:

within Sharp Wave Ripple (SWR) complexes (sub-second bursts of high frequency oscillation of up to 220Hz, see [Buzsáki 2006](#)). SWR sequences can proceed in both a forward and a backward direction, the latter more prominently after the animal collects a reward ([Ambrose et al. 2016](#)). Replays were initially linked to memory consolidation, following the influential hypothesis that the hippocampus may be specialized for the fast learning of episodic memories and may replay experiences off-line to train and consolidate cortical semantic memories ([McClelland et al. 1995](#)). More recent findings support the hypothesis that replays are also involved in prospective forms of cognition. For example, when animals rest between goal-directed spatial navigation episodes, replays are preferentially directed toward known goal sites and are predictive of future choices, suggesting a role in planning ([Pfeiffer and Foster 2013](#)). Furthermore, replays are not limited to the verbatim recollection of spatio-temporal trajectories that the animal has experienced, but can also generalize to novel trajectories or novel combinations of already experienced trajectories ([Gupta et al. 2010](#)), as well as to unexplored spaces in which reward delivery has been observed ([Olafsdottir et al. 2015](#)) or novel environments before they are visited, i.e., preplay ([Dragoi and Tonegawa 2011](#)).

The second form of IGSs is “theta sequences” — time-compressed trajectories that can be decoded in the hippocampal theta rhythm of rodents engaged in behavioural tasks ([Foster and Wilson 2007](#)). Within each theta cycle (7-12Hz), short sequences of place cells (four to six on average) fire with very precise temporal dynamics: each cell fires at a specific phase of the theta rhythm, which changes cycle after cycle (i.e., phase precession) while preserving the sequential order at the population level and the forward direction. Theta sequences are formed very rapidly ([Feng et al. 2015](#)) and often act as a “moving window”, coding for a (forward) sequence of spatial positions loosely centred on the moving animal. The fact that theta sequences (often) include place cells that correspond to (have their “true” place field in) a future position of the animal’s trajectory has motivated the influential proposal that theta sequences afford prospective coding and the prediction of upcoming locations ([Lisman and Redish 2009](#)). Notably, during difficult decisions, theta sequences can support the prospective coding of behavioural plans (e.g., trajectories that lead to preferred goal sites, see [Wikenheiser and Redish 2015](#)) and choice alternatives (e.g., branches of a T-maze, see [Johnson and Redish 2007](#)), possibly implementing a serial deliberation between them. This latter example refers to the vicarious trial and error (VTE) behaviour of rodents at decision points in T-mazes: in early trials before they accumulate sufficient knowledge about the reward location, rodents stop and repeatedly look to the left and right as if they are deliberating between the alternatives ([Tolman 1938](#)). During VTE behavior, hippocampal theta sequences “sweep forward” serially in the two branches of the maze (while the animal remains at the decision point). This suggests that the animal is performing a “search through mental information space” ([Redish 2016](#)).

In summary, sequential neuronal activity in the hippocampus is observed both at a behavioral timescale (while the animal visits successive locations during navigation and receives external stimuli), and at faster timescales (when theta and SWR sequences run in an internally-generated mode). To explain this finding, it has been proposed that support for a broader range of detached cognitive functions stems from the *internalization* of the stimulus-tied hippocampal sequences (and associated phenomena such as theta rhythms) that initially supported overt spatial navigation. Thus, after internalization, hippocampal sequences have a “dual use” and can operate in both a stimulus-tied mode and an internally-generated mode, the latter possibly supporting a wide range of cognitive operations ([Buzsáki et al. 2014](#); [Pezzulo et al. 2014](#)). The possible functions of IGS (internally generated sequence) are various and still under investigation, and include memory consolidation (e.g., forming declarative memories, training an internal model), prediction and planning (e.g., preparing a route to a goal location), and the covert “what if” evaluation of possible action sequences (possibly in combination with other brain structures such as the ventral striatum).

The “dual use” may be conceptualized in terms of internal generative models for PP, which act as “sequence generators” for sequences of spatial locations (or more generally for sequences of events

which may not be navigational) and can form different functional networks with various brain areas (e.g., the entorhinal cortex, the prefrontal cortex and the ventral striatum), depending on task demands, see Pezzulo et al. 2017). These internal models are learned while an animal navigates an environment (though they may be preconfigured to some extent, see Section 6). During navigation, the models are engaged by external stimuli (conveyed to the hippocampus mainly through the entorhinal cortex) and can support the estimation of the animal's spatial position and produce short-range predictions. However, learning the internal models amounts to partially or fully internalizing the agent-environment dynamics, such that the same models can also be spontaneously reactivated (by tapping the self-sustaining internal dynamics of the model) in the partial or almost total absence of external stimuli (e.g., during sleep). This “tapping” can be either intentional, as in the below example of epistemic actions (see Section 4), or non-intentional, as in the case of replay of experience (e.g., for rodents, spatial trajectories and other events) during sleep.

How can this internally-generated mode be useful? An authoritative view is that the replay of spatial trajectories in rodents is useful for aggregating a series of *episodic* memories (temporarily stored in the hippocampus) into a *semantic* internal model in the cortex. This is supported by recent machine learning advancements which suggest that off-line experience replay significantly improves learning (e.g., by removing undesired correlations, Kumaran et al. 2016). There may be additional benefits if one thinks about the hippocampus in terms of an internal model rather than merely as a “storage” of episodic memories. Theoretical considerations suggest that when an internal model is spontaneously engaged in the absence of external stimuli, it can produce “unbiased” resamples of its content (in the case of IGSs, one might say that it would produce samples of trajectories based on the model's prior probability distribution (Buesing et al. 2011)). However, in practice, there will often be some external input or bias to this process. For example, the representation of a desired goal location such as the “home” location of the animal (possibly stemming from the prefrontal cortex) can influence this “re-sampling” process and bias the resampled sequences towards the goal location, possibly supporting planning function. This can be explained using the mechanisms of active inference (or the related framework of *planning-as-probabilistic-inference* (Botvinick and Toussaint 2012)), if one considers that goal-representation acts as a sort of constraint (or allegorically, a sort of attractor state) that funnels the resampling process. The same process can be used to repeatedly resample past experience for memory consolidation or for cognitive map formation, with the possibility to “bias” the sampling in a way that (for example) over-represents rewarded experiences (Kumaran et al. 2016). Importantly, the generative processes described here do not consist in the verbatim recollection of past episodes (as suggested by the term “experience replay”) but have constructive elements. They therefore permit (for example) recombination or interpolation from past experience (Gupta et al. 2010).

This example has illustrated that (hippocampal) neuronal dynamics can operate in a dual mode: one stimulus-tied and one internally-generated. This finding suggests the presence of an internal model that internalized agent-environment dynamics and is able to reproduce them spontaneously, in the absence of stimulus. Although this hypothesis remains to be fully empirically tested, it is exemplificative of a possible pathway from action-perception to detached cognition via internal modeling of the kind used in PP.

Although the evidence I reviewed comes from animal studies and touches only a limited set of detached operations — spatial memory and planning — the scope of IGS and related mechanisms may extend well beyond this. It has been suggested that mechanisms analogous to IGS may support more advanced human abilities including imagination, prospection and “mental time travel” to the past and the future, since these “detached” activities also recruit shared brain structures including the hippocampus and the medial temporal lobe (Buckner and Carroll 2007; Schacter and Addis 2007; Suddendorf 2006). One theoretical proposal bridging these seemingly disconnected fields is that “mechanisms of memory and planning have evolved from mechanisms of navigation in the physical world” and “the neuronal algorithms underlying navigation in real and mental space are fundamen-

tally the same” (Buzsáki and Moser 2013, p. 130). This would suggest that navigation and reasoning in arbitrary domains (“mental spaces”) may be based on the same mechanisms that support overt spatial navigation.

4 Epistemic Actions Can Be Executed both Externally and Internally

The second example concerns *epistemic actions*, which can be executed both externally (e.g., through overt exploration of the environment) and internally (e.g., by using a generative model to simulate the outcome of a series of actions and “gather evidence” in favour of a select few before a choice is made). Recent theoretical and empirical studies have shown formal similarities between these two forms of “information foraging” (Hills et al. 2015; Pezzulo et al. 2013), both of which may be invoked in the face of exploration–exploitation dilemmas, or when collecting information (and thereby reducing uncertainty) prior to a choice is more cost-effective than taking action based on current knowledge.

Both overt exploration (exemplified by searching for external cues before making a choice (Friston et al. 2015) and covert mental exploration (exemplified by rodent *vicarious trial and error* behaviour at decision points (Pezzulo et al. 2016a) have been recently modelled using PP. Importantly, both appeal to the same concept of *epistemic value*, which is an integral part of the active inference scheme. In this scheme, an agent’s plans must balance extrinsic value (e.g., reaching goals) and epistemic value (e.g., reducing uncertainty about the goal location); the latter can gain prominence over the former in circumstances where uncertainty is too high.

It is tempting to speculate, given the formal analogy between overt and covert forms of exploration and epistemic action, that some covert mental operations result from the internalization of mechanisms that balance overt exploration (i.e., exploring novel action possibilities) and greedy exploitation (i.e., selecting the most rewarding action found thus far) in conditions of uncertainty or risk. One might therefore use the internal model to consider and evaluate hypotheses in one’s mind (or to “collect more evidence” for and against each hypothesis) until one is either confident about one’s decision or decides that it is not worth investing further cognitive effort in that task. This captures the trade-off between exploring novel action possibilities and selecting the most rewarding action found thus far. Cognitive neuroscience is starting to scrutinize some of the brain mechanisms underlying exploration–exploitation and cost-benefits computations, including the balance between deliberative and habitual forms behavior (Daw et al. 2005; Redish 2016; Pezzulo et al. 2013) and the trade-offs between the costs of increasing attention demands (or exerting cognitive control over a task) versus the benefits in terms of increased reward (Shenhav et al. 2013). These trade-offs can be conceptualized using hierarchical PP architectures, in which (for example) deliberative mechanisms can supersede and contextualize habitual forms of behavior. Under PP, habitual forms of behavior would be selected when engaging the full deliberative system is not cost-effective, for example, when the animal is sufficiently confident that the environment has not changed, and so repeating a previously successful action is likely to result in a higher pay-off than exploring new opportunities (Pezzulo et al. 2015).

I have thus far focused on an intentional kind of epistemic action that consists in “tapping” or “interrogating” a generative model in order to probe hypotheses or collect evidence. However, this may be one instance of a more general cognitive mechanism that permits one to exert control over one’s own mental processes (as opposed to control over the external world); in other words, a mechanism that sees “thinking as the control of imagination” (Pezzulo and Castelfranchi 2009). This perspective implies that the concepts of “actions” and “skills” are extended beyond those that require the expression of overt behaviour to also include mental operations that have no immediate external referent. In a similar vein, Metzinger 2017 discusses how mental operations can have *epistemic goal-states* (e.g., “Knowing what the sum of 2 + 3 is”). One can also imagine other kinds of mental operations that are controlled towards some desired end-state. For example, an interior designer can move or change furniture pieces in her mind until she reaches a configuration that fits the style of the house; or an

animal can mentally resolve a competition between affordances, or plan to create new affordances, before acting (Pezzulo and Cisek 2016). Here, again, PP permits the identification of a crucial feature of these mental activities: the fact that they are actively controlled *towards* a desired goal state — where achieving the goal state has epistemic value.

As briefly mentioned above, the functional organization of action in PP (specifically, in active inference), as opposed to other schemes such as optimal control theory, revolves around achieving goal states using prediction and error correction mechanisms. In active inference, goals control action and perception engenders a cascade of predictions that are hierarchically decomposed down to set points for peripheral reflex arcs, which steer bodily movements. The same scheme can be adopted in a more internalized way without arc reflexes, if one allows an active inference agent to express goal states that concern his own mental states or “beliefs” (where the term “belief” is used in the technical sense of probability theory, not in the sense of classical propositional attitudes, and may denote for example a Gaussian probability distribution, defined by the two parameters of *expectation* or mean and *precision* or inverse variance). One example of an epistemic internal goal state is “having highly precise beliefs about the value of choice offers” when one needs to gather new evidence before making a decision about an investment. Another example is “having highly precise beliefs about the best placement of furniture pieces in this house” if one wants to design a fancy house layout. Yet another example is “having highly precise beliefs about the best way home” during route planning. In all these examples, an agent can execute a mental action to change his or her belief state³, and to make some of his or her beliefs very precise before making a choice.

It is worth noting that, in active inference, the precision of all the relevant beliefs (e.g., about the agent’s current and goal locations) is always optimized before a choice. This optimization is considered to be a standard aspect of active inference (or free energy minimization), not a form of meta- or cognitive control. However, there may exist (mental) operations that override or finesse the default optimization mechanisms of active inference, which would provide “mental actions” a truly causal role in the architecture of PP (see Metzinger 2017 for a comprehensive discussion). These mental actions may be cast within a Bayesian learning or active inference scheme, too. For example, one can use priors about the precision that a belief or a set of beliefs needs to have before a choice, or one can monitor precision levels, until one has a sufficient “sense of confidence” (Meyniel et al. 2015), i.e., a sufficiently high likelihood that one’s inferences are correct. In other words, the selection of a mental action — and the solution of “decide now vs. collect more evidence” (or optimal stopping) problems — may rest on the precision-modulation of internal epistemic states (e.g., raising priors on expected precision or confidence before a choice). This is analogous to the way raising the precision of an internal belief makes it a strong goal representation that generates a cascade of predictions, which in turn enslave overt action.

From this perspective, the achievement of epistemic goal states (and the resolution of epistemic uncertainty) may be seen as a form of cognitive control over one’s own mental activity by using monitoring, error correction and precision modulation mechanisms that are analogous to overt action control (Pezzulo 2012) in order to control epistemic behavior and attain sufficient confidence in one’s choices. Some examples are: continuing to mentally compare the pros and cons of various invest-

³ While both mental action and perceptual inference change the agent’s belief state, there are significant differences between the two. Perceptual inference is implemented under a predictive coding scheme; it uses prediction errors to change the agent’s belief state as a function of new evidence and prediction error but does not include action selection or the active search for new evidence. Conversely, a mental action can be conceptualized as an epistemic action under the active inference scheme, which generalizes predictive coding to also include action selection and planning. Like other (overt) epistemic actions, mental actions result from the deliberate choice to search for new information (or reconsider old information) for epistemic purposes, e.g., to intentionally reduce uncertainty before a difficult choice. For example, if a conference has two parallel symposia and one is uncertain about which one to attend, he can execute an epistemic action externally (e.g., spend some time reading the talk abstracts) or internally, as a mental action (e.g., explicitly recall past examples of talks by the same speakers). Epistemic (or mental) actions are selected as part of a plan whose benefits (e.g., reading all talk abstracts until one is very confident about the best symposium) and costs (e.g., the cognitive costs of abstract reading) are considered and compared with those of alternative action plans (e.g., drinking another coffee and then going directly to a random symposium). The selection of a mental action thus complies with the same formal principles that regulate the balance of intrinsic (epistemic) and extrinsic (economic) value in active Inference (Friston et al. 2015).

ments (or reading business webpages) until one is confident enough, continuing to imagine moving furniture (or actually moving it) until one is happy with final configuration, or striving to remember past travels (or consulting a GPS navigator) until one is certain about a travel plan. The computational efficiency and empirical validity of these or alternative schemes, and their relations to meta-cognition and cognitive control, remain to be assessed. Furthermore, it remains to be studied whether mental action selection obeys cost-benefit considerations, permitting one to trade off the benefits of extra information and extra confidence against the cognitive and temporal costs of achieving these epistemic goals (Shenhav et al. 2013).

5 From Homeostatic Drives to More Abstract Goal States

My third example concerns the detachment of goal states from homeostatic drives. In active inference, one can describe adaptive control loops by starting from cybernetic error-correction mechanisms (Butz 2016; Pezzulo et al. 2015; Seth 2013). To illustrate the concept, one can start with a homeostatic drive — such as a felt need for glucose — that produces (interoceptive) prediction errors. These errors, in turn, engender autonomic responses but also a sophisticated (crossmodal) generative model that produces a cascade of (exteroceptive and proprioceptive) prediction errors. The latter engage an action pattern — such as locating and consuming an apple — that suppresses all the prediction errors, including the initial interoceptive prediction error (by restoring homeostasis), thus terminating the process.

However, not all adaptive actions are initiated (and controlled) by current needs and interoceptive prediction errors. The fact that one can buy food even when one is not hungry exemplifies the human ability to set and achieve goals in open-ended ways. In other words, there is often a strong *functional* dependence between homeostatic imperatives (e.g., to be satiated) and goal states that drive adaptive action (e.g., finding and then consuming food), but the *causal* (or proximal) coupling between the two can be sometimes loosened. In other words, an interoceptive prediction error (signaling e.g., low glucose levels) is not always required to initiate active inference and control loops for food consumption.

To understand how this may be possible, one needs to understand the aforementioned cybernetic scheme in more detail, in particular, its anticipatory aspects. In PP, adaptive action is realized by a generative model that encodes contingencies across interoceptive, exteroceptive and proprioceptive modalities, e.g., between glucose levels, the visual appearance of apples, and the actions required to secure them. The predictive capabilities of the internal model permit going beyond feedback-based error correction, to steering a series of anticipatory regulatory (or *allostatic* (Sterling 2012)) loops. For example, one can stop eating an apple predictively (i.e., by using the internal model to predict that eating a certain amount will restore glucose levels) rather than reactively (i.e., only after receiving a signal that the glucose level is actually restored). This is more adaptive given that generating the latter signal may take too much time. Moreover, one can use predictions rather than just feedback to select and regulate action. For example, one can decide on an action in anticipation of a predictable need (be satiated) rather than waiting for an interoceptive error signal (for hunger). Another example of (implicitly) anticipatory process during regulatory eating loops is salivation, which prepares resources to digest a to-be-eaten food (Pavlov and Thompson 1902). The theory of *allostasis* (Sterling 2012) encompasses many more examples of anticipatory regulatory mechanisms, which involve (for example) hormonal processes that mobilize resources in anticipation of a need, and which up- or down-regulate the whole system rather than using fixed set points as would be suggested by the idea of homeostasis. All these examples illustrate that even the (relatively) simple regulation of drive states can be largely anticipatory rather than just reactive. The neuronal PP architecture supporting the highly integrative functions required for allostasis is necessarily hierarchical and includes important hubs that combine interoceptive, exteroceptive and proprioceptive modalities (e.g., the insula, see Craig 2015)

During learning and development, the internal model can increase its scope and internalize drive-based regulatory loops to generate (for example) goal-representations and plans for “eating”. It can then initiate a plan for “searching for a restaurant” or “buying food” in an internally-generated mode (i.e., based on goals) rather than in a stimulus-driven mode (i.e., only after feeling hunger). This can be done using a hierarchical PP model that progressively learns regularities at increasingly deeper levels and longer temporal timescales — such as the relations between the act of ordering food in a restaurant and the integrity of the internal milieu, as measured by low prediction error of interoceptive signals (Pezzulo et al. 2015). These internal models permit an organism to anticipate needs rather than merely reacting to them, and to prepare to satisfy a drive (e.g., hunger) before there is an interoceptive error signal. From this perspective, the role of a higher-level cognitive goal like “buying food” would be to produce a sort of anticipatory error-signal, which triggers error-correction actions (for food consumption) before a lower-level drive system produces an interoceptive error-signal (e.g., loss of glucose), which would be more dangerous. The proximal mechanisms for producing goal-related error signals and for monitoring goal achievement error signals may be borrowed from more primitive mechanisms that monitor reward achievement (Montague 2006).

This example illustrates that a goal-based mechanism for action selection provides some detachment from immediate needs and homeostatic drives; in other words, goal-directed (intentional) action can distally relate to basic drives but also acquire autonomy from them. The act of finding a restaurant before one is hungry retains the full intentional and adaptive character of eating an apple when one is hungry; while the latter is driven by interoceptive stimuli, the former is internally-generated (goal-driven) but still adaptive, because the model has internalized a basic homeostatic loop (and the subsequent corrective actions). The power of this hierarchical PP scheme rests on the fact that it allows animals to control and produce effects in the external world and invent “cognitive goals” in open-ended ways, which go well beyond the satisfaction of the homeostatic drives that (often) originated them.

The question of how much cognitive goals can “diverge” from simpler physiological imperatives is still open. In principle, the fact that goal states are learned by internalizing (and predicting) allostatic loops should prevent a radical divergence between the two. Furthermore, in the PP hierarchy envisaged here, the “higher” layers that encode more cognitive goals (like finding a restaurant) remain to some extent linked to the “lower” layers that implement more basic allostatic loops. Prediction errors need to be minimized at all levels, and so even if the “cognitive goal” of eating at a good restaurant has been achieved, one can change restaurant if (after a while) the simpler drive of “being satiated” is not achieved — at least if one has enough time and money — which would also change the “model” of the restaurant. Finally, it is not necessary that higher layers supersede all the operations performed by lower layers. For example, in evaluating how good a food is, one can rely on (relatively higher) cognitive representations of the quality of a restaurant — for example, whether it was positively reviewed in a gourmet magazine — but also engage a (relatively lower) interoceptive simulation that provides anticipated feelings of taste.

However, there may be cases where cognitive goals truly diverge from physiological imperatives. My examples regarded very simple cases of goal states (like finding a restaurant) for which one can reconstruct a plausible causal history back to physiological states (hunger). However, even these (apparently) simple goal states have aspects that originate from cultural dynamics and may not be easily reducible to homeostatic imperatives — and sometimes may run against them. This is even more evident in sophisticated goals such as pursuing an ascetic ideal. Although one may link social and cultural practices to the usual imperatives of survival and reproduction, the ways proximal mechanisms (goal achievement) link these domains are not always easy to reconstruct. Finally, it is important to consider that there is a strong habitual component in human behavior (e.g., buy food at my usual supermarket, or go to a restaurant every Friday); while habits may originate from the routinization of goal-directed control (Pezzulo et al. 2015; Friston et al. 2016b), they do not retain its flexibility (e.g., they can be insensitive to changes in interoceptive state) and thus may become maladaptive. These ex-

amples illustrate that in humans, the relations between sophisticated goals and simpler physiological imperatives may be multifarious — but at least hierarchical PP modeling offers some guiding principles for studying them.

6 Open Questions and Interrogatives

I started with the problem of ‘scaling-up’ action-oriented theories of cognition to account for ‘higher’ cognitive phenomena (such as imagery, navigation, and so on). I provided three examples of such ‘higher’ cognition (in spatial navigation, mental actions, and the creation and attainment of cognitive goals) and discussed them in terms of detached actions — where the detachment process rests on the construction of generative PP models, which permit the internalization of action-environment dynamics. My proposed view of detached cognition as internalized PP has several implications, but also raises numerous interrogatives, which I summarize below schematically:

- It follows from this framework that, although detached cognition is free from some of the constraints of action-perception cycles (e.g., the necessity to support real time action), it may preserve key situated and embodied aspects of overt action control, such as the “natural statistics” of control tasks (e.g., navigation or grasping). One example is the fact that hippocampal theta and SWR sequences run at a faster timescale (thus without the temporal constraints of action-perception cycles), yet their content (in terms of place cells and their relative temporal arrangement) can be the same or very similar. Visual and motor imagery studies similarly suggest important similarities at both functional and neuronal levels between perceptual (or motor) experiences and their covert counterparts (imagery), despite the fact that the latter are free from some constraints (e.g., one can imagine legendary animals “in the mind’s eye” and in some cases one can imagine faster than one can see [Jeannerod 2006](#); [Pylyshyn 2001](#)). Which constraints are shared and which are not (and why) across overt and covert cognitive operations remains to be fully assessed.
- In active inference, bodily processes are part of the inference and contribute to prediction error (or free energy) minimization. One example is active sensing: one can use hand or eye movements to infer (or reduce surprise about) the position and shape of an object ([Friston et al. 2012a](#)). Under certain conditions, other bodily movements may be part of covert cognitive operations, too. For example, eye movements upwards may contribute to our understanding of a story about a person who climbs a building ([Spivey and Geng 2001](#)). Alternatively, some aspects of bodily movements may be internalized, too, and become part of generative models that can run covertly in an internally-generated way; this may be the case, for example, in head direction cells in rodents ([Buzsáki et al. 2014](#)).
- From a neurophysiological viewpoint, an intriguing question is how the same neuronal population can operate in two modalities: stimulus-tied and internally-generated⁴. From the PP perspective, this question is also interesting as it may help shed light on how the brain implements internal generative models. One domain where this question can be addressed is hippocampal processing. In the hippocampus, learning is extremely rapid. This has led to the proposal that it may rest on preconfigured neuronal dynamics: pre-coded sequences of neuronal activities that may serve as “dynamical templates”, permitting the learning of new experiences “in one shot”

⁴ Usually, “stimulus-tied” refers to a process that is initiated by a stimulus from the outside world (e.g., a visual stimulus). In the case of homeostatic regulation and interoceptive loops, the same idea maps to the availability of bodily sensations and feelings (e.g., feeling hungry) rather than truly “external” stimuli. However, the very idea of a process that is initiated by stimuli (either perceptual stimuli or bodily sensations and feelings) is too reactive for PP theories. Active inference is guided (proactively) by exteroceptive, proprioceptive and proprioceptive predictions that obey high-level beliefs and homeostatic imperatives, rather than being governed (reactively) by stimuli — except in special cases such as habits ([Friston et al. 2016b](#)). From this perspective, active inference has always “internally generated” components — especially at high hierarchical levels, which can embed self-sustained dynamics that continuously engender predictions, see e.g., ([Friston et al. 2012c](#)). The distinction between “stimulus-tied” and “internally generated” still makes sense in this context if one considers that in the latter case external stimuli are partially or entirely missing.

with successive refinement of what was learned on the basis of additional experience (Dragoi and Tonegawa 2011). A more comprehensive proposal is that (unlike most current machine learning approaches) this form of learning may consist of aligning internal spontaneous brain dynamics (and rhythms) to external environmental dynamics, and once this is done, brain rhythms spontaneously “replay” agent-environment interactions. This would mean that internal models do not become detachable only after learning but are able to run in a spontaneous mode from the beginning (Buzsáki 2006). The idea that it is possible to *model* agent-environment dynamics by *aligning* internal and external brain dynamics is also appealing from an enactivist viewpoint (see Section 7).

- The second example, on “epistemic action”, suggests that we should expand the concept of an “action” beyond bodily movements that achieve an externally defined goal state, to also encompass (covert) mental actions that have no immediate external correlates but maintain action intentionality. This extended notion of action is fully compatible with PP. Indeed, in PP an action is directed towards a goal state, but the notion of a goal state does not have to be restricted to the exteroceptive or proprioceptive domains (as for physical actions such as grasping), but can encompass an informational state of the system (e.g., the precision of a belief state). This is most evident if one considers that, in PP, actions have both pragmatic (or extrinsic) and epistemic (or intrinsic) components, with the latter having the goal of changing informational state (e.g., searching for a cue). This notion of an epistemic action in PP adequately covers those “mental” actions whose goal is changing the agent’s belief state (and not, technically speaking, the state of the external environment or the agent’s physical internal environment, such as being hungry).
- One interesting implication is that mental actions that resolve uncertainty would have the same intentionality as external exploratory actions. In other words, information foraging is the same when performed in the mind or in the external world. A second interesting implication is that “mental actions” can be conceptualized as actions that change the cognitive (belief or goal) state of an agent that is involved in (for example) action planning or control, and need not to operate on a distinct (amodal) cognitive substrate. Thus, for example, it would not be necessary to postulate two distinct cognitive representation of state uncertainty, one that is controlled using “mental actions” and one that influences action selection.
- The third example, on drives and goals, provides an example of detachment that rests on the construction of a hierarchical generative model rather than a strict dual use of the same parts of the model, as in the case of hippocampal sequences. In such cases, the internalization of a homeostatic loop and the ensuing regulatory plans requires building additional (hierarchically higher) components of the model. While the fact that cognitive goals can be engaged even when there is no sensed interoceptive loop suggests that the hierarchically higher goals enjoy some degree of detachment, it is unclear whether and in which cases this detachment of the “higher part” from the “lower part” can be complete, especially since prediction error has to be minimized across the whole hierarchy, not just in a part of it.

7 Relations with other Approaches Including Embodied and Enactivist Views

This perspective has some similarities with, but also important differences to other proposals, which I briefly summarize below.

7.1 Action-Oriented (AOR) Framework

Compared to most proposals advanced within the AOR or motor cognition frameworks, here the emphasis is not just on the reuse of motor predictions outside motor control loops, but the engagement

of generative models in an internally-generated mode, which may be a broader phenomenon. In other words, the off-line reuse of the motor system's predictive abilities may be just one of the mechanisms permitting a biological organism to temporarily disengage from the perceptual-motor loop and engage in detached forms of cognition.

7.2 Cortical Recycling and Neural Reuse

The ideas of internalization and “dual use” have some relation to theories of “cortical recycling” (Dehaene and Cohen 2007) and of “neural reuse” (Anderson 2010), which focus on the exaptation or recycling of neuronal resources that then acquire novel functions. For example, a brain area adapted for perception might be exapted to also recognize letters. However, implicit in the idea of “dual use” is the assumption that covert cognitive abilities (e.g., planning) remain connected to the overt processes (e.g., spatial navigation) that scaffolded them, because they use a common generative model. In other words, these abilities are not just connected by their ontogenetic or phylogenetic history (e.g., recycling), but continue to share a generative model, which can operate in two dynamic modalities (stimulus-tied vs. internally-generated). The possibility of operating in two modalities is intrinsic to the notion of a generative model, in which the internally-generated mode corresponds to the generative process of “imagining” or “hallucinating” patterns such as images, faces, video frames, etc. (Hinton 2007).

7.3 Dual-Process Theories

The notion of “dual use” is not the same as dual-systems theories in cognitive science (such as the idea of two separate systems of thought, one reflexive and one deliberative (Kahneman 2011)). Nor does the distinction between stimulus-tied and internally-generated processes map to the distinction between habitual and goal-directed control in dual-process theories of reinforcement learning (Daw et al. 2005). Stimulus-tied modalities reflect a process occurring at the same timescale as the action-perception cycle. This process can fully incorporate external stimuli independently of whether goal-directed action planning or stimulus-response is implemented (the latter, in active inference, applies only in rare cases such as in the presence of habits). Rather, internally-generated refers to processes that are outside the action-perception cycle (e.g., the covert replay of spatial trajectories while the animal sleeps and is deprived of external sensations).

Dual-process theories in reinforcement learning assume that goal-directed action and habits depend on segregated neuronal and computational processes, and that they compete to control behaviour (although the possibility that they may also cooperate has been sometimes recognized). Within active inference, however, goal-directed actions and habits are better conceptualized within a hierarchical scheme in which the higher layers that implement goals can contextualize lower levels that implement less flexible responses (Pezzulo et al. 2015). The result is a continuum between goal-directed and habitual action that depends on the relative weight assigned to the different layers. Habits can arise in this scheme too, when the lower layers acquire sufficient precision to become essentially impermeable to the influences of the higher layers. As a further development of this view in the context of policy selection, one can consider that habitual policies can arise from the self-observation of goal-directed action planning when there is no (residual) ambiguity (Friston et al. 2016b). In this scheme, one would initially select policies in a goal-directed manner, and successively (when there is no ambiguity over time) develop a habitual policy: a “copy” of the most-often selected goal-directed policy, which can be selected in a stimulus-based manner rather than using deliberation and expected free energy

minimization. It is worth noting that in this scheme, habitual policies are not learned in parallel with goal-directed action (as assumed typically in dual theories) but only afterwards⁵.

7.4 Perceptual Symbol System Theory

The PP-inspired view of detached cognition sketched here connects quite well with the most developed conceptual framework for embodied cognition: perceptual symbol system theory (PSS) (Barsalou 1999). In PSS, experiences are internalized to form embodied concepts (or “perceptual symbols”), whose re-enactment produces a “simulator” that steers “situated simulations”. Here, one can consider that “perceptual symbols” link to specific (unimodal) elements of a generative model and “simulators” link to multiple interconnected elements that form a multimodal concept (e.g., the concept of a dog generates multimodal predictions regarding what a dog looks like, how the bark sounds and the anticipated softness of touching a dog). A “simulation” refers to the generative process of the generative model, which produces (or “hallucinates”) observations that are compatible with a given simulator or a combination of multiple simulators, much like deep (generative) neural networks are used to generate exemplars in machine learning (Hinton 2007). In PSS, however, a simulation is always “situated”: the (prior) information encoded in the simulator is combined with various contextual elements that are present at the moment a person instantiates a simulation. This implies that a person would produce different “situated simulations” of an airplane if he is flying or at home, if he is happy or worried, or if he is engaged in a memory task (e.g., recalling names of parts of an airplane) or imagining a future flight. This situatedness (or context-sensitivity) is a hallmark of human cognition and is currently beyond what current machine learning techniques can do; perhaps it would require embodying a PSS into an agent that dwells in realistic environments and has a rich set of personal experiences. Despite these limitations, some key ideas of PSS may be explained (or implemented) using the usual constructs of PP; a situated simulation might construct perception by generating and predicting exteroceptive observations (predictive coding), guide action by generating proprioceptive predictions (active inference), and scaffold emotional experience by generating and regulating interoceptive states (interoceptive inference or embodied predictive coding; Barrett and Simmons 2015; Pezzulo 2013; Seth 2013).

7.5 Enactivism and the Relations between Internal Modeling and Representation

Most of the ideas I discussed in this article would also lend themselves quite naturally to an enactivist perspective. This is consistent with previous observations that PP (and in particular active inference) has enactivist elements (Allen and Friston 2016; Friston et al. 2012b; Bruineberg et al. 2016). This seems *prima facie* surprising, given that active inference includes the (cybernetic) notion that adaptive control requires an *internal model* of the environment, and the idea of an internal model is closely related to the idea of internal representation, which is antithetic to enactivism. However, in active inference, internal modeling is instrumental to accurate goal-directed action control, over and above representation (which is not the case in all theories of PP and perceptual predictive coding). Priors play the dual role of hypotheses (in perceptual processing) and goals (in action control). However, in most practical cases, the latter goal-oriented role is more fundamental, because there are some prediction errors, such as those generated by homeostatic processes, which cannot be minimized by “changing one’s mind” but require taking action (e.g., eating or drinking). Accordingly, the brain develops internal models and generates predictions to satisfy the agent’s goals (or to maintain allostasis) rather than to maintain an accurate internal representation of the external environment *per se*. In other words, the

⁵ The rationale for this is that the full power of active inference is only necessary under risk (e.g., many-to-one relations between outcomes and internal states) and ambiguity (i.e., many-to-one relations between internal states and outcomes). When there is no ambiguity, then the external stimuli are usually sufficient to afford a successful state-action policy of the kind commonly used in fields like reinforcement learning. Therefore, developing a habitual policy would amount to a sort of “simplification” of the internal model, which is just another way to reduce free energy, see Friston et al. 2016b.

success criteria for internal models of agent-environment dynamics are accurate prediction and goal achievement, not accurate mirroring of an external reality. In most practical applications, a model can afford good prediction even if it is sketchy and does not capture the full complexity of environmental dynamics. This is evident if one looks at published studies and compares the agent's generative model with the "true" generative process (aka the "real" environmental dynamics). In summary, active inference can be conceived of as the synthesis of two ideas: that "the brain is for prediction" and "the brain is for action". The focus on the latter, action-based and embodied aspects of brain function (which is not mandatory in other PP approaches) relaxes representational aspects of internal modeling.

An even more nuanced view of the relations between internal modeling and representation emerges if one assumes that developing an internal model boils down to aligning (or synchronizing) pre-existent brain dynamics and rhythms to environmental dynamics, as discussed above in relation to hippocampal processing. This may be not entirely satisfactory from an enactivist viewpoint, though, as it still requires postulating that internal models are *within* the brain. Alternatively, one can consider that even if an internal model is required for control, the internal model is not *within* the brain; rather, *the brain-body-environment system as a whole* implements an internal model⁶. For example, a robot may produce efficient locomotion by aligning internal dynamics (e.g., rhythmic behaviour produced by central pattern generators for locomotion) and external dynamics (a treadmill moving at a certain speed), while also exploiting some aspects of its embodiment (e.g., the design of its legs which may afford correct posture) to simplify control (Pfeifer and Bongard 2006). This echoes the claim that "the system" that operates locomotion is not reducible to a brain controller, and hence one need not postulate that internal models (or representations) are *within* the brain. While this latter argument is credible for tasks that require on-line engagement with the external environment (including other agents), as in the walking robot example, it is less clear whether it is sufficient for implementing higher cognitive skills that may require the *detachment* of internal generative models from the rest of the system (e.g., from online environmental dynamics). For example, it is unclear how exactly the walking robot described above may form (or select among) locomotion plans. In other words, if internal models are formed by brain-body-environment systems, they may not be detachable from on-line interactions, or they may require additional elements to be detachable such as the internal emulation of environmental dynamics. It thus remains an open question whether treating the brain-body-environment system as an internal model would be sufficient to explain the kind of phenomena I have discussed here: for example, hippocampal internally generated sequences and their roles in memory and planning; mental actions; and detached goal processing.

8 Conclusions

The key constructs of PP (e.g., prediction and prediction error, generative model, and precision) are increasingly used to explain cognitive phenomena of various complexity, ranging from action-perception loops to interoception and emotion, decision-making, planning, and beyond. However, the mere application of the same principles to several domains leaves room for different interpretations. Embodied theories of cognition tend to assume interdependence between action-perception loops and higher cognitive domains, yet it is unclear how the latter may have originated from the former.

I have discussed three examples that illustrate a general principle: cognitive (or covert) mental activities may result from an internalization process which engages brain circuits and internal generative models originally used for overt behaviour (e.g., goal-directed spatial navigation, epistemic foraging in the external environment, or acquiring food to satisfy a currently felt hunger). Generative models implied in

⁶ Technically speaking, internal states of a dynamic system can be distinguished from external states by appealing to the statistical concept of the "Markov blanket" that separates them. In this scheme, internal states model and act on external state to preserve the system's integrity (Friston 2013). However, this formulation does not prescribe where the boundary between internal and external states should be located, and whether the "Markov blanket" separates (for example) the brain from the rest of the system, or the brain-body from the rest of the system, etc. One can also appeal to the idea that there are multiple, nested "Markov blankets"; see Allen and Friston 2016.

PP may permit us to internalize (or to use fancy words, *em-brain* or *cognitivize*) key aspects of agent-environment interactions, including interoceptive loops as in the case of allostasis. This would permit the use of internal models in a “dual mode”, stimulus-tied vs. internally-generated (or spontaneous). The former mode is associated with (overt) action-perception cycles and the latter with (covert) cognitive processing which is detached from the here-and-now and can thus support, for example, future-oriented (prospective) and past-oriented (retrospective) forms of higher cognition. These examples illustrate a gradualist PP perspective in which higher cognitive abilities are distinct from action-perception loops because they run (covert or overt) PP processes, not because they are implemented in distinct modules.

Interestingly, the three examples illustrate that internalization and dual use may be implemented in various ways. In the spatial navigation example, internalization has a clear neurophysiological connotation: the same neuronal circuit (involving the hippocampus but also other brain areas) can operate in two modes, which correspond to two distinct dynamic regimes or brain rhythms (Buzsáki 2006; Buzsáki et al. 2014). In the epistemic action example, internalization refers to a functional principle — mental actions can achieve the same epistemic goals as external exploration actions — but it is currently unclear whether overt and covert forms of information foraging use shared neuronal circuits. Finally, in the example of drives and goals, internalization rests on the construction of a PP hierarchy (e.g., a hierarchy of drives and goals) that can be successively engaged in a stimulus-tied mode (search for food when hungry) and an internally-generated mode (search for food when not hungry but anticipate hunger), with the latter resting on cognitive goals that enjoy some detachment from drive states. This latter example suggests that some cognitive operations (e.g., the processing of abstract goals) can rest on hierarchical elaborations of action-perception (or interoceptive) loops, but also that the latter can be engaged when necessary during abstract goal processing (e.g., for anticipating the taste of a food).

The idea of internalization is not novel, but PP offers a conceptual framework that facilitates discussions and empirical validation, and generalizes and extends several distinct proposals within a mechanistic and biologically-grounded scheme. I focused on three examples that have been recently characterized in PP terms. However, I consider them to be illustrations of a phenomenon that may be much more general. For example, as discussed above, theories of AOR and motor cognition have provided other useful examples of the off-line reuse of motor predictions in action understanding or simulation (Jeannerod 2006). Furthermore, one can construct the formation of self-models as an internalization of the “self” as the center of predictions and experience (Hohwy and Michael forthcoming), or the construction of an epistemic agent model (Metzinger 2015). Or one can think of *somatic markers* in terms of an internalization of evaluative processes, which permits the running of “what if” loops (Damasio 1994). Finally, possible extensions of the same framework to social contexts, cultural and linguistic practices remain to be investigated (Clark 2016; Pezzulo et al. 2016b).

The extent to which one can describe higher forms of cognition in terms of the internalization of the process described here remains to be assessed, but using PP as a process model may help in charting this territory. It is also important to clarify that internalization does not exclude other ways to implement higher cognition; for example, this model is not necessarily antithetic to the idea that part of cognition is “externalized” or off-loaded to the environment, as in the case of using a computer (or even an abacus) to do maths, or rotating Tetris pieces to aid in deciding where to place them (the latter was indeed an early example of “epistemic action” in the literature, cf. Kirsh and Maglio 1994). A complete theory should encompass these and other possible roots to higher cognition.

References

- Allen, M. & Friston, K. J. (2016). From cognitivism to autopoiesis: Towards a computational framework for the embodied mind. *Synthese*, 1–24.
- Ambrose, R. E., Pfeiffer, B. E. & Foster, D. J. (2016). Reverse replay of hippocampal place cells is uniquely modulated by changing reward. *Neuron*, 91 (5), 1124–1136.
- Anderson, M. L. (2010). Neural reuse: A fundamental organizational principle of the brain. *Behavioral and Brain Sciences*, 33 (04), 245–266. <https://dx.doi.org/10.1017/S0140525X10000853>.
- Barrett, L. F. & Simmons, W. K. (2015). Interoceptive predictions in the brain. *Nature Reviews Neuroscience*, 16, 419–429.
- Barsalou, L. W. (1999). Perceptual symbol systems. *Behavioral and Brain Sciences*, 22, 577–600.
- Botvinick, M. & Toussaint, M. (2012). Planning as inference. *Trends Cogn Sci*, 16 (10), 485–488. <http://dx.doi.org/10.1016/j.tics.2012.08.006>.
- Bruineberg, J., Kiverstein, J. & Rietveld, E. (2016). The anticipating brain is not a scientist: The free-energy principle from an ecological-enactive perspective. *Synthese*, 1–28.
- Buckner, L. B. & Carroll, D. C. (2007). Self-projection and the brain. *Trends Cogn Sci*, 11 (2), 49–57. <https://dx.doi.org/10.1016/j.tics.2006.11.004>.
- Buesing, L., Bill, J., Nessler, B. & Maass, W. (2011). Neural dynamics as sampling: A model for stochastic computation in recurrent networks of spiking neurons. *PLoS Comput Biol*, 7 (11), e1002211.
- Butz, M. V. (2016). Towards a unified sub-symbolic computational theory of cognition. *Frontiers in Psychology*, 7, 925.
- Buzsáki, G. (2006). *Rhythms of the brain*. New York: Oxford University Press.
- Buzsáki, G. & Moser, E. I. (2013). Memory, navigation and theta rhythm in the hippocampal-entorhinal system. *Nature Neuroscience*, 16 (2), 130–138.
- Buzsáki, G., Peyrache, A. & Kubie, J. (2014). *Cold Spring Harbor symposia on quantitative biolog*. Emergence of cognition from action (pp. 41–50). Cold Spring Harbor Laboratory Press.
- Clark, A. (2016). *Surfing uncertainty: Prediction, action, and the embodied mind*. New York: Oxford University Press.
- (2015). Embodied prediction. In T. K. Metzinger & J. M. Windt (Eds.) *Open MIND*: 7(T). Frankfurt am Main: MIND Group. <https://dx.doi.org/10.15502/9783958570115>.
- Clark, A. & Grush, R. (1999). Towards a cognitive robotics. *Adaptive Behavior*, 7 (1), 5–16.
- Cotterill, R. (1998). *Enchanted looms: Conscious networks in brains and computers*. Cambridge University Press.
- Craig, A. D. (2015). *How do you feel? An interoceptive moment with your neurobiological self*. Princeton University Press.
- Cruse, H. & Schilling, M. (2015). The bottom-up approach: Benefits and limits. In T. K. Metzinger & J. M. Windt (Eds.) *Open MIND*: 9(R). Frankfurt am Main: MIND Group. <https://dx.doi.org/10.15502/9783958570931>. <http://open-mind.net/papers/the-bottom-up-approach-benefits-and-limits2014a-reply-to-aaron-gutknecht>.
- Damasio, A. R. (1994). *Descartes' error: Emotion, reason and the human brain*. New York: Grosset/Putnam.
- Daw, N. D., Niv, Y. & Dayan, P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nature Neuroscience*, 8 (12), 1704–1711. <https://dx.doi.org/10.1038/nn1560>.
- Dehaene, S. & Cohen, L. (2007). Cultural recycling of cortical maps. *Neuron*, 56 (2), 384–98.
- Diba, K. & Buzsáki, G. (2007). Forward and reverse hippocampal place-cell sequences during ripples. *Nature Neuroscience*, 10 (10), 1241–1242.
- Donnarumma, F., Costantini, M., Ambrosini, E., Friston, K. & Pezzulo, G. (2017). Action perception as hypothesis testing. *Cortex*. <http://dx.doi.org/10.1016/j.cortex.2017.01.016>
- Dragoi, G. & Tonegawa, S. (2011). Preplay of future place cell sequences by hippocampal cellular assemblies. *Nature*, 469 (7330), 397–401.
- Feng, T., Silva, D. & Foster, D. J. (2015). Dissociation between the experience-dependent development of hippocampal theta sequences and single-trial phase precession. *The Journal of Neuroscience*, 35 (12), 4890–4902.
- Foster, D. & Wilson, M. (2006). Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature*, 440, 680–683.
- Foster, D. J. & Wilson, M. A. (2007). Hippocampal theta sequences. *Hippocampus*, 17 (11), 1093–1099.
- Friston, K. (2010). The free-energy principle: A unified brain theory? *Nat Rev Neurosci*, 11 (2), 127–138. <https://dx.doi.org/10.1038/nrn2787>.
- Friston, K. J. (2013). Life as we know it. *J R Soc Interface*, 10 (86), 20130475. <https://dx.doi.org/10.1098/rsif.2013.0475>.
- Friston, K. J. & Frith, C. (2015a). A duet for one. *Consciousness and Cognition*.

- Friston, K. J. & Frith, C. D. (2015b). Active inference, communication and hermeneutics. *Cortex*, 68, 129–143. <http://dx.doi.org/10.1016/j.cortex.2015.03.025>.
- Friston, K. J., Adams, R. A., Perrinet, L. & Breakspear, M. (2012a). Perceptions as hypotheses: Saccades as experiments. *Front Psychol*, 3, 151. <https://dx.doi.org/10.3389/fpsyg.2012.00151>.
- Friston, K., Samothrakis, S. & Montague, R. (2012b). Active inference and agency: Optimal control without cost functions. *Biol Cybern*, 106 (8-9), 523–541. <https://dx.doi.org/10.1007/s00422-012-0512-8>.
- Friston, K. J., Shiner, T., FitzGerald, T., Galea, J. M., Adams, R., Brown, H., Dolan, R. J., Moran, R., Stephan, K. E. & Bestmann, S. (2012c). Dopamine, affordance and active inference. *PLoS Comput Biol*, 8 (1), e1002327. <https://dx.doi.org/10.1371/journal.pcbi.1002327>.
- Friston, K. J., Schwartenbeck, P., FitzGerald, T. A., Behrens, T. & Dolan, R. J. (2013). The anatomy of choice: Active inference and agency. *Front Hum Neurosci*, 7, 598. <https://dx.doi.org/10.3389/fnhum.2013.00598>.
- Friston, K. J., Rigoli, F., Ognibene, D., Mathys, C., FitzGerald, T. & Pezzulo, G. (2015). Active inference and epistemic value. *Cogn Neurosci*, 6, 187–214. <https://dx.doi.org/10.1080/17588928.2015.1020053>.
- Friston, K. J., FitzGerald, T., Rigoli, F., Schwartenbeck, P. & Pezzulo, G. (2016a). Active inference: A process theory. *Neural Computation*.
- Friston, K. J., FitzGerald, T., Rigoli, F., Schwartenbeck, P., O’Doherty, J. & Pezzulo, G. (2016b). Active inference and learning. *Neuroscience & Biobehavioral Reviews*, 68, 862–879.
- Gallagher, S. (2005). *How the body shapes the mind*. Oxford.
- Grush, R. (2004). The emulation theory of representation: Motor control, imagery, and perception. *Behavioral and Brain Sciences*, 27 (03), 377–396.
- Gupta, A. S., van der Meer, M. A. A., Touretzky, D. S. & Redish, A. D. (2010). Hippocampal replay is not a simple function of experience. *Neuron*, 65 (5), 695–705. <https://dx.doi.org/10.1016/j.neuron.2010.01.034>.
- Harnad, S. (1990). The symbol grounding problem. *Physica D: Nonlinear Phenomena*, 42 (1-3), 335–346.
- Hesslow, G. (2002). Conscious thought as simulation of behaviour and perception. *Trends in Cognitive Sciences*, 6, 242–247.
- Hills, T. T., Todd, P. M., Lazer, D., Redish, A. D. & Couzin, I. D. (2015). Exploration versus exploitation in space, mind, and society. *Trends in Cognitive Sciences*, 19 (1), 46–54.
- Hinton, G. E. (2007). To recognize shapes, first learn to generate images. *Progress in Brain Research*, 165, 535–547.
- Hohwy, J. (2013). *The predictive mind*. Oxford University Press.
- Hohwy, J. & Michael, J. (forthcoming). Why should any body have a self? In F. de Vignemont & A. Alsmith (Eds.) *The body and the self, revisited*. Cambridge, MA: MIT Press.
- Jeannerod, M. (2006). *Motor cognition*. Oxford University Press.
- Johnson, A. & Redish, A. D. (2007). Neural ensembles in CA3 transiently encode paths forward of the animal at a decision point. *J Neurosci*, 27 (45), 12176–12189. <https://dx.doi.org/10.1523/JNEUROSCI.3761-07.2007>.
- Kahneman, D. (2011). *Thinking, fast and slow*. Macmillan.
- Kirsh, D. & Maglio, P. (1994). On distinguishing epistemic from pragmatic action. *Cognitive Science*, 18, 513–549.
- Kumaran, D., Hassabis, D. & McClelland, J. L. (2016). What learning systems do intelligent agents need? Complementary learning systems theory updated. *Trends in Cognitive Sciences*, 20 (7), 512–534.
- Lisman, J. & Redish, A. D. (2009). Prediction, sequences and the hippocampus. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 364 (1521), 1193–1201.
- McClelland, J. L., McNaughton, B. L. & O’Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychol Rev*, 102 (3), 419–457.
- Metzinger, T. (2015). M-autonomy. *Journal of Consciousness Studies*, 22 (11-12), 270–302.
- (2017). The problem of mental action. Predictive control without sensory sheets. In T. Metzinger & W. Wiese (Eds.) *Philosophy and predictive processing*. Frankfurt am Main: MIND Group.
- Meyniel, F., Schlunegger, D. & Dehaene, S. (2015). The sense of confidence during probabilistic learning: A normative account. *PLoS Comput Biol*, 11 (6), e1004305.
- Montague, R. (2006). *Why choose this book? How we make decisions*. EP Dutton.
- O’Keefe, J. & Dostrovsky, J. (1971). The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Research Volume*, 34, 171–175.
- Olafsdottir, H. F., Barry, C., Saleem, A. B., Hassabis, D. & Spiers, H. J. (2015). Hippocampal place cells construct reward related sequences through unexplored space. *Elife*, 4, e06063.
- Pavlov, I. P. & Thompson, W. H. (1902). *The work of the digestive glands*. Charles Griffin.

- Pezzulo, G. (2012). An active inference view of cognitive control. *Frontiers in Theoretical and Philosophical Psychology*, 3, 478.
- (2013). Why do you fear the bogeyman? An embodied predictive coding model of perceptual inference. *Cognitive, Affective, and Behavioral Neuroscience*, 14 (3), 902–11.
- Pezzulo, G. & Castelfranchi, C. (2009). Thinking as the control of imagination: A conceptual framework for goal-directed systems. *Psychological Research*, 73 (4), 559–577.
- Pezzulo, G. & Cisek, P. (2016). Navigating the affordance landscape: Feedback control as a process model of behavior and cognition. *Trends Cogn Sci*, 20 (6), 414–424. <https://dx.doi.org/10.1016/j.tics.2016.03.013>.
- Pezzulo, G. & Rigoli, F. (2011). The value of foresight: How prospecting affects decision-making. *Front. Neurosci*, 5 (79), 79.
- Pezzulo, G., Barsalou, L. W., Cangelosi, A., Fischer, M. H., McRae, K. & Spivey, M. (2011). The mechanics of embodiment: A dialogue on embodiment and computational modeling. *Frontiers in Psychology*, 2 (5), 1–21.
- Pezzulo, G., Rigoli, F. & Chersi, F. (2013). The mixed instrumental controller: Using value of information to combine habitual choice and mental simulation. *Front Psychol*, 4, 92. <https://dx.doi.org/10.3389/fpsyg.2013.00092>.
- Pezzulo, G., van der Meer, M. A. & Lansink, C. S. and P. (2014). Internally generated sequences in learning and executing goal-directed behavior. *Trends in Cognitive Sciences*, 18 (12), 647–657. <https://dx.doi.org/10.1016/j.tics.2014.06.011>.
- Pezzulo, G., Rigoli, F. & Friston, K. J. (2015). Active inference, homeostatic regulation and adaptive behavioural control. *Prog Neurobiol*, 134, 17–35. <https://dx.doi.org/10.1016/j.pneurobio.2015.09.001>.
- Pezzulo, G., Cartoni, E., Rigoli, F., Pio-Lopez, L. & Friston, K. J. (2016a). Active inference, epistemic value, and vicarious trial and error. *Learn Mem*, 23 (7), 322–338. <https://dx.doi.org/10.1101/lm.041780.116>.
- Pezzulo, G., Iodice, P., Donnarumma, F., Dindo, H. & Knoblich, G. (2016b). Avoiding accidents at the champagne reception: A study of joint lifting and balancing. *Psychological Science*.
- Pezzulo, G., Kemere, C. & Van der Meer, M. (2017). Internally generated hippocampal sequences as a vantage point to probe future-oriented cognition. *Annals of the New York Academy of Sciences*. <https://dx.doi.org/10.1111/nyas.13329>.
- Pfeifer, R. & Bongard, J. (2006). *How the body shapes the way we think: A new view of intelligence*. Cambridge, MA: MIT Press.
- Pfeiffer, B. E. & Foster, D. J. (2013). Hippocampal place-cell sequences depict future paths to remembered goals. *Nature*, 497, 74–79.
- Piaget, J. & Cook, M. (1952). *The origins of intelligence in children*. International Universities Press New York.
- Pirolli, P. & Card, S. (1999). Information foraging. *Psychological Review*, 106 (4), 643.
- Pylyshyn, Z. W. (2001). Visual indexes, preconceptual objects, and situated vision. *Cognition*, 80 (1-2), 127–158.
- Rao, R. P. & Ballard, D. H. (1999). Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field effects. *Nat Neurosci*, 2 (1), 79–87. <https://dx.doi.org/10.1038/4580>.
- Redish, A. D. (2016). Vicarious trial and error. *Nature Reviews Neuroscience*, 17 (3), 147–159.
- Schacter, D. L. & Addis, D. R. (2007). The cognitive neuroscience of constructive memory: Remembering the past and imagining the future. *Philos Trans R Soc Lond B Biol Sci*, 362 (1481), 773–786. <https://dx.doi.org/10.1098/rstb.2007.2087>.
- Seth, A. K. (2013). Interoceptive inference, emotion, and the embodied self. *Trends in Cognitive Sciences*, 17 (11), 565–573.
- Shenhav, A., Botvinick, M. M. & Cohen, J. D. (2013). The expected value of control: An integrative theory of anterior cingulate cortex function. *Neuron*, 79 (2), 217–240.
- Spivey, M. J. & Geng, J. J. (2001). Oculomotor mechanisms activated by imagery and memory: Eye movements to absent objects. *Psychol Res*, 65 (4), 235–241.
- Sterling, P. (2012). Allostasis: A model of predictive regulation. *Physiology & Behavior*, 106 (1), 5–15.
- Stoianov, I., Genovesio, A. & Pezzulo, G. (2016). Prefrontal goal-codes emerge as latent states in probabilistic value learning. *Journal of Cognitive Neuroscience*, 28 (1), 140–157.
- Suddendorf, T. (2006). Foresight and evolution of the human mind. *Science*, 312, 1006–1007.
- Tolman, E. C. (1938). The determiners of behavior at a choice point. *Psychological Review*, 45 (1), 1.
- Varela, F. J., Thompson, E. T. & Rosch, E. (1992). *The embodied mind: Cognitive science and human experience*. Cambridge, MA: MIT Press.
- Wikenheiser, A. M. & Redish, A. D. (2015). Hippocampal theta sequences reflect current goals. *Nature Neuroscience*, 18 (2), 289–294.