

Title:

Embodied Skillful Performance: Where the Action Is

Authors:

Inês Hipólito 1, 2, 3, 4

Manuel Baltieri 5

Karl Friston 4

Maxwell J. D. Ramstead 4, 6, 7

Affiliations:

1. Department of Philosophy, University of Wollongong, Wollongong, Australia.
2. Institute of Psychiatry, Psychology and Neuroscience, King's College London, London, United Kingdom.
3. Mind, Brain Imaging and Neuroethics, Institute of Mental Health Research, University of Ottawa, Ottawa, Canada.
4. Wellcome Centre for Human Neuroimaging, University College London, London, United Kingdom.
5. Lab for Neural Computation and Adaptation RIKEN Center for Brain Science Wako, Saitama, Japan.
6. Division of Social and Transcultural Psychiatry, Department of Psychiatry, McGill University, Montreal, Quebec, Canada.
7. Culture, Mind, and Brain Program, McGill University, Montreal, Quebec, Canada.

Abstract:

When someone masters a skill, their performance looks to us like second nature: it looks as if their actions are performed smoothly without explicit, knowledge-driven, online monitoring of their performance. Contemporary computational models in motor control theory, however, are *instructionist*: that is, they cast skillful performance as a knowledge-driven process, one that is driven by explicit motor representations of the action to be performed skillfully, which harness instructions for performance. Optimal control theory, a popular representative of such approaches, casts skillful performance as the execution of *motor commands*, the deliverances of a motor control system implemented by separable forward and inverse models that work in tandem with a state estimator to control the motor plant. These models rest on the principle that motor control is realized by the concerted action of separate modular subsystems, which transform an explicit motor representation into a sequence of physical movements. This paper aims to show the limitations of such instructionist approaches to skillful performance. More specifically, we address the question of whether the assumption of modular knowledge-driven motor control in optimal control theory (based on motor commands computed by separable state estimators, forward models, and inverse models) is warranted. The first section of this paper examines the instructionist assumption, according to which skillful performance consists in the execution of instructions invested in motor representations. The second and third sections characterize the implementation of motor representations as motor commands, with a special focus on formulations from optimal control theory. The final sections of this paper examine

predictive coding and active inference – behavioral modeling frameworks that descend, but are distinct, from optimal control theory – and argue that the instructionist assumption is ill-motivated in light of new developments in motor control theory, which cast motor control and motor planning as a form of (active) inference.

Keywords:

Optimal control theory, instructionism, motor representation, action-oriented representation, active inference, skillful performance

Acknowledgements:

We are grateful to Alex Kiefer, Michael Kirchhoff, and Carlotta Pavese for helpful comments and discussions that contributed to the ideas in this paper. The authors of this article were supported by the Social Sciences and Humanities Research Council of Canada (MJDR), by the University of Wollongong (IH), and by a Wellcome Trust Principal Research Fellowship (Ref: 088130/Z/09/Z) (KJF).

1. Introduction

Expert performance dazzles us. The performance of a dance, of a musical piece, or of martial arts brings before us a display of human skills that, from a cognitive perspective, can only result from extensive practice. As opposed to bare movements, such as breathing and blinking, skillful performances are intelligent bodily activities, which harness knowledge about how to perform certain movements expertly. This knowledge, however, is not always ready-to-hand in an explicit fashion, if at all; and indeed, explicit conscious appraisal of one's performance while it is still ongoing often leads to 'choking' (Cappuccio et al., 2019). Thus, whatever kind of knowledge drives action cannot be of the explicit sort.

This action-driving form of knowledge is not limited to that which would inform the practical performance itself; skillful action is accountable to the practices within which these skills are cultivated. That is to say, the norms that govern specific cultures of practice also determine whether a skill is truly mastered and enter into the guidance and execution of the skillful performance. Accordingly, beyond merely accomplishing the physical act entailed by skillful performance, mastering a skill requires a bodily doing that is culturally embedded, situated, and intelligible within a meaningfully structured social context (Hasselberger, 2018; Hutto, 2005; Veissière et al., 2020; Hutto et al., in press).

What makes skillful performance so challenging to study – in cognitive science – is that skillful performances, as fine-grained bodily responses to salient features of a dynamically changing situation, can be described in terms of norms, knowledge, and expertise. Creatures interact with their context via intelligent behavioral adjustments, which entails varied acts of cognition, such as intending, perceiving, engaging with others in the social world, but also attending to this or that, deliberating, speaking, and so on. All of these actions are performed relative to sets of culturally sanctioned standards of practice, which must be enacted to a large extent through extensive training. This, combined with the expertise embodied in the smooth and skillful execution of a motor task, suggests peculiar knowledge in the generation of skillful performance.

Skillful performance thus stands in a seemingly paradoxical relation to knowledge: it both requires it and is confounded by it. Given its exquisite sensitivity to norms and context, and given the expertise that it requires, skillful performance is (or at least seems to be) guided by knowledge that becomes internalized through practice. However, explicit use of knowledge also seems to hamper expert performance. What is the relation between knowledge of skillful performance? What kind of knowledge, if any, guides skillful performance?

To address the epistemic and normative aspects of skillful performance, the position in the study of motor control that we will label *instructionism* (Wheeler & Clark 1999) casts skillful performance in terms of *explicit instructions*, that is, *forms of knowledge that directly guide performance* and that are *harnessed in separable structures that are internal to the performing agent* (Jeannerod, 1997, 2006; Jankovic, 2019; Pavese, 2019; Piñeros Glasscock, 2019; Stanley & Williamson, 2017; Pacherie, 2017). The instructionist assumption says that skillful performance is enabled by (what we will call) *motor representations*, which harness knowledge about how a specific skillful performance is to be executed in the form of *instructions* for movement. Instructionism, then, is the view that skillful performance depends

on the capacity of an agent to *represent to itself* explicitly the procedure to be accomplished as a set of motor instructions – and to execute those instructions for movement accordingly.

This construct of motor representation has been cashed out in different, sometimes overlapping ways. In the philosophy and cognitive science literatures, we find flavors of this construct variously formulated as “practical representations” (Pavese, 2018), “action-based ways of thinking” (Peacocke, 1986), “ability-entailing concepts” (Stanley, 2011), “executable concepts” (Pacherie, 2011), “genic representation” (Wheeler & Clark, 1999), “action-oriented representations” (Clark, 1997), and so on. What these constructs have in common is that they operationalize the kind of motor knowledge at play in the execution of skillful motor action. Some of these accounts assume that motor knowledge is harnessed in internal structures that encode explicit instructions for movement, and they will be our focus here.

In computational neuroscience, skillful performance has usually been studied under the rubric of *optimal control theory* (Stengel, 1994; Gregory, 2018; Anderson & Moore, 1990), with models often conforming to a *separation principle* (Baltieri & Buckley, 2018). This is the modularist assumption (Fodor, 1975) according to which motor control is realized by concerted processes performed by separable, modular subsystems. According to optimal control theory, skillful performance – indeed, all motor control – is realized computationally by three separate modules: the inverse model (or optimal control), forward model, and state estimator (Drayson, 2018; Jeannerod, 2018; Levy, 2017; Mylopoulos and Pacherie, 2017; Fridland 2015, 2017; Friston, 2011). Optimal control theory is instructionist, in that it posits that skillful performance is realized through the construction and execution of an explicit *motor command*, which harnesses knowledge about (instructions for) skillful, knowledge-driven motor task execution. Thus, on this model of motor control, the so-called forward model and optimal controller work together to select an optimal action, based on a value function specified in terms of desired states; where the motor command is specified in terms of instructions for movement formulated in an intrinsic frame of reference (i.e., formulated in terms of the states of motor effectors, such as stretching and compressing of muscle fibers).

The aim of this paper is to discuss critically the limitations of instructionist models of skillful performance. More specifically, we target the theoretical and empirical plausibility of separable, modular forward and inverse models and estimators responsible for the selection of actions based on a (value) function of future states, as postulated by optimal control theory. The first section of the paper characterizes the instructionist assumption, which casts skillful performance as being based in the construction and execution of motor representations. The following two sections characterize the implementation of motor representations as motor commands, focusing on computational models from optimal control theory. The final sections of this paper leverage work in predictive coding and active inference – behavioral modeling frameworks that inherit but are distinct from optimal control theory – to argue that instructionist models of motor control are ill-motivated.

2. The instructionist model of skillful performance

In this section, we examine the commitments of instructionism. Instructionist models define motor control of the kind involved in skillful performance as the execution of a set of

instructions for movements to be executed according to a prespecified method or procedure. A *motor representation* is defined as a structure internal to an agent that encodes or otherwise harnesses a set of explicit instructions for movement, the execution of which leads to skillful performance. As we will see, such a motor representation represents the specific manner in which a task is to be accomplished.

How to make sense to this? What does it mean for a thing to represent some state of affairs? It is common in the philosophy of mind to argue that representations involve modes of presentation (Frege, 1892; Millikan, 1997). This construct of mode of presentation has two main components: a representation presents some state of affairs (1) as being so-and-so (2) from a specific vantage point. For instance, when I visually perceive the presence of a red apple, I perceive it from a certain *point of view* (i.e., from my visual vantage point), precisely *as being a red apple* (i.e., as opposed to perceiving it as being, say, a fruit or as being a red object). To represent a state of affairs thus entails that we represent it in a perspectival way as being so-and-so; which is equivalent to saying that representations, essentially, must have a mode of presentation (Burge, 2009, 2010). This entails that, if there exist motor or practical representations, there must also exist a motor or practical mode of presentation.

The modes of presentation at play in perception, thought, and action involve a set of (perceptual, conceptual, and motor or practical) *abilities* that constitute a *motor or practical perspective* (Pavese 2019; Burge, 2009, 2010). Pavese's (2019) discussion of representations situates what she calls practical representations (which we equate to motor representations as defined above) with respect to other kinds – perceptual and conceptual representations. The different varieties of representation differ in the manner in which they enable agents to represent states of affairs. Consider, e.g., the nature of perspectives that are involved in the perceptual representation of a situation. On this account, perceptual abilities (e.g., being able to discriminate between a middle C and a D sharp) constitute a perspective from which one can perceive states of affairs in the world; in this case, a musical state of affairs about the key of a song. To be endowed with such perceptual abilities enables an agent to *track states of affairs* in the world from a given perceptual perspective opened by these abilities (Dretske, 1988; Millikan, 1984; Fodor, 1987). Conceptual representations, similarly, are related to the conceptual abilities with which agents represent states of affairs to themselves conceptually (Laurence & Margolis, 1999; Machery, 2009; Margolis & Laurence, 2014; Peacocke, 1992; Prinz, 2004). To represent some state of affairs conceptually thus entails the existence of a conceptual perspective, itself rooted in the conceptual abilities of the agent.

Importantly, this account allows us to fix the *content* of a representation, namely, the state of affairs that the representation is *about*, i.e. that which is disclosed by the relevant set of (perceptual and conceptual) abilities with which an agent is endowed – and which constitute the perspective from which it can represent that content. In the perceptual and conceptual cases, what is represented is the state of affairs that can be represented as being so-and-so thanks to the perspective that is opened by the perceptual and conceptual abilities with which an agent is endowed; i.e., the state of affairs that is perceived or that is entertained in thought or predicated, respectively.

Pavese (2019) extends this line of reasoning to practical representation (and which we extend to motor representation). Similarly, to perceptual and conceptual varieties, practical representations also represent by virtue of a set of *motor or practical abilities* that constitute a

perspective from which state of affairs in the world is represented practically, in a format amenable to motor control. Practical abilities are defined as abilities to execute an action in a prespecified and typified manner. The content of a practical representation is a *method*: a specific sequence of physical movements to be carried out by the agent (Wolpert 1997; Girard 1989; Pavese 2019, 2015). To be more precise, a method decomposes a particular task to be executed into component actions, perhaps nested the ones within the others, that when orchestrated bring about the desired outcome (Pavese, 2019, 2015; Mylopoulos & Pacherie 2017). Thus, to represent the world from the perspective provided by practical abilities means to represent a task as having to be accomplished practically in a prespecified manner, i.e., according to the method or procedure by which the content of the representation – the task – is presented. The distinctive feature of practical representation is their ‘direction of fit’: they function to make the state of affairs in the world fit with the prescriptions harnessed in the practical representation (Pavese, 2020). Whereas perceptual and conceptual abilities have a world-to-mind direction of fit, practical representations have a mind-to-world fit, which is what gives such representations their practical aspect.

Mylopoulos and Pacherie (2017) provide a definition of *motor representations* that dovetails nicely with Pavese’s (2019) account of practical representations and computational neuroscience research in motor control (Jeannerod, 1997, 2006). In sum, they argue: (1) that motor representations represent objects and situations in terms of their *properties relevant for action*, in a *proprietary format* specified in terms of an *intrinsic frame of reference* – defined, e.g., by the state of motor effectors, muscle fiber extension and contraction, etc.; (2) that these motor representations are informed by or contain implicitly some knowledge about the body’s *biomechanical and kinematic constraints*; (3) and that motor representations – at least usually – serve the execution of *transitive movements*, specified in terms of an *extrinsic* frame of reference (i.e., a representation of states of affairs that is ‘objective’ in three-dimensional space rather than body-dependent).

The broad strokes of this definition seem common to most specific accounts of motor representation. For instance, on Pavese’s (2019) account, motor commands (which, as we will see below, implement motor or practical representations in optimal control theory) represent the procedure or method according to which a task is to be accomplished, and are informed by a sensorimotor mapping from the actions being generated to their sensory consequences, satisfying condition (2). Moreover, they represent the method of task execution in a format that can both be used by the motor system to generate a motor action – i.e., in an intrinsic frame of reference, satisfying condition (1) – and also in a format that is sensitive to online, real time sensory feedback – i.e., in a manner that renders it responsive to outcomes specified in an extrinsic frame of reference, satisfying condition (3) of the definition just discussed.

Pacherie (2018) notes that motor representations meet criteria for representationality as set out by Bermudez (1998): they have correctness or satisfaction conditions; they have a structure that exhibits and leverages some form of compositionality (i.e., evinces identifiable constituent or elementary units); and they also have a “grammar” that regulates the assembly of the constituent units into a coherent pattern. In cognitive science, this has led to the investigation of principles common to all skills, premised on the idea that what is thus common must be some set of representational processes. This view is labeled *intellectualism* (Stanley & Williamson 2017) and can be seen as the broader rubric under which falls our target in this

article, namely, *instructionism*. At the root of such unifying models of skill is the instructionist assumption, which would allow for the construction of a general theory of skill, with epistemic attributes such as generativity, abstract rules or norms, and patterns of learning (Christensen et al. 2019; d'Avella et al. 2016; Sutton 2018).

Finally, we distinguish two kinds of instructionism (Wheeler & Clark, 1999; Wheeler, 2005), one strong and one weak. *Strong* instructionism is the claim that neural representations (in this case, motor representations) completely specify, on their own, the specific movements to be executed by an agent. We will see that this assumption is prevalent in many versions of motor control theory (e.g., Jeannerod, 1997, 2006). The *weak* version of instructionism is the more modest claim that, among the many dynamically coupled systems that generate skillful performance (e.g., an able body, a normal ecological backdrop of cultural practices and standards, and so on) one kind stands out: structures internal to an agent that are responsible for encoding information that can be interpreted as explicit instructions for action, given a background of ecologically normal processes that enable them to play this role (Clark, 1997; Engel et al., 2013).

On this more modest account, motor representations would play in the generation of behavior a role analogous to that of genes in the generation of phenotypic traits (Wheeler & Clark 1999). It is well established (e.g., Goodwin, 1994; Kelso, 1995; Varela, Thompson, & Rosch 1991; Thompson, 2007) that genes are able to code for proteins in the context of a set of factors that are causally involved in gene expression, but that do not themselves code for proteins (e.g., epigenetic transcription factors, the overall healthy and normal functioning of the cell, that cell's being embedded in an organism, etc.) (Hipólito and Martins, 2017). Analogously, the weak instructionist framework for motor representation says that skillful performance is the result of an orchestrated process spanning components in the brain, body, and world, but that of these components, some special structures in the brain play the specific, explanatorily irreducible role of encoding explicit instructions for motor performance. Note, en passant, the conformity of this definition of representation with the definition of motor representation by Mylopoulos and Pacherie (2017) that was discussed above. In what follows, we will argue that neither kind of instructionism is warranted.

3. From motor representations to motor commands

An appropriate scientific representational theory of motor action must elucidate both the kind of content in which motor representations traffic and, crucially, how such content is supposed to causally guide the generation of skillful performance – lest the story have no explanatory bite. Mylopoulos and Pacherie note that a scientifically respectable theory of motor action “cannot provide a full account of purposive action without appealing to motor representations and without explaining how intentions interface with motor representations.” (2017, p. 334). Computational models of motor control must explain the manner in which motor representations are able to play the role of interface between the conative states of an agent (that is, desires and intentions to perform some task) and the motor performance.

Pavese (2019) argues that the construct of a *motor command*, which is widely used in the study of motor control, implements the construct of practical (or motor) representations in computational models of motor control. On this model, motor tasks are realized through a process involving “a series of sensorimotor transformations that map the intentions of the agent together with visual and other sensory information about the location of the targeted objects [...] and the location of the limbs into a series of motor commands” (Pavese, 2019, p. 791). On this view, a motor command is a practical or motor representation that enables the transformation from conative states or intentions of a motor agent (i.e., the agent’s intention to perform a task according to a prespecified method) to the actual motor performance itself (i.e., to the sequence of muscle movements that together comprise the skillful action).

On Pavese’s (2019) denotational model, the *content* of a motor command is the task to be performed itself; a view which finds echoes in related theories of motor representation (e.g., Wolpert 1997; Girard 1989). More precisely, the content of a motor command is the task outcome, what the task is meant to accomplish; e.g., moving one’s body to some location in space. The motor command thus comprises the specification of the outcome of a task in an *external frame of reference* (i.e., in terms of movement in three dimensional space). A motor command is thus the *output* of a (conative) system responsible for motor planning.

We have discussed what the contents of motor or practical representations are: they represent a specific *method* or *procedure*, which is defined as the explicit specification of movements in three dimensional space (i.e., limb movements prespecified by a method or procedure, and harnessed as instructions for movement in an intrinsic frame of reference) that lead to some desired task outcome. We also examined how such practical representations get their content through their coupling to those practical abilities that open up a practical or motor perspective. The *mode of presentation* of a motor command is the *prespecified method* according to which the task is to be carried out. Thus, motor commands are also the *inputs* of the system that controls motor actions (Fridland, forthcoming). They stand as an intermediary between the conative systems of the motor agent (intention and desire) and the motor system responsible for carrying out the actual motor performance that ends up being executed.

Crucial to note is that, in order to play the intermediary role of informing the motor plant about what movements it must execute, motor commands must be generated via the inversion of a process mapping consequences in an extrinsic frame of reference, in which the desired movement is specified in terms of a task outcome in external coordinates (e.g., moving my finger to a point in three dimensional space), from an intrinsic frame of reference, specified in terms of muscle movements. This entails an *inverse inference problem*, which requires working back from the desired sensory consequences (e.g., desired visual and proprioceptive sensory feedback that confirms “my finger is now pressing the left button”) to a specification of their motor cause in an intrinsic frame of reference (i.e., a set of muscle activations that can generate such desired consequences). In other words, given some goal state that is specified in terms of extrinsic coordinates (and given conative states like desires and intentions), the problem to solve is the generation of a sequence of muscle movements, explicitly specified intrinsically in terms of stretching and compressing of muscle fibers. This has been called the “interface challenge” (Butterfill & Sinigaglia, 2014). In other words, how are motor representations implemented such that they can realize or cohere with the intentions of an agent while also instructing motor performance?

4. Motor commands and their representational role in optimal control theory

In this section, we examine how motor representations are implemented as motor commands in computational models of motor control from optimal control theory. We will see that the instructionist assumption that motor behavior is underwritten by the construction and execution of motor representations that are implemented in the brain as motor commands is, as it turns out, a pervasive one in studies of motor behavior.

This inverse inference discussed in the previous section – to wit, the problem of inferring how to specify muscle movements in an intrinsic frame of reference that bring about a goal state specified in an extrinsic frame of reference – is a nontrivial one, which has been addressed and finessed by optimal control theory. A general schema as how motor control is implemented in optimal control theory is depicted in Figure 1.

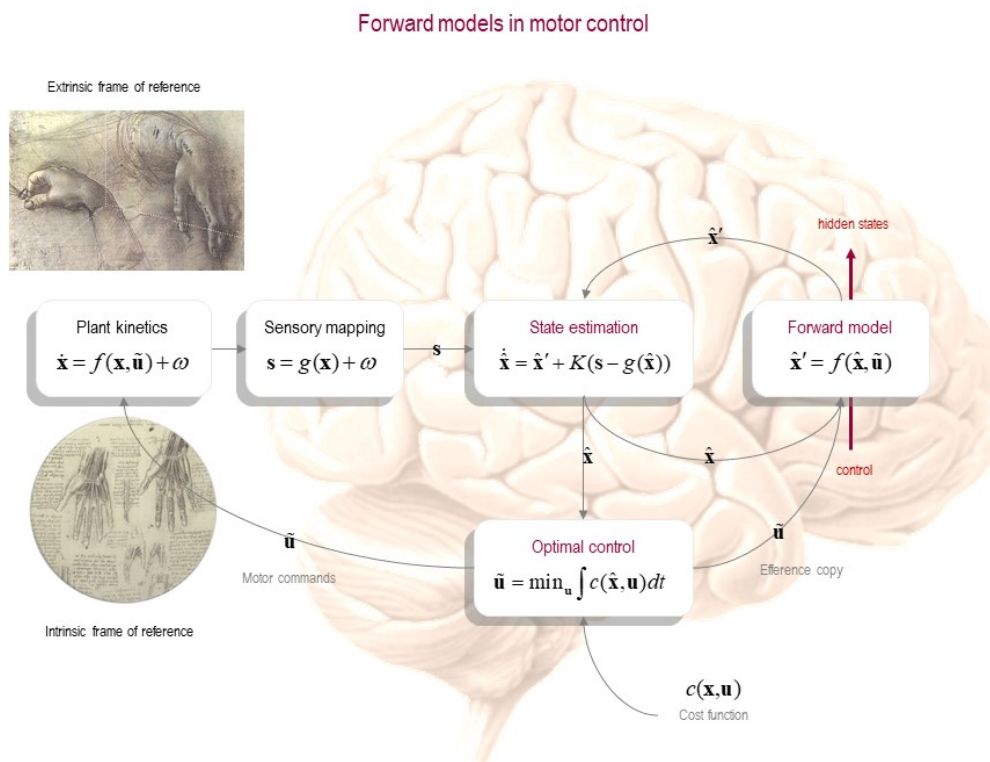


Figure 1. A computational model of optimal control. This figure presents a schematic of the computational architecture that underwrites optimal control theory. Note the separate optimal control or inverse model, state estimator, and forward model and the use of a cost function by the optimal control. Adapted with permission from Friston (2011).

In optimal control theory (Wolpert, 1997; Kawato, 1999; Todorov, 2004; Scott, 2004), there are four main components at play in the generation of motor action: the motor plant, the state estimator, the forward model, and the optimal control (also called the inverse model). The

motor control scheme functions, heuristically, as follows. The core of the model is the *optimal controller*, which tackles the inverse problem that was just discussed (hence, its other name, the inverse model). The optimal control is a mapping from desired trajectories, specified in extrinsic coordinates, to muscle movements (i.e., to changes in muscular states specified in terms of intrinsic coordinates). The optimal control selects an action based on the minimization of a *cost function*: the action that is selected is the one that leads to outcomes associated with the lowest cost or, equivalently, that leads to the most valuable states. The output of the optimal control is a *motor command*, which in our reading is a kind of practical representation, as discussed above.

Once an action is selected by the optimal control – i.e., once the optimal control has constructed a motor command – the latter is sent to the *motor plant* for execution. The motor plant is the physical motor system (e.g., a limb) that executes the task to be performed; it carries out the movement prescribed by the motor command, which contains a specification of the muscle movements needed to realize the task outcome (a representation of the method, in the parlance of practical representation theory). Thus, the optimal control generates motor commands, which implements a specific method or procedure as specified in terms muscle movements in an intrinsic frame of reference (the motor command). It follows that the motor command qualifies as a motor or practical representation in the sense discussed above.

Physical movements of the motor plant, in turn, generate sensory information. This information is conveyed to a *state estimator*, via a sensory mapping. The function of the state estimator is to infer in what state the system finds itself, given its sensory feedback. The state estimator, technically speaking, comprises a probabilistic mapping from hidden parameters and states (i.e., hidden causes) to sensory observations; and its inference process inverts this mapping, to infer the most probable hidden cause, given available sensory data.

As the motor command is being relayed to the motor plant, a copy of the motor command, known as an efference copy, is sent to a *forward model*. Actions have sensory (e.g., visual, and proprioceptive) consequences; and accordingly, the function of the forward model is to improve the execution of action by helping to finesse the inferences of the state estimator. Forward models do this by converting the (efference) copy of the motor command generated by the optimal control into a prediction of its sensory consequences, which can be discounted in state estimation. In effect, the state estimator uses information, pooled from the motor plant (via the sensory mapping) and the forward model, to form a prediction error: it compares the sensory outcome predicted by the forward model with the actual sensory data that is received from the motor plant. It uses this error to finesse its posterior state estimates. Of note is that, in optimal control schemes, this prediction error is not typically represented in the model explicitly with a distinct variable or parameter; in Figure 1, it is denoted as the update term $\mathbf{s} - \mathbf{g}(\mathbf{x})$ weighted by the Kalman gain K . Finally, posterior state estimates are used to guide the process of action selection that is carried out by the optimal control; which brings us to where we began.

The standard approach to computational models of separable subsystems is based on linear quadratic gaussian (LQG) control (Stengel, 1994). LQG-based models focus especially on formulations of perception and action in terms of (Bayesian) inference on the hidden states of the environment and on (deterministic) optimal control of a motor system (i.e., the body). Following this architecture, perception is often implemented using Kalman filters or similar

Bayesian methods for estimation; while action is modelled as a process of feedback control based on linear quadratic regulators. The applications of the LQG framework in optimal motor control are ubiquitous, but often only implicit, with a few major exceptions more directly advocating its use in cognitive (neuro)science (Todorov and Jordan, 2002; Todorov, 2004; McNamee and Wolpert, 2019). This concludes our heuristic description of motor control as it is implemented in optimal control theory.

5. The instructionist assumptions of optimal control theory

The formulation of sensorimotor control in terms of optimal control theory heavily hinges on two different, but highly interconnected, assumptions: (1) the central specification of descending motor commands, and their (efferent) copies, in the form of detailed low-level instructions for control of the motor plant, which is specified in terms of an intrinsic frame of reference (i.e., extension and contraction of muscle fibers), and (2) a separation of forward and inverse models, operating on complementary aspects of action planning and execution.

As highlighted in the previous section, the constructs of motor commands and their efference copies are typically used in frameworks focusing on the computational role of various components (the state estimator, forward and inverse models) derived from (optimal) control theoretic approaches to the problem of motor control. In this light, motor commands are cast as the product of an optimal control (or inverse model), which builds accurate action policies based on internal models of the biomechanical and kinematic properties of an agent's musculoskeletal system (the sensory mapping). While forward models are thought to emulate the mechanical properties of a body and its interactions with an environment, once a certain action policy is implemented, inverse models are normally portrayed as inverting these cause-effect relationships to form plans over future actions, based on state estimators (also called comparator models) that combine internal simulations of agent-environment couplings and desired target states.

The presence of these two models, forward and inverse, then naturally introduces the idea of different frames of reference over which internal models must operate: an intrinsic one, specified in terms of musculoskeletal properties of the body (e.g., muscle fibers), and an extrinsic, movement-based one, characterizing the external features of motor programs (e.g., hand position); see Friston, (2011) for a discussion of these ideas in the literature. In particular, a forward model takes a system from an intrinsic to an extrinsic frame, predicting the effects of different movements using musculoskeletal plans specified by neural activity, and essentially translating motor commands into actions on the world and their consequences. On the other hand, an inverse model builds motor commands by inverting this causal chain. The inverse model first leverages a value function of states, to form a mapping from desired target states in an extrinsic frame of reference (i.e., in a coordinate system based on external consequences of movements) to a set of intrinsic coordinates in the space of muscle fiber activations; and then maps these activations to a set of neural activation patterns in the motor system that are capable of generating the appropriate and desired muscle activations. From a more mechanistic perspective, frameworks based on optimal control theory are sometimes characterized in terms of “force control,” stressing the idea that, in these models, motor

commands specify actions in the form of muscle forces and joint torques (Hollerbach, 1982; Ostry and Feldman, 2003).

This architecture based on a dual frame of reference rests on the assumption, central to optimal control theory, that *value* (valuable states) is what *causes* behavior. As we have discussed, in models from optimal control theory, sequences of actions are selected according to a value function of states. This means that actions are selected by the optimal control that maximize the value of – or, equivalently, minimize the cost or risk associated with – future outcomes, defined in terms of desirable states.

A second major assumption in computational models of optimal control for action is their (often implicit) reliance on a sequential, modular architecture of perception-cognition-action, notably described as the “sense-model-plan-act” paradigm (Brooks, 1991) or the “classical sandwich” of cognition (Hurley, 2001); see Baltieri and Buckley (2018) for discussion. On this conception, action, perception, and cognition are depicted as separate processes, working relatively independently with specialized kinds of representations (practical, perceptual, or conceptual, respectively) based on different mechanistic and neurophysiological (e.g., localized) implementations (Wolpert and Kawato, 1998). This is a classical idealization of the sensorimotor loop, in which perception is portrayed as a bottom-up or feed-forward process with the primary goal of receiving information through the senses in order to build internal representations of the surrounding environment (Marr, 1982). Action is then cast as a process of deriving appropriate motor commands based on the outcomes of cognitive internal manipulations, such as thinking and planning.

This notion of separable subsystems has its roots in the classical hypothesis of the modularity of the mind (Fodor, 1983) and often constitutes one of the underlying assumptions in various applications of optimal control theory to the study of cognitive agents (Wolpert, 1997; Wolpert and Kawato, 1998); see Baltieri and Buckley (2018) and George and Sunny (2019) for some reviews. On the modularist view, more ‘peripheral’ components of cognitive systems, i.e., those subserving action and perception (but according to some, perhaps also some of “central processing”) are implemented as separable modules, working independently to transform sensations incoming through input interfaces (perception) into internal models, used to plan actions executed via output layers (motor control, behavior). The information content of each specialized module is encapsulated (i.e., its flow is restricted to the module), and the kinds of computations it performs is specialized as well; an idea closely related to the concept of cognitive impenetrability typically discussed in the context of perceptual processes (Pylyshyn 1999; Coltheart, 1999; Barrett and Kurzban, 2006; Raftopoulos, 2019).

In summary, motor control schemes in optimal control theory are *instructionist*, as we described the notion in the opening sections. This can be seen from the modular architecture in these schemes, which is based on separable forward-inverse models, estimators, and on the use of value functions to select actions. This architecture for motor control is used to compute motor commands, which implement the construct of motor representation: they harness explicit motor instructions, canvassed in a proprietary format that the motor plant can use to guide the execution of action (i.e., specified in an intrinsic frame of reference), so obtain desired states specified in extrinsic coordinates. We now critically examine this assumption.

6. Less control, more action: From optimal control to predictive coding and active inference

6.1. From forward-inverse models and cost functions to generative models

The optimal control approach has been repeatedly challenged over the years, with work questioning its neurophysiological plausibility (Ostry and Feldman, 2003; Latash et al. 2010), its computational scheme of forward and inverse models with separate roles (Adams et al. 2013, Clark 2015a; Pickering and Clark, 2014), its reliance on cost functions, and its claims regarding optimality expressed in terms of the value of states (Friston 2011; Friston et al (2012); Pezzulo et al. 2015).

The account of separable, modular perceptual and motor subsystems, in particular, has recently been suggested to reflect a classical result in the control theory literature, where modular regulators are defined using the “separation principle” (Baltieri and Buckley, 2018). In control theory, this principle describes a set of necessary and sufficient conditions for the independent optimization of the two main components of a device regulating a system in the presence of uncertainty: a paired state estimator and forward model, and a (deterministic) controller. Under the assumptions of the separation principle, teleological behavior can be cast as a sequential process of *optimal* estimation, combining state estimation and forward models, perhaps followed by a phase where internal world (forward and inverse) models are refined and used for off-line planning. This leads to an *optimal* control stage, where actions are produced by an inverse model using accurate estimates of the current state of a system. An intrinsic assumption of optimal control approaches based on the separation principle is thus that sensorimotor control is orchestrated mainly by two separate modules: a combined state estimator/forward model and control/inverse model. The assumptions behind the separation principle in control are, however, rather strict and include, for instance, the presence of linear dynamics, and the plausibility of using quadratic cost functions representing uncertainty with Gaussian noise. As previously suggested, some of these assumptions can be easily violated when applied to biological systems (Todorov, 2005; Baltieri and Buckley, 2018).

Perhaps the most important shortcoming of this approach comes from the fact that its formulation expresses the neutrality, or lack of *dual effects*, of motor signals (Bar-Shalom and Tse, 1974). In practice, this means that the canonical controls generated by LQG models cannot affect a system’s levels of uncertainty in the future, i.e., actions can only be instrumental, and have no epistemic effect on future state estimates – with a possible exception to this account found in the optimal feedback control extension of the model by Todorov and Jordan (2002). In accordance with the differences in terms of epistemic actions, approaches based on the separation principle have variously been addressed also as adaptive (as opposed to dual) controllers (Kappen, 2011), or feedback (as opposed to closed-loop) methods (Bar-Shalom and Tse, 1974).

An alternative approach can however be found in frameworks such as *active inference* (Friston et al., 2012; Friston et al. 2017). In these approaches, some of the assumptions that underwrite the separation principle are dropped in favor of a more cohesive and unifying perspective on forward and inverse models (Baltieri and Buckley, 2018); see also George and Sunny (2019). Active inference comprehensively challenges the optimal control theoretic

approach to sensorimotor behavior, highlighting some of the limitations associated with such schemes based on value functions (Friston, Adams, & Montague, 2012; Friston, 2011). First, there is good reason to believe that behavior cannot be specified by a single number – here, the single number or scalar that is tracked by the value function. Indeed, the physics of flow shows that motion in a biologically realistic state space irreducibly includes two orthogonal kinds of motion: an irrotational (or curl-free) component and a solenoidal (or divergence-free) component. Heuristically, the irrotational component is what allows the flow to climb a gradient towards more valuable or probable states; while the solenoidal component specifies a flow around an isoprobability contour, where all states entered have an equal value or probability. The irrotational component contributes the appetitive, motivated aspect to behavior, getting the agent closer to desired states or observations; whereas the solenoidal component describes behavior that does not aim directly at need satisfaction (e.g., circling around a prey, walking or simply trembling). Value functions – and indeed any motor scheme based on scalar value functions – are not up to the task of modelling behavior because, by construction, they can only account for irrotational, gradient destroying, value maximizing aspect of flow.

In a nutshell, active inference (Friston, 2020) says that action and perception are in the service of maximizing not a value function of states, but a functional of beliefs about states (known as variational or expected free energy). Active inference models question the role of inverse models, previously claimed to be physiologically unrealizable (Ostry and Feldman, 2003) and computationally intractable (Adams et al., 2013). Active inference replaces value functions and solutions to optimal control problems – formulated as motor commands based on dynamic programming methods – with *priors* (or Bayesian beliefs). That is, active inference replaces the inverse-forward model pair with a single forward model (a *generative model*) that encompasses probabilistic beliefs about expected sensory consequences of action. Rather than using a separate inverse model to infer the most appropriate course of action, active inference schemes use Bayesian inference techniques to invert the generative model in order to select action policies.

Active inference eliminates recourse to explicit value functions (Friston, Adams, & Montague, 2012; Friston, 2011). Instead of selecting actions using a (value) function of states, active inference models directly construct a prior preference over sensory outcomes or observations, which is used to guide motor control in a feedback-sensitive, online fashion, in an extrinsic frame of reference. Technically, active inference extends popular predictive coding models used in neuroscience, where perception is cast in terms of prediction error minimization (Rao and Ballard, 1999). Active inference extends this account to model motor control and explains action selection by appealing to the minimization of divergence between predicted (c.f., desired) sensory data and actual sensory consequences, e.g., in visual and proprioceptive modalities. Crucially, action is modelled as an extrinsic frame of reference (e.g., “my hand is over there”) and the forward model generates predictions of sensory consequences in an intrinsic frame (e.g., “this is what I would feel and see if my hand is over there”). The idea, then, is that rather than select an explicit motor command, the agent infers what it is doing, under prior beliefs that are realized autodidactically (see Friston, 2010). Crucially, this brings perception and action together in the same functional frame – and also explains some of the similarities between the functional architecture of sensory and motor cortices (Adams et al.,

2013). While this move from a problem of control to one of inference – in terms of predictive coding – does not make the problem mathematically easier in and of itself (Friston, 2011), it provides a hypothesis about the computational architecture that underwrites action selection, grounded in neurophysiological evidence and consistent with the literature on predictive coding models for other sensory modalities.

In this light, the active inference approach stands in stark contrast to optimal control accounts described earlier, where forward and inverse models are seen as distinct functional units with perception and action lying at the two opposite ends of a chain of sequential processing (cf. the classical sandwich of cognition). Active inference, instead, posits that the functions of inverse models are absorbed into the inversion of forward models, now building actions by inverting a hierarchical generative model, where motor commands become proprioceptive predictions – and corollary discharge becomes exteroceptive predictions.

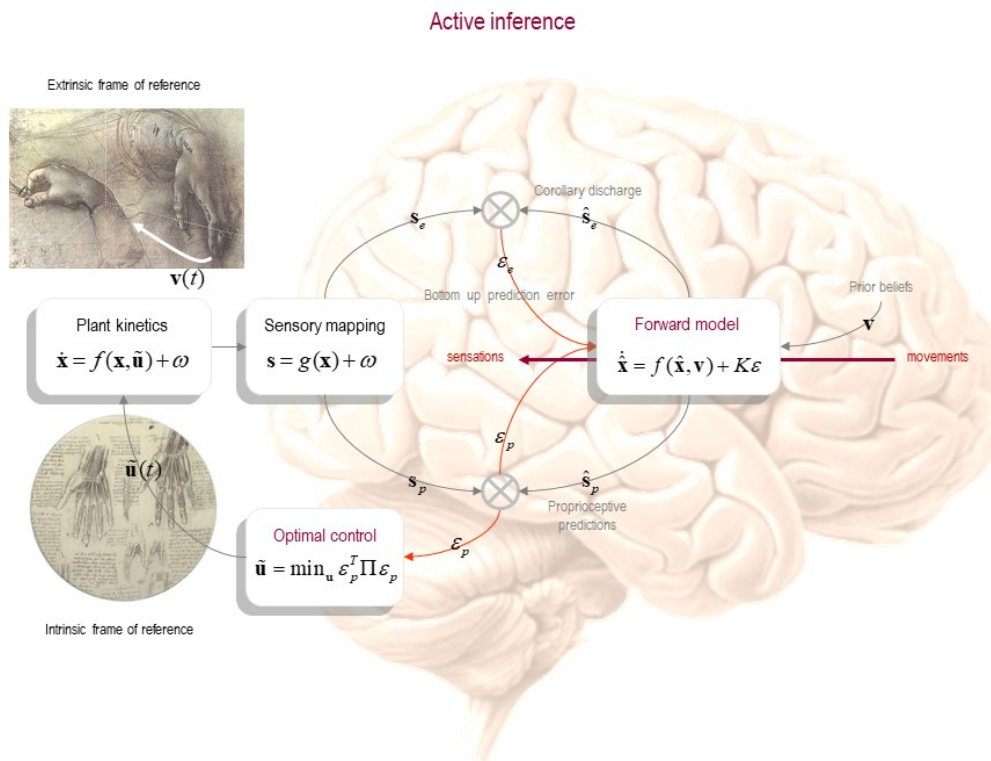


Figure 2. Motor control in active inference. This figure presents the computational architecture that underwrites active inference. Note that the cost function has been replaced with proprioceptive prediction-error based control and that the separate inverse-forward models and state estimator have been merged into a single forward (generative) model. From Friston (2011)

6.2. From motor commands to proprioceptive predictions

A second important move afforded by active inference is the replacement of motor commands in the form of accurate motor plans in intrinsic (bodily) coordinates, considered to be unrealistic due to the required specificity of a plan and the huge number of degrees of freedom of the neuromuscular system, with predictions about proprioception (Ostry and Feldman, 2003; Adams et al., 2013). This implicitly solves some of the main issues with models relying on the inversion of the many-to-one mapping from a high-dimensional intrinsic frame of reference to a low-dimensional external, movement-based, coordinate system. In practice, this summarizes the problem of motor redundancy (see Latash (2012)) – and dissolves Bernstein's problem (Bernstein, 1967) – where several combinations of different muscle activations can lead to the same final goal: think for instance of an arm reaching task and the virtually infinite number of possible arm trajectories that could satisfy a given final goal in the form of a target location.

In active inference, following models of predictive coding for perception as inference, action planning is described in terms of inverting a generative model (i.e., mapping from consequences to causes) via the inclusion of a proprioceptive modality, and an ensuing minimization of proprioceptive prediction errors. While this proposal provides an alternative, arguably more parsimonious, alternative to inverse models, it only apparently solves the least problematic aspect of instructionist models: the inversion of the process generating musculoskeletal motor plans from patterns of neural activity. The hard part still consists of ultimately explaining action execution. To solve this problem, active inference replaces the value function with prior beliefs about what an agent is doing; in other words, a hypothesis that best explains the sensory evidence at hand. This construction inherits directly from 19th-century ideomotor theory put simply: the best explanation for my sensations is that I am walking and when walking, I expect these sensations. When peripheral reflexes resolve prediction errors at the level of the spinal cord, these explanations become self-fulfilling prophecies – and my prior beliefs about walking are realized, in an embodied and enactive fashion.

Active inference proposes an account of action selection that is consistent with some ideas of the mechanical description of motor actions provided in threshold or referent control (previously also known as the “equilibrium-point hypothesis” or “virtual trajectories control hypothesis”) (Feldman, 2015). Similarly to this framework, active inference suggests that, rather than encoding muscle forces or joint torques, descending motor signals act as reference points, setpoints or thresholds for stretch receptors, in order to create movement as a “chain of reflexes” (Adams et al., 2013). Unlike referent control however, active inference commits to the idea that such thresholds can be interpreted directly in terms of proprioceptive predictions of the target state, as opposed to thresholds “lambda” typical of referent control models (Feldman, 2015).

In active inference, proprioceptors become perception-action units whose combined functions for perception and action are controlled by their (Kalman) gain or precision (Adams et al., 2013). This has two deep ramifications for motor control. First, in active inference, classical motor command and efference copy constructs of optimal control theory become redundant; and second, control assumes a dual role in active inference schemes, reflective of the dual role of action itself. The former point speaks to the idea that frameworks based on optimal control and the separation principle typically require (efference) copies of motor commands (forces and torques) to be passed from an inverse to a forward model, such that

predictions generated by forward models can discount the effects of one's own actions on one's perception of the world. While in robotics and control theory, this is classically solved by the presence of an efference copy of motor signals sent to the estimator (Kawato, 1999) that is known to the engineer/roboticist. In neurobiology, the role of this copy is hotly debated (Bridgeman, 2007; Feldman, 2009; Adams et al., 2013; Feldman, 2016). Thus, for principled reasons, active inference avoids the requirement for a controller to send an efference copy to the estimator and forward model. This is due to the fact that forward connections already denote prediction errors in their mappings from prior beliefs about expected limb trajectories to their (proprioceptive) sensory outcomes. Active inference thus softens the lines between perception and action, reconciling Helmholtz' account of perception as unconscious inference and Sherrington's description of movement as a chain of reflexes, by expressing sensorimotor control as an inseparable problem of prediction error (or free energy) minimization. Note that the notion of efference copy can be more gracefully framed in terms of corollary discharge; namely, descending predictions of the consequences of action in other sensory modalities (Sperry, 1950; Von Helmholtz, 1867; Wurtz, 2008). This is an integral part of active inference as high-level constructs generate predictions at lower levels.

Further, by building a framework that takes advantage of simple, lower-level motor functions, which are increasingly recognized as being more than simplistic, pre-programmed reflexes (Bizzi et al., 2000; Buhrmann and Di Paolo, 2014; Weiler et al., 2019), active inference introduces, at a computational level, an account of the dual effects of action at different levels. On a short spatiotemporal scale (action execution), one finds an implicit account expressed in terms of *variational* free energy (or prediction error) minimization, constrained by the dual role of proprioception in predictive coding models with reflex arcs (Friston et al., 2010). On longer time scales (such as those involved in action planning), on the other hand, a more explicit account of this exploration/exploitation problem emerges with the minimization of *expected* free energy that underwrites prior beliefs about action, and the emergence of epistemic and instrumental imperatives (Friston et al., 2017).

7. Motor control as interactive engagement with sensorimotor contingencies

Let us take stock of what has been said so far. We started from the observation that the most popular models in the field of motor control studies make an *instructionist assumption*. In instructionist models, skillful performance is explained by appealing to the construction and execution of motor commands. That is to say, these models posit motor or practical representations, which harness knowledge about how a specific skillful performance is to be executed in the form of explicit motor instructions that is specified in terms of an intrinsic (muscle-based) frame of reference. We then reviewed new frameworks in the study of motor control – namely, active inference and predictive coding – which undermine the instructionist assumption. We saw that, in these frameworks, nothing like an explicit motor command ever needs to be computed, which undermines even the weak version of instructionism (Wheeler & Clark, 1999). Where does this leave us in terms of a positive proposal? What is motor control if it does not consist in the skillful execution of motor commands?

Broadly speaking, active inference offers a model of motor control as a process of online, real-time motor adaptation to an environment; what has been cast in terms of *attunement* between and environment and its denizens (Bruineberg et al., 2014; Anderson 2017; Ramstead et al., 2019). The tight and reciprocal reconnection between perception and action in the active inference framework resonates deeply with several key ideas developed within *embodied* and *enactive* approaches to cognition and agency (Newen et al. 2018; Gallagher 2020; Ramstead et al., 2019). In particular, the inescapable codependence between action and perception in active inference coheres nicely with one brand of enactive-embodied cognition, namely, *sensorimotor approaches* to the study of cognition (Engel, Friston, & Kragic, 2015; Engel et al., 2013; Di Paolo et al., 2017; Gallagher 2020). According to the sensorimotor approach, cognition is a process of interactive engagement with the world, based on smooth online coping with minimal models enabling agents to interact with salient features of their environment (Engel, Friston, & Kragic, 2015). On this account, perception, cognition, and action are premised on the recognition of, and interaction with, *sensorimotor contingencies*, defined as a series of invariant correlations describing the relations between sensory and motor modalities (Noë, 2004). Perception is thus only appropriately defined for agents actively interacting with their milieu, when the world is dynamically coupled to an agent (Di Paolo et al., 2017); rather than on the “classical sandwich” of cognition (Hurley, 2001), which casts motor control in terms of sequential perception, planning, and action. On this account, perception and action are cast as the *mastery* of sensorimotor contingencies.

Importantly, as suggested by Di Paolo and colleagues (2017), this view reflects a spectrum of related ideas, which includes simple open-loop sensorimotor correlations, closed-loops ones, regularities given a goal, and optimal sets of regularities according to a certain performance metric. Some of these ideas speak to a move back to dynamical models of cognition, based on writing down equations of motion rather than symbolic computation; ideas that go back to ecological psychology (Gibson 1979), and which speak to the “lawful linkages between sensory and motor systems” advocated by Varela et al. (1991) and the “subjective physics” of perception (Brette 2013). When they are situated in the context of biological systems and their biomechanical constraints, sensorimotor contingencies may be seen in terms of “synergies,” capturing the attunement of different muscle groups to specific tasks engaged by an agent (Latash, 2008). Thus, instead of constructing elaborate instructions harnessed in motor representations, motor control deploys smooth real-time adaptation to the salient aspects of a situation, leveraging the biophysics of interacting physical bodies.

In active inference, a similar account emerges once we consider non-modular approaches to cognition, combining predictive approaches to perception, dynamic reflex arcs, and mechanisms for planning over expected future outcomes (Parr et al., 2018). As previously suggested, for instance, by Brette (2013) and Di Paolo and colleagues (2017), the idea of sensorimotor contingencies is well captured by simple relationships between proprioceptive sensations and motor actions. We further suggest that the predictive role of proprioception – advocated in active inference – extends causally linear account of motor control (such as the one by Brette 2013), which tend to focus only on the contingency between new actions and their proprioceptive consequences (i.e., new action \rightarrow new proprioceptive state). Active inference proposes a complementary view, where predictions of expected proprioceptive states are not just seen as passive reactions to new motor signals, but as also triggering adjustable,

dynamic reflex arcs to generate new actions (new proprioceptive state \rightarrow new action \rightarrow new proprioceptive state \rightarrow new action \rightarrow ...). The temporal depth of this model confers a more active, anticipatory role to proprioception, now seen in a causally circular model of sensorimotor control, in line with the enactive and embodied approach of Di Paolo and colleagues (2017); where action is informed by perceptual processes and perception is itself an active process of an agent engaging with the world (Baltieri 2017).

Furthermore, active inference can capture and generalize other aspects of the sensorimotor account, including, for instance, the trade-off between exploration and exploitation. In sensorimotor contingency theory, this trade-off is used to characterize sensorimotor regularities in terms of equilibrium solutions in a dynamical system analysis of an agent/environment coupled system (Maye & Engel, 2013; Di Paolo et al., 2017).

Ultimately, by invoking a formulation in terms of random dynamical systems (rather than deterministic ones, as in Di Paolo et al. 2017), active inference offers a more general and direct explanation for sensorimotor invariants in terms of nonequilibrium steady states. Technically, in active inference, the irrotational and solenoidal components of the solution flow (see section 6.1) characterize the behavior of a dynamical system in terms of components that increase/decrease value (irrotational) and maintain constant value over a trajectory of isoprobability in the phase space (solenoidal). Unlike the standard approach to the exploration-exploitation dilemma based on value functions, this formulation can define steady states in the form of trajectories rather than fixed points; and in doing so, can better capture the idea of sensorimotor *invariants* in terms of patterns whose value (of being sensorimotor contingencies functionally useful to achieve a goal) is fixed over a trajectory (e.g., simply breathing).

8. Conclusion

This paper aimed to critically discuss the limitations of instructionist approaches to skillful performance and also to assess what kind of knowledge (if any) is involved in motor control. The instructionist assumption is that according to which skillful performance is, at bottom, driven by motor representations that harness instructions about how to perform a given task. We examined the manner in which motor representations are operationalized as motor commands in optimal control theory. We asked whether the assumption of modular knowledge-driven motor control in optimal control theory, which is based on a modular architecture implementing separable state estimators, forward models, and inverse models, is warranted, and concluded that it is not. We argued that the new behavioral modeling tools and strategies from generalizations of optimal control theory – namely, active inference – show that the instructionist assumption is ill-motivated.

References

- Adams, R. A., Shipp, S., & Friston, K. J. (2013). Predictions not commands: active inference in the motor system. *Brain Structure and Function*, 218(3), 611-643.
- Ao, P. (2004). Potential in stochastic differential equations: novel construction. *J. Phys. Math. Gen.* 37, L25–L30.
- Anderson, B., & Moore, J. B. (1990). *Optimal control: linear quadratic methods*. Prentice-Hall, Inc.
- Anderson, M. L. (2017). *Of Bayes and bullets: An embodied, situated, targeting-based account of predictive processing*. Johannes Gutenberg-Universität Mainz.
- Baltieri, M., & Buckley, C. L. (2017). An active inference implementation of phototaxis. In *Artificial Life Conference Proceedings*. MIT Press.
- Baltieri, M., & Buckley, C. L. (2018). The modularity of action and perception revisited using control theory and active inference. In *Artificial Life Conference Proceedings*. MIT Press.
- Baltieri, M., & Buckley, C. L. (2019). Generative models as parsimonious descriptions of sensorimotor loops. *Behavioral and Brain Sciences*, 42, E218.
- Bar-Shalom, Y., & Tse, E. (1974). Dual effect, certainty equivalence, and separation in stochastic control. *IEEE Transactions on Automatic Control*, 19(5), 494-500.
- Barrett, H. C., & Kurzban, R. (2006). Modularity in cognition: framing the debate. *Psychological review*, 113(3), 628.
- Brette, R. (2013). Subjective physics. arXiv preprint arXiv:1311.3129.
- Bridgeman, B. (2007). Efference copy and its limitations. *Computers in biology and medicine*, 37(7), 924-929.

Bizzi, E., Tresch, M. C., Saltiel, P., & d'Avella, A. (2000). New perspectives on spinal motor systems. *Nature Reviews Neuroscience*, 1(2), 101-108.

Brooks, R. A. (1991). New approaches to robotics. *Science*, 253(5025), 1227-1232.

Bruineberg, J., & Rietveld, E. (2014). Self-organization, free energy minimization, and optimal grip on a field of affordances. *Frontiers in human neuroscience*, 8, 599.

Buhrmann, T., & Di Paolo, E. A. (2014). Spinal circuits can accommodate interaction torques during multijoint limb movements. *Frontiers in computational neuroscience*, 8, 144.

Cappuccio, M. L., Kirchhoff, M. D., Alnajjar, F., & Tani, J. (2019). Unfulfilled Prophecies in Sport Performance: Active Inference and the Choking Effect. *Journal of Consciousness Studies*, 27(3-4), 152-184.

Clark, A. (2015a). *Surfing uncertainty: Prediction, action, and the embodied mind*. Oxford University Press.

Clark, A. (2015b). Radical predictive processing. *The Southern Journal of Philosophy*, 53, 3-27.

Clark, A. (1997). *Being There*. Cambridge, MA: MIT Press.

Coltheart, M. (1999). Modularity and cognition. *Trends in cognitive sciences*, 3(3), 115-120.

Di Paolo, E., Buhrmann, T., & Barandiaran, X. (2017). *Sensorimotor life: An enactive proposal*. Oxford University Press.

Drayson, Z. (2018). The realizers and vehicles of mental representation. *Studies in History and Philosophy of Science Part A*, 68, 80-87.

Engel, A. K., Maye, A., Kurthen, M., & König, P. (2013). Where's the action? The pragmatic turn in cognitive science. *Trends in cognitive sciences*, 17(5), 202-209.

Engel, A. K., Friston, K. J., & Kragic, D. (Eds.). (2015). *The pragmatic turn: Toward action-oriented views in cognitive science* (Vol. 18). MIT Press.

Feldman, A. G. (2015). Referent control of action and perception. *Challenging Conventional Theories in Behavioral Neuroscience*.

Frege, G. (1892). Über sinn und bedeutung. *Zeitschrift für Philosophie und philosophische Kritik*, 100, 25-50.

Fridland, E. (2017). Skill and motor control: intelligence all the way down. *Philosophical Studies*, 174(6), 1539-1560.

Fridland, E. (forthcoming) The Nature of Skill: Functions and Control Structures Ellen Fridland. in Fridland, E., & Pavese, C. (forthcoming) *Routledge Handbook on Skill and Expertise*. Routledge.

Fridland, E., & Pavese, C. (forthcoming) *Routledge Handbook on Skill and Expertise*. Routledge.

Fridland, E. R. (2015). Skill, nonpropositional thought, and the cognitive penetrability of perception. *Journal for General Philosophy of Science*, 46(1), 105-120.

Friston, K. J., Daunizeau, J., Kilner, J., & Kiebel, S. J. (2010). Action and behavior: a free-energy formulation. *Biological cybernetics*, 102(3), 227-260.

Friston, K. (2011). What is optimal about motor control?. *Neuron*, 72(3), 488-498.

Friston, K., Samothrakis, S., & Montague, R. (2012). Active inference and agency: optimal control without cost functions. *Biological cybernetics*, 106(8-9), 523-541.

Friston, K., Adams, R., & Montague, R. (2012). What is value—accumulated reward or evidence?. *Frontiers in neurorobotics*, 6, 11.

Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P., & Pezzulo, G. (2017). Active inference: a process theory. *Neural computation*, 29(1), 1-49.

Gallagher, S. (2020). *Action and Interaction*. Oxford University Press.

Glick, E. (2015). Practical modes of presentation. *Noûs*, 49(3), 538-559.

Gregory, J. (2018). *Constrained optimization in the calculus of variations and optimal control theory*. Chapman and Hall/CRC.

George, N., & Sunny, M. M. (2019). Challenges to the modularity thesis under the Bayesian brain models. *Frontiers in human neuroscience*, 13.

Hay, N., Stark, M., Schlegel, A., Wendelken, C., Park, D., Purdy, E., ... & George, D. (2018, February). Behavior Is Everything: Towards Representing Concepts with Sensorimotor Contingencies. In *AAAI* (pp. 1861-1870).

Hipólito, I., Martins, J. (2017) Mind-life continuity: a qualitative study of conscious experience. *Progress in Biophysics and Molecular Biology*. Vol 131, pp. 432-444.

- Hollerbach, J. M. (1982). Computers, brains and the control of movement. *Trends in Neurosciences*, 5, 189-192.
- Hutto, D. D. (2005). Knowing what? Radical versus conservative enactivism. *Phenomenology and the Cognitive Sciences*, 4(4), 389-405.
- Hutto, D., Gallagher, S., and Ilundain-Agurruza, J. Hipólito, I. (in press). Culture in Mind - An Enactivist Account: Not Cognitive Penetration But Cultural Permeation. In L. J. Kirmayer, S. Kitayama, C. M. Worthman, R. Lemelson, & C. A. Cummings (Eds.), *Culture, mind, and brain: Emerging concepts, models, applications*. New York, NY: Cambridge University Press. (with Hutto, D., Gallagher, S., and Ilundain-Agurruza, J.)
- Jankovic, M. (2019). Ascribing practical knowledge. *Linguistics and Philosophy*, 1-29.
- Jeannerod, M. (Ed.). (2018). *Attention and performance XIII: motor representation and control*. Psychology Press.
- Jeannerod, M. 1997. *The cognitive neuroscience of action*. Oxford, UK: Blackwell Publishers, Inc..
- Jeannerod, M. 2006. *Motor Cognition: What actions tell the self*. New York, NY: Oxford University Press.
- Kappen, H. (2011). Optimal control theory and the linear Bellman equation. In D. Barber, A. Cemgil, & S. Chiappa (Eds.), *Bayesian Time Series Models* (pp. 363-387). Cambridge: Cambridge University Press.
- Latash, M. L. (2008). *Synergy*. Oxford University Press.
- Latash, M. L. (2012). The bliss (not the problem) of motor abundance (not redundancy). *Experimental brain research*, 217(1), 1-5.
- Latash, M. L. (2020). On Primitives in Motor Control. *Motor Control*, 1(aop), 1-29.
- Latash, M. L., Levin, M. F., Scholz, J. P., & Schöner, G. (2010). Motor control theories and their applications. *Medicina*, 46(6), 382.
- Levy, N. (2017). Embodied savoir-faire: knowledge-how requires motor representations. *Synthese*, 194(2), 511-530.
- Maye, A., & Engel, A. K. (2013). Extending sensorimotor contingency theory: prediction, planning, and action generation. *Adaptive Behavior*, 21(6), 423-436.

- McNamee, D., & Wolpert, D. M. (2019). Internal models in biological control. *Annual review of control, robotics, and autonomous systems*, 2, 339-364.
- Millikan, R. G. (1997). Images of identity: In search of modes of presentation. *Mind*, 106(423), 499-519.
- Mylopoulos, M., & Pacherie, E. (2017). Intentions and motor representations: The interface challenge. *Review of Philosophy and Psychology*, 8(2), 317-336.
- Newen, A., De Bruin, L., & Gallagher, S. (Eds.). (2018). *The Oxford handbook of 4E cognition*. Oxford University Press.
- Noë, A. (2004). *Action in perception*. MIT press.
- O'Regan, J. K., & Noë, A. (2001). A sensorimotor account of vision and visual consciousness. *Behavioral and brain sciences*, 24(5), 939-973.
- Ostry, D. J., & Feldman, A. G. (2003). A critical evaluation of the force control hypothesis in motor control. *Experimental brain research*, 153(3), 275-288.
- Parr, T., & Friston, K. J. (2018). The discrete and continuous brain: from decisions to movement—and back again. *Neural computation*, 30(9), 2319-2347.
- Pavese, C. (forthcoming) Practical Representation. in Fridland, E., & Pavese, C. *Routledge Handbook on Skill and Expertise*.
- Pavese, C. (2019). The psychological reality of practical representation. *Philosophical Psychology*, 32(5), 784-821.
- Perry, J. (2019). *Frege's Detour: An Essay on Meaning, Reference, and Truth*. Oxford University Press.
- Pezzulo, G., Rigoli, F., & Friston, K. (2015). Active Inference, homeostatic regulation and adaptive behavioural control. *Progress in neurobiology*, 134, 17-35.
- Pickering, M. J., & Clark, A. (2014). Getting ahead: forward models and their place in cognitive architecture. *Trends in cognitive sciences*, 18(9), 451-456.
- Piñeros Glasscock, J. S. (2019). Practical Knowledge and Luminosity. *Mind*.
- Pylyshyn, Z. (1999). Is vision continuous with cognition?: The case for cognitive impenetrability of visual perception. *Behavioral and brain sciences*, 22(3), 341-365.
- Prosser, S. (2019). Shared modes of presentation. *Mind & Language*, 34(4), 465-482.

- Stanley, J., & Williamson, T. (2017). Skill. *Nous* 51:4 713-726.
- Rao, R. P., & Ballard, D. H. (1999). Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature neuroscience*, 2(1), 79-87.
- Raftopoulos, A. (2019). Cognitive Penetrability and the Epistemic Role of Perception (pp. 223-250). Palgrave Macmillan, Cham.
- Ramstead, M. J., Kirchhoff, M. D., & Friston, K. J. (2019). A tale of two densities: Active inference is enactive inference. *Adaptive Behavior*, 1059712319862774.
- Scott, S. H. (2004). Optimal feedback control and the neural basis of volitional motor control. *Nature Reviews Neuroscience*, 5(7), 532-545.
- Stelmach, G. E., & Diggles, V. A. (1982). Control theories in motor behavior. *Acta psychologica*, 50(1), 83-105.
- Stengel, R. F. (1994). Optimal control and estimation. Courier Corporation.
- Suzuki, K., Schwartzman, D. J., Augusto, R., & Seth, A. K. (2019). Sensorimotor contingency modulates breakthrough of virtual 3D objects during a breaking continuous flash suppression paradigm. *Cognition*, 187, 95–107.
- Todorov, E. (2004). Optimality principles in sensorimotor control. *Nature neuroscience*, 7(9), 907-915.
- Todorov, E. (2005). Stochastic optimal control and estimation methods adapted to the noise characteristics of the sensorimotor system. *Neural computation*, 17(5), 1084-1108.
- Todorov, E., & Jordan, M. I. (2002). Optimal feedback control as a theory of motor coordination. *Nature neuroscience*, 5(11), 1226-1235.
- Varela, F. J., Thompson, E., & Rosch, E. (2016). *The embodied mind: Cognitive science and human experience*. MIT press.
- Veissière, S. P., Constant, A., Ramstead, M. J., Friston, K. J., & Kirmayer, L. J. (2019). Thinking through other minds: A variational approach to cognition and culture. *Behavioral and Brain Sciences*, 1-97.
- Verdejo, V. M. (2019). In defence of the shareability of Fregean self-thought. *Acta Analytica*, 34(3), 281-299.

Weiler, J., Gribble, P. L., & Pruszynski, J. A. (2019). Spinal stretch reflexes support efficient hand control. *Nature neuroscience*, 22(4), 529-533.

Wilkinson, S. (2019). Distinguishing volumetric content from perceptual presence within a predictive processing framework. *Phenomenology and the Cognitive Sciences*, 1-10.

Wolpert, D. M. (1997). Computational approaches to motor control. *Trends Cognit. Sci.* 1, 209–216. doi: 10.1016/S1364-6613(97)01070-X

Wolpert, D. M., and Kawato, M. (1998). Multiple paired forward and inverse models for motor control. *Neural Netw.* 11, 1317–1329. doi: 10.1016/S0893-6080(98)00066-5

Bernstein, N.A., 1967. *The co-ordination and regulation of movements*. Pergamon Press, Oxford.

Sperry, R.W., 1950. Neural basis of the spontaneous optokinetic response produced by visual inversion. *Journal of comparative and physiological psychology* 43, 482-489.

Von Helmholtz, H., 1867. *Handbuch der physiologischen Optik*. Voss.

Wurtz, R.H., 2008. Neuronal mechanisms of visual stability. *Vision Res* 48, 2070-2089.