

Bayesian Sensorimotor Psychology

Michael Rescorla

Abstract: *Sensorimotor psychology* studies the mental processes that control goal-directed bodily motion. Recently, sensorimotor psychologists have provided empirically successful Bayesian models of motor control. These models describe how the motor system uses sensory input to select motor commands that promote goals set by high-level cognition. I highlight the impressive explanatory benefits offered by Bayesian models of motor control. I argue that our current best models assign explanatory centrality to a robust notion of mental representation. I deploy my analysis to defend *intentional realism*, to rebut *eliminativism* and *instrumentalism* regarding mental representation, and to explore the relation between intentionality and normativity.

1. Motor Control and Mental Representation

We routinely achieve our goals by moving our bodies. The mental processes that generate these bodily motions are already quite non-trivial for everyday actions, let alone for highly skilled activities such as playing a musical instrument or hitting a tennis ball across the net. For

I am greatly indebted to Tyler Burge, Bence Nanay, and two anonymous referees for this journal for helpful feedback on earlier drafts of this paper. I also thank participants at a UC Working Group on Perception, especially Jonathan Cohen and Matthew Fulkerson, where I presented a version of this material. Finally, I thank Andy Clark, Kevin Falvey, Stephen Scott, Emanuel Todorov, and Daniel Wolpert for helpful discussion.

Address for correspondence: Department of Philosophy, University of California, Santa Barbara, CA, 93106, USA

Email: rescorla@philosophy.ucsb.edu

example, suppose I resolve to pick up a nearby book. Achieving this goal requires that my motor system estimate the book's location (so that my hand reaches to the right place), shape (so that my hand grips the book), and weight (so that I exert enough force to lift the book). Even assuming suitable estimates of these properties, choosing appropriate motor commands remains a non-trivial task. There are indefinitely many different ways I might move my body so as to pick up the book. (Try it!) Somehow, my motor system must quickly choose from among these indefinitely many alternatives. How does the motor system so reliably and so effortlessly select motor commands that promote my goals? How does it transform sensory inputs into suitable motor outputs? Such questions are the research focus of *sensorimotor psychology*, which studies the mental processes that control goal-directed bodily motion.¹

This paper investigates the philosophical foundations of sensorimotor psychology. My main thesis is that our best current sensorimotor psychology assigns a central explanatory role to *mental representation*. The science describes how representational mental states mediate between sensory input and bodily motion. It thereby presupposes the traditional picture of the mind as a representational organ. Sensorimotor psychologists elaborate that picture into well-confirmed, mathematically rigorous models of motor control.

Sections 2 and 3 discuss how current sensorimotor psychology uses *Bayesian decision theory* to model the motor system. Section 4 argues that the resulting Bayesian models assign a pivotal role to mental representation. Based on my analysis, I conclude that representational properties are scientifically indispensable aspects of mentality. To bolster my position, I critique philosophical treatments that seek to eliminate mental representation from scientific discourse

¹ 'Sensorimotor psychology' is my name for a field typically called 'motor control' by its practitioners. Many of these practitioners are institutionally affiliated with engineering, computer science, physiology, or neuroscience rather than psychology. I contend that their research is thoroughly *psychological*, by virtue of its heavy reliance upon mental representation. The label 'motor control' is regrettably ambiguous between the *study* of various mental, neural, and physiological processes versus the processes themselves.

(Section 5) or that construe representational locutions in purely instrumental terms (Section 6). I conclude by exploring the relation between representation and normativity (Section 7).

My underlying goal is to bring sensorimotor psychology into contact with traditional philosophy of mind. Contact benefits both disciplines. It benefits philosophy of mind by placing some familiar notions, such as *representational content*, in a broader scientific context. It benefits sensorimotor psychology by clarifying the field's conceptual foundations. I hope to enrich some familiar philosophical apparatus with empirical details *and* to show how that apparatus illuminates scientific research into motor control.

2. Bayesian Modeling of Motor Control

The motor system operates under conditions of considerable uncertainty. There are at least three main sources of uncertainty:

- *Ambiguity of sensory input*. Sensory stimulations underdetermine their environmental causes. For instance, I may receive retinal stimulations ambiguous between two possibilities: that the perceived object is convex and that light comes from overhead; or that the perceived object is concave and that light comes from below.
- *Noise*, i.e. corruption by random errors. Neural noise plagues the transmission of signals from sensory organs to the brain and from the brain to muscles.
- *Dynamic environment*. The environment (including the subject's own body) constantly changes, often in ways highly relevant to the motor task. Much of the variation arises from unpredictable factors, including motor noise, fatigue, injury, or external disruption.

All three factors arise even in relatively simple motor tasks (e.g. reaching towards a target), and they ramify in more complex tasks (e.g. returning an incoming tennis ball across the net). How does the motor system overcome uncertainty to achieve desired results?

Contemporary researchers address this question by deploying *Bayesian decision theory*, the standard mathematical model of reasoning and decision-making under uncertainty. The most successful models feature two main elements: *estimation* and *control*.

2.1 Estimation

On a Bayesian approach, the motor system estimates environmental state through a statistical inference. The core notion is *subjective probability* $p(h)$, where h is an hypothesis that reflects a possible state of the world: a possible hand position; or a possible shape of some distal object; or a possible weight of some distal object; and so on. Intuitively, $p(h)$ is “degree of belief” in h . A *hypothesis space* H contains all hypotheses h relevant to the motor task. The motor system maintains a probability assignment over H , regularly updated in light of motor commands and sensory input.

Motor efference copy plays a crucial role. For every motor command u , the motor system transmits a copy of u back to the Bayesian estimator. The estimator deploys a *forward model*, encoding the likely effect of u upon environmental state. In Bayesian terms, the forward model encodes a conditional probability $p(h | h^*, u)$: degree of belief in h given that h^* currently obtains and given motor command u . In effect, the forward model mirrors relevant dynamics of the environment (Grush, 2004). For example, the forward model employed during estimation of hand position mirrors the dynamics of the human arm (Wolpert, Ghahramani, and Jordan, 1995).

Execution of motor commands is noisy. For that reason, a forward model by itself is inadequate: errors quickly accumulate, rendering the state estimate useless. Sensorimotor processing corrects the forward model's initial prediction by exploiting *sensory input* (e.g. visual or proprioceptive input). *Bayes's Rule* dictates that, when one receives sensory input e , one should update the *prior probability* $p(h)$ by replacing it with the *posterior probability* $p(h | e)$: the degree of belief in h given e . *Bayes's Theorem* states that

$$p(h | e) = \eta p(h) p(e | h),$$

where the *prior likelihood* $p(e | h)$ is degree of belief in e given h , and where η is a normalizing constant to ensure that probabilities sum to 1. Bayes's Rule and Bayes's Theorem jointly specify how to reallocate probabilities across the hypothesis space in light of sensory input.

During a typical motor task, the motor system issues multiple motor commands and receives multiple sensory signals. Bayesian decision theory yields precise algorithms that dictate how to update the probability assignment in light of these developments. The algorithms embody a recursive strategy along the following lines:

- (1) ***Initial state estimate***. The forward model receives two inputs: a state estimate, and efference copy of the latest motor command. It converts these inputs into a prediction of the subsequent state resulting from the motor command.
- (2) ***Corrected state estimate***. The Bayesian estimator corrects the feedforward prediction from step (1), using sensory feedback, the prior likelihood, and Bayes's Rule.
- (3) ***Return to step (1)***, using as inputs the estimate from step (2) and efference copy of the latest motor command.

The motor system maintains a running *probabilistic state estimate* (an assignment of probabilities to hypotheses), updated in rough accord with steps (1)-(3). In some Bayesian

models, although not all, the motor system also maintains a privileged *non-probabilistic state estimate* suitably related to the probabilistic state estimate.²

Bayesian sensorimotor psychology develops these ideas into empirically well-confirmed models that describe how the motor system forms probabilistic state estimates. The estimates may concern either one's own body (e.g. current hand position) or the external environment (e.g. the location, shape, size, and weight of an external object). The resulting models help explain how sensorimotor processing surmounts the three sources of uncertainty enumerated above:

- *Ambiguity of sensory input.* The prior probability treats certain environmental conditions as likelier than others. Bayesian estimation uses the prior to discriminate among hypotheses that are equally compatible with sensory input. For example, the motor system maintains a prior that assigns higher probability to convex face shapes than concave face shapes (Hartung, Schrater, Bühlhoff, Kersten, and Franz, 2005). It uses this prior when presented with face shapes, assigning higher posterior probability to convexity even when retinal input is ambiguous between concavity and convexity.
- *Noise.* Under appropriate assumptions, the posterior probability has lower variance than either the prior probability or the prior likelihood (Bays and Wolpert, 2007). So Bayesian updating can reduce uncertainty arising from noise.
- *Dynamic environment.* A forward model generates real-time predictions about the expected consequences of motor commands. Sensory correction mitigates errors arising from motor noise, external interference, and other factors.

For more detailed overviews of Bayesian sensorimotor modeling, see (Franklin and Wolpert, 2011), (Shadmehr and Mussa-Ivaldi, 2012), (Wolpert, 2007), (Wolpert and Landy, 2012).

² The *Kalman filter* implements recursive Bayesian estimation when the priors are Gaussian and the dynamical system is linear. The Kalman filter encodes a Gaussian probability distribution by its median and its covariance matrix. However, nothing about the Bayesian framework *per se* requires non-probabilistic state estimates.

An experiment due to Saunders and Knill (2004) illustrates the empirical support for Bayesian sensorimotor models. Subjects in a virtual reality set-up reached quickly towards a target. Experimenters surreptitiously altered visual feedback regarding finger position and/or velocity. By perturbing the subject's virtual fingertip (actually a small sphere), they generated conflict between visual input and other information sources. Subjects responded by altering finger trajectory. Observed trajectories conformed to a Bayesian model describing estimation of finger position and velocity. Alternative theoretical frameworks have great difficulty explaining the observed trajectories. For example, consider a *homing model*, on which the motor system attempts to keep finger motion directed towards the target. A *rotation perturbation* shifts the virtual fingertip's position but not its motion relative to the target. The rotation perturbation generates no change in relative motion, so the homing model predicts no trajectory change. Yet finger trajectory changes markedly, disconfirming the homing model. See Figure 1 for details.

INSERT FIGURE 1 ABOUT HERE

Bayesian modeling of mental activity raises some pressing questions: Where do prior probabilities and prior likelihoods come from? To what extent do they derive from genetic endowment? To what extent, and through what mechanisms, does experience shape them? See (Clark, 2013a) for helpful discussion. For present purposes, we may set these issues aside.

2.2 Control

Motor control faces a *redundancy* problem (Bernstein, 1967). Consider a relatively simple motor goal: moving one's right hand to some location. There are infinitely many paths carrying the

hand from one location to another; infinitely many trajectories along each path; multiple joint angle combinations that achieve each trajectory; and multiple muscle activations that yield each joint angle combination. The motor system must choose from among these infinitely many redundant options. Despite the large space of options, actual observed behavior is confined to a narrow subspace. Of course, there is trial-to-trial variation. But we reliably execute sensorimotor tasks in relatively stereotyped fashion (Todorov, 2004). How does the motor system solve the redundancy problem? And why do the resulting solutions fall within a comparatively narrow range of theoretical possibilities?

The most successful models address these questions through a framework called *optimal control* (Todorov and Jordan, 2002). A *cost function* $c(u, h)$ assigns a cost to issuing motor command u when hypothesis h obtains. The cost function typically has two components:

- The first component penalizes deviation from the task goal (e.g. distance between actual hand position and desired hand position).
- The second component reflects additional task-independent desiderata. For example, it may penalize energy expenditure, or it may reward smooth motion.

A motor command is *optimal* when it minimizes expected total costs. The mathematical details here are rather complex, and we need not explore them. The key point is that, under suitable assumptions, a state estimate and a cost function determine a unique optimal motor command. At each stage of the motor task, a *controller* converts the current state estimate and the cost function into the optimal motor command. Intuitively, the motor system selects the best motor command given its estimate of environmental conditions. An optimal control model of a motor task yields a sequence of optimal motor commands, thereby predicting an average trajectory for the task.

Trial-to-trial variation from the average trajectory arises through several sources, including noise, fatigue, and external interference. Whenever a deviation occurs, the motor system faces a choice: correct the deviation or ignore it. Correction expends energy and generates further noise. The optimal policy is to correct a deviation from the average trajectory only if the deviation is relevant to the task goal. Todorov (2004) calls this *the minimal intervention principle*. To illustrate, suppose one's goal is to maintain a certain hand position. One can maintain this position while shoulder, elbow, and wrist joint angles vary considerably. These variations are task-irrelevant, so the optimal policy is to leave them uncorrected.

Many areas of human motor control --- including reaching, grasping, bimanual control, speech, and object manipulation --- display much lower variation along task-relevant dimensions than task-irrelevant dimensions (Todorov and Jordan, 2002). Task-constrained variability is precisely what optimal control predicts, because the minimal intervention principle dictates that one should correct only those deviations that bear upon the task goal. Thus, optimal control explains complex patterns of uniformity and variation in human motor control. Those same patterns elude rival theoretical frameworks. To illustrate, consider *the desired trajectory hypothesis*, according to which sensorimotor processing selects a desired limb trajectory and then tries to implement that trajectory. Assuming the desired trajectory hypothesis, one would expect sensorimotor processing to correct deviations by adjusting motor organs back to pre-selected trajectories. Yet numerous studies demonstrate a pervasive bias towards task-specific correction of motor perturbations. At every stage, the motor system selects motor commands that promote the overall task goal, without trying to maintain some pre-selected trajectory.

An experiment by Liu and Todorov (2007) illustrates the explanatory power of optimal control. Experimenters instructed subjects to reach for a target within a certain timeframe. In

some cases, the target jumped during the reaching motion. When the target jumped early in the motion, subjects easily changed course to reach it. When the target jumped late in the motion, subjects undershot the target *even though there still was ample time to reach it*. Almost all extant theories find undershooting difficult or impossible to explain, but optimal control explains it quite easily. The intuitive idea is that the cost function balances several desiderata: *accuracy* (hitting the target); *stability* (stopping once the target is reached); and *energy expenditure minimization*. Rapid last-minute course corrections consume substantial energy. They also generate considerable noise, thereby imperiling accuracy. Finally, they undermine stability by threatening to keep the hand moving when it should stop. Thus, attempting a full last-minute course correction is not optimal. Liu and Todorov codified these intuitive formulations through an optimal control model that captures observed hand trajectories quite well. The model predicts that subjects undershoot to a lesser degree when stability is relatively less important (i.e. when the cost function assigns relatively less weight to stability). Liu and Todorov confirmed this prediction by comparing two scenarios. In the first scenario, subjects were instructed to slow down their hand motion and touch the target gently. In the second, subjects were not so instructed but were instead allowed to hit the target hard. Thus, stability was an important desideratum in the first scenario but not the second. Undershooting occurred in both scenarios, but much more so in the first. So accuracy increased when stability became less important --- as predicted by the optimal control model.

2.3 Motor Control as Unconscious Inference and Decision-making

Helmholtz (1867) hypothesized that perception involves an ‘unconscious inference’ from sensory stimulations to perceptual states. Contemporary *perceptual psychology* uses Bayesian

modeling to develop a broadly Helmholtzian conception (Geisler and Kersten, 2002), (Knill and Richards, 1996). The basic idea is that the perceptual system executes an unconscious *statistical* inference from sensory input (e.g. retinal stimulations) to perceptual estimates of environmental state (e.g. estimates of the shapes, sizes, locations, and colors of distal objects). Researchers have used the Bayesian framework to explain numerous perceptual illusions and constancies. For philosophical analysis of Bayesian perceptual psychology, see (Clark, 2013a) and (Rescorla, 2015).

The optimal control framework extends Helmholtz's conception to motor control, positing unconscious inferences that estimate environmental state and unconscious decisions that select motor commands. Figure 2 summarizes the basic explanatory template. Researchers have applied this template, or variants on it, to numerous motor tasks.

INSERT FIGURE 2 ABOUT HERE

Mental activity depicted between the dotted lines in Figure 2 is *subpersonal*. The agent herself does not compute a sensory estimate, nor does she issue a motor command to her motor organs. Those tasks are instead performed by her mental subsystems. In rare cases, agents monitor and select detailed aspects of limb trajectories. But agents do not usually consider such details --- we are busy with more pressing matters. Agents typically choose a fairly abstract goal (e.g. grasping an object with one's hand), relying on the motor system to convert the goal into motor commands. Even when agents attentively guide limb trajectories, many details regarding those trajectories and the muscle activations that cause them are not consciously accessible (Pacherie, 2008). Likewise, many detailed aspects of subpersonal sensorimotor estimation are

not consciously accessible. For example, subjects in the Saunders and Knill experiment did not notice the perturbation depicted in Figure 1, yet sensorimotor processing took the perturbation into account. Thus, the motor system converts task goals into motor commands through mental processes that lie outside the agent's conscious purview. Practical reasoning interfaces with bodily motion only as mediated by these unconscious processes.

Higher-level mental activity influences subpersonal sensorimotor processing by setting a task goal. In rare cases, the task goal dictates relatively detailed limb trajectories. In more typical cases, it specifies a desideratum without dictating detailed trajectories to achieve the desideratum. For example, I may seek to move my hand to some location; or to move my hand to some location and then another distinct location; or to grasp some visually perceived object; or to hit an incoming ball with a racket; and so on. Each task goal induces a different cost function. A few exceptions aside (Trommershäuser, Maloney, and Landy, 2003), sensorimotor psychology does not explicitly model the mental processes that determine tasks goals. Most models simply take the task goal as exogenously determined by higher-level mental processes.

How does sensorimotor estimation relate to perception? The estimates deployed during motor control are subpersonal, whereas perception yields a personal-level percept that can serve as input to reasoning, planning, and decision-making by the individual. According to the 'two visual systems' hypothesis (Milner and Goodale, 1995), there exist distinct mental subsystems that process visual input for distinct purposes (motor control vs. conscious perception) and that are largely functionally independent from one another. However, this hypothesis is controversial (Briscoe, 2008), (Schenk, Franz, and Bruno, 2011). For present purposes, we need not address the controversy. I focus on the mental processes that underwrite motor control, without settling how exactly those processes relate to perceptual processing.

Current Bayesian models of motor control are often highly idealized. For example, Bayesian inference and expected cost minimization are computationally intractable in many motor tasks. We would eventually like to delineate less idealized models that posit only tractable computations. AI, engineering, and robotics offer various techniques for constructing tractable approximations to idealized Bayesian computation. How to apply these techniques to psychological modeling is a topic of active research.

3. Alternative Viewpoints

Having canvassed optimal control models of goal-directed bodily motion, let us consider some alternative viewpoints. I first examine several popular objections to the optimal control framework (Section 3.1). I then consider two rival frameworks: *the equilibrium point hypothesis* (Section 3.2); and *active inference* (Section 3.3).

3.1 Objections to Optimal Control

A recurring worry about Helmholtzian unconscious inference is that it requires an ‘inner homunculus.’ If we postulate unconscious inferences, mustn’t we also postulate a homunculus who executes the inferences? A version of this worry applies to Bayesian sensorimotor models. Don’t these models require a homunculus who consults Bayesian norms?

I respond that nothing about Bayesian sensorimotor psychology mandates an inner homunculus. Bayesian sensorimotor modeling requires that the motor system *conform* (at least approximately) to Bayesian norms. It does not require that the motor system *represent* Bayesian norms. For example, it does not require that the motor system represent Bayes’s Rule. There is

no reason why the motor system must consult Bayes's Rule in order to conform (or approximately conform) to Bayes's Rule. Modern computer science provides an existence proof. We can easily program a computer so that it conforms to Bayesian norms, including Bayes's Rule and expected cost minimization. Yet the computer need not somehow represent or consult those norms during its computations. The computer does not have an inner homunculus. This analogy shows that Bayesian sensorimotor models do not require a homunculus.

A second objection condemns Bayesian modeling as *ad hoc* curve fitting. Critics complain that a Bayesian model is a mere 'just so story' depicting how some behavior *might* have arisen (Bowers and Davis, 2012). With enough ingenuity, one can choose priors and a cost function to generate desired behavioral data. Since one can adjust these parameters at will, the Bayesian framework is vacuous and unfalsifiable.

This objection may have merit when applied to Bayesian modeling of certain mental phenomena. It is less convincing when applied to Bayesian modeling of the motor system, where the priors and costs usually reflect independently motivated physical, biological, and psychological constraints:

- The prior usually mirrors law-like or statistical regularities in the environment (e.g. that convex face shapes are much more common than concave face shapes).
- The forward model mirrors environmental dynamics (e.g. dynamics of the human arm).
- The task-dependent component of the cost function reflects the task goal itself.
- The task-independent component of the cost function often reflects well-motivated desiderata, such as energy expenditure minimization.

Of course, fitting a Bayesian model to the data always requires some curve-fitting. We must set various free parameters (e.g. variances of probability distributions). However, the model's main

qualitative predictions often do not depend upon these parametric choices. For example, the qualitative explanation summarized in Figure 1 is highly robust to parametric variation.

Most importantly, optimal control has the unifying power characteristic of good explanation. It explains task-constrained variability in diverse motor tasks: reaching, speech, object manipulation, etc. It also explains motor behavior resulting from diverse experimental interventions, including perturbed visual feedback on bodily state (Figure 1), perturbed target location (Section 2.2), and altered task instructions (Section 2.2). There are many similar results along these lines that I have not discussed. This remarkable unifying power establishes that optimal control offers genuine explanations, not just post hoc redescriptions of the data.

Another common complaint is that Bayesian sensorimotor psychology does not say how the nervous system implements Bayesian activity (Bowers and Davis, 2012), (Loeb, 2012). However, a similar complaint applies to many areas of cognitive psychology. The complaint strikes me as misguided. As Marr (1982) argues, one may legitimately describe mental activity at an abstract level that prescind from neural implementation details. Certainly, any *complete* Bayesian sensorimotor psychology must illuminate neural implementation. How does neural activity encode (or approximately encode) a probability assignment? How do neurophysiological processes implement (or approximately implement) Bayes's Rule? The literature offers some tentative suggestions (Pouget, Ma, Beck, and Latham 2013). Even lacking well-confirmed neurophysiological underpinnings, Bayesian models can explain numerous motor phenomena.

A final objection maintains that optimal control cannot accommodate various motor tasks. Sensorimotor psychologists have only used optimal control to model a fairly limited array of tasks, such as reaching and pointing. Can we extend the framework to more sophisticated

tasks, such as walking, writing, or tying one's shoes? Opponents contend that this will prove difficult or even impossible (Bowers and Davis, 2012), (Friston, 2011).

I favor a 'wait and see' attitude here. Predicting in advance whether a scientific framework will accommodate recalcitrant phenomena is a risky endeavor. For example, satisfactory treatment of the three-body problem within Newtonian physics took more than two centuries. At present, we cannot say how well optimal control will handle more sophisticated motor tasks.³ But we *can* say that it offers excellent explanations for some relatively simple motor tasks. No other known framework approaches it in explanatory power. I now substantiate this assessment by considering two prominent rival frameworks.

3.2 The Equilibrium Point Hypothesis

A long tradition in neuroscience seeks to explain motor control without invoking state estimation by the motor system. One manifestation is *the equilibrium point hypothesis* (EPH), espoused by Feldman (2011) and Latash (2010). According to EPH, voluntary bodily motion occurs when motor organs transit through a series of 'equilibrium points' along a desired trajectory. The motor system uses proprioceptive feedback to adjust resting lengths of muscles, whose springlike properties ensure that organs move through desired equilibrium points. Motor control does not involve sensory estimation of environmental state. It does not involve complex computations of any kind. It involves peripheral reflex responses to proprioceptive feedback.

Although EPH still finds some proponents, there are good grounds to find it unsatisfying.

I mention three particularly important points:

³ Friston (2011) claims that optimal control cannot model tasks, such as walking, that involve cyclic movement. Yet a suitably general version of the optimal control framework can in fact model cyclic movement, as Friston (2011, p. 494) basically concedes. For cost minimization models of robotic walking, see (Erez and Todorov, 2012), (Srinivasan and Ruina, 2006). These models conclusively demonstrate that one can model cyclic movement within an optimal control framework. (Thanks to Emanuel Todorov and Daniel Wolpert for discussion of these issues.)

- As Shadmehr (2010) notes, EPH predicts that subjects with proprioceptive loss cannot make normal voluntary movements. This prediction is false. In contrast, optimal control correctly predicts that voluntary motion can proceed absent proprioceptive feedback, because subpersonal state estimation can exploit visual feedback and efference copy.
- As Figure 1 illustrates, there is compelling evidence that the motor system continually alters trajectories based upon subpersonal estimates of bodily state. This evidence conflicts with EPH, which rejects anything like subpersonal estimation.
- EPH assumes that the motor system pursues some desired trajectory. But how does the motor system determine a desired trajectory? Proponents of EPH do not try to answer this question. They simply assume an exogenously fixed trajectory (Rosenbaum, 2002). They thereby leave unexplained a huge, vital component of the overall phenomenon we wish to explain: the mental transition from goals to motor commands. In contrast, optimal control explains how the motor system computes a trajectory in real time. Thus, optimal control offers a large increase in explanatory power.

EPH offers many insights. It shows that surprisingly sophisticated motor activity can arise from the interaction between motor reflexes and proprioceptive feedback. Proponents of optimal control should try to accommodate these insights. In that spirit, Shapiro and Körding (2010) propose a hybrid approach incorporating ideas from both traditions. Viewed as a *competitor* to optimal control, EPH is not satisfactory.

3.3 Active Inference

Friston (2011) advocates a version of Bayesian sensorimotor psychology that dispenses with cost functions. According to Friston, the motor system maintains a prior probability over possible

trajectories. Intuitively, the motor system expects motor organs to follow some desired trajectory. Using Bayesian inference, the motor system predicts proprioceptive consequences of the expected trajectory. If motor organs deviate from the expected trajectory, then proprioceptive prediction error results. Motor reflexes suppress proprioceptive prediction error, steering motor organs towards the expected trajectory. Bodily motion does not reflect an attempt by the motor system to minimize costs. Rather, it reflects an ongoing effort to minimize sensory predictive error. Hohwy describes the view as follows: ‘Action therefore does not come about through some complex computation of motor commands that control the muscles of the body. In simple terms, what happens instead is that the muscles are told to move as long as there is prediction error’ (2014, p. 82). Friston calls this approach *active inference*. See (Clark, 2013b) and (Hohwy, 2014) for philosophical analysis of active inference.

Although active inference employs Bayesian apparatus, the underlying picture differs considerably from optimal control. In many respects, the picture more closely resembles EPH. Like EPH, active inference rejects complex computation of motor commands. Like EPH, active inference holds that interaction between motor reflexes and proprioceptive feedback impels motor organs towards desired trajectories.

Friston (2011, p. 492) writes: ‘Perhaps the most definitive argument in favor of active inference, as a normative model of motor control, is that prior beliefs about behavior emerge naturally as top-down or empirical priors during hierarchical perceptual inference. This contrasts with optimal control, which, at the end of the day, still has to explain how cost functions themselves are optimized. In short, active inference eliminates the homunculus implicit in cost functions.’ Friston does not say how hierarchical perceptual inference generates an expected trajectory. Despite what he intimates, his model does *not* explain how expected trajectories come

to be expected. His model simply takes expected trajectories as exogenously determined, even for relatively simple reaching motions. Thus, his model does not explain why the motor system follows one trajectory rather than another when pursuing a task goal.⁴ By assuming an exogenously determined prior over trajectories, Friston leaves unexplained how the motor system selects a trajectory. Like proponents of EPH, Friston does not even try to explain a huge, vital component of the mental transition from goals to motor commands.⁵

Optimal control explains how the motor system determines a trajectory in real time. Contrary to what Friston suggests, the explanation requires no homunculus. True, the explanation assumes a cost function with exogenously determined properties. But scientific theories always presuppose exogenously determined explanantia. The question is whether those explanantia subserve good explanations. By assuming an exogenously determined but well-motivated cost function, optimal control explains numerous features of the resulting bodily motions. To illustrate, consider the Liu and Todorov (2007) experiment, where different instructions cause different degrees of undershooting. Optimal control readily explains this phenomenon: different instructions cause the motor system to instantiate a different cost function, thereby causing a different sequence of optimal motor commands. Active inference does not explain the variation in undershooting, because it does not explicitly model causal influences upon the putative prior over trajectories.

Active inference also seems unable to explain task-constrained variability in motor behavior (Scott, 2004). According to Friston, the motor system strives to eliminate prediction errors that arise when the actual trajectory deviates from some expected trajectory, whether or

⁴ In certain passages, Friston appears to reject the very idea of goal-directed bodily motion. But I think it more charitable to interpret him as countenancing goal-directed bodily motions and as trying to elucidate the subpersonal mental processes that produce those motions. Hohwy (2014, pp. 81-92) pursues an interpretation along these lines.

⁵ My argumentation in this paragraph is heavily influenced by unpublished work by Emanuel Todorov and Tom Erez, who argue along similar lines. Needless to say, the blame for any mistakes or infelicities lies entirely with me.

not the deviations are task-relevant. This framework provides no principled reason to expect task-constrained variability. In contrast, optimal control offers a principled explanation for task-constrained variability: namely, that the optimal policy is to correct only task-relevant deviations. By invoking expected cost minimization, optimal control explains a crucial phenomenon that Friston does not even try to explain. So optimal control once again offers decisive advantages over active inference.

4. Representation as Explanatorily Central

I now discuss how sensorimotor psychology illuminates mental representation.

Philosophers and scientists use the phrase ‘mental representation’ in many different ways. The type of mental representation that concerns me involves *representational content* (sometimes called *intentional content*). In many important cases, a mental state has a content that represents the world as being a certain way. We can ask whether the world is indeed that way. These states are *semantically evaluable* with respect to such properties as truth, accuracy, and fulfillment (Fodor, 1987, pp. 10-11), (Searle, 1983, pp. 10-13). To illustrate:

- Beliefs are evaluable as true or false. For example, my belief *that John has gone fishing* is true iff John has gone fishing.
- Perceptual states are evaluable as accurate or inaccurate. For example, my perceptual experience *as of a red sphere standing before me* is accurate only if a red sphere is standing before me.
- Intentions are evaluable as fulfilled or unfulfilled. For example, my intention *to donate money to Oxfam* is fulfilled iff I donate money to Oxfam.

In all three cases, a mental state is associated with some condition under which the state achieves ‘representational success’ (truth, accuracy, or fulfillment). Intuitively, the state achieves representational success when it correctly represents reality. Following Burge (2010), I use *veridicality* as an omnibus term to describe representationally successful mental states. Burge (2010) regards *truth* and *accuracy* as distinct species of veridicality, but I remain neutral on this question. The crucial point is that many important mental states are associated with *veridicality-conditions*: conditions for veridical representation of reality.

There are many conflicting philosophical theories of representational content. The basic idea behind most theories is to posit abstract entities --- *representational contents* --- related in some systematic way to veridicality-conditions. Popular candidates include Fregean thoughts, Russellian propositions, and sets of possible worlds. I will not wade into these turbulent waters. For our purposes, the differences among philosophical theories of representational content do not matter. What matters is a core doctrine shared by virtually all such theories: many important mental states are evaluable as veridical or non-veridical.

What is it for mental states to have representational content? Various philosophers have tackled this question, trying to elucidate representation in non-representational terms. I will not take sides here. I am not addressing how mental states come to have representational content. My concern is the more basic thesis that many mental states *have* representational content, in the minimal sense that they are associated with conditions for veridical representation. As Bermúdez (1995) stresses, both personal-level mental states and subpersonal mental states can have representational content in this minimal sense.

I contend that our current best sensorimotor psychology assigns a central role to representational content. The science posits *intentions* that have fulfillment-conditions (e.g. an

intention to move my hand to a location, or an intention to grasp some perceived object). It also posits *subpersonal state estimates*: probability assignments to hypotheses with accuracy-conditions (e.g. hypotheses that represent hand position, or hypotheses that represent the location, shape, or weight of a perceived object). Mathematically precise generalizations describe how these mental states interact with one another and with sensory input to yield bodily motion. Thus, sound philosophical discussion of motor control must acknowledge the indispensable role that representational mental states play in mediating between sensory input and motor output.

I will develop my position by considering intentions (Section 4.1) and subpersonal state estimates (Section 4.2). I will then discuss how mental representation informs the explanatory generalizations offered by sensorimotor psychology (Section 4.3).⁶

4.1 Intention

Philosophical discussion of practical reason assigns a pivotal role to *intentions* (Anscombe, 1957), (Davidson, 1980). Philosophers vigorously debate the nature of intentions, but most discussants agree that intention has the following three properties:

- (a) It is a personal-level state, i.e. a state of the individual rather than her subsystems.
- (b) It represents some goal of the individual.
- (c) It has a fulfillment-condition, which is satisfied if the intention causes bodily motion that achieves the relevant goal.

⁶ Grush (2004) also discusses the relation between sensorimotor psychology and mental representation. He focuses especially on forward models and the Kalman filter. I place much more emphasis than Grush does upon Bayesian underpinnings and upon veridicality-conditions. I also deploy the science to rebut various views (e.g. eliminativism and instrumentalism) that Grush does not explicitly consider.

I claim that sensorimotor psychology presupposes mental states with these properties. I call the mental states ‘intentions,’ but I am not concerned with how closely they resemble intentions as standardly construed by philosophers. The key point is that they have properties (a)-(c).

As Shadmehr, Smith, and Krakauer observe, ‘[m]otor control is the study of how organisms make accurate goal-directed movements’ (2010, p. 89). Our topic is *goal-directed* bodily motion, i.e. motion in pursuit of a goal. Research proceeds from the assumption that people routinely achieve their goals by moving their bodies. We want explain how people do this so well. In some cases, an individual may not consciously choose the relevant goal. While immersed in my work, for example, I may thoughtlessly scratch my cheek without forming any conscious intention to do so. However, many core cases involve goals consciously pursued *by the individual*, not merely her subsystems (Pacherie, 2000), (Jeanerrod, 2006, pp. 1-21). To pursue a goal, the individual must mentally represent the goal as one to be pursued. She must instantiate a personal-level mental state that is fulfilled when resulting bodily motions achieve the goal. In other words, she must instantiate a mental state with properties (a)-(c).

For this reason, sensorimotor psychology presupposes intentions. Researchers want to explain how mental activity converts intentions into bodily motions that fulfill or approximately fulfill those intentions. We cannot even specify the explanandum unless we presuppose personal-level mental states with fulfillment-conditions.

Intention plays a particularly central role within optimal control modeling. The main idea behind the optimal control framework is that motor activity chooses commands suited to advance some goal. This idea underlies the framework’s distinctive explanations, including the explanation for task-constrained variability. In many cases (though perhaps not all), the relevant

goal is set by higher-order cognition. In any such case, there must be a personal-level mental state that represents the goal, i.e. an intention.

More concretely, suppose we want to model some reaching task using optimal control. Then we must isolate a cost function appropriate to the task. The cost function enshrines a task goal (e.g. *reach to this location and then that location*). When deciding which cost function to attribute to the motor system, we evaluate what goal the subject is pursuing. In other words, we identify the intention operative during the motor task. Identifying the operative intention is easier in some cases than others. In a typical experiment, we simply tell the subject what goal to pursue. We assume that the subject understands our instructions and seeks to obey them. In more naturalistic settings, identifying the subject's intentions may be less straightforward. In all cases, we explicitly or implicitly rely upon our general folk psychological capacity to discern one another's mental states. We deploy this capacity to identify the operative intention, thereby constraining the class of plausible cost functions. Thus, applying an optimal control model to a specific individual requires us to consider the individual's intentions.

A good illustration is the Liu and Todorov (2007) experiment, where instructions influence degree of undershooting. The optimal control explanation for this phenomenon hinges upon a contrast in the cost function, reflecting the relative importance of stability. But what causes the contrast in the cost function? And how do we decide which cost function to attribute in which scenario? Intention is key here. The subject understands our instructions and forms a suitable intention, which causes the motor system to instantiate a suitable cost function. We decide which cost function to attribute by evaluating what the subject intends (i.e. whether she aims to hit the target gently, or whether she merely aims to hit it). Our explanatory paradigm presupposes personal-level intentions that mediate between instructions and cost functions.

Most intentions influence behavior only as mediated by further intentions. My intention to donate money to Oxfam causes bodily motion only as mediated by an intention to sign my name to a check, an intention to place the check in the mailbox, and so on. As this example illustrates, only relatively ‘low-level’ intentions can serve as input to subpersonal sensorimotor processing. Which intentions can interface directly with subpersonal processing? Where should we draw the line between the personal-level processes that produce intentions and the subpersonal processes that convert intentions into bodily motions? Philosophers standardly tackle these questions by distinguishing among intentions. For example, Searle (1983) distinguishes between *prior intentions* and *intentions-in-action*, while Mele (1992) distinguishes between *distal intentions* and *proximal intentions*. The details vary considerably, but the basic idea is to identify relatively low-level intentions that can initiate a bodily motion, sustain the motion until completion, and guide the motion as it unfolds (Pacherie, 2006). There is room here for fruitful interchange between philosophers and scientists. Such interchange might clarify how personal-level decision-making and subpersonal motor activity relate. For present purposes, we need not settle precisely which intentions interface directly with subpersonal processing.⁷

4.2 Subpersonal State Estimation

Intuitively, we often achieve our goals because we have true beliefs. The models canvassed in Section 2 extend this intuitive idea to the subpersonal realm. Researchers postulate a *probabilistic state estimate*: a probability assignment to hypotheses. The science remains fairly

⁷ Pacherie (2008) distinguishes between *distal intentions*, *proximal intentions*, and *motor intentions*. Motor intentions are non-propositional representational mental states that represent the task goal and that figure in motor processing. She offers a detailed theory of the ‘intentional cascade’ from distal to proximal to motor intentions. Similarly, Butterfill and Sinigaglia (2014) argue that we should recognize non-conceptual *motor representations* of the task goal, distinct from personal-level intentions. I am sympathetic to the thesis that subpersonal, non-propositional, non-conceptual representations of the task goal mediate between personal-level intentions and subpersonal cost assignments. Officially, though, I am not committed to that thesis.

neutral regarding the nature of ‘hypotheses,’ but it makes one crucial assumption: each hypothesis is accurate or inaccurate, depending on the environment. Thus, each hypothesis has an accuracy-condition. For example, an hypothesis that represents some object as having a certain weight is accurate only if the object has that weight. The Bayesian estimator begins with priors that mirror law-like or statistical regularities in the environment (e.g. that certain shapes are likelier than others). Efference copy and sensory input cause reallocation of probabilities, conforming roughly with Bayesian norms. Assuming favorable environmental conditions, the resulting state estimate assigns high probability to accurate or approximately accurate hypotheses. The motor system selects a motor command apt to advance the task goal, assuming that high probabilities attach to approximately accurate hypotheses. The motor system thereby converts sensory inputs into bodily outputs apt to advance the task goal.

This explanation turns upon the notion of *accurate representation*. To explain why motor behavior yields (approximately) intended results, we cite high probability assignments to (approximately) accurate subpersonal hypotheses. We presuppose subpersonal hypotheses with accuracy-conditions, and we cite satisfaction of those accuracy-conditions to explain why bodily motions fulfill intentions.⁸

Sensorimotor psychology also studies cases where motor behavior does not yield desired results. There are diverse possible sources of failure: neural noise, muscle fatigue, internal malfunction, external interference, and so on. One particularly notable source of failure is *misestimation*, i.e. assignment of overly high probabilities to inaccurate hypotheses. Figure 1

⁸ Nanay (2013) postulates *pragmatic representations* that mediate between sensory input and motor output. Pragmatic representations represent action-relevant features of the environment. For example, when I pick up a book, my motor system deploys pragmatic representations of the book’s location, shape, size, and weight. Pragmatic representations have accuracy-conditions, and they are typically unconscious. Thus, they share many important features with the probabilistic state estimates emphasized by my account. One difference is that Nanay’s pragmatic representations do not seem to be probabilistic in nature. As indicated in note 2, I am open to the possibility that motor control deploys non-probabilistic state estimates, but nothing about the Bayesian framework per se requires such items.

illustrates how misestimation can stymie consummation of the task goal. To explain why the subject's virtual finger misses the target, we cite an inaccurate subpersonal estimate of current finger state (or overly high probability assignments to inaccurate subpersonal hypotheses). Our explanation presupposes subpersonal hypotheses with accuracy-conditions.

There are important connections here between sensorimotor psychology and perceptual psychology. In a *perceptual illusion*, misleading sensory input causes an inaccurate percept. Analogues to many well-known perceptual illusions --- including the *Ebbinghaus illusion* (Franz, Bühlhoff, and Fahle, 2003) and the *Müller-Lyer illusion* (Bruno and Franz, 2009) --- arise for sensorimotor estimation. Consider *the hollow face illusion*, in which a concave mask appears convex despite several visual cues that it is concave. A comparable illusion afflicts motor control: subjects trying to touch the nose on the mask fail systematically, due to misestimation of the mask's shape (Hartung, Schrater, Bühlhoff, Kersten, and Franz, 2005).⁹ There are at least two potential contrasts between perceptual illusions and corresponding motor illusions:

- (a) A perceptual illusion concerns a personal-level percept. A motor illusion concerns a subpersonal estimate. How the personal-level percept and the subpersonal estimate relate to one another is an important open question.
- (b) Introspectively, it seems clear that perception produces a determinate non-probabilistic percept. In contrast, we do not know whether the motor system produces a non-probabilistic state estimate. It may simply deploy an updated probability assignment.

⁹ Proponents of the 'two visual systems' hypothesis frequently claim that motor control resists various well-known perceptual illusions, including the Ebbinghaus, Müller-Lyer, and hollow face illusions. Numerous papers by Franz and his colleagues, including the papers cited above, cast such claims into doubt. Carefully designed experiments that match the perceptual task to the motor task seem to eradicate the alleged asymmetry between perception and motor control. One apparent exception is the *size-weight illusion*: perception dramatically overestimates the weight of smaller objects, but the motor system is somewhat biased in the opposite direction (Brayanov and Smith, 2010).

Still, there is a fundamental affinity between perceptual illusion and motor illusion. In both cases, misleading sensory input causes sensory misestimation of environmental conditions.

Bayesian perceptual psychology offers satisfying explanations for numerous perceptual illusions (Rescorla, 2015). The same basic template can also explain numerous motor illusions. For example, the hollow face illusion arises because the prior assigns high probability to convex face shapes (Hartung, Schrater, Bühlhoff, Kersten, and Franz, 2005). A concave mask flouts this prior, thereby inducing Bayesian misestimation of shape. More generally, Bayesian estimation deploys priors to overcome the inherent ambiguity of sensory input. No matter how apt the priors, there are possible circumstances that flout them. Thus, the strategy through which sensory estimation handles ambiguous sensory input ensures that sensory misestimation will occur under certain conditions. Misestimation in unfavorable circumstances is a corollary of successful estimation in favorable circumstances. (Cf. Burge, 2010, pp. 87-98.)

4.3 Representational Explanation in Sensorimotor Psychology

I submit that representation is firmly embedded within our current best sensorimotor psychology. The science explains how motor activity causes bodily motions that fulfill (or approximately fulfill) intentions. The central explanatory strategy is to illuminate why subpersonal estimation routinely assigns high probability to accurate (or approximately accurate) hypotheses. The science also illuminates cases where intentions go unfulfilled due to subpersonal misestimation, i.e. overly high probability assignments to inaccurate hypotheses. These explanations presuppose personal-level intentions with fulfillment-conditions and subpersonal hypotheses

with accuracy-conditions. The aims, methods, and results of our current best sensorimotor psychology embody a representationalist viewpoint.¹⁰

Optimal control modeling describes how the controller converts state estimates and cost assignments into bodily motions. A state estimate allocates probabilities to hypotheses, where each hypothesis h is accurate only if the environment meets certain conditions. A cost assignment allocates costs to pairs (u, h) , where u is a motor command and where h once again has an accuracy-condition. Thus, state estimates and cost assignments are defined over representationally contentful hypotheses. The hypotheses might represent possible hand positions, or possible shapes of some distal object, and so on. Mathematically precise *ceteris paribus* generalizations describe how a state estimate and a cost assignment jointly determine a specific motor command. The generalizations characterize state estimates and cost assignments by describing hypotheses in representational terms --- *as* representations of specific hand positions, or specific distal shapes, and so on. The generalizations describe how mental states *as characterized partly through environmental conditions those states represent* interact with one another and with sensory input to cause motor commands.

Current sensorimotor psychology also describes formation of subpersonal state estimates in great detail. It cites priors, the forward model, sensory input, and efference copy to explain how the motor system estimates environmental conditions. For example, Saunders and Knill (2004) present a Bayesian model that describes estimation of finger state based upon visual feedback and efference copy. The model delineates precise rules for reallocating probabilities over a hypothesis space whose component hypotheses represent possible finger positions and velocities. Detailed explanatory generalizations describe how the current state estimate *as described representationally* influences the next state estimate *as described representationally*.

¹⁰ In (Rescorla, 2015), I develop a kindred representationalist analysis of Bayesian perceptual psychology.

Hence, representational mental states figure crucially as both *explanantia* and *explananda* within our current best sensorimotor psychology. To explain bodily motion, the science cites representational properties of mental states, including intentions, subpersonal state estimates, and subpersonal cost assignments. To explain a state estimate with certain representational properties, the science cites representational properties of an earlier state estimate.¹¹

Admittedly, sensorimotor psychologists do not use locutions such as ‘representational content’ or ‘veridicality-condition.’ Nevertheless, such locutions illuminate the explanatory structure of optimal control sensorimotor modeling. Here again there is room for fruitful exchange between philosophers and scientists. Tools from philosophy of mind can help sensorimotor psychologists better understand the conceptual foundations of their own research.

4.4 Intentional Realism

I favor a broadly scientific realist viewpoint: explanatory success is a *prima facie* guide to truth. From a scientific realist viewpoint, sensorimotor psychology supports *intentional realism*, i.e. realism towards mental representation. Our current best sensorimotor psychology offers successful explanations for diverse bodily motions. The explanations attribute veridicality-conditions at both the personal level and the subpersonal level. There is no evident way to forego these intentional attributions while preserving the explanatory benefits that they provide. Thus,

¹¹ Consider two thinkers X_1 and X_2 who are duplicates save that they inhabit distinct, qualitatively indistinguishable spatial environments. Suppose that each thinker’s motor system forms a subpersonal non-probabilistic estimate of hand position. Call the estimates E_1 and E_2 . The estimates have different accuracy-conditions, because each estimate denotes a different location. For many explanatory purposes, one may want to “hive off” the external denotation, isolating context-free representational properties shared by X_1 and X_2 . I believe that this is how current models proceed. The models cite context-free representational properties, abstracting away from contextually-determined differences between subjects. The models individuate subpersonal estimates not through contextually-determined veridicality-conditions, but through context-free representational properties that *contribute to* veridicality-conditions. Developing my analysis would require extensive inquiry into various subtle issues involving context, content, and psychological explanation. For present purposes, it does not matter *which* representational properties of mental states figure within sensorimotor modeling. What matters is simply that the science assigns a crucial explanatory role to *some* representational properties of mental states.

current sensorimotor psychology assigns a central explanatory role to representational content. We should embrace representation as a genuine, scientifically indispensable aspect of mentality.

Intentional realism is a fairly popular position, espoused by Burge (2010), Fodor (1987), Peacocke (1992), and many others. However, it has encountered substantial opposition over the past few decades. To bolster my intentional realist perspective, I will now examine opposing *eliminativist* (Section 5) and *instrumentalist* (Section 6) positions.

5. Eliminativism

Eliminativists seek to expunge representationality (or *intentionality*) from scientific psychology. In this spirit, Quine bemoans ‘the baselessness of intentional idioms’ (1960, p. 221). He maintains that, when we are ‘venturing to formulate the fundamental laws of a branch of science,’ we should reject intentionality in favor of the ‘the physical constitution and behavior of organisms.’ Similarly, Stich wants to ‘banish talk of content in scientific settings.’ He dismisses ‘intentional locutions’ as ‘not the sort of locutions we should welcome in serious scientific discourse’ (1991, p. 240). Both philosophers think that scientific psychology should eschew intentional discourse. Churchland (1981) also inclines towards this position, although officially he only rejects personal-level propositional attitudes. Quine and Stich explicitly aim to purge scientific psychology of *all* intentional attribution, including intentional attributions to subpersonal states. They hold that rigorous theorizing about the mind should reject representational content *in general*.

Eliminativists defend their position through various arguments, which I will not try to canvass or rebut. Burge (2010, pp. 296-298) contends that the arguments are unconvincing, and I

agree. Burge (2010) also argues that eliminativism conflicts with perceptual psychology, a conclusion that I likewise defend in (Rescorla, 2015). For present purposes, my main point is that eliminativism conflicts with our current best sensorimotor psychology. Bayesian models of motor control reveal a rich structure of representational mental states that mediate between sensory stimulation and bodily motion. The science posits personal-level intentions with fulfillment-conditions. It also posits subpersonal hypotheses with accuracy-conditions. Rigorous explanations describe how representational mental states interact with one another and with sensory input to yield bodily motion.

Eliminativists often claim that folk psychology is ill-suited for incorporation into mature science. It may be useful for daily life, but it does not meet proper scientific standards of clarity, precision, rigor, or depth. For example, Churchland (1981, p. 75) critiques folk psychology as a ‘stagnant or degenerating research program’ unlikely to inspire fruitful theories of the mind.

I disagree. To delineate an optimal control model of an actual human subject, we must identify the relevant cost function. As I argued in Section 4.1, we identify the cost function in a typical motor task by deploying general folk psychological capacities to discern personal-level propositional attitudes (especially intentions). Only by deploying those capacities can we isolate the task goal enshrined by the cost function. Thus, even though an optimal control model describes subpersonal rather than personal-level mental activity, we apply the model to an actual subject only by supplementing it with broadly folk psychological analysis of personal-level mental activity. We cannot preserve the science in anything resembling in its current form if we utterly reject all common sense discourse about personal-level propositional attitudes. The science may not vindicate *every* aspect of folk psychology. But it requires us to posit personal-level states that share key features with intentions, as characterized by folk psychology.

There is a second important respect in which optimal control builds upon folk psychology. As Davidson (1980) emphasizes, Bayesian decision theory refines ordinary belief-desire explanation. Thus, optimal control sensorimotor modeling refines common sense psychological discourse. It transmutes folk psychology into *ceteris paribus* explanatory generalizations as mathematically precise and empirically fruitful as one could desire. Granted, an optimal control model differs in important ways from folk psychological description. Most notably, an optimal control model describes subpersonal activity, while folk psychology describes personal-level mental activity. Nevertheless, optimal control lends scientific rigor to a broadly folk psychological explanatory template.

5.2 Eliminating Intentionality?

Eliminativists usually insist we can eschew intentional locutions while retaining any explanatory benefits offered by intentional psychology. They maintain that we can offer equally good explanations cast in non-intentional terms. Generally speaking, though, one cannot radically alter how a science taxonomizes its subject matter while preserving the science's explanatory successes. Why expect that we can transfigure the taxonomic scheme employed by current sensorimotor psychology while retaining its explanatory benefits, any more than we can eschew talk about *gravity* or *genes* while retaining the benefits offered by current physics and current biology? Only a detailed demonstration should convince us that we can purge sensorimotor psychology of representationality while preserving its explanatory achievements. Eliminativists have not even begun to provide such a demonstration.

Field concedes that we may need to cite representational properties when 'explaining behavior described in an intentional way ("murdered his wife")' (2001, p. 77), or when

explaining mental states under intentional descriptions. But he questions whether we must cite such properties when explaining behavior described non-intentionally, as when we try to answer the question ‘Why did her arms move in that way?’ (p. 77). He suggests that we can forego representational locutions when explaining bodily motions under non-intentional descriptions.

As we have seen, there is abundant evidence that sophisticated mental processes mediate between sensory feedback and bodily motions. If we want to explain why someone moved her arm a certain way, we must model these mental processes. The most successful models postulate probabilistic updating over hypotheses individuated through their representational properties. Sensorimotor psychologists postulate mental states with representational properties because doing so yields excellent explanations for bodily motion, just as physicists postulate gravity because doing so yields excellent explanations for celestial and terrestrial motion. *Pace* Field, our best science of motor control assigns an essential role to representational properties when explaining bodily motions under non-intentional descriptions.

Over the past century, numerous scientists have tried to explain bodily motion in entirely non-representational terms. This anti-representationalist program persists among many contemporary proponents of *embodied cognition* (Chemero, 2009) and *dynamical systems theory* (Kelso, Dumas, and Tognoli, 2013). These scientists commonly motivate their approach through physicalist or naturalist rhetoric much like that advanced by eliminativist philosophers. Anti-representationalist research has proved unable to explain the complex ways that sensory input influences bodily motion. Indeed, anti-representationalists do not usually even try to explain the vast range of behavioral data that the models canvassed in Section 2 easily accommodate.

5.3 The Computational Theory of Mind

Many philosophers hold that cognitive science should assign a central role to *non-representational computational description*. The most familiar version of this view emphasizes *formal syntactic manipulation*, epitomized by a Turing machine manipulating stroke marks on a machine tape. Stroke marks are *formal syntactic items* that do not have inherent representational import. Fodor (1987) applies this picture to the mind. He models mental activity as Turing-style computation over formal syntactic mental items. Other proponents of the formal syntactic picture include Chalmers (2011), Cummins (1989), Field (2001), Pylyshyn (1984), and Stich (1983). Proponents vary in their attitude towards intentionality. At one extreme, Fodor (1987) espouses intentional realism. He wants to reserve a central explanatory role for representational content *in addition to* formal mental syntax. At the opposite extreme, Field and Stich deploy the formal syntactic picture to support eliminativism. They recommend that we eschew intentional psychology in favor of formal syntactic computational modeling.

Formal syntactic description is supposed to be *multiply realizable* in Putnam's (1967) sense. Physical systems with radically different physical properties can instantiate the same formal syntactic properties. For example, a carbon-based creature and a silicon-based creature might satisfy the same formal syntactic description. Formal syntactic description prescind from both neural *and* representational properties of mental activity.

Our current best sensorimotor psychology does not feature formal syntactic description along these lines. The science describes how efference copy and sensory feedback *as characterized neurophysiologically* combine with a state estimate *as characterized representationally* to yield a new state estimate *as characterized representationally*. It describes how a subpersonal cost assignment *as characterized representationally* combines with the state estimate *as characterized representationally* to yield a motor command *as characterized*

neurophysiologically. Formal syntactic individuation of mental states plays no role. There may be many excellent reasons to postulate formal syntactic mental items, but our current best sensorimotor psychology postulates no such items. The science does not employ formal syntactic descriptions that prescind from both representational and neural details. It individuates mental states in representational terms *as opposed to* formal syntactic terms.¹²

Field (2001, pp. 72-82, pp. 153-156) proposes a version of Bayesian modeling on which subjective probabilities attach to formal syntactic items individuated without regard to meaning or content. He claims that this framework can preserve any explanatory benefits offered by intentional explanation. However, he does not indicate anything resembling a detailed Bayesian model of specific psychological phenomena. He simply asserts that Bayesian models defined over formal syntactic items offer the same explanatory advantages as Bayesian models defined in representational terms.

The details of Field's account call into question whether his favored formal syntactic models can match the explanatory benefits of representationally-specified models. According to Field, there is no viable interpersonal notion of type-identity for mental representation tokens (2001, p. 75, fn. 3). As he puts it, 'the notion of type-identity between tokens in one organism and tokens in the other is not needed for psychological theory, and can be regarded as a meaningless notion' (2001, p. 57, fn. 32). Thus, Field's formal syntactic taxonomic scheme cannot type-identify the mental states of distinct subjects. Field's proposal departs markedly from actual sensorimotor psychology, which routinely type-identifies mental states of distinct subjects. To illustrate, consider two subjects Y_1 and Y_2 executing the same reaching task. In principle, Y_1 and Y_2 can instantiate precisely the same Bayesian sensorimotor computations. They

¹² In (Rescorla, 2012; forthcoming), I argue that we can model the mind as a computational system while eschewing any appeal to formal mental syntax. On this view, computational models of the mind can individuate mental states through their representational properties rather than through any alleged formal syntactic properties.

can receive the same sensory input and efference copy, yielding the same state estimate, interacting with the same cost assignment to yield the same motor command. Bayesian sensorimotor psychology can offer a unified psychological explanation for the bodily motions of Y_1 and Y_2 . In contrast, Field envisions a disunified treatment. He recognizes no useful sense in which Y_1 and Y_2 are psychological duplicates.

In practice, there are always differences between two subjects. People have different bodies, possibly yielding slightly different forward models. They receive slightly different sensory feedback, possibly yielding different state estimates. There may be other differences. Sensorimotor modeling often ignores these differences, just as physical modeling often ignores the fact any two objects have slightly different masses, slightly different forces acting upon them, and so on. Even when a sensorimotor model registers minor psychological differences between subjects, it simultaneously highlights various uniformities and similarities: that the task goal is the same, or that one state estimate is similar to another, and so on. Sensorimotor psychology can offer a *relatively* unified treatment even when Y_1 and Y_2 differ in various ways. Yet Field must treat all such cases in equally disunified fashion.

Sensorimotor psychology offers unified or relatively unified explanations for the bodily motions of diverse subjects. Field vows instead to explain those bodily motions in highly disunified fashion. So Field's favored approach seems unlikely to preserve the explanatory benefits offered by current sensorimotor psychology.

6. Instrumentalism

Instrumentalists regarding intentionality acknowledge that intentional psychology is predictively successful, but they question whether mental states *really* represent. Dennett (1987) develops a broadly instrumentalist approach. According to Dennett, theorists who attribute representational properties to mental states are not *literally* asserting that mental states have representational properties. They are merely adopting the ‘intentional stance.’ Hornsby (2000) and McDowell (1994) agree with Dennett regarding intentional attribution to the *subpersonal* level, although they favor a literal construal of intentional attributions to *personal*-level states. As Hornsby puts it, ‘[a]dopting the [intentional] stance towards... subsystems... is a matter of treating them *as if* they had some of the intentional properties that persons have’ (2000, p. 20). Similarly, McDowell writes that ‘the content-involving truth at the “sub-personal” level is irreducibly metaphorical’ (1994, p. 197). Dennett, Hornsby, and McDowell share an instrumentalist attitude towards subpersonal intentional states.

Sensorimotor psychology casts doubt upon instrumentalism regarding intentionality, including instrumentalism directed solely towards subpersonal representational states. The science describes numerous subpersonal states in representational terms. Contrary to what Dennett, Hornsby, and McDowell suggest, the science does not advance these intentional attributions in a metaphorical or ‘as if’ fashion. It advances them as true (or approximately true) descriptions. Talk about subpersonal mental representation within sensorimotor psychology is no more metaphorical than talk about gravity within physics or talk about genes within biology.

Colombo and Seriès (2012) advance an opposing instrumentalist analysis. They hold that ‘Bayesian models should be understood as no more than toolboxes for making predictions and systematizing data’ (p. 714). The apparatus of forward models, priors, state estimates, and cost

assignments is nothing but a useful device for predicting bodily motions. We should not construe Bayesian sensorimotor models as true (or approximately true) descriptions of mental activity.

I think that Colombo and Seriès seriously underrate the explanatory aspirations and achievements of sensorimotor psychology. The science does not merely aim to provide systematic tabulations of behavioral data. It aims to explain human motor performance. Why does the subject follow this trajectory rather than that one? Why does she successfully execute the task goal in these conditions but not those? Why does motor behavior vary less along task-relevant dimensions than task-irrelevant dimensions? Bayesian sensorimotor psychology illuminates these and many other questions. The science purports to offer, and in many cases *does* offer, satisfying explanations for the target phenomena. Talk about forward models, priors, state estimates, and cost assignments is not just a useful predictive device. Sensorimotor psychologists postulate these theoretical entities for their explanatory power, just as physicists postulate gravitational forces and biologists postulate genes for their explanatory power. In all cases, the resulting explanatory power provides good reason to believe that the postulated entities exist. Instrumentalism towards sensorimotor psychology is no more justified than instrumentalism regarding physics, biology, or any other successful science.

Readers may complain that I have applied scientific realism to sensorimotor psychology in an overly crude fashion. Even committed scientific realists must acknowledge that some explanatorily successful scientific models posit non-existent entities, such as frictionless planes or massless strings. Explanatory success does not mandate a realist attitude towards *all* entities postulated by a scientific model. But then why should we adopt a realist attitude towards the subpersonal estimates postulated by sensorimotor psychologists?

I agree that we should not believe some entity exists simply because it figures in an explanatorily successful theory. For example, we should not believe that frictionless planes exist simply because explanatorily successful models invoke them. However, we should believe that an entity exists when it figures in our best explanations *in a seemingly indispensable way*. We are confident that, in principle, we could expunge frictionless planes from physics while preserving any important explanatory achievements. There is no evident way to preserve those achievements without citing gravity. That is why physics provides reason to believe in gravity but not frictionless planes.

A similar diagnosis applies to sensorimotor psychology. Existing models are highly idealized, neglecting various information-processing constraints that more realistic models would take into account. Moreover, as Colombo and Seriès (2012) note, researchers often choose the details of a model (e.g. the precise form of the priors and the cost function) based mainly on mathematical convenience rather than psychological realism. However, we feel confident that more realistic models would preserve the explanatory achievements of current models. In contrast, there is no evident way to preserve those explanatory achievements while eschewing subpersonal state estimation. For example, the explanation summarized by Figure 1 hinges upon the claim that visual feedback influences a subpersonal estimate of finger state. Similarly, the optimal control explanation of task-constrained variability requires (some approximation to) subpersonal expected cost minimization grounded in (some approximation to) subpersonal Bayesian estimation. These explanations assign an essential role to subpersonal hypotheses with accuracy-conditions. Apparently, then, subpersonal representational mental states figure indispensably within our best explanations for important motor phenomena. This provides strong reason to believe that subpersonal representational mental states exist.

To appreciate the advantages of realism over instrumentalism, let us compare possible intentions I_1, \dots, I_r . For example, these might be intentions to reach one's right hand to r different locations. Each intention I_n induces a distinct mapping M_n from sensory input to motor commands. The mapping reflects how sensorimotor processing transforms sensory input into motor commands, assuming that the subject currently seeks to execute intention I_n . The mapping is stochastic, rather than deterministic, due to noise and other interfering factors.

Realists can offer a unified explanation for the diverse mappings M_1, \dots, M_r . Each intention I_n causes a subpersonal cost assignment, which combines with the current state estimate to yield the next motor command. Varying the intention while holding fixed the Bayesian estimator (including the priors and the forward model) alters the mapping M_n from sensory input to motor commands, in a way made precise by the Bayesian model. So realists can offer a systematic account that reveals the diverse mappings M_1, \dots, M_r as manifestations of a single underlying Bayesian estimator. Instrumentalists cannot offer a systematic explanation along these lines. They cannot say that the diverse mappings M_1, \dots, M_r reflect a uniform Bayesian estimator, because they dismiss as metaphorical all talk about subpersonal Bayesian estimation. Can instrumentalists offer an alternative explanation, perhaps grounded in neurophysiological or other non-intentional properties? Given present scientific knowledge, the suggestion is idle speculation. We have no idea how to explain the mappings within a unified framework, save as manifestations of subpersonal estimation and control.

7. Normative Constraints on Intentional Attribution

Many philosophers view intentional psychology as an inherently normative enterprise. Quine (1960) anticipates this view. Davidson (1980) develops it at length. According to Davidson, a *constitutive ideal of rationality* informs how we interpret the mental states and speech acts of others. We can ascribe representational content to a creature's mental states only when we depict the creature as largely conforming to norms of logic, probability theory, and decision theory. Davidson recognizes that people deviate from rational norms due to self-deception, akrasia, and various other factors. But he insists that these deviations occur against a baseline of conformity. Good interpretation should not impute too much irrationality to the interpreted subject. In that respect, normative evaluation constrains psychological description. Dennett (1971) espouses a similar position, as do many other philosophers.

The research canvassed in Section 2 embodies a broadly Davidsonian methodology. Researchers adopt a two-step approach: first, construct a normative model describing how an optimal Bayesian decision-maker would proceed; second, fit the normative model as well as possible to the data. We begin with Bayesian norms governing estimation and control, as encapsulated by Figure 2. To apply those norms to a specific motor task, we specify details left open by the norms: the cost function, the forward model, the priors, and so on. By specifying these details, we construct a detailed normative model of the task. Our model yields *ceteris paribus* generalizations relating sensory input, mental activity, and behavior. We evaluate through experimentation how well the generalizations describe actual humans. Hence, the basic explanatory strategy is to use Bayesian normative models as descriptive psychological tools. This explanatory strategy presupposes that the motor system largely conforms (at least approximately) to Bayesian norms. Deviations can arise due to malfunction, external

interference, and various other factors. If our current best sensorimotor psychology is on the right track, then those deviations occur against a baseline of conformity.

Davidson emphasizes beliefs, desires, and other personal-level propositional attitudes. In contrast, sensorimotor psychology emphasizes subpersonal activity. The science extends standard Bayesian norms (e.g. Bayes's Rule and expected cost minimization) from the personal level to the subpersonal level. It treats motor processes as unconscious inferences and decisions that largely conform (approximately) to appropriate Bayesian norms.¹³

Philosophers sometimes suggest that normative constraints undermine the scientific pretensions of intentional psychology. Science aims to describe how the world is, not how it should be. But then what business does science have with rational norms? Setting aside this worry, normative constraints on intentional description may still seem to conflict with proper scientific methodology. How can a true science base itself on dogmatic allegiance to the rationality of mental activity? Davidson, Dennett, Quine, and many others argue in this vein.¹⁴

Bayesian sensorimotor psychology casts doubt upon all such arguments. The field seems as scientific as one could desire. It is mathematically rigorous, empirically well-confirmed, and explanatorily powerful. Yet it employs the broadly Davidsonian methodology highlighted above: fit behavior as well as possible to normative models. This methodology, which assigns rational norms an essential constraining role, allows us to connect behavioral data with the explanatorily fruitful apparatus of forward models, Bayesian estimation, cost minimization, and so on. Rather than *undermining* scientific rigor, normative constraints *promote* scientific explanation of bodily

¹³ See (Rescorla, 2013) for further discussion of how Bayesian cognitive science relates to Davidson's philosophy, especially his emphasis upon rationality as a constitutive ideal.

¹⁴ Davidson, citing normative constraints on intentional ascription, argues that intentional generalizations are always *ceteris paribus*. On that basis, he draws an invidious distinction between psychology and physics. However, as Fodor (1987) notes, all special sciences require *ceteris paribus* generalizations. Reliance on such generalizations does not distinguish intentional psychology from biology, geology, or other special sciences.

motion. They guide us towards mathematical structures of Bayesian decision theory instantiated by the motor system. Normative constraints need not flout proper scientific methodology.

According to Dennett (1971, p. 99), '[i]ntentional psychology is vacuous as psychology because it presupposes and does not explain rationality or intelligence.' As Fodor (1981, pp. 100-123) observes, it is unclear why Dennett thinks that rationality presuppositions render intentionality psychology vacuous. Dennett's worry seems to be that rationality presuppositions are just heuristic idealizations, so that a theory based upon them cannot be entirely factual. In any case, I disagree with Dennett's assessment. Intentional psychology may look vacuous if one confines attention to folk psychological platitudes, but there is nothing vacuous about Bayesian sensorimotor psychology. Researchers develop the Bayesian framework by providing detailed models of specific motor tasks. Non-vacuity stems from mathematically precise generalizations, patterned after Bayesian norms. The generalizations yield substantive predictions, which we compare with behavioral data. When the data confirm the generalizations, we attain satisfying explanations of bodily motion. This methodology presupposes that motor processes largely conform (approximately) to Bayesian norms.

The presupposition is not an unwarranted dogma or a heuristic idealization. It is a *working hypothesis*. Researchers elaborate the hypothesis into well-confirmed models of specific motor tasks. By analogy: Newtonian physicists take as a working hypothesis that physical systems conform to Newton's laws, and they elaborate that hypothesis into well-confirmed models of specific systems (pendulums, orbiting planets, etc.). In each case, we adopt the working hypothesis because it has produced some striking explanatory successes. In each case, the working hypothesis enables fruitful theorizing. Thus, normative presuppositions of Bayesian sensorimotor modeling generate no evident pressure towards eliminativism or instrumentalism.

Scientifically-minded philosophers often regard intentionality and normativity quite warily. Although relatively few authors overtly endorse eliminativism or instrumentalism, many authors downplay both intentional content and rational norms when limning the foundations of scientific psychology (e.g. Chalmers, 2011; Cummins, 1989; Egan, 2010). I have argued that intentionality and normativity figure irreplaceably within explanatorily powerful theories of motor control. The theories postulate representationally contentful mental states that causally interact in approximate accord with Bayesian norms. If these theories lie anywhere close to the truth, then intentional description guided by normative evaluation provides unsurpassed insight into the etiology of goal-directed bodily motion. We should embrace intentionality and normativity as indispensable contributors to any complete scientific psychology.

Word count: 14,077

References

- Anscombe, E. 1957: *Intention*. Oxford: Blackwell.
- Bays, P., and Wolpert, D. 2007: Computational principles of sensorimotor control that minimize uncertainty and variability. *Journal of Physiology*, 578, 387-396.
- Bermúdez, J. L. 1995: Nonconceptual content: from perceptual experience to subpersonal computational states. *Mind and Language*, 10, 333-369.
- Bernstein, N. 1967: *The Coordination and Regulation of Movements*. Oxford: Pergamon.
- Bowers, J., and Davis, C. 2012: Bayesian just-so stories in psychology and neuroscience. *Psychological Bulletin*, 138, 389-414.

- Brayanov, J., and Smith, M. 2010: Bayesian and 'anti-Bayesian' biases in sensory integration for action and perception in the size-weight illusion. *Journal of Neurophysiology*, 103, 1518-1531.
- Briscoe, R. 2008: Another look at the two visual systems hypothesis: the argument from illusion studies. *Journal of Consciousness Studies*, 8, 3-62.
- Bruno, N., and Franz, V. 2009: When is grasping affected by the Müller-Lyer illusion?: a quantitative review. *Neuropsychologia*, 47, 1421-1433.
- Burge, T. 2010: *Origins of Objectivity*. Oxford: Oxford University Press.
- Butterfill, S., and Sinigaglia, C. 2014: Intention and motor representation in purposive action. *Philosophy and Phenomenological Research*, 88, 119-145.
- Chalmers, D. 2011: A computational foundation for the study of cognition. *The Journal of Cognitive Science*, 12, 323-357.
- Chemero, A. 2009: *Radical Embodied Cognitive Science*. Cambridge: MIT Press.
- Churchland, P. 1981: Eliminative materialism and the propositional attitudes. *Journal of Philosophy*, 78, 67-90.
- Clark, A. 2013a: Expecting the world: perception, prediction, and the origins of human knowledge. *Journal of Philosophy*, 110, 469-496.
- Clark, A. 2013b: Whatever next? predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, 36, 181-204.
- Colombo, M., and Seriès, P. 2012: Bayes in the brain --- on Bayesian modeling in neuroscience. *British Journal for the Philosophy of Science*, 63, 697-723.
- Cummins, R. 1989: *Meaning and Mental Representation*, Cambridge: MIT Press.
- Davidson, D. 1980: *Essays on Actions and Events*. Oxford: Clarendon Press.

- Dennett, D. 1971: Intentional systems. *Journal of Philosophy*, 68, 87-106.
- Dennett, D. 1987: *The Intentional Stance*. Cambridge: MIT Press.
- Egan, F. 2010: Computational models: a modest role for content. *Studies in History and Philosophy of Science*, 41, 253-259.
- Erez, T., and Todorov, E. 2012: Trajectory optimization for domains with contacts using inverse dynamics. *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 4914-4919.
- Feldman, A. 2011: Space and time in the context of the equilibrium-point theory. *Wiley Interdisciplinary Reviews: Cognitive Science*, 2, 287-304.
- Field, H. 2001: *Truth and the Absence of Fact*. Oxford: Oxford University Press.
- Fodor, J. 1981: *Representations*. Cambridge: MIT Press.
- Fodor, J. 1987: *Psychosemantics*. Cambridge: MIT Press.
- Franklin, D., and Wolpert, D. 2011: Computational mechanisms of sensorimotor control. *Neuron Review*, 72, 425-442.
- Franz, V., Bühlhoff, H., and Fehle, M. 2003: Grasp effects of the Ebbinghaus illusion: obstacle avoidance is not the explanation. *Experimental Brain Research*, 149, 470-477.
- Friston, K. 2011. What is optimal about motor control?. *Neuron*, 72, 488-498.
- Geisler, W., and Kersten, D. 2002: Illusions, perception, and Bayes. *Nature Neuroscience*, 5, 508-510.
- Grush, R. 2004: The emulation theory of representation: motor control, imagery, and perception. *Behavioral and Brain Sciences*, 27, 377-442.
- Hartung, B., Schater, P., Bühlhoff, H., Kersten, D., and Franz, V. 2005: Is Prior Knowledge of Object Geometry Used in Visually Guided Reaching?. *Journal of Vision*, 5, 204-214.

- Helmholtz, H. von. 1867: *Handbuch der Physiologischen Optik*. Leipzig: Voss.
- Hornsby, J. 2000: Personal and sub-personal: a defense of Dennett's early distinction. *Philosophical Explorations*, 1, 6-24.
- Hohwy, J. 2014: *The Predictive Mind*. Oxford: Oxford University Press.
- Jeanerrod, M. 2006: *Motor Cognition*. Oxford: Oxford University Press.
- Kelso, J., Dumas, G., and Tognoli, E. 2013: Outline of a general theory of behavior and brain coordination. *Neural Networks*, 37, 120-131.
- Knill, D., and Richards, W. (eds) 1996: *Perception as Bayesian Inference*. Cambridge: Cambridge University Press.
- Kroliczak, G., Heard, P., Goodale, M. A., and Gregory R. L. 2006: Dissociation of perception and action unmasked by the hollow face illusion. *Brain Research*, 1080, 9-16.
- Latash, M. 2010: Motor control: in search of physics of the living systems. *Journal of Human Kinetics*, 24, 7-18.
- Liu, D., and Todorov, E. 2007. Evidence for the flexible sensorimotor strategies predicted by optimal feedback control. *The Journal of Neuroscience*, 27, 9354-9368.
- Loeb, G. 2012: Optimal isn't good enough. *Biological Cybernetics*, 106, 757-765.
- Marr, D. 1982: *Vision*. New York: Freeman.
- McDowell, J. 1994: The content of perceptual experience. *Philosophical Quarterly*, 44, 190-205.
- Mele, A. 1992: *Springs of Action*. Oxford: Oxford University Press.
- Milner, A., and Goodale, M. 1995: *The Visual Brain in Action*. Oxford: Oxford University Press.
- Nanay, B. 2013: *Between Perception and Action*. Oxford: Oxford University Press.
- Quine, W. V. 1960: *Word and Object*. Cambridge: MIT Press.

- Pacherie, E. 2000: The content of intentions. *Mind and Language*, 15, 400-432.
- Pacherie, E. 2006: Towards a dynamic theory of intentions. In S. Pockett, W. P. Banks, and S. Gallagher (eds). *Does Consciousness Cause Behavior? An Investigation of the Nature of Volition*. Cambridge: MIT Press.
- Pacherie, E. 2008. The phenomenology of action: a conceptual framework. *Cognition*, 208, 179-217.
- Peacocke, C. 1992: *A Study of Concepts*. Cambridge: MIT Press.
- Pouget, A., Beck, J., Ma., W. J., and Latham, P. 2013: Probabilistic brains: knowns and unknowns. *Nature Neuroscience*, 16, 1170–1178.
- Putnam, H. 1967: Psychophysical predicates. In W. Capitan and D. Merrill (eds), *Art, Mind, and Religion*. Pittsburgh: University of Pittsburgh Press.
- Pylyshyn, Z. 1984: *Computation and Cognition*. Cambridge: MIT Press.
- Rescorla, M. 2012: How to integrate representation into computational modeling, and why we should. *Journal of Cognitive Science*, 13, 1-38.
- Rescorla, M. 2013: Rationality as a constitutive ideal. In E. Lepore and K. Ludwig (eds), *A Companion to Davidson*. Malden: Wiley-Blackwell.
- Rescorla, M. 2015: Bayesian perceptual psychology. In M. Matthen (ed.), *The Oxford Handbook of the Philosophy of Perception*. Oxford: Oxford University Press.
- Rescorla, M. Forthcoming: From Ockham to Turing --- and back again. In A. Bokulich and J. Floyd (eds), *Turing 100: Philosophical Explorations of the Legacy of Alan Turing*. Berlin: Springer.
- Rosenbaum, D. 2002: Motor control. In H. Pashler and S. Yantis (eds), *Stevens' Handbook of Experimental Psychology*, vol. 1, 3rd ed. New York: Wiley.

- Saunders, J., and Knill, D. 2004: Visual feedback control of hand movements. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 24, 3223-3234.
- Schenk, T., Franz, V., and Bruno, N. 2011: Vision-for-perception and vision-for-action: which model is compatible with the available psychophysical and neuropsychological data?. *Vision Research*, 51, 812-818.
- Scott, S. 2004: Optimal feedback control and the neural basis of volitional motor control. *Nature Reviews Neuroscience*, 5, 532-546.
- Searle, J. 1983: *Intentionality*. Cambridge: Cambridge University Press.
- Shadmehr, R. 2010: From equilibrium point to optimal control. *Motor Control*, 14, e25-e30.
- Shadmehr, R., and Mussa-Ivaldi, S. 2012: *Biological Learning and Control*. Cambridge: MIT Press.
- Shadmehr, R., Smith, M., and Krakauer, J. 2010: Error correction, sensory prediction, and adaptation in motor control. *Annual Review of Neuroscience*, 3, 89-108.
- Shapiro, M., and Körding, K. 2010: Looking for synergies between the equilibrium point hypothesis and internal models. *Motor Control*, 14, e31-e34.
- Srinivasan, M., and Ruina, A. 2006: Computer optimization of a minimal biped model discovers walking and running. *Nature*, 439, 72-75.
- Stich, S. 1983: *From Folk Psychology to Cognitive Science*. Cambridge: MIT Press.
- Stich, S. 1991: Narrow content meets fat syntax. In B. Loewer and G. Rey (eds), *Meaning in Mind: Fodor and His Critics*. Cambridge: Blackwell.
- Todorov, E. 2004: Optimality principles in sensorimotor control. *Nature Neuroscience*, 7, 907-915.
- Todorov, E., and Jordan, M. 2002: Optimal feedback control as a theory of motor

coordination. *Nature Neuroscience*, 5, 1226-1235.

Trommershäuser, J., Maloney, L., and Landy, M. 2003: Statistical decision theory and selection of rapid, goal-directed movements. *Journal of the Optical Society of America*, 20, 1419-1433.

Wolpert, D. 2007: Probabilistic models in human sensorimotor control. *Human Movement Science*, 26, 511-524.

Wolpert, D., Ghahramani, Z., and Jordan, M. 1995: An internal model for sensorimotor integration. *Science*, 269, 1880-1882.

Wolpert, D., and Landy, M. 2012: Motor control is decision-making. *Current Opinion in Neurobiology*, 22, 1-8.

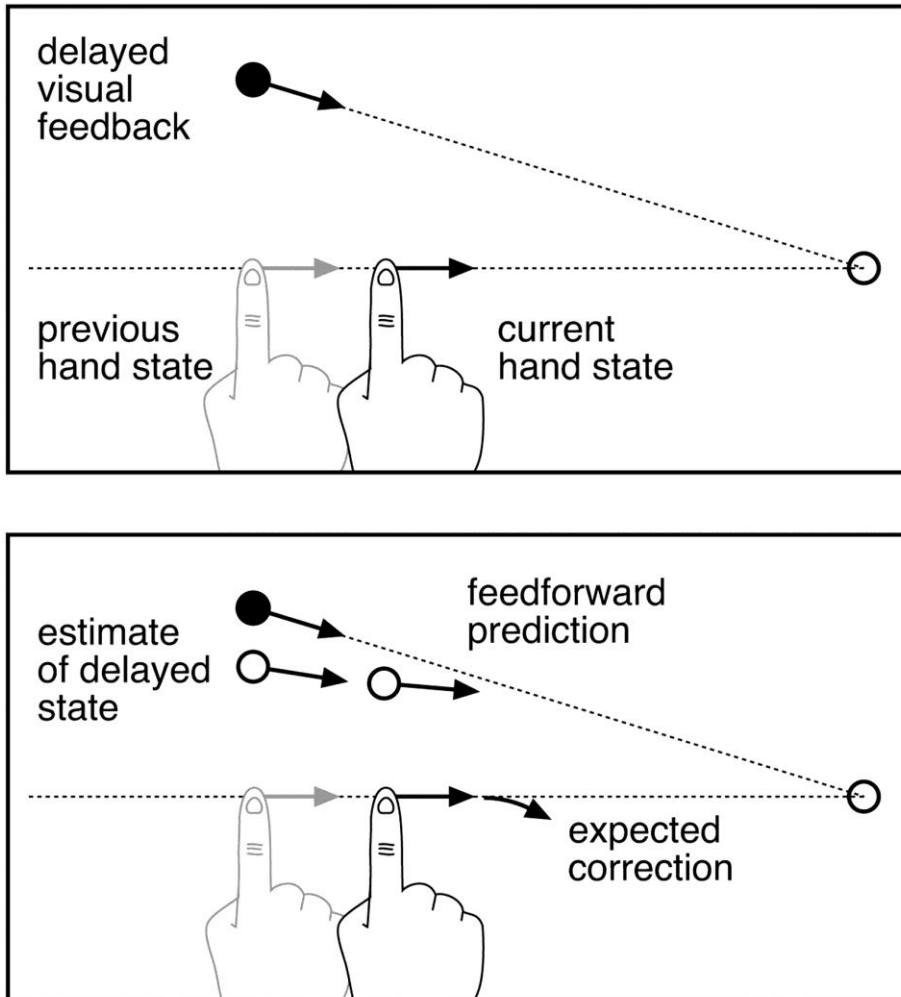


Figure 1. The open circle in the top panel is the target location. The black circle is the virtual fingertip, whose movement is controlled by the subject's hand motion. The subject's real hand is not visible. Due to sensory delay, visual feedback always reflects an earlier state of the finger. In a *rotation perturbation* (top panel), visual space is rotated around the target location, so that the virtual fingertip has a different location than the subject's real hand but still moves towards the target location. Subjects do not consciously notice the perturbation, which occurs while an occluder obscures the virtual fingertip. Contrary to the homing model's predictions, subjects systematically alter course downwards, thereby missing the target location. The lower panel illustrates how Bayesian modeling explains this phenomenon. On a Bayesian approach, the motor system estimates *finger state* (position and velocity). The estimator receives delayed visual feedback, which it uses to make a corrected estimate of delayed state. The estimator then uses the forward model and efference copy of motor commands issued during the delay period to make a feedforward prediction of current state. The resulting mistaken state estimate causes the motor system to "correct" finger trajectory downwards, away from the target location. (From Saunders and Knill, 2004; Figure 6. Reproduced with permission of the Society for Neuroscience.)

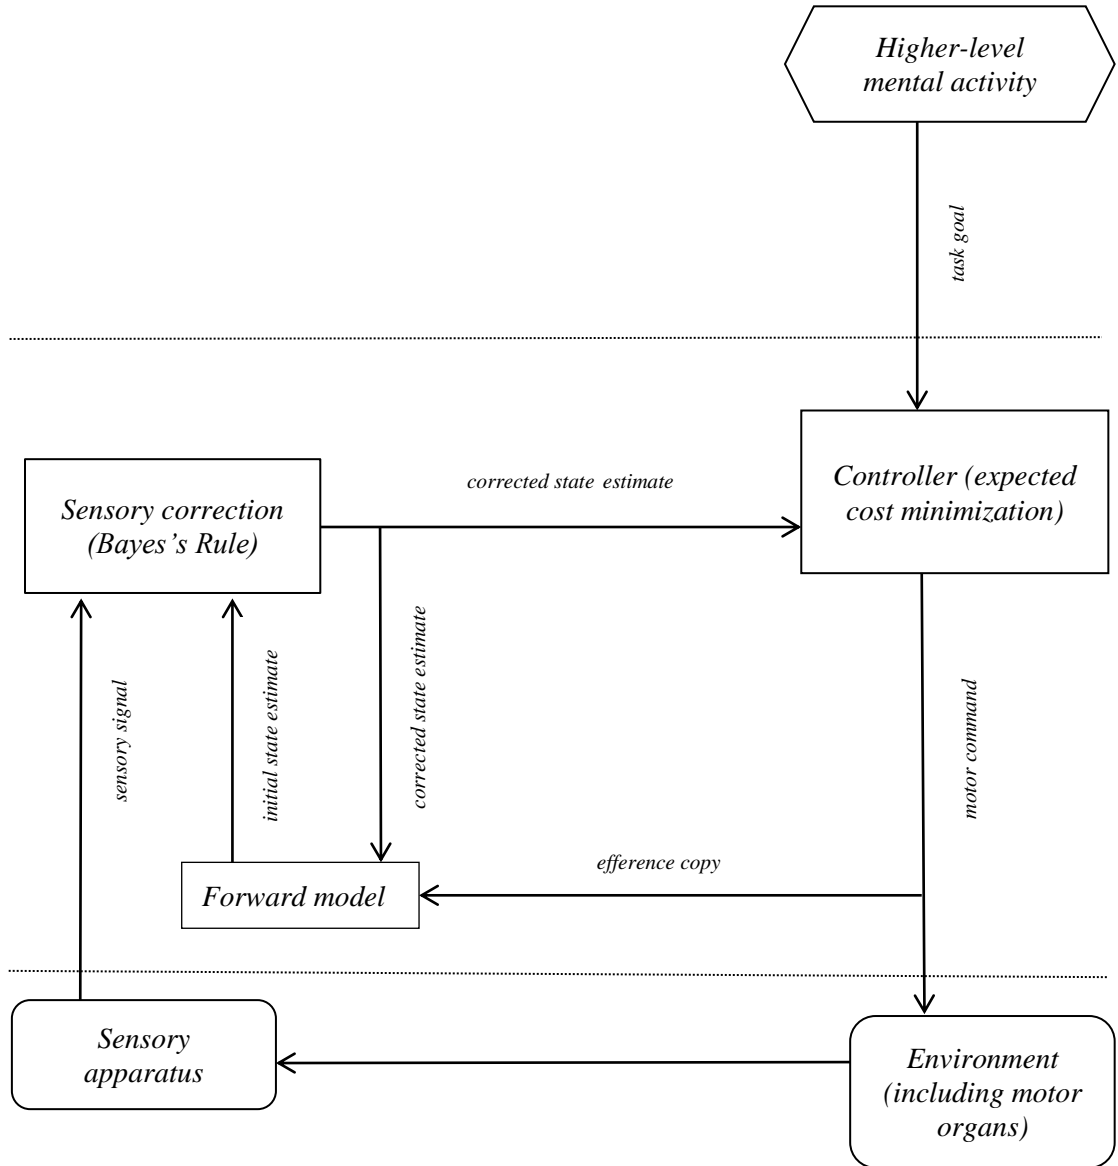


Figure 2. A template for Bayesian models of sensorimotor control. Some models vary the template somewhat. For example, the Saunders and Knill (2004) model handles sensory delay by transmitting the initial state estimate rather than the corrected state estimate to the controller.