Motor Computation

Michael Rescorla

Enjoying a bravura performance by a professional athlete or virtuoso musician, we marvel at the performer's skilled bodily motion. We are less apt to appreciate that relatively humdrum activities—such as walking, talking, riding a bicycle, tying shoelaces, typing, using silverware, or pouring milk into a cup without spilling—already require impressive control over one's motor organs. The ease with which we execute these activities belies their difficulty, as evidenced by our current inability to build robots that match human performance. During the course of each day, a typical adult achieves myriad goals through an extraordinary range of dexterous bodily motions. How do humans manage to achieve their goals by exerting such refined control over their motor organs?

Sensorimotor psychology, the scientific study of motor control, emerged from Hermann von Helmholtz's (1867) pioneering investigations and assumed its modern form in the work of Nikolai Bernstein (1923; 1930; 1967). Building on insights of Helmholtz and Bernstein, contemporary sensorimotor psychologists have convincingly established that human motor control involves sophisticated unconscious computations that mediate between cognition and bodily motion. The present entry will discuss basic aspects of motor computation, along with some implications for philosophy of mind.

1. From intentions to motor commands

Suppose I form an intention to do something—say, to push an elevator button with my right index finger. My motor system must transform the intention into a sequence of suitable motor commands. How does this transformation work? How does my motor system select commands that promote my goal? Bernstein first highlighted a *redundancy problem* that bedevils the transformation of intentions into motor commands. There are innumerable ways the motor system might achieve some goal. For example, there are many trajectories my finger might take to reach the elevator button and many muscle activations that would achieve a given trajectory. The motor system must choose rapidly from among these infinitely many options.

Bernstein decisively advanced research into the redundancy problem with a pivotal discovery: the movement details through which an agent completes a motor task vary considerably across trials in certain characteristic ways. In particular, performance across trials varies far less along task-relevant dimensions than task-irrelevant dimensions. For example, suppose the task is to aim a laser pistol at a target. Joint configurations in the arm can fluctuate widely even as the agent maintains a fixed aim. Scholz, Schöner, and Latash (2000) found that fluctuations in joint angle that affect how well the pistol aims at the target are much smaller than fluctuations that have no such effect. In a similar vein, Todorov and Jordan (2002) studied a task where the subject moved her hand through a sequence of five widely spaced targets arranged in a plane. Hand trajectories varied much more between the targets than near the targets, reflecting the fact that hand position between targets was task-irrelevant. The scientific literature conclusively shows that a large disparity between task-relevant variation and task-irrelevant variation occurs within a wide range of motor activities (Todorov and Jordan, 2002), such as postural control, walking, talking, skiing, writing, reaching, and bimanual coordination. Any adequate theory of motor control must explain the disparity.

2. Optimal feedback control

Sensorimotor psychologists have explored various theoretical frameworks for explaining motor control (Rosenbaum, 2002). The most empirically successful is *optimal feedback control* (OFC). This framework uses *optimal control theory*, a mathematical approach to decision-making that has been extensively developed in engineering and statistics (Stengel, 1986). The basic idea behind OFC is that the motor system selects motor commands that are in some sense optimal (or near-optimal) with respect to one's current goal. Researchers begin with a normative model delineating how an idealized decision maker would accomplish the goal. They then investigate how well the normative model fits actual human performance.

The central construct of optimal control theory is a *cost function*, which measures the desirability of issuing a motor command under the assumption that certain environmental conditions obtain. The cost function quantifies various performance criteria. Some of the performance criteria are task-independent (e.g. minimizing expenditure of energy). Some are task-dependent and reflect the goal being pursued (e.g. that I move the tip of my index finger to a specific location; or that I walk to the other end of a room; or that I pick up a nearby ball). Optimal motor commands are those that minimize expected costs. Thus, "optimality" is relative to a cost function. An *optimal controller* selects motor commands that are optimal given the current cost function. Depending on the details, more than one motor command may be optimal.

OFC uses optimal control theory to illuminate how the motor system solves Bernstein's redundancy problem (Scott, 2004; Todorov, 2004; Todorov and Jordan, 2002). The crucial insight is that the motor system selects certain motor commands over the infinitely many alternatives because the selected commands are "better" according to well-defined performance criteria. The selected commands minimize (or nearly minimize) expected costs. OFC elaborates this insight into detailed normative models describing how the human motor system *should* choose motor commands (given certain performance criteria). These models are pressed into service as psychological descriptions of human motor computation, yielding testable predictions. In many cases, researchers have validated the predictions. A well-confirmed optimal control model of a motor task provides satisfying explanations for key aspects of human performance.

A crucial plank of OFC is the *minimal intervention principle*, articulated by Todorov (2004). When I pursue a goal, factors such as motor noise or external interference frequently perturb my bodily trajectory. It is not a good idea to correct each perturbation, because every correction expends energy. An optimal controller corrects only those perturbations that are task-relevant. Hence the minimal intervention principle, which enjoins: "make no effort to correct deviations away from average behavior unless those deviations interfere with task performance" (Todorov, 2004, p. 911). The minimal intervention principle directs the controller to tolerate perturbations so long as they do not interfere with the task.

The minimal intervention principle is a norm. It decrees how motor processing *should* operate. OFC postulates that there are numerous cases where the motor system conforms (at least approximately) to this norm. In such cases, the motor system pursues a task goal rather than enforcing a predetermined movement plan that effectuates the goal. As the task unfolds, the motor system chooses motor commands that optimally (or near-optimally) promote the goal. Perturbations from the average trajectory may occur, but the motor system only corrects perturbations that impede task completion. Motor activity optimizes relatively abstract performance criteria, rather than implementing a pre-specified sequence of bodily movements. OFC develops these ideas into mathematically detailed, empirically successful models of specific motor tasks (Haith and Krakauer, 2013; Wolpert and Landy, 2012).

By invoking the minimal intervention principle, OFC achieves several notable explanatory successes:

- OFC explains the striking disparity noted above between task-relevant and task-irrelevant variation across trials. If a controller corrects task-relevant deviations but leaves task-irrelevant deviations uncorrected, then task-irrelevant deviations accumulate so that bodily movements vary far more along task-irrelevant dimensions.
- OFC illuminates how the motor system responds to experimentally-induced perturbations of bodily motion. The minimal intervention principle dictates that the motor system should preferentially correct task-relevant perturbations, and this is indeed what happens in numerous motor tasks (Crevecoeur, Cluff, and Scott, 2014). For example, Nashed, Crevecoeur, and Scott (2012) studied subjects performing two slightly different tasks: reaching rapidly either to a small circular target or to a long horizontal bar. In random trials, an unexpected mechanical perturbation disrupted hand trajectory horizontally. In the circular target task, the motor system responded by correcting course back towards to the target. Corrective motions began as early as 70 ms after the perturbation—strong evidence that they were automatic and not generated voluntarily. In the horizontal bar task, the motor system did not correct the perturbation, because it could still accomplish the goal (reaching the horizontal bar) without any course correction. Thus, the task goal decisively influenced how the motor system responded to mechanical perturbations. Perturbations were corrected only when they affected the goal, as dictated by the minimal intervention principle.

These successes exemplify the fruitful nexus between normative evaluation and psychological explanation made possible by OFC.

OFC differs significantly from virtually all rival theories of motor control. Rival theories usually enshrine the *desired trajectory hypothesis*, which postulates a rigid division between motor planning and motor execution: the motor system first chooses a detailed movement plan, and it then tries to implement the chosen plan. From the perspective of OFC, this rigid division is non-optimal. No matter how sensible some predetermined detailed movement plan may initially seem, enforcing it will waste energy by correcting task-irrelevant deviations. According to the minimal intervention principle, the motor system should not solve Bernstein's redundancy problem in advance of executing the motor task. It should instead solve the redundancy problem *on-line*, selecting motor commands in furtherance of the task goal as the task unfolds. The desired trajectory hypothesis does not explain why motor processing preferentially corrects task-relevant perturbations, nor does it explain why performance varies more along task-irrelevant dimensions than task-relevant dimensions. Thus, the desired trajectory hypothesis looks unpromising. Nevertheless, it has recently attracted high profile advocates in the scientific (Friston, 2011) and philosophical (Clark, 2015; Hohwy, 2014) communities. These authors do not say how they hope to explain the asymmetry between task-relevant and task-irrelevant variation—a fundamental aspect of motor control that looks incompatible with their favored approach and that OFC easily explains.

3. Estimating environmental state

Environmental conditions are highly relevant to task performance. If my goal is to walk across the room and pick up a ball, then relevant factors include the ball's location, its size, its weight, the presence of any obstacles, my own current bodily configuration, and so on. To choose appropriate motor commands, the motor system takes these and other factors into account. However, motor processing cannot directly access environmental conditions. It can access the environment only by way of sensory stimulations, including stimulations of the retina, the inner ear, muscle spindles, and so on. The motor system must use sensory stimulations to *estimate* environmental state.

Here we encounter a striking commonality between motor control and perception. I can consciously perceive certain properties of my environment, such as the shapes, colors, sizes, and locations of nearby objects. But perceptual processing cannot directly access these environmental properties. The perceptual system must estimate environmental conditions based upon sensory stimulations. Helmholtz (1867) postulated that it does so through an *unconscious inference*. In response to proximal sensory input, the perceptual system forms a “best hypothesis” regarding which environmental conditions caused that proximal input. Contemporary sensorimotor psychologists adapt Helmholtz’s approach to explain how motor processing estimates environmental conditions. Sensorimotor psychology postulates that the motor system executes an unconscious inference from sensory stimulations to estimates of environmental state. The motor system chooses “best hypotheses” regarding the environmental causes of sensory stimulations. It consults the chosen hypotheses when selecting motor commands that promote the current task goal.

Recently, perceptual and sensorimotor psychologists have elucidated unconscious inference in terms of *Bayesian decision theory*, which is a mathematical framework for modeling reasoning and decision-making under uncertainty. Bayesian agents handle uncertainty by assigning *subjective probabilities* to hypotheses. The probabilities are “subjective” because they reflect the agent’s own psychological degree of confidence rather than objective probabilities out in the world. A Bayesian agent begins with initial subjective probabilities (called *priors*) and then revises those initial probabilities in light of new evidence. Precise Bayesian norms dictate how the agent should revise her probabilities in light of new evidence.[1] As applied within sensorimotor psychology, the basic idea is that the motor system assigns probabilities to hypotheses regarding environmental conditions that are relevant to the current task. Ongoing sensory stimulation causes constant revision of these probabilities, in rough conformity to Bayesian norms. Hence, the motor system estimates environmental conditions through an unconscious Bayesian inference. When transforming intentions into motor commands, motor processing selects commands that minimize expected cost—i.e. the cost one expects to incur, given current probabilities. Bayesian modeling of motor estimation has proved extremely successful and is an important component of the overall OFC framework (Shadmehr and Mussa-Ivaldi, 2012; Wolpert, 2007).

A pervasive challenge facing sensorimotor estimation is *sensory delay*. Motor control requires rapid on-line selection of motor commands, but sensory signals take a while to reach the brain. To overcome sensory delay, the motor system anticipates how its own commands are likely to impact environment state (Wolpert and Flanagan, 2009). For each motor command sent to the musculature, the brain produces an *efference copy*. Using efference copy, the motor system predicts likely consequences of its own commands (Miall and Wolpert, 1996). Basically, the motor system begins with its current estimate of environmental state and then extrapolates forward using efference copy. The result: a new environmental state estimate that is correctable by future sensory signals but that can guide motor control until such signals arrive.

4. Challenges for optimal feedback control

[1] For general discussion of Bayesian modeling in cognitive science, see (Colombo and Hartmann, 2017).

OFC has been successfully applied to various relatively simple motor tasks, such as reaching, pointing, aiming, and so on. It has also been successfully applied to at least one fairly complex task involve the whole body (Stevenson, et al., 2009): balancing on a snowboard. An important agenda item for sensorimotor psychology is to apply OFC to more complex real-world tasks, such as riding a bicycle or playing a musical instrument.

A significant hurdle here is that Bayesian inference and expected cost minimization are, in general, computationally intractable. Aside from a few simple cases, a physical system with limited computational resources cannot update probabilities in precise conformity to Bayesian norms, nor can it minimize expected cost with complete precision. A version of this problem arises for all fields that employ Bayesian decision theory, including engineering, statistics, and artificial intelligence. The standard solution is to investigate algorithms through which a physical system can efficiently *approximate* Bayesian inference and expected cost minimization. Sensorimotor psychologists have explored computationally tractable approximation schemes tailored to the human motor system, with excellent results (e.g. Li, Todorov, and Pan, 2004; Todorov, 2009). Future research will presumably deploy these and perhaps other as-yet-undiscovered approximation schemes so as to model a range of motor tasks.

Another important agenda item is to illuminate the *neurophysiological processes* through which the motor system implements (or approximately implements) computations postulated by OFC models. How exactly does the brain encode an assignment of subjective probabilities to hypotheses? How does it encode a cost function? Through what neural processing does it update subjective probabilities and select optimal (or near-optimal) motor commands? We do not know the answers to these questions, although recent research offers some intriguing suggestions (e.g. Denève, Duhamel, and Pouget, 2007; DeWolf and Eliasmith, 2011).

With so many questions left unanswered, OFC at present hardly constitutes a complete theory of human motor computation. Even in its current incomplete state, OFC offers powerful explanations for a range of motor phenomena.

5. Motor control as unconscious inference and decision-making

I favor a broadly *scientific realist* perspective: when a scientific theory is explanatorily successful, this gives us a *prima facie* reason to accept that the theory is at least approximately true. I apply the scientific realist perspective to sensorimotor psychology: OFC models are explanatorily successful, far more so than rival theories, so we have strong *prima facie* reason to regard them as at least approximately true. I do not say that we should regard a successful OFC model as true in every detail. For example, current models often use priors that embody a highly simplified dynamics for the human body. They do this for reasons of analytical tractability, not because there is any independent reason to think that the motor system employs these precise priors. A more accurate model would presumably use more psychologically realistic priors. While we should not accept that current OFC models are precisely true in every detail, we should accept in broad strokes the picture of motor computation that they embody. More specifically:

- The motor system assigns subjective probabilities to hypotheses regarding environmental conditions. It updates those probability assignments in response to sensory stimulations and efference copy. Transitions between probability assignments conform at least approximately to Bayesian norms.
- When pursuing a motor task, the motor system encodes a cost function that reflects the goal being pursued. As the task unfolds, the motor system selects motor

> commands that are near-optimal in light of the cost function and current subjective probabilities.

I defend my realist perspective at length in (Rescorla, 2016; forthcoming). For an opposing *instrumentalist* viewpoint on Bayesian modeling, see (Colombo and Seriès, 2012). For helpful comparison of the realist and instrumental viewpoints, see (Sprevak, 2016).

Assuming a realist interpretation of OFC modeling, motor control results from subpersonal mental processes similar in key respects to high-level conscious inference and decision-making. The processes fall under (and approximately conform to) the same Bayesian norms that govern inference and decision-making by agents. However, the processes are executed not by the person herself but rather by her subsystems. They are *subpersonal*. Except in very unusual cases, the *person* does not consciously choose a detailed movement plan. She is not aware of specific motor commands relayed to her musculature. She does not consciously access the environmental state estimates that inform selection of motor commands. She simply sets a goal (e.g. *pick up that nearby ball*) and lets her motor system do the rest.

One important consequence of the realist viewpoint is that volitional bodily motion results from highly sophisticated computations executed by the motor system. The computations are subpersonal and inaccessible to consciousness, but they approximately implement personal-level rational norms that have been extensively studied within engineering, statistics, robotics, and artificial intelligence. Unconscious Bayesian inference and decision-making mediate the transition from intentions to motor commands.

6. Motor learning

Improvements in motor performance are crucial to the refined control that we exert over our bodies. When I master a new motor skill, such as playing the clarinet or hitting a ball with a tennis racket, many changes occur in my motor processing. Even everyday activities such as walking or talking require extensive practice for their mastery. Compensatory adjustments in motor control are also needed when my body changes over the short-term (e.g. through fatigue or injury) or the long-term (e.g. through development or aging) and when the external environment changes (e.g. walking on pavement versus walking through mud). Any change in motor control that improves task performance is called *motor learning*. This rubric covers the acquisition of new motor skills and also improved performance in previously learned activities (Haith and Krakauer, 2013).

Scientific research mainly studies a specific type of motor learning called *adaptation*. During adaptation, the motor system corrects a disruption of some previously mastered activity. In a seminal experiment, Helmholtz instructed subjects to reach to a target, which they did quite easily. He then equipped the subjects with prism lenses that shifted the visual field to the left or right. Subjects initially missed the target, but they quickly adapted and learned to reach the target. Upon removal of the prism lenses, subjects once again failed at the reaching task, although they quickly re-adapted. Helmholtz's prism experiment vividly illustrates the speed and efficiency with which the motor system responds to ever-changing conditions (in this case, a perturbation of visual input). It is well-established that adaptation mainly involves subpersonal processes rather than conscious correction by the subject herself. Although subjects sometimes realize that conditions have changed and try to compensate accordingly, adaptation by the motor system proceeds in relative independence from high-level conscious thought (Mazzoni and Krakauer, 2006; Shadmehr and Mussa-Ivaldi, 2012, pp. 187-192).

Sensorimotor psychologists have extensively investigated the subpersonal mental activity that underlies adaptation. A key finding is that adaptation, like motor control itself, involves sophisticated computations that draw upon available sensory information. When sensory prediction error occurs (e.g. my hand does not reach the visual target as expected), the motor system cannot directly pinpoint what caused the error. Instead, it must estimate what caused the error. Bayesian decision theory provides models that dictate how to solve this estimation problem (Shadmehr and Mussa-Ivaldi, 2012, pp. 192-212). Researchers have applied Bayesian modeling to several adaptation paradigms, sometimes with striking explanatory success.

To illustrate, suppose we perturb hand movements during a reaching task by applying an external force field to the hand. Motor performance quickly adapts, so that the subject reaches the target despite the perturbing force field. Surprisingly, though, force field adaptation does not merely change the mapping from sensory estimates to motor commands. It also changes subpersonal sensory estimates themselves (Haith, Jackson, Miall, and Vijayakumar, 2009). Repeated exposure to an external force field on the right hand causes a shift in visual and proprioceptive estimates of the hand's position. Due to this shift, estimates of hand position become markedly less accurate. Thus, force field adaptation generates a *sensorimotor illusion*. The illusion is puzzling, because visual and proprioceptive input agree even in the presence of the external force field. Why should visual and proprioceptive estimation change when no discrepancy between them occurs?

A plausible answer is that the sensorimotor illusion reflects a "credit assignment" problem. Sensory prediction error might arise from perturbed motor execution (e.g. an external force field), perturbed sensory input (e.g. prism lenses), or some combination thereof. The Bayesian strategy is to divide the credit for prediction error between motor execution and sensory estimation, taking into account the reliability of all relevant information sources (Haith, Jackson, Miall, and Vijayakumar, 2009). A Bayesian estimator will attribute the prediction error partly to sensory miscalibration, even though the error actually arises entirely from an external force field. Thus, the sensorimotor illusion does not manifest some underlying defect in sensorimotor processing. Rather, it reflects the motor system's ongoing effort to estimate environmental conditions based upon ambiguous sensory cues.

This example illustrates two important points: first, motor learning involves sophisticated subpersonal computations that share many notable properties with high-level reasoning; second, Bayesian modeling illuminates these computations by isolating norms to which they (approximately) conform.

7. Mental representation as explanatorily central

Sensorimotor psychology has significant implications for longstanding debates about the nature of the mind.

Philosophers traditionally regard the mind as a *representational organ*. On the traditional picture, mental states represent the world, and representational properties of mental states are vital for understanding how the mind works. Contemporary philosophers often develop this picture by linking representation to *veridicality-conditions*—conditions for veridical representation of the world. To illustrate:

- Beliefs are the sorts of things that can be true or false. My belief *that Emmanuel Macron is French* is true if Emmanuel Macron is French, false if he is not.
- Perceptual states are the sorts of things that can be accurate or inaccurate. My perceptual experience *as of a red sphere* is accurate only if a red sphere is before me.

- Intentions are the sorts of things that can fulfilled or thwarted. My desire *to eat chocolate* is fulfilled only if I eat chocolate.

Beliefs have truth-conditions, perceptual states have accuracy-conditions, and intentions have fulfillment-conditions. Truth, accuracy, and fulfillment are species of veridicality. It is widely agreed that veridicality-conditions are vital for understanding belief, perception, intention, and many other mental states (Burge, 2010, p. 9; Fodor, 1987, pp. 10-11).

The traditional emphasis upon mental representation encounters periodic resistance among philosophers and psychologists. Some authors castigate mental representation as a scientifically unrespectable notion that should be expunged from serious theorizing (e.g. Chemero, 2009; Quine, 1960; Skinner, 1938; van Gelder, 1992). In response, *representationalists* maintain that representational discourse is both legitimate and scientifically indispensable. Fodor (1975; 1987) develops the representationalist viewpoint. He argues that current cognitive science offers impressive representational explanations whose benefits are not replicable within a non-representational framework. Burge (2010) argues similarly, focusing especially upon the role that mental representation plays within perceptual psychology.

Sensorimotor psychology provides strong support for representationalism. The science seeks to explain how motor activity transforms intentions into motor commands that promote fulfillment of those intentions. An intention's fulfillment-condition plays a key role in explaining the motor commands that it triggers. In particular, the fulfillment-condition helps explain why a particular cost function is operative in a given motor task. For example, if I intend to move my index finger to some target, then the cost function assigns lower cost to outcomes where my finger reaches the target. If I intend to move my hand through five widely space targets, then the cost function assigns lower cost to outcomes where my hand moves through those targets. More generally, the intention operative in a motor task helps determine the cost function that the motor system employs when computing minimal (or near-minimal) expected costs. This explanatory strategy presupposes personal-level intentions with fulfillment-conditions. We cite the condition under which an intention is fulfilled (e.g. that my finger reaches the target) to explain which cost function the motor system deploys.

The science also invokes veridicality-conditions when characterizing subpersonal computations. According to current Bayesian models, the motor system estimates environmental conditions by updating subjective probabilities $p(h)$ assigned to hypotheses h. Each hypothesis h represents the environment as being a certain way. For example, h might represent the size of an object, or it might represent the current configuration of one's body. h is veridical just in case the environment is as h represents it. The cost function $c(h, u)$ measures the desirability of issuing motor command u in situations where h is veridical. Detailed explanatory generalizations describe **(a)** how sensory input and efference copy influence reallocation of subjective probabilities over hypotheses **(b)** how subjective probabilities along with the cost function inform selection of motor commands. Generalizations of types (a) and (b) cite veridicality-conditions. To illustrate, suppose the cost function rewards placing my index finger at a target. Then a different motor command will result when the Bayesian estimator treats it as likely that my finger is left of target than when the Bayesian estimator treats it as likely that my finger is right of target. A type (b) generalization rigorously captures this pattern by citing veridicality-conditions of hypotheses, e.g. by adducing whether hypothesis h represents my finger as left or right of target. In this manner, our current best sensorimotor psychology assigns a central explanatory role to representational aspects of subpersonal motor computation.

Over the past century, many scientists have tried to explain motor control in non-representational terms (e.g. Chemero, 2009; Kelso, Dumas, and Tognoli, 2013). These attempts have proved far less explanatorily successful than representational theories. In particular, anti-representationalists have not successfully explained Bernstein's fundamental observation that motor performance displays more variability along task-irrelevant dimensions than task-relevant dimensions. Indeed, it is unclear whether one can so much as state this explanandum in non-representational terms, since the task-relevant/task-irrelevant distinction presupposes a *goal* represented and pursued by the subject.

Anti-representationalists often insist that, even when a scientific explanation invokes mental representation, we can recast the explanation in non-representational terms while replicating any benefits that it offers. They claim that we can purge representational locutions from cognitive science without explanatory loss (e.g. Field, 2001; Stich, 1983).

I think that we should regard such claims quite warily. We cannot usually purge a science of its central notions while preserving its explanatory achievements. Physicists cannot renounce forces while offering anything resembling the explanations offered by Newtonian physics. Biologists cannot renounce genes while preserving anything resembling modern genetics. Similarly, anti-representationalists have given us little reason to suspect that we can purge cognitive science of representational mental states while preserving its explanatory achievements. For further defense of a representationalist perspective on sensorimotor psychology, see (Rescorla, 2016). For a kindred representationalist analysis of perceptual psychology, see (Rescorla, 2015).[2]

[2] I have argued that explanations offered by sensorimotor psychology cite representational properties of mental states. However, some representational properties are far more relevant than others to the explanation of bodily motion. Context-dependent representational properties are usually much less relevant than context-invariant representational properties. For example, suppose I perceive a marble on the floor and form an intention to grab *that*

8. Future directions

We are only beginning to understand motor control in humans and other animals. To conclude, I highlight a few areas where further scientific or philosophical inquiry is needed. In several cases, these areas are under active investigation.

Computational architecture. The motor system and the perceptual system both estimate environmental state. How does sensorimotor estimation relate to perceptual estimation? This question has generated considerable controversy (e.g. Briscoe, 2008; Milner and Goodale, 1995; Schenk, Franz, and Bruno, 2011). A less widely discussed question concerns the extent to which high-level cognition influences motor control. Intentions crucially influence motor processing by influencing the cost function. Beliefs, desires, and other mental states influence the cost function *by way of* influencing intentions. To what extent can higher-level mental states influence motor activity without mediation by intentions? For example, to what extent can conscious beliefs influence subpersonal estimation by the motor system? Such questions relate to longstanding debates about the *modularity* of perception sparked by Fodor's (1983) famous discussion.

Format and content. I have argued that motor control features subpersonal representational mental states: probability assignments to hypotheses that have veridicality-conditions. We would like to understand the format and content of the hypotheses. Are they propositional? Conceptual? Do they involve something like imagistic or map-like representation? How do they bear upon standard philosophical theories of representational content, such as Fregean thoughts or Russellian propositions? Further investigation would illuminate the

---

marble. My intention is fulfilled only if I grab that particular marble. My intention's fulfillment-condition depends upon the specific marble represented by the intention. But the specific marble does not seem explanatorily relevant to explaining which motor commands my motor system issues. The marble could have been replaced by a qualitatively indistinguishable duplicate, and this change does not seem relevant to explaining my bodily motion. Thus, not *all* aspects of the intention's fulfillment-condition seem relevant to explaining my bodily motion. The intention has context-invariant representational properties that, in conjunction with context, help determine its fulfillment-condition. The context-invariant representational properties are what seem relevant to explaining the bodily motion. These issues merit extensive further discussion. For present purposes, I must set them aside.

underpinnings of sensorimotor psychology. It would also enhance philosophical discussion of representational content with an expanded diet of scientifically important examples.

Subpersonal subjective probability. What is it to assign a subjective probability to a hypothesis? This question is basic for Bayesian decision theory. Philosophers have made several attempts to answer the question in a non-trivial way, without much visible success (Erikkson and Hájek, 2007). Existing proposals tend to emphasize sophisticated personal-level activities, such as gambling or linguistic communication. The motor system does not engage in such activities. Thus, it seems doubtful that existing proposals shed much light on subjective probability as it figures in motor computation. A major philosophical task is to elucidate the subjective probabilities employed by the motor system and to clarify how exactly they resemble the personal-level subjective probabilities emphasized by traditional philosophical inquiry. Can we provide an analysis of subjective probability that applies equally well to the personal and subpersonal levels?

Intention. The nature of intention is a central topic for philosophy of mind and philosophy of action. Discussion usually focuses upon the role that intention plays within practical reasoning. The basic strategy is to explore how intention interfaces with belief, desire, planning, instrumental reasoning, and other high-level facets of human psychology. A complementary strategy, pursued recently by Pacherie (2000; 2006), is to explore how intention interfaces with motor control. Fulfilling an environment-directed intention requires a capacity to control one's motor organs in appropriate ways. Accordingly, one might hope to illuminate environment-directed intentions by studying how they engage subpersonal motor processing. A potential benefit of this complementary strategy is that sensorimotor psychology is far better developed than the science of high-level propositional attitudes, so that it arguably provides a sounder basis for philosophical inquiry.

These are just a few research avenues suggested by sensorimotor psychology. There are many additional avenues that I have not addressed. Given the central role that motor control plays within our mental lives, the foundations and ramifications of sensorimotor psychology merit thorough scrutiny by the philosophical community.

Acknowledgments

I am grateful to Matteo Colombo and Mark Sprevak for comments that improved this entry.

Works Cited

Bernstein, N. 1923. “Studies on the Biomechanics of Hitting Using Optical Recording.” *Annals of the Central Institute of labor* 1: 19-79.

---. 1930. “A New Method of Mirror Cyclographie and its Application Towards the Study of Labor Movements During Work on a Workbench.” *Hygiene, Safety and Pathology of Labor* 56: 3–11.

---. 1967. *The Coordination and Regulation of Movements*. Oxford: Pergamon.

Briscoe, R. 2008. “Another Look at the Two Visual Systems Hypothesis: the Argument from Illusion Studies.” *Journal of Consciousness Studies* 8: pp. 3-62.

Burge, T. 2005. “Disjunctivism and Perceptual Psychology.” *Philosophical Topics* 33: pp. 1-78.

---. 2009. “Five Theses on *De Re* States and Attitudes.” In *The Philosophy of David Kaplan*, eds. J. Almog and P. Leonardi. Oxford: Oxford University Press. pp. 246-324.

---. 2010. *Origins of Objectivity*. Oxford: Oxford University Press.

Chemero, A. 2009. *Radical Embodied Cognitive Science*. Cambridge: MIT Press.

Clark, A. 2015. *Surfing Uncertainty*. Oxford: Oxford University Press.

Colombo, M., and Hartmann, S. 2017. “Bayesian Cognitive Science: Unification and Explanation.” *The British Journal for the Philosophy of Science* 68: pp. 451-484.

Colombo, M., and Seriès, P. 2012. “Bayes in the Brain --- on Bayesian modeling in Neuroscience.” *British Journal for the Philosophy of Science* 63: pp. 697-723.

Crevecoeur, F., Cluff, T., and Scott, S. 2014. “Computational Approaches for Goal-Directed Movement Planning and Execution.” In *The Cognitive Neurosciences*, eds. M. Gazzaniga and G. Mangun. Cambridge: MIT Press.

Denève, S., Duhamel, J.R., and Pouget, A. 2007. “Optimal Sensorimotor Integration in Recurrent

Cortical Networks: A Neural Implementation of Kalman Filters." *The Journal of Neuroscience* 27: 5744-5756.

DeWolf, T., and Eliasmith, C. 2011. "The Neural Optimal Control Hierarchy for Motor Control." *The Journal of Neural Engineering* 8: 065009.

Erikkson, L., and Hájek, A. 2007. "What Are Degrees of Belief?" *Studia Logica* 86: pp. 183-213.

Field, H. 2001. *Truth and the Absence of Fact*. Oxford: Oxford University Press.

Fodor, J. 1975. *The Language of Thought*. New York: Thomas Y. Crowell.

---. 1983. *The Modularity of Mind*, Cambridge: MIT Press.

---. 1987: *Psychosemantics*. Cambridge: MIT Press.

Friston, K. 2011. "What Is Optimal about Motor Control?". *Neuron* 72: pp. 488-498.

Haith, A., Jackson, C., Miall, C., and Vijayakumar, S. 2009. "Unifying the Sensory and Motor Components of Sensorimotor Adaptation." In *Advances in Neural Information Processing Systems* 21, eds. D. Koller, D. Schuurmans, Y. Bengio, and L. Bottou: pp. 593-600.

Haith, A., and Krakauer, J. 2013. "Theoretical Models of Motor Control and Motor Learning." In *Progress in Motor Control VII: Neural Computational and Dynamic Approaches*, eds. M. Richardson, M. Riley, and K. Shockley. Springer: New York. pp. 7-28.

Helmholtz, H. von. 1867. *Handbuch der Physiologischen Optik*. Leipzig: Voss.

Hohwy, J. 2014. *The Predictive Mind*. 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: pp. 120-131.

Knill, D., Bondada, A., and Chhabra, M. 2011. "Flexible, Task-Dependent Use of Sensory Feedback to Control Hand Movements." *The Journal of Neuroscience* 31: pp. 1219-1237.

Li, W., Todorov, E., and Pan, X. 2004. "Hierarchical Optimal Control of Redundant Biomechanical Systems." *Proceedings of the 26th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*. San Francisco.

Mazzoni, P. and Krakauer, J. 2006. "An Implicit Plan Overrides an Explicit Strategy During Visuomotor Adaptation." *Journal of Neuroscience* 26: pp. 3642-3645.

Miall, R. C., and Wolpert, D. 1996. "Forward Models for Physiological Motor Control." *Neural Networks* 9: pp. 1265-1279.

Milner, A., and Goodale, M. 1995. *The Visual Brain in Action*. Oxford: Oxford University Press.

Nashed, J., Crevecoeur, F. and Scott, S. 2012. "Influence of the Behavioral Goal and Environmental Obstacles on Rapid Feedback Responses." *Journal of Neurophysiology* 108: pp. 999-1009.

Pacherie, E. 2000. "The Content of Intentions." *Mind and Language* 15: pp. 400-432.

---. 2006. "Towards a Dynamic Theory of Intentions." In *Does Consciousness Cause Behavior? An Investigation of the Nature of Volition*, eds. S. Pockett, W. Banks, and S. Gallagher. Cambridge: MIT Press. pp. 145-167.

Palmer, S. 1999. *Vision Science*. Cambridge: MIT Press.

Peacocke, C. 1992. "Scenarios, Concepts and Perception." In *The Contents of Experience*, ed. T Crane. Cambridge: Cambridge University Press. pp. 105-135.

---. 1993. "Externalist Explanation." *Proceedings of the Aristotelian Society* 93: pp. 203-230.

Quine, W. V. 1960. *Word and Object*. Cambridge: MIT Press.

Rescorla, M. 2015. "Bayesian Perceptual Psychology." In *The Oxford Handbook of the*

Philosophy of Perception, ed. M. Matthen. Oxford: Oxford University Press.
---. 2016. “Bayesian Sensorimotor Psychology,” *Mind and Language* 31: pp. 3-36.
---. Forthcoming. “A Realist Perspective on Bayesian Cognitive Science.” In *Inference and Consciousness*, eds. T. Chan and A. Nes. Routledge.
Rosenbaum, D. 2002. “Motor Control.” In *Stevens' Handbook of Experimental Psychology*, vol. 1, 3rd ed., eds. H. Pashler and S. Yantis. New York: Wiley. pp. 315-340.
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: pp. 812-818.
Scholz, J., Schöner, G., and Latash M. 2000. “Identifying the Control Structure of Multijoint Coordination During Pistol Shooting. *Experimental Brain Research.* 135: pp. 382–404.
Scott, S. 2004. “Optimal Feedback Control and the Neural Basis of Volitional Motor Control.” *Nature Reviews Neuroscience* 5: pp. 532-546.
Shadmehr, R., and Mussa-Ivaldi, S. 2012. *Biological Learning and Control*. Cambridge: MIT Press.
Shadmehr, R., and Wise, S. 2005. *The Computational Neurobiology of Reaching and Pointing*. Cambridge: MIT Press.
Skinner, B. F. 1938. *The Behavior of Organisms*. New York: Appleton-Century-Crofts.
Sober, S. J., and Sabes, P. N. 2005. “Flexible Strategies for Sensory Integration During Motor Planning.” *Nature Neuroscience* 8: pp. 490–497.
Sprevak, M. 2016. “Philosophy of the Psychological and Cognitive Sciences.” In *The Oxford Handbook of Philosophy of Science*, ed. P. Humphrey. Oxford: Oxford University Press.
Stengel, R. 1986. *Optimal Control and Estimation*. New York: Dover.
Stevenson, I., Fernandes, H., Vilares, I., Wei, K., and Körding, K. 2009. “Bayesian Integration and Non-Linear Feedback Control in a Full-Body Motor Task.” *PLOS Computational Biology* 5: e1000629.
Stich, S. 1983. *From Folk Psychology to Cognitive Science*. Cambridge: MIT Press.
Todorov, E. 2009. “Stochastic Optimal Control and Estimation Methods Adapted to the Noise Characteristics of the Sensorimotor System.” *Neural Computation* 17: 1084-1108.
Todorov, E., and Jordan, M. 2002. “Optimal Feedback Control as a Theory of Motor Coordination.” *Nature Neuroscience* 5: pp. 1226-1235.
Todorov, E. 2004. “Optimality Principles in Sensorimotor Control.” *Nature Neuroscience* 7: pp. 907-915.
van Gelder, T. 1992. “What Might Cognition Be, If Not Computation?” *Journal of Philosophy* 92: pp. 345–381.
Wolpert, D. 2007. “Probabilistic Models in Human Sensorimotor Control.” *Human Movement Science* 26: pp. 511-524.
Wolpert, D., and Flanagan, J. R. 2009. “Forward Models.” *The Oxford Companion to Consciousness*. Oxford: Oxford University Press. pp. 294-296.
Wolpert, D., and Landy, M. 2012. “Motor Control is Decision-making.” *Current Opinion in Neurobiology* 22: pp. 1-8.