Draft for *The Routledge Handbook of Bodily Awareness*, eds. Adrian Alsmith and Matthew Longo:
Please don't cite, as revisions may follow

Bodily Skill

Joshua Shepherd¹

Abstract

To a first approximation, 'bodily skill' refers to the capacity to successfully utilize the body in the world to achieve goals. But the body is complex, and bodily skill manifests in many different ways. Further, work on bodily skill spans the philosophy of mind, action, and cognitive science, as well as the sciences of motor control and perception. This chapter aims to provide an overview of recent themes and key ideas. First, we review work on the nature of skill as such. Second, we discuss ways theorists have discussed the implementation of skill in the human body, and a theoretical difficulty associated with explanations of implementation. Third, we discuss the importance of the idea of a hierarchical representational architecture for understanding skill's implementation, and various ways this architecture can be made more sophisticated. Fourth, we discuss work that attempts to understand skill execution on-the-fly, and in particular, work that attempts to understand the connection between intentions and motor representations, and the connection between implicit sensorimotor adaptation and explicit planning and executive control. Fifth, we discuss work on the models that assist fine-grained skill execution, as well as the idea that even at rapid, fine-grained motoric levels, one finds intelligent systems and mechanisms supporting skill.

My charge in this chapter is to offer an overview of 'bodily skill.' I will take my topic to be distinct, and more abstract, than that of bodily skills. There are a number of capacities that, when well-functioning, may qualify as bodily skills in some sense. We have perceptual skills, attentional skills, skills related to memory, imagination, executive control, and so on. This chapter is about such skills only indirectly. For I will focus on the issue of how many of the capacities that make up the living body come together to enable and generate skilled behavior. More specifically still, I will tend to focus on the generation of skilled action.

The chapter is organized like so. In section one, I discuss the nature of skill as such. In section two, I discuss ways theorists have discussed the implementation of skill in the human body, and a theoretical difficulty associated with explanations of implementation. In section three, I discuss the importance of the idea of a hierarchical representational architecture for understanding skill's implementation, and various ways this architecture can be made more sophisticated. In section four, I discuss work that attempts to understand skill execution on-the-fly, and in particular, work that attempts to understand the connection between intentions and motor representations, and the connection between implicit sensorimotor adaptation and explicit planning and executive control. In section five, I discuss work on the models that assist fine-grained skill execution, as well as the idea that even at rapid, fine-grained motoric levels, one finds intelligent systems and mechanisms supporting skill.

¹ The author gratefully acknowledges two sources of support. First, funds from European Research Council Starting Grant 757698, awarded under the Horizon 2020 Programme for Research and Innovation. Second, a fellowship from the Canadian Institute for Advanced Research's Azrieli Global Scholar programme, and the Brain, Mind, and Consciousness program.

1 The nature of skill

Let us begin with an abstract question: what is skill as such? Within the philosophy of psychology, we can discern two families of views.

One seeks to characterize skill in terms of successful action. We might call these success views. Shepherd (2021, 125), for example, argues that skill within some domain of action (e.g., skill at basketball, or skill at chess, or skill at architecture, see Shepherd 2020a) should be understood as the capacity to produce high success rates at goals central to that domain, where the success-rates are understood as explained by the agent's control over her behavior (including her action), and the agent's facility at constructing good plans for action within that domain. Such an account suggests much work in understanding the nature of control over behavior, and in understanding the agent's planning facility, but ultimately the account ties skill to a set of cases that involve some degree of success. Similarly, Fridland (2020, 251) outlines a view of skill on which skills are functions from intentions to successful actions, where the successes are explained by practice-honed control structures that implement the agent's intentions. Again, such an account suggests a focus on control structures, the influence of practice, and the relation of intentions to the execution of action. Again, ultimately skill is yoked to a set of cases of successful action.

A different family of views place less emphasis on control and success, and more on knowledge. We might call these knowledge views. So, for example, Stanley and Williamson (2017) think of skill as a special kind of disposition. Skill at some particular action, for them, is a disposition to form knowledge apt for guiding the action. We do have here an emphasis on guidance or control. But explicating control is much less important for Stanley and Williamson. They wish to connect their account of skill to broader currents in epistemology, and to argue that the direct manifestation of skill is not the action, but is rather the knowledge state the agent forms. (For a different implementation of a similar approach to skill, see Stanley and Krakauer 2013.)

Carlota Pavese (2016a, 2016b) offers a different kind of knowledge view, on which skill is closely tied to an agent's knowledge of how to act. For Pavese, an agent is skilled at an action if and only if the agent knows not only how to perform that action, but knows how to perform that action sufficiently well. (What it is to know how well is a complicated issue, see Pavese 2017). Thus the content of the knowledge, and in some ways the structure of the agent's knowledge (Pavese 2019), is critical to Pavese's understanding of skill.

I leave the philosophical disputes between such accounts aside in this chapter, and I emphasize two things both families share. First, both families recognize that skill comes in degrees – one can be more or less skilled – and that excellence at action execution is in some ways central to the nature of skill. It is important to emphasize this point, for the following reason. The term 'skill' is often used to refer to action abilities that are easily learned in the lab, and that are widely shared. In this sense, skill refers to general competence. Often in psychology, what philosophers call skill is referred to as expertise. This expertise, or excellence, usage of the term skill draws an important distinction between, for example, the weekend shinny (pickup hockey) player and the great Ryan O'Reilly, captain of the St. Louis Blues.

There is no strong need to police usage of the term skill, but it is important to keep the distinction in mind. Consider, for example, a recent debate about the human ability to visually track incoming objects (in, for example, sports like baseball, cricket, or tennis). It was once widely thought that the ball moved too fast for conscious or executive control processes to have any impact on tracking and responding in such sports. But when Mann et al. (2013) brought world-class cricket players into experimental conditions, they found something different. Highly skilled batters could do things that novices could not, namely, couple head movement to ball movement, and make predictively accurate saccades that mirrored the motion of the ball (see Shepherd 2015 for discussion of the relevance of this work to ‘conscious control,’ and see Mann et al. 2019 for recent work on this). The difference between high levels of skill and mere competence, then, may frequently be important for our understanding of system capacities and limitations, as well as for the ways various capacities are interrelated by processes of learning, planning, and executive control.

A second commonality between success views and knowledge views of skill is a tendency towards abstraction. But of course we not only want answers to questions about the form of skill, we want to know how – by means of what physical and psychological mechanisms – agents like human beings manage to produce skilled behaviors. I turn, then, to discussion of skill’s implementation in the human body (and mind), and to a difficulty we all face in understanding it.

2 The implementation of skill, and a difficulty

In a recent discussion, Mylopoulos and Pacherie’s (2020) note that one hallmark of skilled behavior is its flexibility – the skilled agent (within limits, of course) is able to adjust to differences and difficulties in her environment and body, and to produce successful action come what may. How is such flexibility achieved? Mylopoulos and Pacherie offer three aspects to consider. First, skilled agents display sensory sensitivity – they closely track immediate sensory changes. Second, skilled agents display situational sensitivity – they track the nature of the circumstances in which they act, and how these circumstances evolve. Third, skilled agents display strategic sensitivity – they track relationships between the situation, the sensory environment, and their own intentions and goals, and their sense of how these relationships are unfolding informs the way that they implement their plan (by way of, e.g., motor commands).

But how are sophisticated sensitivities like these expressed in the nervous system? Work on this issue is, at the moment, fairly fragmented. Some authors focus on the importance of broader scale capacities to skill (and to action more broadly). One can find, for example, discussions of the role of consciousness (Montero 2016, Shepherd 2015), attention (Wu 2016), motor acuity (Stanley and Krakauer 2013), cognitive control (Christensen et al. 2016), practical reasoning (Shepherd 2018, Fridland 2021), motor control (Fridland 2017), imagination (Briscoe 2008), memory (De Brigard 2019, Christensen et al. 2019), and more, for skilled behavior. Some of these discussions intersect, but in many places the focus remains somewhat siloed.

Other authors focus on the importance of different kinds of representational states to skill (and to action more broadly). One can find, for example, discussions of different sorts of visual representations (Milner and Goodale 2008, Briscoe and Schwenkler 2015, Shepherd 2016),

bodily representations (Alsmith 2019), multi-format representations (Shepherd 2018), metacognitive representations (Mylopoulos and Pacherie 2020), practical representations (Pavese 2019), motor representations (Pacherie 2008, Mylopoulos and Pacherie 2018, Brozzo 2017), motor imagery (Jeannerod 1995, Nanay 2020), as well as discussions of different sorts of representational models that might inform skill (Christensen 2020, Dogge et al. 2019), and discussion of specialized notions like basic action concepts (Schack and Frank 2020), structured action representations (Mylopoulos and Pacherie 2017), and motor primitives (Latash 2020).

We can see, now, that any treatment of how skilled behavior is produced and organized faces a difficulty. In order to understand the relative importance or unimportance of any of these discussions of various capacities and representations (to say nothing of discussions of various processes and systems), we have to have some sense of how they might come together in a comprehensive model of skill. And this is difficult for several reasons.

First, skilled behavior is developed and exercised within some domain of activity (Shepherd 2021). But there are an indefinite number of domains, and no a priori guarantee that skill will look the same across domains as disparate, for example, as chess, jousting, lecturing, tennis, water polo, sculpture, and so on.

Second, bodily skill centrally involves the body. But, as Bernstein (1996) memorably catalogues, the human musculo-skeletal system (to say nothing yet of the nervous system) is an extraordinarily sophisticated machine. What are the features that bodily skill shares in very disparate action domains? How many different ways might there be to construct a skill? These are difficult but important questions for the skill theorist.

Third, bodily skill depends heavily on a range of psychological capacities and processes. But these processes are subtle and the right way to identify and individuate them, and to explain how they relate to one another and what roles they have in the production of action, raises a host of issues for science and philosophy, about which disagreement abounds.

The central difficulty here, then, is not unique to skill, but it is arguably more pronounced. Skilled action enforces a kind of organization on a wide array of topics within cognitive psychology and neuroscience. The difficulty is how to integrate and unify, in a theoretically illuminating way, the plurality of contributors to skilled behavior.

3 A convergence on hierarchical architecture

Some progress regarding the difficulty can be made by noting an important convergence in work on skilled action. Multiple discussions of skill emphasize the centrality of hierarchy in the psychological architecture that supports skilled behavior. Hierarchy is invoked to explain features of specific action types, for example, speech control (Perkell et al. 1997), typing (Logan and Crump 2011), and reaching (Kim et al 2020). It is also invoked (by almost everyone) to explain general features of psychological architecture (Lashley 1951, Rosenbaum et al. 2007, Grafton and Hamilton 2007, Badre 2008, Pacherie 2008, Shepherd 2015, Christensen et al. 2016, Mylopoulos and Pacherie 2018). As Yokoi and Diedrichsen explain, ‘the core idea of a multi-level action hierarchy is that each level combines elements in the hierarchically lower level, facilitating

the transition between lower-level elements. Each hierarchically higher level would lose some of the fine temporal details encoded at the lower level and, hence, represent the action at a more abstract level' (2019, 1).

How is the hierarchy organized? Theorists offer different taxonomies, and I cannot cover all of them here. But let us look, briefly, at a recent proposal due to Wayne Christensen (2020), which builds upon earlier influential work (Christensen et al. 2016). Christensen proposes an interactive multi-level architecture (IMLA). Towards the top of the hierarchy, Christensen includes three systems capable of impacting and interacting within working memory. There is, at the top, a general reasoning system. One level down, he posits a situated action system, 'which controls action in relation to situation and task structure, such as using a screw driver to screw in a screw while attaching a leg of an Ikea chair' (2020, 10). This system interacts in significant ways with an integrative motor system, and the joint operations of these two systems are responsible for a good deal of action execution. Importantly, however, all three systems exhibit rich interaction and functional integration, and all three are capable of having direct impact upon conscious awareness and conscious cognitive control. At the bottom of the hierarchy Christensen posits a modular motor system, which is akin to the dorsal stream system responsible for fine-grained action planning and execution, and familiar from Milner and Goodale's work (e.g., Goodale and Milner 1992, Milner and Goodale 2008).

We can fruitfully compare the working memory-involving portion of Christensen's proposal with a recent architecture proposed by Badre and Nee (2018). Their hierarchy is more directly tied to neural organization, and in particular a rostrocaudal functional organization found within the frontal cortex. A basic idea that Badre and Nee share with many neuroscientists goes as follows.

[T]he common element has been that rostral frontal areas are involved in more abstract forms of control than more caudal areas. This putative rostrocaudal abstraction gradient has been theorized to support a hierarchical processing architecture of the frontal lobes, wherein abstract goals are actively translated into movements via a rostral-to-caudal flow of processing, and rostral areas of frontal cortex influence and organize processing in posterior areas. (Badre and Nee 2018, 270)

Badre and Nee propose an interesting update on this general architecture. First, they posit three functional subdivisions within frontal cortex. At the most caudal, they posit a sensory-motor control system responsible for execution of movements. There is likely further hierarchical organization within this system, structured by the abstraction of various movement representations. And one can see a rough accord with Christensen's integrative motor system. Second, they posit a contextual control system located within the mid-lateral prefrontal cortex. This system is responsible for relating ongoing representations of the context to goals and action execution possibilities. Again, one can see a rough accord with Christensen's situated action system. Third, they posit a system for schematic control at the most rostral, within the right lateral prefrontal cortex. Here there is a possible divergence from Christensen. As Badre and Nee explain their proposal, schematic control has a 'generality beyond only

temporal or episodic signals’ (2018, 179). The notion of a schema in play here is akin to a knowledge structure, or model, ‘that organizes many lower-order features and their relationships’ (179).

Importantly, for Badre and Nee this three-fold division is not strictly hierarchical. Rich functional integration is posited, and the most rostral region responsible for schematic control is not, strictly speaking, the top of the hierarchy. They place the anterior mid-DLPFC, which (speaking roughly here) sits at the bottom of schematic control, and at the border with contextual control, at the top of the hierarchy. Within schematic control, right lateral prefrontal cortex performs a domain-specific function analogous to that of the more caudal regions. Its function is to apply schematic knowledge to cognitive control tasks.

The idea that, within a generally hierarchical architecture, one might find significant further interlevel structure, is promising. A similar idea has been recently taken up by Mylopoulos and Pacherie (2020), who argue (drawing on Cushman and Morris 2015, and see Shepherd 2021b for a development of this idea along slightly different lines) that action control can be understood as heterarchical in an important sense. They argue that heterarchy should not be understood as a feature of the representations guiding action, but rather as a feature that emerges from learning, or skill development.

[A] representational hierarchy is in principle compatible with a control heterarchy. For instance, Cushman and Morris (2015), from whom we borrow the expression “heterarchical control”, propose that whether or not one relies on automated control processes (what they call “habitual control”) or on more cognitive control processes (what they call “planned control”) is not necessarily a matter of the level of abstraction in the hierarchy of action specification. Rather, this depends on the amount of variability present at a given level in terms of the action types that are available to serve the corresponding goals within the given context and on the amount of experience the agent has with the task as specified at that level. (2020, 5)

The connection to learning is important, for it seems clear that one cannot understand the psychological architecture that supports bodily skill as static. This architecture is dynamic, and is constantly impacted by practice and feedback and learning of many sorts (see Krakauer et al. 2019 for a review). As they develop skills, agents are simultaneously developing models of a task space (or an action domain), as well as progressively building and refining a number of important representational structures. Throughout the learning process, high-level conceptualization processes interact closely with more fine-grained sensorimotor adaptation, motor chunking and parsing, and motor sequencing processes (Krakauer et al. 2019).

One interesting recent illustration of this was provided by Popp et al. (2020). They asked participants to execute finger tapping sequences on a keyboard, in two different conditions. One condition had participants chunk (by asking them to remember and practice) the sequence of finger movements in a way that accorded efficiently with basic physical constraints on the action. A second condition had the participants chunk the movements inefficiently. They found that

even though the movement sequences were the same, the inefficient chunking group performed worse initially, and after ten days of practice, and that the inefficient chunking structure – the inefficient cognitive representations with which they began – were to blame. (Interestingly, with enough practice participants would overcome the deficit, in effect by solving the problems in their own way – Popp et al. report that ‘participants found idiosyncratic, but inter-individually stable, ways of performing each sequence’ (1449).

One key feature of work on skilled action is the way that practice makes the representational structures that guide action more efficient. Thomas Schack and colleagues (see, e.g., Schack 2004, Schack and Mechsner 2006, Schack 2012, Schack and Frank 2020) have used structural dimensional analysis applied to what they call basic action concepts to better understand how. A general finding in this region is that experts have very different basic action concepts, and differently organized networks of basic action concepts, than non-experts.

Of course possession or development of a useful representational structure is one thing, and using that structure in a skillful way is another. How do agents manage to put representations like basic action concepts to work in guiding on-the-fly activity?

4 Representations, systems, and interfaces

Consider what Butterfill and Sinigaglia (2014) have called the interface problem. This is a problem about how intentions and motor representations coordinate to jointly guide on-the-fly action. On the initial specification of the problem, there is a mystery involved in how they do so. For intentions and motor representations, arguably, represent an action’s outcome in distinct representational formats: a propositional format and a motoric format. Representational formats are understood, roughly, as representational systems that contain proprietary representational primitives, and distinct combinatorial rules for how the primitives interact (Shepherd 2018). Think, analogously, of different languages, or of the difference between a picture and a linguistic description of a picture. What kinds of interaction between distinct representational systems are possible, and how does the mind manage them? From one perspective, the process(es) can seem computationally difficult, and possibly very expensive. Perhaps an understanding of how the mind manages such coordination will help illuminate on-the-fly skilled action.

Butterfill and Sinigaglia’s initial specification of the problem generated a wave of work in the philosophy of action and psychology, with a number of different ‘solutions’ proposed (Sinigaglia and Butterfill 2015, Mylopoulos and Pacherie 2017, Burnston 2017, Brozzo 2017, Ferretti and Zipoli Caiani 2018, Shepherd 2018, Shepherd 2019, Christensen 2020). Mylopoulos and Pacherie (2017) argue that intentions take executable action concepts as contents, and that these contents can trigger appropriate motor representations via links between executable action concepts and motor schemata. Burnston (2017) argues that intentions do not determine the contents of motor representations, but can bias the values motor representations take along certain parameters. Shepherd (2018, 2019) argues that interaction between representational formats within practical reasoning is the norm, and that intentions can sometimes specify action in a motoric format that motor representations take on board. Christensen (2020) motivates and

emphasizes mechanisms that use metarepresentational rules for developing links between different representational systems.

One interesting feature of this work is the way that it draws on a similar – but arguably importantly distinct – problematic within literatures on motor control and sensorimotor adaptation. From the perspective of this work, widespread sharing of representational resources is the norm, and the issue is rather to understand what representational systems, and what representations, do important work.

Recent work on sensorimotor adaptation suggests that motor learning in some contexts requires coordination between at least two sensorimotor learning systems that operate in parallel (Taylor and Ivry 2011, McDougle et al. 2015, McDougle et al. 2016, Krakauer et al. 2019, Albert et al. 2020). An ‘explicit system’ is sensitive to task success, and failures at this level can generate explicit changes in strategy by the agent. By contrast, an ‘implicit’ system is sensitive to sensory prediction error, and failures at this level lead to small changes to the implementation of an action (e.g., the direction of a reach). Kim et al. (2020) argue that this makes good functional sense:

The computational goal of the implicit component of adaptation is to keep the motor system well calibrated, ensuring that motor commands are faithfully implemented (Miyamoto et al. 2020). This system appears impervious to information about task success. In contrast, the computational goal of a strategic process such as aiming is to ensure that the appropriate action is selected to optimize performance. For this to occur, it is essential to monitor how well a selected action achieves its desired goal. (83)

Even so, interactions between the implicit and explicit systems raise open questions. As Krakauer et al. (2019) note, it is not clear how far work on sensorimotor learning with well-established actions like reaching can generalize to skill learning for novel tasks. And Hadjiosif and Krakauer (2020) have recently questioned whether the implicit/explicit distinction simplistically obscures an underlying reality that is ‘more graded and granular’ (4). It might be that full answers to questions about skill learning require a merging of explanations between issues that animate neuroscientists and issues that animate philosophers.

5 Intelligence all over

We have seen that bodily skill is underwritten by a sophisticated representational hierarchy that enables massive and fruitful interaction between a wide range of psychological capacities and sub-systems. One point that emerges from recent work on bodily skill is that processing up and down this hierarchy operates in a sophisticated way. To use Ellen Fridland’s (2017) memorable phrasing, it’s intelligence all the way down. Or, to put the point in a slightly different way, a perspective on which practical reasoning is the seat of the agent’s intelligence appears to be false. The agent is comprised of a range of intelligent processes, directed towards different aspects of skill learning and skill execution.

One striking illustration of this comes from work on the long-latency stretch reflex, a reaction to perturbation that occurs roughly 50-100ms afterwards (Pruszynski and Scott 2012). This reflex occurs so quickly that agents are rarely aware of the processing that drives it. But the reflex is intelligently integrated into what the agent is doing, and, as Forgaard et al. (2021) recently note, the stretch reflex is differentially influenced by an agent's intention, including intended movement distance, and by verbal instruction given to the agent.

Weiler et al. had participants perform actions that required different motions of the wrist in order to reach a target. Perturbations to the arm generated differential, but appropriate, long-latency reflexes in either the wrist flexor or wrist extensor muscles. Weiler et al. comment:

[L]arge long-latency stretch responses were always evoked in the triceps, but flexibility evoked in the WF when wrist flexion assisted the reaching movement and in the WE when wrist extension assisted the reaching movement and that this flexibility required no training or practice. Thus, both voluntary actions and rapid feedback responses possess the ability to flexibly coordinate the movement of multiple arm joints during reaching, possibly because they both have access to a similar—or the same—internal model that accounts for the arm's dynamics. (545)

Weiler et al. (2018) claim that 'the neural circuits that generate the long-latency stretch response use an internal model that accounts for biomechanical complexities (e.g., interaction torques) associated with controlling a multi-joint linkage like the arm' (545). Regarding this claim, Maeda et al. (2020) offer fascinating evidence in favor of the idea that an internal model of the arm is utilized by the system that generates long-latency reflexes. (For elaboration of the optimal feedback control theory that has been influential in understanding how motor coordination might be computationally achieved, see Todorov and Jordan 2002, Diedrichsen et al. 2010.) They consistently perturbed participant arms in a passive condition to generate a reflex. Then, they had participants perform reaching actions. These reaching actions displayed signatures of learning from the passive reflex condition. In other words, learning at the level of passive responses transferred to the planning of voluntary motor responses.

Interestingly, this model of the arm appears to go beyond arm-muscle reflexes, and to flexibly integrate feedback from cutaneous receptors in the hand, as well as to direct feedback responses in conditions of fine grip control (Crevecour et al. 2016, Forgaard et al. 2021). So, for example, Hernandez-Castillo et al. (2020) applied stimulus to the fingertips, mimicking an object slipping out of hand, and also applied stimulus to the shoulder, jostling the arm. Stimulus to the fingertips generated a long-latency reflex at the shoulder – the reflex aimed to move the arm in the direction of the falling object. Further, the long-latency reflex at the shoulder was flexibly changed depending on whether stimuli were delivered to the fingertips alone, or to the fingertips and shoulder. In a review, Forgaard et al. note that 'these results indicate that afferent feedback from the fingers can be quickly integrated with proprioceptive feedback from the arm to direct rapid, behaviorally relevant responses' (4).

In summary, understanding what Krakauer (2019) has called the ‘intelligent reflex’ is interesting on its own. The broader point, however, is that bodily skill is subserved by a range of systems, and systems that operate primarily outside the purview of the agent’s awareness display flexibility and task-appropriate intelligence.

The point here is not the relatively crude one that unconscious functionality drives much of the agent’s skill. In fact, much recent work demonstrates the opposite – cognition and strategic planning are rife throughout motor learning tasks (Krakauer et al. 2019), and cognitive control is critical for skill (Christensen et al. 2016). This idea has to be maintained at the same time as the idea that a range of mechanisms outside the direct purview of strategic planning and cognitive control demonstrate flexible intelligence as well.

References

- Albert, S. T., Jang, J., Haith, A. M., Lerner, G., Della-Maggiore, V., Krakauer, J. W., & Shadmehr, R. (2020). Competition between parallel sensorimotor learning systems. *bioRxiv*.
- Alsmith, A. J. (2019). Bodily structure and body representation. *Synthese*, 1-30.
- Badre, D. (2008). Cognitive control, hierarchy, and the rostro–caudal organization of the frontal lobes. *Trends in cognitive sciences*, 12(5), 193-200.
- Badre, D., & Nee, D. E. (2018). Frontal cortex and the hierarchical control of behavior. *Trends in cognitive sciences*, 22(2), 170-188.
- Bernstein, N. A. (1996). *Dexterity and its development*. Psychology Press.
- Briscoe, R. E. (2008). Vision, action, and make-perceive. *Mind & Language*, 23(4), 457-497.
- Briscoe, R., & Schwenkler, J. (2015). Conscious vision in action. *Cognitive Science*, 39(7), 1435-1467.
- Brozzo, C. (2017). Motor intentions: How intentions and motor representations come together. *Mind & Language*, 32(2), 231-256.
- Burnston, D. C. (2017). Interface problems in the explanation of action. *Philosophical Explorations*, 20(2), 242-258.
- Butterfill, S. A., & Sinigaglia, C. (2014). Intention and motor representation in purposive action. *Philosophy and Phenomenological Research*, 88(1), 119-145.
- Christensen, W. (2020). The skill of translating thought into action: framing the problem. *Review of Philosophy and Psychology*, 1-27.
- Christensen, W., Sutton, J., & McIlwain, D. J. (2016). Cognition in skilled action: Meshed control and the varieties of skill experience. *Mind & Language*, 31(1), 37-66.
- Christensen, W., Sutton, J., & Bicknell, K. (2019). Memory systems and the control of skilled action. *Philosophical Psychology*, 32(5), 692-718.

- Crevecœur, F., Thonnard, J. L., Lefèvre, P., & Scott, S. H. (2016). Long-latency feedback coordinates upper-limb and hand muscles during object manipulation tasks. *Eneuro*, 3(1).
- Cushman, F., & Morris, A. (2015). Habitual control of goal selection in humans. *Proceedings of the National Academy of Sciences*, 112(45), 13817-13822.
- De Brigard, F. (2019). Know-how, intellectualism, and memory systems. *Philosophical Psychology*, 32(5), 719-758.
- Diedrichsen, J., Shadmehr, R., & Ivry, R. B. (2010). The coordination of movement: optimal feedback control and beyond. *Trends in cognitive sciences*, 14(1), 31-39.
- Dogge, M., Custers, R., & Aarts, H. (2019). Moving forward: On the limits of motor-based forward models. *Trends in Cognitive Sciences*, 23(9), 743-753.
- Ferretti, G., & Caiani, S. Z. (2019). Solving the interface problem without translation: The same format thesis. *Pacific Philosophical Quarterly*, 100(1), 301-333.
- Forgaard, C. J., Reschechtko, S., Gribble, P. L., & Pruszynski, J. A. (2021). Skin and muscle receptors shape coordinated fast feedback responses in the upper limb. *Current Opinion in Physiology*.
- Fridland, E. (2017). Skill and motor control: Intelligence all the way down. *Philosophical Studies*, 174(6), 1539-1560.
- Fridland, E. (2021). Skill and strategic control. *Synthese*, 1-28.
- Fridland, E. (2021). The nature of skill: functions and control structures. In E. Fridland and C. Pavese, *The Routledge Handbook of Philosophy of Skill and Expertise*. Routledge, Oxon, 245-257.
- Goodale, M. A., & Milner, A. D. (1992). Separate visual pathways for perception and action. *Trends in neurosciences*, 15(1), 20-25.
- Grafton, S. T., & Hamilton, A. F. D. C. (2007). Evidence for a distributed hierarchy of action representation in the brain. *Human movement science*, 26(4), 590-616.
- Hadjiosif, A. M., & Krakauer, J. W. (2020). The explicit/implicit distinction in studies of visuomotor learning: Conceptual and methodological pitfalls. *European Journal of Neuroscience*.
- Hernandez-Castillo, C. R., Maeda, R. S., Pruszynski, J. A., & Diedrichsen, J. (2020). Sensory information from a slipping object elicits a rapid and automatic shoulder response. *Journal of neurophysiology*, 123(3), 1103-1112.
- Jeannerod, M. (1995). Mental imagery in the motor context. *Neuropsychologia*, 33(11), 1419-1432.
- Kim, H. E., Avraham, G., & Ivry, R. B. (2020). The psychology of reaching: action selection, movement implementation, and sensorimotor learning. *Annual review of psychology*, 72.
- Krakauer, J. W. (2019). The intelligent reflex. *Philosophical Psychology*, 32(5), 822-830.
- Krakauer, J. W., Hadjiosif, A. M., Xu, J., Wong, A. L., & Haith, A. M. (2019). Motor learning. *Comprehensive Physiology*, 9(2), 613-663.

- Lashley, K. S. (1951). *The problem of serial order in behavior* (Vol. 21). Oxford, United Kingdom: Bobbs-Merrill.
- Latash, M. L. (2020). On primitives in motor control. *Motor control*, 24(2), 318-346.
- Logan, G. D., & Crump, M. J. (2011). Hierarchical control of cognitive processes: The case for skilled typewriting. In *Psychology of learning and motivation* (Vol. 54, pp. 1-27). Academic Press.
- Maeda, R. S., Gribble, P. L., & Pruszynski, J. A. (2020). Learning new feedforward motor commands based on feedback responses. *Current Biology*, 30(10), 1941-1948.
- Mann, D. L., Spratford, W., & Abernethy, B. (2013). The head tracks and gaze predicts: how the world's best batters hit a ball. *PloS one*, 8(3), e58289.
- Mann, D. L., Nakamoto, H., Logt, N., Sikkink, L., & Brenner, E. (2019). Predictive eye movements when hitting a bouncing ball. *Journal of vision*, 19(14), 28-28.
- McDougle, S. D., Bond, K. M., & Taylor, J. A. (2015). Explicit and implicit processes constitute the fast and slow processes of sensorimotor learning. *Journal of Neuroscience*, 35(26), 9568-9579.
- McDougle, S. D., Ivry, R. B., & Taylor, J. A. (2016). Taking aim at the cognitive side of learning in sensorimotor adaptation tasks. *Trends in cognitive sciences*, 20(7), 535-544.
- Milner, A. D., & Goodale, M. A. (2008). Two visual systems re-viewed. *Neuropsychologia*, 46(3), 774-785.
- Montero, B. G. (2016). *Thought in action: Expertise and the conscious mind*. Oxford University Press.
- Mylopoulos, M., & Pacherie, E. (2017). Intentions and motor representations: The interface challenge. *Review of Philosophy and Psychology*, 8(2), 317-336.
- Mylopoulos, M., & Pacherie, E. (2019). Intentions: The dynamic hierarchical model revisited. *Wiley Interdisciplinary Reviews: Cognitive Science*, 10(2), e1481.
- Nanay, B. (2020). Motor imagery and action execution. *Ergo, an Open Access Journal of Philosophy*, 7.
- Nanay, B. (2020). Motor imagery and action execution. *Ergo, an Open Access Journal of Philosophy*, 7.
- Pacherie, E., & Mylopoulos, M. (2020). Beyond automaticity: The psychological complexity of skill. *Topoi*, 1-14.
- Pacherie, E. (2008). The phenomenology of action: A conceptual framework. *Cognition*, 107(1), 179-217.
- Pavese, C. (2016). Skill in epistemology I: Skill and knowledge. *Philosophy Compass*, 11(11), 642-649.
- Pavese, C. (2016). Skill in epistemology II: Skill and know how. *Philosophy Compass*, 11(11), 650-660.
- Pavese, C. (2017). Know-how and gradability. *The Philosophical Review*, 126(3), 345-383.

- Pavese, C. (2019). The psychological reality of practical representation. *Philosophical Psychology*, 32(5), 784-821.
- Perkell, J., Matthies, M., Lane, H., Guenther, F., Wilhelms-Tricarico, R., Wozniak, J., & Guiod, P. (1997). Speech motor control: Acoustic goals, saturation effects, auditory feedback and internal models. *Speech communication*, 22(2-3), 227-250.
- Popp, N. J., Yokoi, A., Gribble, P. L., & Diedrichsen, J. (2020). The effect of instruction on motor skill learning. *Journal of Neurophysiology*, 124(5), 1449-1457.
- Pruszynski, J. A., & Scott, S. H. (2012). Optimal feedback control and the long-latency stretch response. *Experimental brain research*, 218(3), 341-359.
- Rosenbaum, D. A., Cohen, R. G., Jax, S. A., Weiss, D. J., & Van Der Wel, R. (2007). The problem of serial order in behavior: Lashley's legacy. *Human movement science*, 26(4), 525-554.
- Schack, T. (2004). The cognitive architecture of complex movement. *International journal of sport and exercise psychology*, 2(4), 403-438.
- Schack, T. (2012). Measuring mental representations. *Handbook of measurement in sport and exercise psychology*, 203-214.
- Schack, T., & Frank, C. (2020). Mental Representation and the Cognitive Architecture of Skilled Action. *Review of Philosophy and Psychology*, 1-20.
- Schack, T., & Mechsner, F. (2006). Representation of motor skills in human long-term memory. *Neuroscience letters*, 391(3), 77-81.
- Shepherd, J. (2015). Conscious control over action. *Mind & language*, 30(3), 320-344.
- Shepherd, J. (2016). Conscious action/zombie action. *Noûs*, 50(2), 419-444.
- Shepherd, J. (2018). Intelligent action guidance and the use of mixed representational formats. *Synthese*, 1-20.
- Shepherd, J. (2019). Skilled action and the double life of intention. *Philosophy and Phenomenological Research*, 98(2), 286-305.
- Shepherd, J. (2021). Skill and sensitivity to reasons. *Review of Philosophy and Psychology*, 1-13.
- Shepherd, J. (2021). *The shape of agency: Control, action, skill, knowledge*. Oxford University Press, USA.
- Sinigaglia, C., & Butterfill, S. A. (2015). On a puzzle about relations between thought, experience and the motoric. *Synthese*, 192(6), 1923-1936.
- Stanley, J., & Krakauer, J. W. (2013). Motor skill depends on knowledge of facts. *Frontiers in human neuroscience*, 7, 503.
- Stanley, J., & Williamson, T. (2017). Skill. *Noûs*, 51(4), 713-726.
- Taylor, J. A., & Ivry, R. B. (2011). Flexible cognitive strategies during motor learning. *PLoS Comput Biol*, 7(3), e1001096.

Todorov, E., & Jordan, M. I. (2002). Optimal feedback control as a theory of motor coordination. *Nature neuroscience*, 5(11), 1226-1235.

Weiler, J., Gribble, P. L., & Pruszynski, J. A. (2018). Rapid feedback responses are flexibly coordinated across arm muscles to support goal-directed reaching. *Journal of neurophysiology*, 119(2), 537-547.

Wu, W. (2016). Experts and deviants: The story of agentic control. *Philosophy and Phenomenological Research*, 93(1), 101-126.

Yokoi, A., & Diedrichsen, J. (2019). Neural organization of hierarchical motor sequence representations in the human neocortex. *Neuron*, 103(6), 1178-1190.