Anti-Intellectualism for the Learning and Employment of Skill

Daniel C. Burnston
Forthcoming in *Review of Philosophy and Psychology*
Penultimate draft. Please cite final version.

Abstract

I draw on empirical results from perceptual and motor learning to argue for an anti-intellectualist position on skill. Anti-intellectualists claim that skill or know-how is non-propositional. Recent proponents of the view have stressed the *flexible* but *fine-grained* nature of skilled control as supporting their position. However, they have left the nature of the mental representations underlying such control undertheorized. This leaves open several possible strategies for the intellectualist, particularly with regards to skill learning. Propositional knowledge may structure the *inputs* to sensorimotor learning, may *constitute* the outcomes of said learning, or may be needed for the *employment* of learned skill. I argue that sensorimotor learning produces multi-scale associational representations, and that these representations are of the right sort to underlie flexible and fine-grained control. I then suggest that their content is vitally *indeterminate* with regards to propositional content attribution, because they exhibit a kind of open-ended structure. I articulate this kind of structure, and use it to respond to the three intellectualist strategies. I then show how the perspective I advance offers insights for understanding both instruction and expert practice.

1. Introduction

In this paper, I will argue in favor of anti-intellectualist views of skill, via analysis of perceptual and motor learning. Recent anti-intellectualist arguments have contended that the fine-grained and flexible nature of skillful control resists explanation in terms of propositional knowledge. However, they have left the nature of the mental representations involved in this kind of control undertheorized. As a result, intellectualists can respond in a number of ways, particularly regarding how skills are *learned*. Propositional knowledge can be vital for skill either as the *input* to learning, as the *outcome* of learning, or in the *employment* of learned abilities.

I will discuss the nature of the mental representations acquired during perceptual and motor learning, as a way of motivating the anti-intellectualist position. There are two key points about these representations: (i) their content *outstrips* any propositional knowledge agents have prior to learning (e.g., via instruction); and (ii) their contents are *structured* in a way that resists describing them as modes of presentation of propositional contents. I will argue that these kinds of representations are part and parcel of the "intelligent but non-propositional mechanisms responsible for skill" (Fridland, 2014, p. 2736).

In section 2, I lay out the dialectic in more detail. In section 3, I discuss the kinds of representations generated by perceptual and sensorimotor learning. Section 4 shows how analysis of these representations can answer the three intellectualist strategies, and section 5 expands the discussion to argue that sensorimotor representations are at work in expert, deliberative performance. Section 6 concludes.

2. Setup

In contrast to traditional anti-intellectualist approaches that reduce skill to automated action or habit (Dreyfus, 2006; Dreyfus & Dreyfus, 2005), recent anti-intellectualists stress the flexible and adaptive nature of skilled performance (Christensen, Sutton, & McIlwain, 2015; 2016; Fridland, 2014, 2017; Sutton, McIlwain, Christensen, & Geeves, 2011). Skills, they argue, involve a form of control that is intelligent and goal-directed, and sensitive to task context, but still is not best understood in terms of propositional knowledge. Two particularly important properties that anti-intellectualists have focused on are the *flexibility* and *fine-grained* nature of skilled action control. Fridland (2015, p. 114) notes that tasks like riding a bike or performing surgery "can vary in an almost infinite variety of ways," but that skilled performance requires "responsiveness to the actual nuances of the very situation in which the skill is performed." Sutton et al. (2019) note that particular instances of, say, shooting a basketball, require fine-grained motor representation, but that any instance will involve a unique location on the court. Thus the "domain size" of controlled action is large.

Anti-intellectualists, however, have not done enough to theorize about the kinds of representations that contribute to non-propositional control. One reason for this is that, since modern anti-intellectualists want to resist the view that skilled behaviors are non-mental or purely automatic, they construe them as interacting with the agent's person-level intentional states (Fridland, 2017; Christensen et al., 2016). This seems to place skilled performers' actions in close relation to their propositional attitudes. As such, it becomes unclear why to prefer anti-intellectualism to a sophisticated intellectualism, or to a "hybrid" view that posits both propositional and non-propositional mechanisms of skill employment (Buskell, 2015; Levy, 2017). As Farkas (2018) notes, anti-intellectualists of the modern sort undertake a difficult task, needing to show that skill is "different from propositional knowledge … but … similar enough to be the right kind of cognitive achievement" (p. 106; cf. Springle, 2019).

Intellectualists have responded in one of three ways, although these are often not clearly distinguished in the literature. The first is to claim that propositional knowledge is the input to the skill learning process, such that learned skills are executions of a previously known content. This argument starts from the fact that skill often results from explicit instruction and practice, and moves to the conclusion that what is learned in the development of skill is how to carry out, motorically, the encoded rules or instructions. As Krakauer puts it, propositional knowledge can be "baked in" to the motor system, such that "propositional knowledge can … be transformed into goal-directed, automatized responses – intelligent reflexes" (2019, p. 825).

The second approach is to argue from semantics to the claim that propositional knowledge is *constitutive* of skill. On this view, know-how or skill just *is* a form of propositional knowledge – as Pavese (2018) puts it, it is "the state of knowing a proposition about how to perform that action under a practical mode of presentation" (p. 803; cf. Stanley 2011). For Stanley (2011), knowing how to perform a task is equivalent to knowing a proposition that, for some w, w is a way one could perform that task. Pavese (2015, 2017, 2018) further cashes out the notion of a practical representation. For Pavese, a practical representation represents a "method" – i.e., a way of performing the task – and this in turn is comprised of a set of motor commands, which translate, "bit by bit" (Pavese, 2017, p. 70), the agent's intentions into a set of elementary operations performable by the motor system. So, skilled knowledge of how to φ, in each instance, consists in the agent's knowledge that some m is a method to φ, where m is represented practically.

The third strategy is to focus on the *employment* of skill. On this kind of view, it is the propositional knowledge of a skilled performer that allows them to *exercise* their skills in certain ways. Montero and Evans (2011) focus on the ability of skilled chess players to give detailed, post-facto explanations of their deliberations to show that skilled performers can control their actions via explicit deliberation. Buskell (2015) draws on Montero to stress that expert practitioners often employ *reminders* to themselves about what to focus on, how to perform in what setting, etc. To those motivated by these kinds of data points, what makes experts experts is their ability to conceptualize their actions. Relatedly, Stanley and Williamson (2017) have recently argued that skill is the *disposition* to have the relevant kinds of propositional attitudes in the right contexts.

Notice that each strategy, in its own way, relies on a kind of *explanatory priority* claim. While there may be non-propositional processes that contribute to the learning of skill, what explains the ability of skilled performers (their know-how) is their having a certain kind of propositional knowledge – i.e., as coming to be able to represent a certain set of propositional contents. It is particularly important to understand this kind of priority on the constitutive strategy. On a view like Pavese's, practical modes of presentation needn't be propositionally structured (they may be, e.g., "imperatival pictures"). But it is the grasping of propositional content *through* the mode of presentation that is constitutive of know-how.

So, the intellectualist has a series of powerful strategies to reply to the anti-intellectualist. In what follows I will discuss in detail the kinds of mental representations that are developed in perceptual and motor learning. I will contend that these representations constitute a kind of associative structure that is indeterminate with regards to propositional contents, and that this indeterminacy is functionally necessary for the kinds of flexible-but-detailed control processes that anti-intellectualists cite. I will then argue that positing these kinds of representations answers all three intellectualist strategies.

3. Perceptual and Motor Learning.

3.1. Assumptions

I begin with a few assumptions. First, in addition to motor control, skill importantly involves perceptual representation, as well as interactions between perceptual and motor systems. It has become widely noted in the action control literature that the motor system is itself intentional – it can represent goals and outcomes (Butterfill & Sinigaglia, 2014; Ferretti & Caiani, 2018). Further, as we will see, interactions between perceptual and motor representation are particularly important in sensorimotor learning.

Second, I will assume that there is an important distinction between "discursive" representations – language like representations comprising amodal concepts – and *dimensional* representations that more closely correspond to perceptual and motor representation. The basic idea here is that perceptual and motor representations have a metric structure that is committal along the dimensions that define their referents. So, color representations are defined in terms of the dimensions of the color space, shape representations along dimensions of orientation, depth, and curvature, etc. (Burnston, 2017a, forthcoming). Motor representations are constituted by dimensions of kinematic action (Burnston, 2017b). For instance, making a grasp with the hand requires wrist flexion, force parameters on each of the muscles, and a launch angle of the elbow, etc. A non-discursive representation of a grasping movement defines a value along each of these parameters.

These assumptions should be agreeable to all parties, given that intellectualists admit that skill involves both motor control and perceptual discrimination, and that they agree that practical modes of presentation need not be discursively structured. In what follows, I describe how these kinds of representations are developed in perceptual and sensorimotor learning.

3.2. Perceptual Learning

Consider the stimuli below, from a highly cited study by Fiser and Aslin.

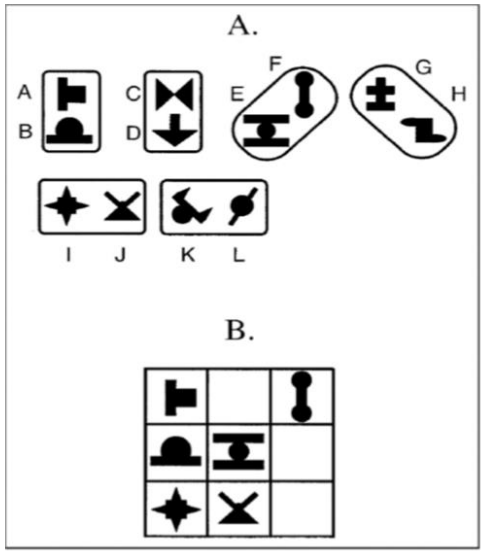

Figure 1. (Left) Perceptual learning stimuli from Fiser and Aslin (2001).

Subjects were told to passively observe a series of stimuli, an example of which is shown in panel B of the figure. Within the stimuli were embedded a series of spatially correlated "base pairs," examples of which are shown in panel A. The two elements of each base pair always occurred in the same configuration relative to each other, but the pair could appear at different locations in the overall stimulus. Importantly, subjects were given no instructions as to what to look for; they were simply told to attend to each successive stimulus.

In a subsequent phase of the experiment, subjects were shown pairs that were either familiar from the training period, or novel. They could successfully recognize base pairs from the familiarization phase at roughly 75% accuracy. So, despite having no instructions regarding what to look for, subjects were able to notice that certain pairs of particular figures regularly co-occurred in relative position to each other, even if their absolute position in the grid was consistently changing.

There is an interesting further complication, however. In the first experiment, the parts of novel pairs might have been seen in the familiarization phase, but had never been seen in the particular location that they were presented in during the test phase. This allowed for spatial location to be used as a cue, in addition to component identity. In a second experiment, this was modified so that individual components of novel pairs could be located in the same location they had occurred in during familiarization – the only difference being that the conjunction of the components in the novel pair had never been seen. Subjects were much worse at the second variety of the task, which suggested that they had encoded the association between components and particular locations, in addition to the spatially invariant correlations between multiple features. Despite this, they still recognized the base pairs from the familiarization phase at a higher-than-baseline rate.

This example exhibits some extremely important properties for the discussion to come. First, the mental representations developed during the task are *multi-scale*. Subjects learn not only what features are present during the familiarization phase, but also the associations between particular features and particular locations, as well as the associations between features comprising base pairs. Second, and relatedly, this multi-scale nature allows for the learned representation to capture patterns of *variance* and *invariance* in the series of stimuli. This is shown in the fact that, during the test condition, subjects are sensitive to both the correlations present in base pairs, regardless of where they occur, *and* to the learned correlations between particular features and particular locations. This is why their performance is hindered, but not back to baseline, by the competing associations.

The ability to represent patterns of variance and invariance also underlies *holistic* effects in perception, and these in turn can importantly influence perceptual *attention*. The following figure shows stimuli from a study by Zang et al. (2016), which is an example from a very wide range of results about *contextual cueing* and visual search.

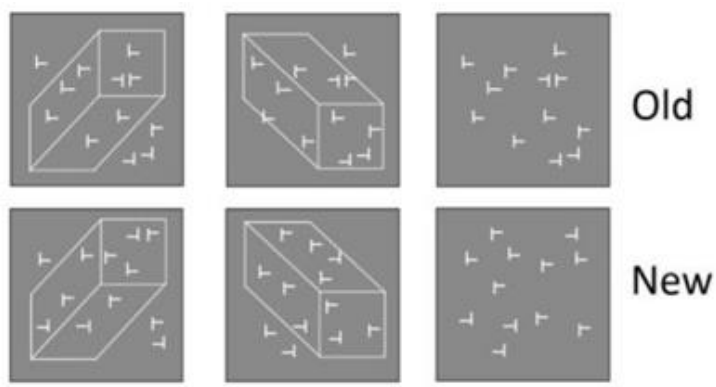

Figure 2. Cuboid stimuli from Zang et al. (2016).

In this kind of task, subjects must isolate one stimulus amongst distractors – in this case, a difficult-to-discern 'T' shape amongst 'L' shaped distractors. All that subjects are instructed to do in the experiments is to look for the 'T'. Again, the important contrast is between a learning condition and a subsequent test condition. Zang et al. showed that changing the *arrangement* of distractors from training to test conditions impedes performance (as shown in the "old" versus "new" panels). Even more striking, they showed that completely *task-irrelevant* aspects of the stimuli can structure visual attention. For instance, if, during the training phase, the stimuli are distributed around a task-irrelevant cuboid, shifting the orientation of the cuboid during the test phase hinders performance. Performance is not impacted, however, if during training, binocular cues are added to the cuboid so that it is appears to be at a different depth plane from the target and distractor stimuli.

The idea here, is that the scene is processed *holistically*. During learning, subjects notice not only the individual stimuli, or only the *task-relevant* stimuli, but pick up on the overall structure of the scene. This representation of the scene then constrains attentional search, so that subsequently, when part of that structure is varied, this variation – despite being task-irrelevant – hinders processing.[1] This sort of holistic processing applies to *categorical* perception as well. Consider the "Greebles" depicted below.

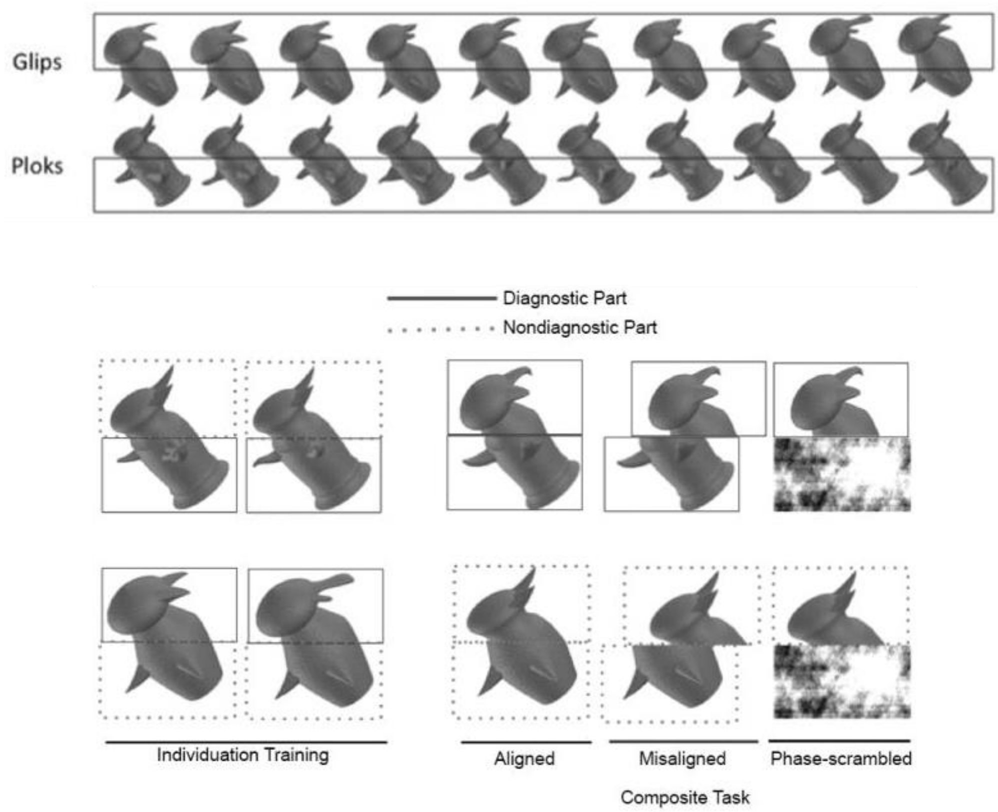

Figure 3. Greeble individuation and matching, from Chua et al. (2015).

Chua et al. (Chua, Richler, & Gauthier, 2015) had subjects learn to identify two categories of Greebles, Glips and Ploks, across a range of exemplars. There was more variation in the top half of Glips than the bottom, and vice-versa for Ploks, so subjects had to learn how to generalize across this variation in recognizing each kind.

[1] On one recent account (Wu, 2011), attention is important for *action selection* – it is what allows us to select the objects in our environments that we will act upon. The view I express here, on which learned perceptual representations structure attentional search, is compatible with this perspective.

In the test condition, subjects had to look at *composite* Greebles – Glip in the top half, and Plok in the bottom half. The task was quite different. They had to look at two composite Greebles and, focusing only on the top half (for Glips) or only on the bottom half (for Ploks), say whether the two were the same or different. Interestingly, what happened in the *task irrelevant* half affected subjects' performance. If the particular task-irrelevant half was equivalent to a half that *had been* part of the diagnostic learning set during the learning phase, it hindered performance. If it had not been seen before, it did not hinder performance. Moreover, if it was spatially displaced relative to the task-relevant half, this effect disappeared, generating no worse performance than a scrambled stimulus (right part of figure 3).

This result shows the properties we have discussed so far. Greebles, as a kind, are processed *holistically*, which is why replacement of a task-irrelevant feature can hinder Greeble perception, and why this effect is ameliorated when the configurational relationship between parts is broken. Moreover, the representation is *multi-scale*, since it applies holistically but also is mediated by particulars – in this case, whether the task-irrelevant half was itself previously relevant as part of a holistically-processed Greeble. Again, it is the multi-scale nature of these associations that explains the pattern of results.

3.3. Sensorimotor learning.

As noted above, most motor learning is in fact *sensorimotor* learning. It is learning that develops and fine-tunes interactions between perception of the environment and the motor system. Consider the figure below, which is from a study on motor learning by Gallivan et al. (Gallivan, Stewart, Baugh, Wolpert, & Flanagan, 2017).

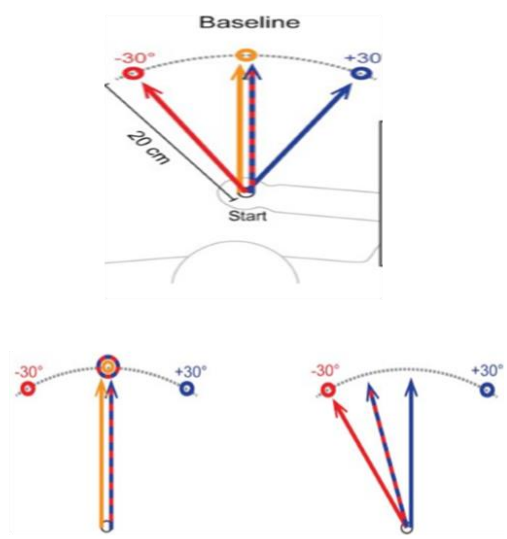

Figure 4. Motor learning study from Gallivan et al. (2017).

In the study, subjects performed a "go-before-you-know" task, in which they had to begin reaching towards a number of simultaneously presented objects *prior* to a cue which instructs them which one is the target. It is well established that in these tasks, independently of any instruction of *how* to perform the task, subjects' initial launch angle is directed towards the spatial average of the targets, as shown in the top panel of the figure. Generally, the ability to compute the average is attributed to the motor system, because it is this system which computes particular launch angles to particular targets (Gallivan, Chapman, Wolpert, & Flanagan, 2018).

Gallivan et al.'s ingenious tweak to this paradigm was to have subjects consistently reach towards a pair of targets but, trial-by-trial, subtly shift the location of the rightmost target so that it was slightly closer to the leftmost one. On each trial, the change in location was so slight that it was below the threshold for conscious discrimination. But, over the course of the trials, the initial launch angle persistently changed to reflect the new spatial average. (The bottom part of the figure shows this contrast.) That is, on the assumption that the motor system is the one that computes the trajectories in general, the most likely hypothesis is that it is this same system that keeps track of the new spatial arrangements of the scene and updates the initial launch accordingly.

As another example, consider Memelink and Hommel's (2005) exploration of the "Simon effect." In the Simon effect, spatial congruency of stimulus-response associations affects behavior – if subjects, for instance, have to press a button on their left when a picture of a horse is presented, they will be quicker to do so if the picture is presented *on their left*.

The point of Memelink and Hommel's task was to show how the Simon effect generalized across tasks, and how it interacted with explicit instruction. They had two tasks, a first task which was designed to elicit a Simon effect, and a secondary transfer task to see if that learned effect carried over. Subjects were instructed to press differently located buttons on a keyboard. So, in one condition, subjects had to press a button higher on the keyboard when they saw a picture of an animal facing up, and a lower button when they saw the picture facing down. This kind of task, which focuses attention spatially, elicited a Simon effect.

In the second, transfer task, subjects simply had to press a higher key when they saw an 'X', and a lower key when they saw an 'O'. If the X was presented higher on the screen, subjects were quicker to press the upward key, and similarly if the 'O' was presented lower on the screen, eliciting a baseline Simon effect. Interestingly, however, this effect was *enhanced* in subjects who had performed the vertical Simon-effect-inducing primary task. This shows that learning in one task transfers over to a secondary task with similar action parameters.

Importantly, these effects outstrip the explicit instructions given to subjects. Memelink and Hommel showed this via a secondary manipulation, which varied the instructions amongst groups of subjects. While in some cases, subjects were instructed in the primary tasks by location (e.g., "press the higher up button"), in some cases they were instructed by a non-spatial feature (e.g., "press the blue button," where the blue button would happen to be higher on the keyboard). The variation in instructions had *no effect* on what occurred in the transfer task. This shows that sensorimotor learning effects occur even with no relevant instructions encoding those aspects of the task.

As a last example, consider an attentional study by Myung, Blumstein, and Sedivy (2006). Myung et al. showed subjects a concurrently presented set of pictures as they listened to spoken words, and tracked their eye movements. Suppose one of the pictures was a piano, and during viewing subjects were presented with the word 'piano'. Unsurprisingly, subjects made saccades immediately to the image of a piano, rather than other objects. More surprisingly, however, when the subjects were presented with 'piano', but *none* of the objects were pianos, subjects performed saccades to a picture of a typewriter. The explanation for this effect is that, while not exact, typewriters share a general set of perceptual attributes and motor affordances with pianos.[2]

These results exhibit similar properties of multi-scale association and holism to those discussed in the previous subsection. In the Gallivan et al. study, subjects perceive, at a subthreshold level, the changes in the target locations, and these minute changes are tracked into updates in reach trajectories in the motor system. This suggests that fine-grained detail is represented in sensorimotor learning. But the Memelink and Hommel results suggest that such learning also occurs at the scale of a *task setting*. So, when subjects are in the same task setting – sitting in front of a keyboard, looking at stimuli on the monitor – in the transfer task as they were in the primary task, the Simon effect exhibited in the primary task carried over, despite the changes in the lower-level stimuli that cued the task. These two effects thus suggest the multi-scale nature of the associations. Holism is shown in the Myung et al. study. While there are differences between typewriters and keyboards, obviously, the fact that they share similar perceptual and motor parameters is taken to underlie the cross-categorical priming effect.

So, in sum, the results shown in perceptual and motor learning exhibit the key properties of multi-scale association and holism. Let us call the results of learning that exhibit these properties "structured sensorimotor representations" (SSRs). In section 4, I will suggest

[2] Levy (2017) considers the case of an expert pianist who sits down at an entirely unfamiliar piano. In this situation, Levy suggests, the pianist must use their propositional knowledge of pianos to inform them what to do in the novel situation. If it is holistic object effects, however, that explain the eye tracking results, and if these are not best read as propositional knowledge (as I will argue in section 4), then the Myung et al. results can be read as an objection to Levy's view that propositional knowledge is required in this case.

that the properties of SSRs are hard for the intellectualist to capture in propositional terms. First, however, I will propose that these kinds of representations are employed in skilled action, and underlie the anti-intellectualist's combined posits of fine-grained and flexible control.

3.4. Employment in skilled action.

Consider what a skilled athlete has to do. They have to recognize the *context* they are in, to an arbitrary degree of detail, determined by their current goals. An ice hockey player has to "know" whether they are in the offensive, defensive, or neutral zone, but within those contexts their precise positioning, as well as the action opportunities afforded to them, will depend on the precise locations of the puck and other players. As they attempt a shot on goal, their movements will need to vary depending on their own position on the ice, as well as that of the goalie and the defenders (cf. Buskell, 2015). As players progress to higher levels, they constantly have to re-learn these affordances as the other players get bigger and more skilled, and the game moves faster.

I suggest that SSRs underlie this type of rich understanding. The ability to recognize one's surroundings and their action affordances, and to flexibly organize one's action in those surroundings to enact one's current goals are part of the basic functioning of sensorimotor control. As Memelink and Hommel note, learning the structure of a task context allows for subsequent actions to be planned and performed quickly and flexibly *within* that structure.[3] They suggest that what subjects learn in their study is the context-bound associations that afford actions in that task environment. These associations contribute to specific actions, but also afford flexibility as to the precise movements that are enacted.

Indeed, studies have shown that, when subjects are performing manual tasks, they can respond flexibly and easily to perturbations of their arms while they reach towards targets (Nashed, Crevecoeur, & Scott, 2014). Moreover, once subjects have experience manipulating a certain kind of object, they react immediately to external forces added to the object, in a way that maintains their goal (Diamond, Nashed, Johansson, Wolpert, & Flanagan, 2015). Finally, Gallivan et al. (Gallivan, Barton, Chapman, Wolpert, & Flanagan, 2015) have shown that the motor system is sensitive, at very fine-grained levels, to *sequences* of goals. If subjects are presented with sequential action targets, where one can be gripped multiple ways and the other requires a specific grip, they will modify their grip of the ambiguous object to match the one required for the other object. These effects suggest that,

[3] Of course, there is a downside to this, given that once a representational structure is learned, it can be hard to operate outside of it. Brownstein and Michaelson (2016) give an amusing example of this phenomenon, in which Major League Baseball all-stars persistently failed to hit pitches from USA softball ace Jennie Finch. What this suggests is a different sensorimotor structure involved in hitting softball pitches, as opposed to baseball ones, despite the very general properties that the two contexts share.

once the properties of objects and task context are learned, actions can be flexibly adapted within those contexts.

Do these principles scale up directly to expert performance? There is certainly piecemeal evidence from many areas that suggests it. Of course, athletes often speak of visualizing outcomes. Chess masters are known to be able to rapidly recognize board configurations (Newen, 2016), and report visualizing the progression of games as the movement of "lines of force" (Chabris & Hearst, 2003). Skillful users of the "mental abacus," who imagine manipulating an abacus in their heads, can win speeded mental calculation competitions, but their acuity is still limited by general structures of visual working memory. Moreover, they perform *gestures* while they calculate, and interrupting those gestures hinders their performance (Frank & Barner, 2012).

Let's discuss one area, athletics, in detail. Psychologists of sport have investigated whether skilled athletes have enhanced abilities in general cognitive faculties, with very mixed results (Voss, Kramer, Basak, Prakash, & Roberts, 2010). However, a range of studies have shown clearly that expert athletes have advanced capacities for learning properties of dynamic visuo-spatial scenes.

For instance, Faubert (2013) compared expert athletes across multiple sports to high-level amateurs (NCAA-level) and non-athletes in their ability to learn a speeded 3-D object tracking task. Subjects viewed a series of spheres which moved in three dimensions, and had to keep track of four of them and report their locations at the end of the trial. If they got them right, the next trial was speeded up, and this continued until they reached a threshold. Expert athletes both performed better within trials, and improved faster across training sessions, compared to amateurs, who in turn were much better than non-athletes. Similarly, Chaddock et al. (Chaddock, Neider, Voss, Gaspar, & Kramer, 2011) had subjects cross streets in a virtual reality environment by walking on a treadmill. Athletes could successfully navigate more complex and quickly moving environments than non-athletes.

The experimenters attribute these effects to greater visual processing speed, increased biological motion recognition, increased abilities for holistic attention, and greater ability to use binocular disparity cues to process depth in a scene. I suggest that the general way to understand the results is to say that expert athletes have highly developed abilities to learn and act within sensorimotor contexts – that is, they are skilled at developing SSRs. Now, skill at a particular sport or task is clearly more specific than the general properties of SSRs evinced in the studies discussed so far. But, there seems to be enough evidence to offer the view that skill learning involves the refinement and sensitization of domain-general learning abilities in specific contexts. I will return to this point in section 5. In the next section, I will assume the perspective based on SSRs to be roughly correct, and argue that, if so, it resists all three intellectualist strategies.

4. SSRs and Intellectualist Strategies.

Let's assume it is the case that SSRs in fact underlie the flexibility and fine-grained control cited by anti-intellectualists. The question is whether the explanatory priority posited by intellectualists can capture that function. Here, I argue that they cannot, because the propositional content of SSRs is *indeterminate*. Not only is it indeterminate, but it is functionally vital that it is so, because that is what allows both fine-grained and open-ended control.

As noted above, SSRs are multiscale and holistic. What this does is allows for the representations to capture both variance and invariance. As we saw in the Fiser and Aslin and Chua et al. studies, learning effects are driven both by particulars, and by correlational and configural relationships that generalize over particulars. This combination is important. The specific particulars that subjects have encountered are fine-grained. But holistic and correlational effects are inherently open-ended; they generalize over space (Fiser and Aslin) and feature particularities (Chua et al.; Memelink and Hommel). Overall spatial layout of scenes is used to flexibly construct particular movements (Gallivan et al.) and to guide attentional search (Zang et al.). The picture that emerges from these properties is that SSRs are *multilevel constraints* on actions, without being determinate representations of them. I now use these properties to reply to each anti-intellectualist strategy.

The first strategy is the input strategy. On this view, subjects receive explicit instruction, and the development of skill involves training perceptual and motor systems to implement those instructions. What is clear in the foregoing examples, however, is that sensorimotor learning considerably *outstrips* the instructions that subjects receive. The Memelink and Hommel study shows that development of associations based on Simon effects occurs despite variation in the stimulus properties described in the instructions. Both the Zang et al. and Chua and Gauthier studies show that that holistic effects occur even when they are task-irrelevant. And the Fiser and Aslin study shows that learning picks up on stimulus configurations when there are no instructions at all, simply through passive viewing. Task configurations, the Gallivan et al. study showed, can affect learning even if subjects are entirely unaware of them. If this is the case, then what subjects learn is not simply a "baking in" of their instructions. And that is a good thing, because the actual situations that skilled performers encounter outstrip, in their variation and detail, the explicit instructions that they receive during training.

The constitutional strategy does not rely on the implausible "baking in" view of explicit knowledge, and admits that there are multiple kinds of representations at work in the exercise of skill. These different forms of representation are taken to be distinct modes of presentation for propositional contents. On this kind of view, what sensorimotor learning does is develop new modes of presentation for propositional contents. The constitutional view is the strongest strategy because it purports to have the resources to capture the fine-grained contents involved.

To assess the constitutional view, it is important to reiterate that intellectualism is a claim about *explanatory priority*. It is knowledge of propositional contents that is supposed to explain the subject's ability to perform skill. On the constitutive view, practical modes of presentation are particular *ways* of representing those contents. While sensorimotor representations might be necessary or important, it is the propositional content that explains how skilled practitioners do what they do. Propositional contents play a role in the psychology of skilled individuals.

To construct the most charitable interpretation of the constitutional view possible, notice how it might incorporate the *multimodal* nature of the learning discussed here. As I argued above, sensorimotor learning spans both perceptual and motor representation. But this can be accounted for by arguing that perception is picture like, and providing a propositional semantics for pictures. A picture's contents, for instance, might just be a long conjunction of propositions (Grzankowski, 2015). Similarly, complex motor representations might be represented as a long list of motor commands. Pavese suggests that practical representations represent propositional content regarding methods, and methods comprise combinations of the elementary operations that the motor system can perform (Pavese, 2017). Note that this allows for the propositional content of a representation to be, potentially, very rich in detail, and to engage with the motor system.

While the constitutional strategy can invoke richer or sparser propositional contents, it is commitment of the view that there must *be some fact of the matter* about the contents. That is, there must be some fact about the contents expressed, which explains the how the representation guides action. To put it another way, the constitutional view posits a *determinate*, if potentially ornate, propositional content, represented via perceptual and/or motor modes of presentation, as explaining skill. For instance, in Pavese's picture, there must be a fact of the matter about what method is driving the motor system, and therefore what the content is of the practical sense that refers to the method. And this content must have the kind of cognitive import that Pavese cites – i.e., it must be what explains skilled performance.

So, if there is no fact of the matter about the determinate propositional contents expressed by SSRs, then the constitutional strategy fails. I contend that SSRs indeed lack such determinate contents. A framework for pursuing this claim is expressed well in Kulvicki's description of "analog" representations. Kulvicki characterizes analog representations as representations that "support a pattern of interaction, specifically open-ended search for content across levels of abstraction" (2015, p. 165), and argues that they do so in virtue of sharing a structure with what they represent – in our case, task contexts. It virtue of this open-endedness, there is no fact of the matter about exactly what contents are represented by the representation – i.e., analog representations are indeterminate. I contend that SSRs exhibit the kind of open-endedness that Kulvicki cites, and that this property undermines the determinateness commitment of the constitutional view.

Kulvicki's primary example is a thermometer. He suggests that a thermometer supports, *as part of its content*, an indeterminate number of abstractions. So, while one *might* take the thermometer to represent the most specific temperature possible (say, 40 degrees), one might also take it to represent any other number of facts – for instance that it is warmer than 10 degrees. Moreover, one can pick any range within the thermometer, and read that it is warmer (/cooler) than *these* temperatures, etc. The fact that the thermometer reading doesn't explicitly privilege one set of abstractions over another is another way to say that it lacks a "canonical decomposition" (cf. Fodor, 2006). Lacking such a decomposition is a property that is often taken as characteristic of iconic representations. Descriptive representations (e.g., "the freezing point of water"), by contrast, do not support such open-ended search, because they fix the level of abstractness at which they are to be analyzed.

In particular, it is the combination of variance and invariance represented in SSRs, as well as their multi-scale nature, that leads to their ability to contribute to both fine-grained and flexible control. Consider the thermometer again. While a given instance of a thermometer reading is fixed at a given point, it is part of a system of representation that supports a range of abstractions, and hence the current reading can be re-grouped into different contrast classes as the need arises (e.g., "what is the precise temperature?" versus "should I bring a sweater?"). So, while the particulars of the representation matter, they don't exhaust its content, which supports an indeterminate search through abstractions.

Multi-scale and holistic representations do this as well. Recall from the Fiser and Aslin study that it is both particular feature instantiations and configural relationships which can be located flexibly around the scene that affect psychological processing. This combination of properties is, I suggest, strictly analogous to the kinds of properties Kulvicki attributes to the thermometer, and we can see this in the other studies as well. While Greebles have a holistic shape that is invariant across feature transformations, those feature particulars still affect attentional patterns. Similarly, the fact that the learned spatial association in Memelink and Hommel's Simon task generalizes across changes in stimuli, and *intersects* with Simon effects present in the second task, shows that they contribute across a range of abstraction. That is, there is not one set kind of stimuli or specific range of situations in which they must be operative. Instead, general configural patterns intersect with particulars at multiple scales to determine performance.

Consider our hockey player again. No two entries into the offensive zone will be quite the same, but positional awareness is a generalizable skill that is instantiated in those specific contexts. Or consider Fridland's bike rider, who must adapt to novel inclines and surfaces, no two of which will be exactly alike. There are an indeterminate number of detailed instantiations, none of which are explicitly named by, but all of which are compatible with, the general representations constituting the skill – i.e., those representations can range across instances. Indeed, I will argue in the next section that this kind of indeterminate structure supports creative and collaborative engagement in skilled practitioners.

This open-ended structure suggests that there is simply no fact of the matter about *the* propositional content expressed by SSRs. So, if intellectualism is committed to determinate contents, then it cannot capture skill if skill is in fact underlain by SSRs. This is the case both for Stanley's general intellectualist view and the more detailed version supplied by Pavese. Even on Pavese's account, there must be a fact of the matter about which method – i.e., what exact combination of (sensory-) motor elements – is driving behavior. If the SSR-based perspective is right, then no such determinate content postulate explains the exercise of skill.

I have argued not only that SSRs exhibit content-indeterminacy, but that this determinacy is *key* to the combined flexibility and fine-grained nature of skilled control. This allows for a response against a common intellectualist strategy. A common tactic for the intellectualist when talking about motor control is to fix propositional content at a certain level of abstraction, and suggest that anything beneath that level is non-representational. Pavese (2018) suggests that methods are represented at the level of "motor schemas," whereas anything more detailed is simply a (perhaps learned) elemental operation of the system. Stanley and Krakauer (2013) suggest that any learning occurring that cannot be captured propositionally is simply an increase in *acuity*, rather than a change of content.

But, if content-indeterminacy in Kulvicki's sense is what is important, then the strategy of fixing propositional content at a specific level of abstraction fails. If you fix the content at one level of abstraction, it will fail to capture more abstract elements of the action context, or more specific on-line adjustments of the motor system (cf. Burnston, 2017b). That is, it will fail to capture the multi-scale aspects of SSRs. Again, this kind indeterminacy is *vital* to capturing the variant and invariant aspects of a task, and it is by no means obvious how it can be captured as a propositional content.

So, here we have an argument that the constitutional strategy fails. And if that is the case, then the employment strategy falls in short order. The employment strategy says that it is skilled agents' propositional knowledge that allows them to *utilize* their skills intentionally. But, given the arguments here, this precisely puts the explanatory priority the wrong way around. If I am right about the structure of SSRs, and it is true that they underlie skill, then it is not propositional knowledge that explains how skill is activated in a given instance. Rather, it is non-propositional knowledge that shows how agents can come up with fine-grained strategies and flexible responses in specific contexts. Recall also Stanley and Williamson's dispositional view, on which skilled performers are disposed to generate propositional knowledge that is distinct from novices. It may well be true that skilled performers have such dispositions, but if the *categorical bases* of those dispositions are SSRs, and if their content is non-propositional, then it is non-propositional content that explains skill. So, the explanatory priority of propositional knowledge posited by the employment strategy fails to hold. I explore this conclusion further in the next section.

5. Instruction, Deliberation, and Structured Knowledge

I have suggested that skilled knowledge is not constituted by propositional contents. This opens up, perhaps, other ways of thinking about instruction and employment. Recall that in section 3.4, I noted that there is a gap between domain-general sensorimotor learning and the learning of specific skills. In this section, I suggest, first, that refined skill involves learning more specific action contexts, and that this process can be mediated, if not fully determined by, instruction. Indeed, I argue that the multi-scale and holistic nature of SSRs account for some obvious facts about instruction (5.1), and can shed light on some fascinating practices of experts performers (5.2).

5.1. Multimodal instruction.

De Vega et al. (2004) give an amusing example that illustrates the lack of resources that even relatively detailed discursive representations offer for motor representation. Consider the sentence "She scratched her back using the floppy disk." They suggest that a sentence like this should be easier to understand than a sentence like "She scratched her back with the thread." Why would this be, despite their overall similarity?

De Vega et al. suggest that discursive representations index *objects* and *roles*. For instance, the sentence determines that the floppy disk should be scratching the back, not the other way around. But this function is different from determining exactly the perceptual and motor consequences of the sentence. Rather, De Vega et al. suggest that the indexed objects are subsequently represented in terms of their *motor affordances*, and that these are what subjects use to determine the details of the action. Given the kind of thing a floppy disk is, and the target of the action, one can "figure out" through motor imagery that the corner is probably the best thing to use to scratch the back. Since a thread gives no such affordances, "She scratched her back with the thread" is harder to understand. They subsequently give empirical evidence for this overall picture, which I won't discuss here. The point is that imagery of acting with motor affordances might play a significant role in action understanding that is not specified by discursive representation.

I suggest that this kind of perspective accounts for something that should be obvious: the *multimodal* nature of instruction. When subjects are actually learning a new skill, their instructional materials simply don't consist in fine-grained descriptions of relevant movements. Rather, instructions work in tandem with external diagrams, emulation, and exemplar-based training to implement skill learning. The best interpretation of this is that instructions are guideposts in a much more complicated process of thinking and learning that is primarily non-propositional. Instructions can, however, direct these processes in specific ways to aid in the learning of specific skills.

In fact, some preliminary work on exactly this kind of dynamic has been done. Kirsh and colleagues (2010, 2011) performed a pilot study on how subjects interpreted instructions for how to create origami sculptures. They both varied the type of instructional materials and filmed the subjects to see how they acted while interpreting the instructions. The results need to be taken with a grain of salt, as the preliminary nature of the study, and concomitant small sample size, restrict us to interpreting statistical trends. However, a few interesting phenomena emerged. First, subjects actually performed better (quicker time to completion; fewer errors) with a combination of diagrams and very short captions than they did with diagrams and longer, more detailed captions. Both were more effective than plain text. This suggests that the discursive representations, like the sentences in De Vega et al.'s example, get subjects' attention focused on the right aspects of the diagrams, so that they can *then* reason through the example in a sensorimotor format.

This is, it turns out, supported heavily by a variety of gestures that subjects performed while doing the task. Subjects would manipulate the actual paper so that it aligned with the visual diagram, both prior to and after making a particular fold. This was done both to check that the starting point and end points of a particular move aligned with the structure depicted in the diagram. Moreover, once the subjects had aligned their paper with the starting point for a move depicted on the diagram, they would often take their other hand and use it to mimic the actual move they were considering, so that they could envision the effect it would have. This, I think, is a clear indicator that what the discursive symbols do is encourage subjects to focus on particular parts of a structure that they can represent in sensorimotor format. That is, instructions do not simply encode a content that the sensorimotor system learns to follow.

If you pick up *any* set of instructional materials for *any* kind of skilled action – say, golf, tennis, or softball – what you will tend to find is a series of descriptions accompanied by a wide range of visual tools, including simplified diagrams and photographs with aids to focus attention on particular aspects. These will often include both visual aids for specific motor actions, as well as diagrams of playing surfaces (golf hole, softball field, etc.) in which to situate the particular actions in the broader strategic aims of the act. I suggest that what this multimodal instruction does is to begin to develop – or to structure practice so that learners can develop – the kind of open-ended structure that characterizes SSRs. In the next subsection, I suggest that, once learned, these multiscale representations are employed by expert performers in the exercise and creative deployment of their skill.

5.2. Experts and Employment

Kirsh (2013) has pursued a series of underappreciated ethnographic studies on renowned choreographers, as well as on the individual and collective practice habits of an expert dance troupe. These studies show (i) the continued use of spatial metaphor for creative action and increased understanding, and (ii) the importance of spatial and visual imagery for both individual and collective practice. Consider the following figure:

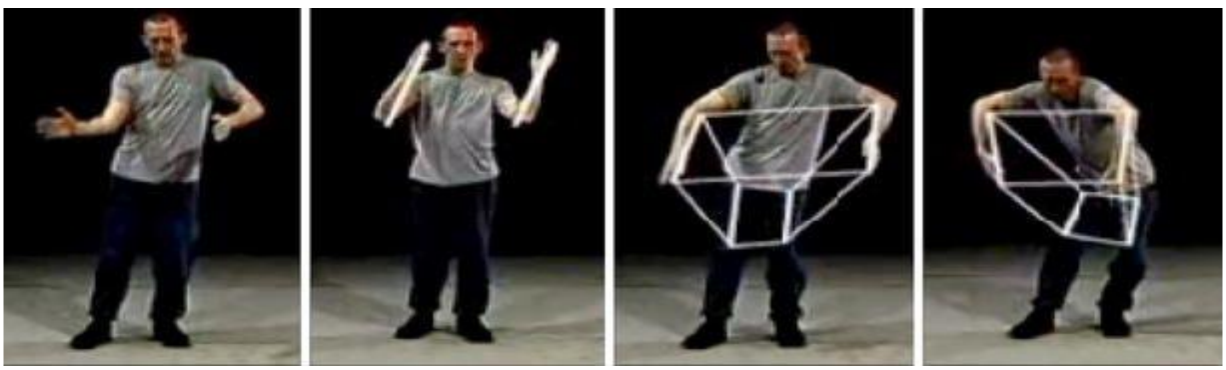

Figure 5. Spatial metaphor in dance instruction, from Kirsh (2010).

The figure depicts a series of instructional practices pioneered by choreographer Bill Forsythe. In the instructional materials Forsythe uses geometrical patterns superimposed on his own body to help convey the kind of movements he wishes his dancers – who are already experts – to perform. Dancers are invited to picture the movements as a series of manipulations performed on these patterns. Certain movements might maintain the volume of the figure, or keep certain edges parallel while modifying others. Intriguingly, Kirsh suggests that these practices *facilitate discourse* about abstract aesthetic ideas like stress and torsion. Once the publicly shared structure is available, and the dancers understand how to manipulate it, they can *then* discuss it with each other, ask questions about the reasoning behind particular moves, etc. This suggests that SSRs are being employed flexibly for deliberation and communication.

As a last example, consider the practice of "marking," exhibited in Figure 6.

Figure 6. Marking, from Kirsh (2013).

When marking, subjects use their body to mimic partial aspects of dances. Marking comes in several forms: it can be done with the whole body, but can also be done with the hands in a mimic of the full-body; it can be done privately and idiosyncratically, but can also be done in conventionalized forms for communication and demonstration. Marking also plays a variety of roles both for an individual and a troupe. For an individual, it serves as an adumbrated form of practice. Indeed, marking has some advantages *over* full-blooded movement for certain aspects of practice – including memory, technique, and timing – although not for the dynamics of the movement. Kirsh says that marking serves, for the individual, as a kind of sensorimotor mnemonic, as a "support structure for imagining the real thing" (2013, p. 13), and notes that one can see less standardized versions of this practice in individuals performing adumbrated baseball swings, tennis forehands, etc. According to the framework being employed here, marking is an anchor for internal simulation of the sensorimotor representations that constitute knowledge of the dance.

Intriguingly, a troupe as a whole often uses marking to work on coordinated movements, to make adjustments to individual movements, and even to situate the dance on a new stage. This suggests that the dancer in the troupe has a multi-level structural representation of the dance, including their individual movements, their movements situated in those of the other dancers in the troupe, and the troupe situated in a physical environment. This is, I suggest, the same kind of multi-level structure exhibited in SSRs. Now, of course, the dancers also *discuss* what they are doing – as well as emulate, gesture, make ostensive reference, etc. However, these emulations and ostensive practices are working *within* a representational structure shared by the troupe. Without that structure there would be no shared representation for the communicative practices to operate within. It is SSRs that enable these practices, not conjunctions of propositions known by the individuals. Or so the anti-intellectualist should say.

6. Conclusion

Intellectualists argue that skilled know-how is propositional. Anti-intellectualists have recently responded by citing the flexible and fine-grained control that skilled practitioners have. In this paper, I have attempted to put the anti-intellectualist point on firmer psychological footing by discussing the kinds of sensorimotor representations developed in learning a skill.

REFERENCES

Brownstein, M., & Michaelson, E. (2016). Doing without believing: Intellectualism, knowledge-how, and belief-attribution. *Synthese, 193*(9), 2815-2836.

Burnston, D. C. (2017a). Cognitive penetration and the cognition–perception interface. *Synthese, 194*(9), 3645-3668. doi:10.1007/s11229-016-1116-y

Burnston, D. C. (2017b). Interface problems in the explanation of action. *Philosophical Explorations, 20*(2), 242-258.

Burnston, D. C. (2020). Fodor on imagistic mental representations. *Rivista internazionale di Filosofia e Psicologia, 11*(1), 71-94

Buskell, A. (2015). How to be skilful: opportunistic robustness and normative sensitivity. *Synthese, 192*(5), 1445-1466. doi:10.1007/s11229-014-0634-8

Butterfill, S. A., & Sinigaglia, C. (2014). Intention and motor representation in purposive action. *Philosophy and Phenomenological Research, 88*(1), 119-145.

Chabris, C. F., & Hearst, E. S. (2003). Visualization, pattern recognition, and forward search: Effects of playing speed and sight of the position on grandmaster chess errors. *Cognitive Science, 27*(4), 637-648.

Chaddock, L., Neider, M. B., Voss, M. W., Gaspar, J. G., & Kramer, A. F. (2011). Do athletes excel at everyday tasks? *Med Sci Sports Exerc, 43*(10), 1920.

Christensen, W., Sutton, J., & McIlwain, D. (2015). Putting pressure on theories of choking: Towards an expanded perspective on breakdown in skilled performance. *Phenomenology and the Cognitive Sciences, 14*(2), 253-293.

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.

Chua, K.-W., Richler, J. J., & Gauthier, I. (2015). Holistic processing from learned attention to parts. *Journal of Experimental Psychology: General, 144*(4), 723.

De Vega, M., Robertson, D. A., Glenberg, A. M., Kaschak, M. P., & Rinck, M. (2004). On doing two things at once: Temporal constraints on actions in language comprehension. *Memory & cognition, 32*(7), 1033-1043.

Diamond, J. S., Nashed, J. Y., Johansson, R. S., Wolpert, D. M., & Flanagan, J. R. (2015). Rapid visuomotor corrective responses during transport of hand-held objects incorporate novel object dynamics. *Journal of Neuroscience, 35*(29), 10572-10580.

Dreyfus, H. L. (2006). Overcoming the myth of the mental. *Topoi, 25*(1-2), 43-49.

Dreyfus, H. L., & Dreyfus, S. E. (2005). Peripheral vision: Expertise in real world contexts. *Organization studies, 26*(5), 779-792.

Farkas, K. (2018). Know-how and non-propositional intentionality. *Non-Propositional Intentionality*, 95.

Faubert, J. (2013). Professional athletes have extraordinary skills for rapidly learning complex and neutral dynamic visual scenes. *Scientific reports, 3*, 1154.

Ferretti, G., & Caiani, S. Z. (2018). Solving the interface problem without translation: The same format thesis. *Pacific Philosophical Quarterly*.

Fiser, J., & Aslin, R. N. (2001). Unsupervised statistical learning of higher-order spatial structures from visual scenes. *Psychological science, 12*(6), 499-504.

Fodor, J. A. (2008). *LOT 2: The Language of Thought Revisited*. Oxford: Oxford University Press.

Frank, M. C., & Barner, D. (2012). Representing exact number visually using mental abacus. *Journal of Experimental Psychology: General, 141*(1), 134.

Fridland, E. (2014). They've lost control: Reflections on skill. *Synthese, 191*(12), 2729-2750.

Fridland, E. R. (2015). Skill, Nonpropositional Thought, and the Cognitive Penetrability of Perception. *Journal for General Philosophy of Science, 46*(1), 105-120.

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

Gallivan, J. P., Barton, K. S., Chapman, C. S., Wolpert, D. M., & Flanagan, J. R. (2015). Action plan co-optimization reveals the parallel encoding of competing reach movements. *Nature Communications, 6*, 7428.

Gallivan, J. P., Chapman, C. S., Wolpert, D. M., & Flanagan, J. R. (2018). Decision-making in sensorimotor control. *Nat Rev Neurosci, 19*(9), 519-534. doi:10.1038/s41583-018-0045-9

Gallivan, J. P., Stewart, B. M., Baugh, L. A., Wolpert, D. M., & Flanagan, J. R. (2017). Rapid Automatic Motor Encoding of Competing Reach Options. *Cell reports, 18*(7), 1619-1626.

Grzankowski, A. (2015). Pictures have propositional content. *Review of Philosophy and Psychology, 6*(1), 151-163.

Kirsh, D. (2010). Thinking with external representations. *Ai & Society, 25*(4), 441-454. doi:10.1007/s00146-010-0272-8

Kirsh, D. (2011). Situating Instructions.

Kirsh, D. (2013). Embodied cognition and the magical future of interaction design. *ACM Transactions on Computer-Human Interaction (TOCHI), 20*(1), 3.

Kulvicki, J. (2015). Analog representation and the parts principle. *Review of Philosophy and Psychology, 6*(1), 165-180.

Levy, N. (2017). Embodied savoir-faire: knowledge-how requires motor representations. *Synthese, 194*(2), 511-530.

Memelink, J., & Hommel, B. (2005). Attention, instruction, and response representation. *European Journal of Cognitive Psychology, 17*(5), 674-685.

Montero, B., & Evans, C. (2011). Intuitions without concepts lose the game: mindedness in the art of chess. *Phenomenology and the Cognitive Sciences, 10*(2), 175-194.

Mylopoulos, M., & Pacherie, E. (2017). Intentions and motor representations: The interface challenge. *Review of Philosophy and Psychology, 8*(2), 317-336.
Myung, J.-y., Blumstein, S. E., & Sedivy, J. C. (2006). Playing on the typewriter, typing on the piano: manipulation knowledge of objects. *Cognition, 98*(3), 223-243.
Nashed, J. Y., Crevecoeur, F., & Scott, S. H. (2014). Rapid online selection between multiple motor plans. *Journal of Neuroscience, 34*(5), 1769-1780.
Newen, A. (2016). Defending the liberal-content view of perceptual experience: direct social perception of emotions and person impressions. *Synthese, 194*(3), 761-785.
Pavese, C. (2015). Practical senses.
Pavese, C. (2017). A theory of practical meaning. *Philosophical Topics, 45*(2), 65-96.
Pavese, C. (2018). On the psychological reality of practical representation. *Philosophical Psychology*.
Springle, A. (2019). Methods, minds, memory, and kinds. *Philosophical Psychology, 32*(5), 635-661.
Stanley, J. (2011). Knowing (how). *Noûs, 45*(2), 207-238.
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.
Sutton, J., McIlwain, D., Christensen, W., & Geeves, A. (2011). Applying intelligence to the reflexes: Embodied skills and habits between Dreyfus and Descartes. *Journal of the British Society for Phenomenology, 42*(1), 78-103.
Christensen, W., Sutton, J., & Bicknell, K. (2019). Memory systems and the control of skilled action. *Philosophical Psychology, 32*(5), 692-718.
Voss, M. W., Kramer, A. F., Basak, C., Prakash, R. S., & Roberts, B. (2010). Are expert athletes 'expert' in the cognitive laboratory? A meta-analytic review of cognition and sport expertise. *Applied Cognitive Psychology, 24*(6), 812-826. doi:10.1002/acp.1588
Wu, W. (2011). Attention as selection for action. *Attention: philosophical and psychological essays*, 97-116.
Wu, W. (2016). Experts and Deviants: The Story of Agentive Control. *Philosophy and Phenomenological Research, 93*(1), 101-126. doi:doi:10.1111/phpr.12170
Zang, X., Geyer, T., Assumpcao, L., Muller, H. J., & Shi, Z. (2016). From Foreground to Background: How Task-Neutral Context Influences Contextual Cueing of Visual Search. *Front Psychol, 7*, 852.