Experts and Deviants: The Story of Agentive Control

WAYNE WU
Carnegie Mellon University

This essay argues that current theories of action fail to explain agentive control because they have left out a psychological capacity central to control: attention. This makes it impossible to give a complete account of the mental antecedents that generate action. By placing attention at the center of the theory of action, and in particular the *intention-attention* nexus, we can recharacterize the functional role of intention in an illuminating way; explicate agentive control so that we have a uniform explanation of basic cases of causal deviance in action as well as other defects of agency (distraction), explain cases of skilled agency and sharpen questions on the role of thought in agency. This provides for a different orientation in the theory of action.

Introduction

What is agency? The received view is the causal theory of action: actions are events caused by an agent's mental states. Consider the primary case of intentional action where intention causes action. This causation should suffice for *agentive control,* for intentional actions are exercises of such control. Agentive control yields phenomena of central philosophical interest: moral, rational, reason-based, skilled, conscious, epistemic and free agency. To understand these specific forms of agency, however, we must understand the core phenomenon, agentive control.

The causal theory, however, fails to explain control. This is the lesson of deviant causal chains. Here is a common frame: a subject *S* intends to perform a felonious act that involves a bodily movement *M*, but the intention so unnerves him that he accidentally realizes *M*. In Roderick Chisholm's (1966) example, a nephew intends to murder his uncle, but this so unnerves him that he drives excessively fast and accidentally kills a pedestrian who turns out to be his uncle. So, his killing of his uncle was caused by his intention, but that killing is not under the nephew's control. This has led to talk of causation *in the right way*, whatever that comes to. Arguably, there is as yet no consensus on a solution to causal deviance in action although everyone agrees that control is disrupted. This means that there is no consensus on an account of agentive control.

The problem with current causal accounts is that they ignore a basic psychological capacity essential to control: *attention*. If attention is critical to control, then these theories are at best incomplete and at worst inaccurate. I argue that integrating attention into the causal theory provides an analysis of the basic case of agentive control that shows what goes wrong in classic examples of causal deviance on the basis of disruption in attentional guidance in control and illuminates specific forms of agency such as skill and expertise. Attention is central to action theory, yet no philosophical account of action has taken it seriously (well, almost no account).

Like traditional accounts of action, I focus on perception-guided bodily actions, but the picture generalizes to all forms of intentional actions, both bodily and mental. Section 2 elaborates the challenge facing the theory of agency while Section 3 provides the basic apparatus for which I have argued elsewhere, and links together notions of intention, attention, automaticity and control, within the context of an informational challenge all agents face. These notions are central to an adequate characterization of action. In section 4, I link intention to attention by showing how attention in action is set by intentions that enables action-relevant coupling of what the agent attends to and how the agent responds to it. This coupling can be seen in patterns of attention in the visual domain. Given this framework, section 5 shows how attention mediated coupling simply goes missing in cases of deviance and thus explains on the basis of a structural feature of agency exactly why agents lose control in those cases. Section 6 explains how attention plays a central role in skill, drawing on ongoing empirical work that demonstrates heightened attention and attention-based coupling in expert versus novice agents in sports. Section 7 discusses the role of cognition in skilled behavior and section 8 draws some final conclusions.

1. The Explanatory Challenge

I contrast two cases: *failed* agency in causal deviance and *skilled and expert* agency (I shall speak of skilled agency for short). The first questions whether the causal theory can explain agentive control; the second, whether it can explain skillful control. Begin with the first. There is no agreed upon solution to the problem of causal deviance.[1] The problem is that the behavior and the effects, even if intended, seem to happen *accidentally*. I would say that they happen *automatically* in a technical sense of "automatic" that I shall define. Because of this, agentive control goes missing in deviance. Of course, sometimes control *needs* to go missing to an extent. This raises the second case, the sublime forms of agency as in skilled athletes and musicians. We marvel at expert performance, but as experts and skilled agents know, getting better requires practice that allows one to relinquish the reins in certain respects. To learn a new and difficult piano passage, one must focus on the notes and on one's playing them, drilling the passage over and over until, when one comes to that passage in a performance, one automatically plays it with little thought. Automaticity then looks to be part of control in some circumstances and (as tied to the accidental) antithetical to it in others. What we need from a theory of action is an account of agentive control that explains why deviance is defective and skill sublime. The account should illuminate the presence of automaticity consistent with agentive control both in its mundane and sublime forms. So, an account of agentive control must find this balance. My suggestion, to be elaborated, is that the nexus between intention and attention is a key factor in explaining this balance.

2. Control, Automaticity and Action Structure: Setting the Stage

The idea of control and automaticity, applied to psychological processes, is common in cognitive science and even in mundane talk as when we say that an agent did X automatically or that she was fully in control. But if one wants an analysis of these notions, little is available (for a brief summary of empirical conceptions, see (Wu 2013)).

[1] For a recent account, see (Shepherd 2014). I will not discuss individual accounts, and many might be right on specific details. But as causal theories aim to identify the mental antecedents of action, it will be enough to show that attention is necessary for action guidance and that it has gone missing in all previous theories.

In psychology, the earliest discussion of control tied it to automaticity and attention although some prominent accounts offered circular analyses (Schneider and Shiffrin 1977). Since the seminal discussions in the 1970s, psychologists have largely eschewed definitions of control and automaticity for more rough and ready characterizations listing typical features of each. This renders the concepts simply catchall notions gesturing at a large set of properties, a motley bunch. We can do better.

Let us begin with a whiff of paradox in the following claims: (a) automaticity and control cannot be jointly instantiated: when one is present, the other is absent (see causal deviance); yet (b) our actions involve both automaticity and control.[2] Often, paradoxes are a result of poorly defined concepts, and one solution is to clean up our notions. Here is my proposal: automaticity and control should be defined relative to a feature of a process (Wu op. cit.). That is, processes are automatic and controlled relative to its various properties, but for any property at a given time, it is either automatic or controlled but not both. On this view, it is a mistake to speak of a process as either automatic or controlled *simpliciter*. By relativizing automaticity and control in this way, we can affirm both (a) and (b): an action can involve automaticity and control simultaneously so long as no one of its properties is automatic and controlled at the same time.

Control is tied to intention: the relevant properties of action are controlled if they are intended *and* the intention plays the appropriate causal role. I will say more about "appropriate causal role" in a moment (this phrase, like "cause in the right way" expresses the condition ruling out causal deviance). For now, let us focus on the connection between control and intention and defer discussing deviance. We are speaking of control as an executive function, and intentions are a paradigm executive mental state. This proposal allows for other notions of control useful to science, but for *agentive* control, the connection to intention--or other action-oriented mental states--is crucial. By "intention" I mean a broad category of action-representing states that play the functional role of generating action, so it includes more fine-grained notions of intention such as Bratman's (1987), primary reasons and belief-desire pairs (Davidson 1980a).

Agentive control is then defined in this way:

> Action *A* is controlled by *S* relative to its having *F* iff *A* has *F* because *S* intended *A* to be *F* and *S*'s intention plays the appropriate causal role.

For example, an action is controlled in respect of its feature of being directed at a target (e.g. a mug) iff *S* acted in that way because *S* so intended, e.g. intended to reach for *that* mug. In this way, talk of controlled features of action can be tied to talk of action as intentional under a description. Let there be some feature *F* of an action such that *the F-ing* describes it as intentional under that description. Then the relevant *F* is controlled presumably because for the action to be intentional under the description, *the F-ing,* the agent intended to F such that the intention played an appropriate causal role. The analysis of control thus makes contact with the causal theory of action. Automaticity, then, is the absence of control:

2 As I read him, Galen Strawson (2003) uses (a) to argue that there is much less mental action than we think there is. For a response, see (Wu 2013).

> Action A is automatic for S relative to its having F iff A's having F is not controlled by S.

At a time, an action's having F cannot be both controlled and automatic, but for an F and a G where $F \neq G$, an action can be both automatic with respect to one, and controlled with respect to the other. The whiff of paradox is dispelled. Note that if we replace talk of control with respect to F with one's intentionally F-ing, then we have the causal theory: one intentionally F's iff one's behaving in that way is caused by ones intention to F (in the right way). Similarly, automaticity is acting in a way not as one intended.

Consider playing middle C on the violin. A novice can be shown how to play middle C, how to hold down the string, how to draw the bow, to stand with violin appropriately under the chin and so on. The teacher will offer corrections but when the student is ready, she can intentionally play C. Her playing that note, her action's having the property of being a playing of middle C, is subject to her control. That is what she intends. But as a beginner, there is much else that is subject to her control or attempts at control: her holding the bow in a certain way, her placing her finger at a specific position, her standing with her foot out a certain distance, etc. The reason that these are controlled is that she explicitly intends to do so, to act in a way that features these properties, and they characterize her action because the intention plays the appropriate role in bringing them about. Moreover, control can fail in that she doesn't end up doing what she intended (e.g. she did not manage to place the finger in the right place). An expert violinist is different in that most of this is done *automatically*. This is one difference between novice and expert, a certain level of automaticity acquired over time.

In other work, I have argued that a *Many-Many Problem* is a necessary feature of all agency, mental or bodily (the idea is inspired by Allport (1987) and Neumann (1987)). That is a strong claim, so for current purposes, I rely on a weaker claim: in the paradigm bodily actions discussed by the causal theory, human agents must solve the Many-Many Problem. As most of these paradigms involved perception-guided bodily actions, I shall focus on them as well (see (Wu 2013) for extending the idea to mental actions). The Many-Many Problem is simply stated: normally, when you survey a scene for action, you confront many possibilities of action in terms of many perceived targets to act on, and for each, many things you can do with it. I will persist in speaking about the Many-Many Problem, but the same issues arise for cases where there is, in some sense, only one thing to act on and only one thing to do, so long as not doing anything is an option. To act, an input must be *coupled* to inform an output else no action happens at all.

The multiple potential targets and responses define a space of possible behaviors: a *behavior space*. In this space, the inputs are identified with each item *that the subject perceives*, so in fact, the inputs are perceptual experiences, individuated in terms of perceived targets (these can be object, feature or space in vision). The responses are those that are available to the subject for each target, and we shall focus on the movements that the subject can make in respect of those targets: grasping, reaching, pointing, kicking (etc.). This defines a behavior space that is a network of different input-output mappings, that is, perception to movement mappings. A *behavior* then is just when one path is actually chosen from this network of action possibilities, an input that is coupled to a movement. The agent responds in a certain way to what she perceives. It does not follow, however, that these behaviors are intentional actions, for some might be habitual, automatic, involuntary, or unintentional actions. The structure of the space is silent on that further issue.

Once this framework is in place, we can ask the following: when an agent acts intentionally, what role does intention play? The Many-Many Problem give us a different way of explicating the causal role of intentions. On the causal theory, intentions are said to cause action, and this is often understood to be one event causing another. This seems to me an incorrect or at least overly narrow account of the causal relation between intention and action. In part, this is because much of what we call intentions are standing states that persist over time rather than events that occur at a time. Right now, you have many intentions, some long term, some short: to go to a meeting, to meet your spouse for dinner, to finish a report due next week, to secure a new contract. These intentions influence your behavior, but it is not the case that for these to influence, they must be events. It is true that for some intentions, an event, namely a decision, leads to its formation, and that we can also actively think of what we intend. But often at the time of action, there isn't another redecision or rearticulation, an intention being activated as an event that causes action. We should open up our theorizing regarding causal possibilities taking the more static conception of intention seriously. Doing so can provide a different account of the causal role of intention that ties it to attention.

Here is the idea: intentions are states that structure how the Many-Many Problem is solved. As there are many options available to an agent, a crucial question is why the agent takes a *specific* path rather than others? The natural idea is that the relevant path in a behavior space is taken because the agent intended to act *in that way*, say to produce a specific response to a specific perceived target. Barring an accidental alignment of the behavior path and intention (cf. causal deviance), the role of intention is to aid appropriate path selection in an agent's behavior space. That is, it helps to solve the Many-Many Problem. But how?

If we think of intentions as enduring structural features of an agent's mind rather than events that occur when they are needed, we can better explain their persistent influence. The reconceptualization of the causal role of intentions in terms of their aiding the solving of the Many-Many Problem relies on treating them as such structural features. How might this proposed causal role be realized? We can draw on the idea of setting weights in a neural network where this setting of weights can be taken to *bias* the mental processes that occur, much as setting the switches in a complicated network of train tracks alters the course of train traffic. In a neural network that realizes the relevant psychological capacities, this biasing can be achieved by strengthening or otherwise altering the connection between processing units making processes relying on biased connections more likely to occur (cf. setting switches so that trains will go along certain tracks rather than others). For example, while unit X is connected to Y it might also be connected to Z. To bias X's connection to Y, the system strengthens or otherwise alters their connection so that X is more likely to lead to Y and not to Z.

Biasing of this sort can occur in a content-dependent way. In a behavior space that consists of two targets A and B and two responses R and S, an intention to R on A leads to biasing the processing that couples A to R (cf. setting points to get a train from New York to head to Boston); an intention to S on B leads to biasing coupling of B to S (cf. setting the tracks so that the Philadelphia train heads to Washington D.C.) and so on. This can be understood as a form of *top-down* modulation of a cognitive state on an agent's responses to what is perceived (we shall return to this idea later).

In this context, we do not think of intentions as events but as a state the subject is in such that the subject is disposed to act as she intends. This disposition will be realized in proclivities to respond to specific stimuli in specific ways. The realization of this dispositional profile can be understood to be (at least in part) the different weights in the underlying network realizing the capacities captured by the behavior space. Change one's intention, namely by changing its content, and the weights shift: intend now to rather than kick the soccer ball, and a different response gets biased; intend now to kick the pot rather than the ball, and a different input gets biased. The point is that we do not need to think of intentions as events to secure this link between different weights in the system realizing our agentive capacities. In what follows, I shall emphasize the role of intentions in biasing the input, what amounts to intentions setting attention. This will be a key component in expertise and in explaining what goes wrong with core cases of causal deviance. Intention-modulated attention is crucial to explaining agentive control.

To solve deviance, we need new ideas. We can treat the present proposal as a hypothesis that provides a conception of intention linking it to other psychological notions such as control, automaticity and attention, links that have not been actively discussed by causal theorists of action although they have been loosely connected by empirical cognitive science. The test for the hypothesis will be whether it can do significant explanatory work. I have proposed a reconceptualization of intention's causal role in light of the Many-Many Problem in terms of biasing. Having done that, the next step is to tie intention to attention.

3. Intention and Attention: Attention-Based Guidance in Action

The causal theory focuses on an agent performing a bodily action in relation to the environment. This fits the Many-Many framework where the subject's experience of the object informs a response. I have spoken of this input-output link as coupling, but it involves selection for action: the subject selects the object for action. "Selection" might suggest to some readers that the agent must do something else (*selecting*) in order to act. In fact, the idea is simpler. We have appropriate selection when a subject's perception of the environment is coupled to and thereby informs the production of a response. That is, selection is just a necessary product of taking a specific path in behavior space. Where we have a specific aspect of the subject's experience operating in this way, we have selection of what is experienced for action. This just is, I claim, a form of attention.

There are too many inchoate and conflicting intuitions about attention and it is hard to identify constraints to keep steady our theorizing. We need to be more systematic about this (see (Wu 2014) for a start; also (Mole 2013)). Many empirical theorists of attention have pointed out that there is no uniform conception of attention. I disagree, and think that William James basically got it right when he wrote:

> Everyone knows what attention is. It is the taking possession by the mind, in clear and vivid form, of one out of what seem several simultaneously possible objects or trains of thought. Focalization, concentration, of consciousness are of its essence. It implies withdrawal from some things in order to deal effectively with others, and is a

> condition which has a real opposite in the confused, dazed, scatterbrained state which in French is called distraction, and Zerstreutheit in German (James 1890).

The critical idea is the "withdrawal from some things in order to deal effectively with others" where we can see this as a description of what happens within a behavior space when a subject traverses a path in performing an action. The idea is that the process must be in that sense *selective*, entailing removal of resources from other options to deal with the selected behavior. But we can converge on the same idea via assumptions in the scientific study of attention where there is a fundamental connection of attention to the idea of a *task*, namely the actions that subjects perform in a laboratory such as visual search, remembering and recalling specific items, tracking, reporting, and so on. These activities are just specific versions of things we do all the time. In studying attention, scientists direct and control how subjects attend by devising a carefully designed task where an item that the subject needs to attend to is made *relevant to performing the task*. This means that the only way for the subject to fulfill the task is to select that target or its properties to guide performance. Thus, if a scientist wants a subject to attend to an object, she makes it the target for tracking, for remembering, for reporting and so on. In designing a task where the object (or property or location) must be selected in this way, the experimenter then takes the subject to be attending to the object (property, location). So, the methodology assumes the following experimental sufficient condition for attention:

> If S selects X to perform a task T, then S attends to X.

One could treat the link as merely evidential: if one has evidence that S selects X to perform a task T, then one has evidence that S attends to X. Of course, that is the first step, but the question is whether there is reason to deny the experimental sufficient condition linking selection directly to attention. It is hard to see what attention would be in these experimental paradigms over and above task selection by the subject (i.e. the subject's perceptual state playing a selective role). The reason why scientists treat evidence for selection of X for task as evidence for attention to X is that they implicitly endorse the experimental sufficient condition.

In the end, it doesn't matter much whether one thinks that such selection is attention. Some might be inclined to call it *subject level task-relevant selectivity*. So, in what follows, it doesn't matter much whether one thinks of coupling as illustrating a family of subject level selective capacities that include attention or, as I do, that these just are forms of attention. The critical point is that these capacities have gone missing in causal accounts of action. Given the previous points, I will treat the selection for task as attention (and for more detailed arguments for this, see (Wu 2014, chap. 1–3)). Furthermore, given the mundane capacities called upon in experimental tasks, say looking, reaching, remembering, reporting and so on, it is a small leap to generalizing from experimental tasks to mundane actions. Thus, we should treat selection for action as attention because our best experimental methodologies of attention in cognitive science assume that and, if James is right, so does our folk theory of attention. So, in coupling an input to response in solving the Many-Many Problem, the subject's selecting X for response suffices for the subject's attending to X. This is just when what the subject perceives informs her response. And of course, such attention can be controlled or automatic (see previous definitions). It can also be conscious or unconscious but let us set that aside (see (Kentridge 2011) for a detailed discussion of unconscious attention).

We can get at the need for this type of subject level selectivity by returning to the cases of mundane bodily action and noting an element that is also largely missing in discussions in action theory: *perception.* Most of the examples of actions discussed by causal theorists are actions in response to the agent's perceived environment. They are perceptually guided: an agent moves in response to what she perceives, whether an elbow to be touched, a rope to let go, a light switch to flick, or a herd of pigs to scare. Yet if one looks at the classic philosophical discussions of action, one will be hard pressed to find systematic discussion of the role of perception.[3] Placing perception into the mix, we see immediately that selectivity is required. Of all that the subject perceives, only a subset is task relevant. Thus, to the extent that the subject's actions are perceptually guided, to the extent that those actions are sensitive to the perceived environment, the subject must respond to what is perceived in a selective way. In the context of the Many-Many Problem, the point is that we cannot act on all of the inputs presented to us by perception. We must select for action. We must attend.

At the same time, intentions influence attention. We do not attend willy-nilly, but attention is coordinated with intention. This flows from our discussion of the Many-Many Problem and the assumption of selection for action as sufficient for attention (that is what input-coupling amounts to). The idea is that in setting intention, one sets the weights that biases which selections are made in action (psychologists speak of *task sets*). So, if one intends to act on X, then X is selected for action; if on Y, then Y is selected. This is part of intention's causal role, one that is driven by the content of intention. The point is that by registering attention in the causal theory, we can expand our conception of intention's causal role, and we have substantial empirical evidence for this role in normal action (for a review, see (Hayhoe and Rothkopf 2011)).

An illustrative study is by Alfred Yarbus (1967) (see Figure 1). Generally, eye movements such as saccadic eye movements are described as *overt attention*. If we think of *covert* attention as attention independent of bodily movement, then covert attention programs overt attention (Armstrong 2011). I prefer to think of this as covert attention programming saccades, so there really isn't a different form of attention (overt), but set that aside. The point is that we can look at patterns of eye movements as expressions of patterns of covert attention. Yarbus asked his subjects to visually examine a painting and to perform several tasks including: (5) to remember the clothes worn; (6) to remember the objects present in the room; and (7) to estimate the time the man had been away. What is striking is that the pattern of eye movements is intelligible given the subject's goals: the eyes linger on (5) the figures; (6) the objects; (7) the faces.

[3] Please choose your favorite classic discussion of action since Anscombe's work (Anscombe 1957). Search in the index or by a search function over an electronic version for "perception". What do you find? Do the same for "attention."

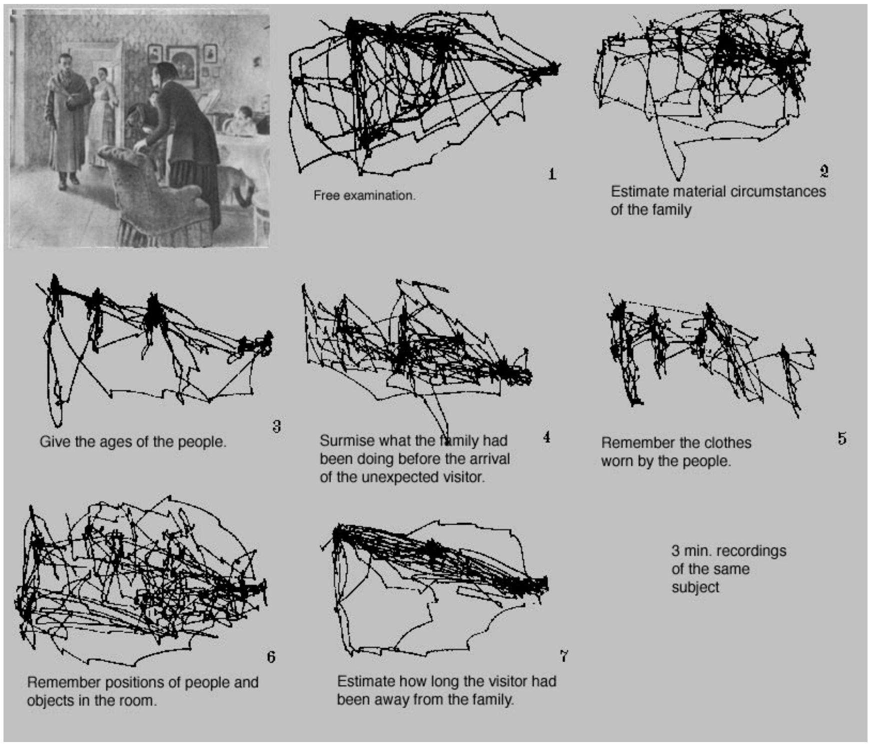

Figure 1. Data from Alfred L. Yarbus, *Eye Movements and Illusion,* 1967, New York: Plenum Press, p. 179, figure 109. This version retrieved from http://en.wikipedia.org/wiki/Eye_tracking. Material used with kind permission from Springer Science and Business Media.

We can think of each panel as an experiment where what is toggled is the intention, given different tasks, while the stimulus, the painting, remains constant. In each case, the subject takes on an intention to follow Yarbus's instructions, say to look at objects, clothes, or the members of the family. With each change in task, the pattern of movements shifts. Since eye movements are driven by covert attention, the pattern of covert attention also shifts, with attention and coupled eye movement changing with different intentions. This illustrates how changing intentions can change visual attention and motor coupling. With each new instruction, different objects, different inputs in the behavior space, become task relevant, and selection for task shifts accordingly.

The experiment also illustrates the difference between automaticity and control. On the one hand, the movements are explained by the specific intentions, movements that amount to visually examining the objects of relevance to the instructions (recall selection for action). This is what the subject presumably intends to do, so that feature of the movements is controlled, namely *looking so as to carry out the instructions* (e.g. to remember objects, to estimate the duration of the man's absence, etc.). Yet at the same time, the precise *patterns* of movements, their order, their kinetics and so forth are automatic, not what the subject intends. We have then in Yarbus's results an illustration of the causal role of intention in influencing attention where shifts of attention exemplify automaticity and control.

The relation between intention and attention has been well-documented in other areas including reading, sports, and mundane activities (see (Hayhoe and Rothkopf 2011); (Schütz, Braun, and Gegenfurtner 2011)). What we can say is that there is good evidence that (a) intentions play a causal role in respect of setting action-relevant attention; (b) that such attention aids task performance (informs it); (c) that this is consistent with the picture from the Many-Many Problem where intention biases coupling; (d) that this biasing does not require that the subject's intention be to explicitly attend in a certain way though it can; (e) as such, such action serving attention need not be intentional (controlled) but can be, and is often, automatic. The fundamental nexus is between intention and attention, illuminated by association with the Many-Many Problem and the contrast between automaticity and control. If the previous considerations are correct, then the intention-attention nexus provides a uniform and central part of agency and as such, we should expect it to be central to our understanding of many aspects of agency. We shall see this nexus as a common point in deviants and experts to be discussed later.

The role of intention can be understood as a form of top-down influence on attention. While I have focused on the functional role of intention on attention, the neural realization of that role is illuminating. One signature of attention on neural activity is *receptive field remapping*. Visual neurons have receptive fields that are tied to the retina in that activation of specific parts of the retina, via stimuli at external locations that project to those parts, will drive the neuron to fire, i.e. generate action potentials. Visual neurons can also be *tuned* to different stimuli, with some stimuli driving a stronger response (preferred) and some a weaker response (less preferred). Now consider two objects in a visual neuron's receptive field, one preferred, the other not. The response of the neuron to two objects is different from its response to each object individually. Essentially, the neural response to the two objects is the weighted average of the response to each object individually. What is striking is that when one of the two objects is rendered task relevant so that the perceiver intends to act on that object, the activity of the neuron can shift as if only the task relevant object is in the receptive field. It is as if the receptive field shrinks around that object (Chelazzi et al. 2001; Chelazzi et al. 1998). Think of this as task-relevant neural selection in respect of objects in the receptive field. This effect has been observed in mid-level visual areas (V4) to high-level visual areas (IT; inferotemporal cortex) in what is known as the ventral visual stream. We can theorize that the animal's intention or similar action-oriented state influences visual attention by influencing visual neural processing underlying visual selectivity.[4]

One model of how conscious experience influences behavior is given by Milner and Goodale's well-known model of the visual system (Milner and Goodale 1995). They hypothesize that the ventral visual stream (including IT), which on their model has the responsibility for realizing conscious visual experience, influences visual-guidance of action through attention. Attention to an object begins with modulations in the ventral stream and feeds back to early visual areas that then aid dorsal visual stream computations that more directly inform action. One way this is done is by identifying relevant targets of action, so a plausible hypothesis is that the receptive field remapping observed in the ventral stream serves this targeting. This suggests a functional role for attention that is directed by intention. Attention identifies intended targets of action so

[4] The main model of such influence is the normalization theory of attention. For a discussion of this model and how to connect it to the role of intention, see my "Shaking the Mind's Ground Floor."

that information from those targets can be used to inform the production of an appropriate response. This would be a neural implementation of the *guidance* role that philosophers have often attributed to intention in order to answer the challenge of causal deviance. I have called this role for attention *coupling* of input to output. Attention guides action in part by selecting objects (or properties) to couple to and inform action (for an argument that attention must influence the dorsal stream in certain actions, see (Wu 2008); for further discussion see (Mahon and Wu Forthcoming)).

The appeal to work in the neuroscience of attention points to the basic picture that underwrites my talk of coupling. What attention is providing to the agent is the relevant target but this involves providing content to systems that process that target in a way that sets parameters for and programs an appropriate response. For example, in visually guided reach, one parameter will be spatial location of the attended target so that an appropriate reach can be programmed to that location (this is part of what the Milner and Goodale account seeks to explain). This information is precisely what is needed to get the hand to where it must be. Normal coupling involves the linking of information of this sort provided at the level of perceived input into the machinery of movement production.

Attention then is a crucial ingredient needed to explain how intention "guides" action. Such guidance marks the goal-directed deployment of attention that is needed in selecting what is perceived to program the motor response. No doubt, there are many empirical issues to be settled, but the point is to see how the empirical issues can illuminate the philosophical issues about control once we have a clearer view of all the mental states and capacities that figure in it. A basic capacity of mind, attention, plays a crucial guidance role: it identifies perceived targets of action to which the agent responds.

4. Causal Deviance and Attentional Guidance

Let us now return to the problem of deviant causal chains, specifically what is called "antecedential" (Brand 1984), "primary" (Mele 1992), and "basic" deviance (Bishop 1989). The idea I pursue is that if the intention-attention nexus is crucial to control, defects in that nexus can illuminate defects of agency as in causal deviance, and indeed other defects as well (e.g. distraction). Here is a recent gloss on the problem: "In all cases of deviance, some control-undermining state or event occurs between the agent's reason states and an event produced by that agent'' (Schlosser 2007). Sometimes philosophers speak of a gap between an intention and the action and this led to one proposal to require that intentions serve as *proximate* cause (Brand 1984). On the view advocated here, there is a gap, but not quite the one that philosophers have thought. The gap results from failure to identify all the psychological components of agentive control.

Let's return to the nephew. What goes wrong is that his intention does not lead to appropriate coupling with the result that his movement fails to be guided by attention to the target. The resulting killing is accidental (automatic, not intended). For there to be an intentional killing, the paradigm case would be that the nephew intends to drive over his uncle, namely *that* pedestrian, requiring coupling of attention to the spatial location of the uncle to inform the driving response. But appropriate coupling is clearly missing despite the causal influence of intention on the process. We have an absence of agentive control because of a disruption in the appropriate intention-attention nexus. Many of the standard cases of deviance, typically involving felons, can be shown to involve disruptions of attention-based coupling. Given that, we can formulate the following sufficient condition for agentive control:

> *S* exerts agentive control in *F*-ing at *t* if *S*'s *F*-ing is a result of an intention to *F* where this intention solves the Many-Many Problem faced by *S* at *t* by directing appropriate attention-mediated coupling.

If you like, talk of "cause in the right way" or "appropriate causal role" is spelled out by delineating more completely the psychological basis of agency, namely the intention-attention nexus. The associated ideas of guidance and control then are cashed out by the mechanisms of biasing and coupling discussed above. The standard cases of deviance fail to meet the antecedent of this condition and hence fail to provide a counterexample to it. If we take the condition here as necessary as well and recall the link between automaticity and control discussed earlier, we can also say the following:

> *S* does not exert agentive control in *F*-ing at *t* (i.e. *F*s automatically) if *S*'s *F*-ing is not a result of an intention to *F* where this intention solves the Many-Many Problem at *t* by directing appropriate attention-mediated coupling.

It turns out then that deviance is quickly addressed in a way that we expected. For causal deviance revealed that the causal theory did not adequately explicate agentive control. Such control reflects a complex psychological capacity of the agent, so that defects in control must be defects in the implementation of that capacity. By highlighting the intention-attention nexus, we illuminate agentive control. On that basis, we can explain how many classical cases of deviance amount to failures of control and how a specification of control avoids those cases.

My definitions suggest that we should say that the nephew's unintended killing is automatic when it is more natural to say that it is merely accidental.[5] The critical point is that the killing is not intended or intentional; specifically, it is not controlled. Since "automatic" is currently used in a technical sense that entails unintended or uncontrolled, the phrasing need not be problematic even though it sounds discordant. To clarify, I claim that technically, automaticity includes accidental behaviors and automatic behaviors, the latter in a colloquial sense that reflects recognition of the behavior's source in repetition and practice (habitual behaviors provide another category of "automatic" behavior). So, the skilled violinist's drawing of her bow in the correct way is automatic in the colloquial sense in that it is honed by endless practice while the nephew's killing is merely accidental in that it doesn't arise because of practice (see below on discussion of skill). Both involve the absence of intention (i.e. are not controlled) for the specific features in question, and both are thereby, in the technical sense, automatic.

Deviance is not the most interesting case of defective agency. A common case is *distracted* agency, but distraction, as James noted, is the opposite of attention and is more common. Here, the intention-attention nexus is again important for we can think of distraction as perceptual selection that is contrary to the intended action, and in that way

5 I owe this observation to a referee and am grateful for the prompt to clarify the terminology.

works against agentive control.[6] This can happen in various ways with varying results: distraction can slow us down, can interfere, or can simply abrogate the current action by initiating a different action. Consider a variation of the nephew case. The murderous nephew drives at high acceleration to his uncle's house but suddenly his attention is *captured* by a man wearing a brightly colored fur coat who just happens to be his uncle. Unfortunately, when driving, one tends to steer in the direction that one is looking (at least this is true for me), and there is then an *automatic* motor response (i.e. unintended) of steering into the man and killing him. This is a kind of distraction that is perhaps too common among drivers leading to unsafe road conditions: distraction leading to veering. In the nephew's case, there is an attention-motor coupling that leads to the desired homicide, but the intention does not contribute to the coupling as its content does not bias attentional processes to shift to the uncle. Rather, the nephew's attention to the uncle was captured by the colorful coat. The upshot of attentional capture is that it is often contrary to the causal influence of one's intentions. So, while there is an attention-motor nexus leading to the desired result, the intention-attention nexus is missing.

Let me deal with one other famous homicidal case, that of Davidson's climber (1980b) who intends to let go of the rope holding his companion and whose intention so unnerves him that he lets go of the rope. What is the contribution of attention? Recall that a Many-Many Problem exists here as well: there are many parts of our body whose movements are subject to control in the sense defined. Use of a specific part of the body requires that it and not another part is selected for motor control. To operate with the hands and not the foot, the current state of the hand must be selected to inform a new state even if it is a simple change such as relaxing the muscles (this is, after all, a targeted relaxing). Here, attention is proprioceptive, selecting a specific body part to inform a response involving that part (see motor control models such as (Wolpert and Ghahramani 2000)). There are two issues then: (a) is this *subject level* selection and (b) is this a form of attention? On my view, which I shall not argue for here, the subject is unconsciously attending to the state of his body, selecting it to inform a response, but I suspect many will find that a controversial claim. As there is no space to address these worries, let me weaken the claim in a way consistent with the previous points.[7]

The weaker claim is that even if one withheld attribution of attention to selection of proprioceptive information, the basic structure of intention-influenced selectivity for action remains. Our intuition is that intentions must play a controlling role in the generation of movement. If this movement requires selective activation of parts of the body in a way sensitive to the current state of that part such as its position, then this selective activation, part of solving a version of the Many-Many Problem, must be informed by intention.[8] That is, even if you treat proprioceptive selection as *sub-personal,* the general framework of a coupling nexus influenced by intention takes hold, modeled on the personal level case we have discussed. But then, the same disruption occurs: normal coupling of a prior bodily state to inform transition to a new bodily state

---

6 Some cases of deviance draw on distraction. See Christopher Peacocke's bank clerk who "absent-mindedly" writes down his name (the wrong name) on a check because of his desire for more money (Peacocke 1979, 55). Here, distraction is contrary to intention and disrupts the normal intention-attention nexus.

7 An argument that these states, in informing action, must be personal level states is given in (Wu 2011). For some remarks in the same direction, see (Burge 2010, 369–76).

8 (Wu 2008) discusses different types of Many-Many Problems

is disrupted. In this case, the bout of nervousness suffices to generate the new bodily state without information from the prior state providing information to program the response. Thus, the general recipe for identifying defects in causal deviance remains. Intentional action requires intention-mediated coupling of selective processes to inform movement, and many standard cases of deviance abrogate agentive control by disrupting those processes. This application to deviance and distraction attests to the explanatory power of attentional selectivity in the theory of action. No doubt, counterexamples will be put forward (philosophy being what it is), but it is important to recognize that much illumination on traditional cases is gained once all the relevant psychological capacities are brought into view.

5. Automatic Attention and Expertise

If we are on the right track, that is, if intention mediated attention is an essential component of agentive control, then focus on attention should illuminate other forms of control, say expertise and skill. This is a *prediction* of the model. Now, the idea of skilled agency is rather broad and it is unclear whether there are general principles that tie together experts in sports, chess, music, dance, the arts, etc. For one thing, in each case, different tasks and goals are performed, different strategies are needed, different capacities are appropriate, different effectors must be called upon etc. There might not be a uniform theory of skill and expertise. Nevertheless, since skills are exercised in agency, the key to understanding them begins with agency *simpliciter*. Correlatively, changes that occur given the acquisition and exercise of skill and expertise must be reflected in changes in core features of agency. We might expect that in many cases, there will be changes at the intention-attention nexus. In this section, I focus on attention in skilled athletes specifically batting in cricket though similar issues arise in other sports (see (Mann, Spratford, and Abernethy 2013) for references). There are differences in how skilled versus novice athletes attend to a scene and pick up information.

Let us consider the case of batsmen in cricket when confronting fast bowlers. The need for selection for action is clear: a batsman is confronted with a variety of stimuli in the visual field but must select just one for current action, namely the ball which can approach at speeds of up to 100 mph (160 kmh). To hit the ball, information about its trajectory and position relative to the batsman's bat must be selected to couple to an appropriate swing. If we describe the batsman's action in terms of the Many-Many Problem, then players of all levels intend to produce the same input-output coupling, namely the ball's trajectory as informing the batting response. So, they should all attend to the ball. If there is a difference in skill, it will be in terms of features of this coupling.

We have already seen from Yarbus's data that eye movements can be automatic yet sensitive to one's goals where such movements are programmed by and indicative of patterns of selective attention. Subjects need not have any idea of the type of movements that they make. Recent experimental work tracking eye movements suggests that there are critical differences at the level of attention between batsmen of different skill levels, from novices to the greatest contemporary batsmen. Land and McLeod (2000) examined three players: two amateurs (strong and weak) and a professional. They showed that there was a consistent pattern in eye movement between the three. Using eye-tracking, they identified an initial fixation at the point of release of the ball, then a *predictive* saccade to where the ball would bounce, and then tracking of the ball after the bounce for about 100-200 milliseconds. This identifies three points of fixation (foveation): at the release point of the ball, at the bounce point, and of the ball itself, at least briefly after the bounce. One thing that is surprising is that the batsmen are not fixating the ball for most of the trajectory pre-bounce (cf. the old coaching adage of *keeping one's eye on the ball*). A key difference between the three players was that in the more skilled hitters, the latency of the first saccade was smaller (i.e. the time that fixation was maintained prior to saccade). One possibility is that the advanced hitters learned to more efficiently extract (select) the information required to make the predictive saccade and that this efficiency of attention contributes to their greater skill (for arguments to this effect, see Land and McLeod, op. cit.). Put another way, they could shift attention more quickly, something that seems advantageous in fast bowling situations. We need not be concerned with the computational role played by this difference in latency, but simply note that it is an alteration in the intention-attention nexus, or more specifically of the deployment of attention.

It is worth noting that phenomenologically, expert hitters claim to be able to see the bat hit the ball, but Land and McLeod did not observe this. Later work by Mann, Spratford and Abernathy (2013) showed that *elite* players, the upper echelon of batters, were able to track the ball after the bounce. Indeed, not only were these players able to keep their eye on the ball, but they were often able to keep the ball in a constant position relative to the head by pursuing the ball via head movement. In this way, they were able to maintain the ball in a constant head-centered egocentric position. It is plausible that by keeping one of the (egocentric) spatial properties of the ball constant, namely its position relative to the head, the batsmen accrue an advantage in how the ball can be struck. The motor system, for example, need not worry about recalculating a response if a critical spatial property is constant throughout the action. What the elite players show is an ability to couple attention to the ball with movement in a unique way. Given that this keeping one's eye on the ball is something that plausibly requires attention and is honed over years of practice, what we see in this highest level of expertise is a very refined input-output coupling.

Similar changes in eye movement, and by implication attention, are seen in athletes in other sports that involve moving objects. Attentional coupling of input to inform a response is common to these cases. In the case of batting, tracking the position of the ball at the time of batting is crucial to successful behavior. The increase in skill level is plausibly tied to improvements in coupling input to output (i.e. hitting the ball). If attention plays a central role in generating action, as I have argued, then in *skilled* action, we might expect that attention not only serves that role but often varies with the level of skill. The studies on athletes support this.

When we consider the *balance* of control and automaticity, it is clear that all players exhibit control in that they intend to hit the visible ball, and when they do so, that feature of their behavior is controlled and intentional. But many of the eye movements that batters make such as the predictive saccades are automatic, something fixed and refined by hours and hours of training and not explicitly intended. Indeed, the batters might never be aware of those movements. Moreover, we can think of this automaticity as freeing up capacities for different sorts of behavior. Put in terms of the Many-Many Problem, what elite players are capable of, on the basis of their input-output coupling profile, are more types of controlled behavioral outputs than that available to less capable players, more ways of hitting the ball that contributes to their unique success (e.g. direction in hitting the ball). This follows from the idea that intentional control (the influence of intention) is limited in terms of its influence. Where one controls some feature *F* of an action, that limits what other features can be controlled. To the extent that the action's being *F* is automatic, new avenues for control are opened (see David Papineau on basic action and skill, (2013)). Not only will the nature of attentional coupling vary between skills, but the behavior space will also vary. This speaks to how there are different ways of balancing automaticity and control in agency, and how practice makes possible not only skilled behavior but nuanced behavior and further possibilities for action.

6. Thought and Skill

Skills highlight a long-standing disagreement about the role of thought in expert action. Hubert Dreyfus has argued that agents can skillfully cope with the world without the intervention of mental representations such as conceptual representations in intention (e.g. (Dreyfus 2002); see recent work on knowledge how and motor control, (Stanley and Krakauer 2013)). For our purposes, the issue is whether intentions play a crucial role in skilled behavior that is at odds with Dreyfus's claims. An important standard here will be that if one disagrees with Dreyfus, one must provide a clear account of *how* thought influences expert action, not merely that one has reason to believe that it does. What we want is an account of the underlying psychological and neural mechanisms that ground such influence.

Let us focus on intentions and attention. We can take the intentions in question to be conceptual in the sense that these states can be the product of explicit practical reasoning, the conclusion of a course of deliberation about what to do. To that extent, having the intention involves the deployment of certain conceptual capacities corresponding to the content of the intention. One cannot intend in this way to do *F* without having the concept of *F*, the very concept that figures in deliberation and corresponding beliefs about *F*. The question we can pose then is this: in skilled agency, how far is the causal reach of intention?

Dreyfus aims to refute the claim that intention permeates *all* forms of skilled agency by providing counterexamples. I suspect that his existential claim is likely correct: there are skilled actions that don't depend on current intentions. Here's a case that I think fits what Dreyfus might have in mind. An expert fencer is standing with her niece in a toy store. At some point, the niece picks up a toy sword and pretends to lunge at her aunt who now automatically, reflexively, perhaps playfully moves in a way that expresses her training in respect of lunges with sharp implements. I don't think we need to postulate an intention to play with her niece here. The fencer acts skillfully but without explicit intention. One might worry that the movement is a mere reflex (as Louise Antony (2002) notes), but we can use the concepts discussed above to show how they reflect a type of intelligence (see especially Fridland's work (2014)).

On Dreyfus's own account, intentions are certainly in the *historical* background of the fencer's acquiring the skill that she automatically deploys. She has intended to learn certain basic moves and over time, has acquired an expertise to deploy those movement capacities in a variety of appropriate situations. That is, there is a history of intentions helping to solve the Many-Many Problem during past practice where her deployment of a parry movement was done intentionally (say x number of times in each practice session). Over time, the input-output couplings might come to be automatically deployed: one parries without having to "think" about it, i.e. to intend to. Bernhard Hommel speaks of intentions as "prepared reflexes" (Hommel 2000), and the idea is suggestive. Intention helps to couple input to output so that the relevant action is efficiently expressed. Accordingly, the fencer's automatic parrying of the niece's lunge is in a way reflexive (better automatic). Yet while intentions are not actively involved, her move counts as intelligent given the diachronic picture of intention-guided practice just noted. Intentions were actively involved in setting up the input-output coupling that is now automatically deployed (these issues are insightfully discussed in Fridland's work, op. cit.). So, we can agree with Dreyfus that in some cases, intentions are not concurrently involved in skilled, intelligent agency as per the standard causal account without losing sight of the influence of intentions over time on current behavior.

The question, then, is how to sharpen any further debate between Dreyfus and others. The core issue of expertise and skill is central to understanding agency in other forms of broader interest to philosophers, say virtue and agency (cf. Aristotle's conception of the practically wise man, especially in John McDowell's characterization of him; see (McDowell 2007) for a recent statement) or epistemic agency in relation to an agent's ability to deliberate selectively only on appropriate reasons (see (Siegel and Silins forthcoming)). As I have emphasized, our accounts of agency should inform any such discussion (see (Wu 2014, chap. 8) for some comments on the epistemic case).

Dreyfus has presented a description of the transition from novice to expert. On the latter, he writes the following:

> The expert not only sees what needs to be achieved: thanks to a vast repertoire of situational discriminations he sees how to achieve his goal. Thus, the ability to make more subtle and refined discriminations is what distinguishes the expert from the proficient performer. Among many situations, all seen as similar with respect to plan or perspective, the expert has learned to distinguish those situations requiring one action from those demanding another. That is, with enough experience in a variety of situations, all seen from the same perspective but requiring different tactical decisions, the brain of the expert performer gradually decomposes this class of situations into subclasses, each of which shares the same action. This allows the immediate intuitive situational response that is characteristic of expertise (op. cit. 371-2).

Much of this has echoes of selection and attention. Dreyfus notes:

> It is crucial that the agent does not merely receive input passively and then process it. Rather, the agent is already set to respond to the solicitations of things. The agent sees things from some perspective and sees them as affording certain actions. What the affordances are depends on past experience with that sort of thing in that sort of situation.

Earlier, I conceded that there is skilled behavior that is not *mere* reflex, exemplifies skill and expertise, and does not depend on current intentions even if it reveals the influence of training and past intentions. But Dreyfus assumes that such cases provide a good model for expert behavior *in general*. The problem is that skillful behavior is almost always generated in light of specific intentions.

Attention is missing in Dreyfus's discussion. An agent can perceive affordances, but a given object affords many different actions. Shift to affordances does not obviate attention: we still have a Many-Many Problem for affordances. Selectivity is needed. What allows for specific selections? On my view, the natural source of biasing is the agent's intention. Consider the fencer again. We could imagine that when she is in an Olympic competition, all her parries and attacks are done simply as a result of built up automatic responses to environmental situations that show some flexibility to new patterns. But emphasis on examples like the toy store case where intention is absence ignores the presence of intention in all normal instances of sporting competition: the fencer is there because she intends to compete, and indeed, intends to win the match. Indeed, she has likely studied her opponent and has formed a plan of attack. It is implausible to take the intention to be causally idle, needed only to get the fencer in front of the opponent. But then what is its causal reach?

We could, for example, speak of the cricket batter as responding to an affordance by seeing what needs to be done given the trajectory of a ball whose velocity approaches 100 miles an hour. If the previous discussion is correct, batting requires selectivity of the relevant affordance in order to respond, and this, I have argued, is attention that is fixed partly by intention. The empirical story to be told here might reach far down indeed, if the empirical work appealed to earlier is correct, down to basic visual processing of objects. In actual expert behavior, intentions modulate the deployment of learned automaticities. It could very well be that paradigm cases of skilled behavior are imbued by conceptual states through and through. This is compatible with saying that the behaviors exhibit quite a bit of automaticity (in that way, they are like reflexes which are paradigmatically automatic), but it also allows that in many ways, they are controlled. What the current debate lacks is recognition of the intention-attention nexus and detailed proposals of underlying computations and mechanisms. The phenomenological points will only get us so far, and to my mind, too often underdetermine the nature of the underlying processes.

We can provide a framework for productive discussion by focusing on the intention-attention nexus in a philosophically and empirically informed way. The philosophical insight sees this nexus as central to agentive control while the empirical insight is to note the physical and computational possibilities. As we saw, the influence of intention can influence basic visual processing of objects at the level of single neurons throughout the ventral visual stream. This raises two questions: (a) how far is the influence of intention in respect of basic processing needed to produce action? and (b) what is the nature of this influence in mechanistic and computational terms? These are questions for cognitive science, both theoretical and empirical, but the answers to them will be relevant to philosophical debates about skill and agency.[9] Still, the basic philosophical task of

---

9 I have focused on the attention (input) side, but similar issues arise on the motor (output) side. On this, see the work of Stanley and Krakauer (2013), Fridland (op. cit.), and Papineau (op. cit.), among others.

providing a theory of agentive control might be largely in place if the previous considerations are correct. At the very least, we have a philosophically motivated hypothesis on what it is for an agent to exert control in light of the lived world.

Conclusion

Where are we? We began with a venerable story of action that focused on understanding the mental causal antecedents of action but curiously failed to deal seriously with attention and, relatedly, perception. Restoring attention to the picture, in the context of spelling out the structure of agentive control within the Many-Many Problem and the role of intention in solving that problem, yields a more complete picture of agency. The test of that picture is in its providing a uniform basis to explain both what happens when agency is defective and what happens when agency is sublime. I have argued that the appeal to such psychological capacities allows us to explain what is wrong in classical causal deviance, in distraction, and what is right in skilled agency. This provides an indirect argument that these capacities do play a critical causal role in agentive control and thus should play a central role in a theory of agency. It makes possible articulation of a set of new questions for the theory of action.[10]

References

Allport, A. 1987. "Selection for Action: Some Behavioral and Neurophysiological Considerations of Attention and Action." In *Perspectives on Perception and Action*, 395–419. Hillsdale, N.J.: Lawrence Erlbaum Associates, Publishers.

Anscombe, G.E.M. 1957. *Intention*. Oxford: Blackwell Publishers.

Antony, Louise M. 2002. "How to Play the Flute: A Commentary on Dreyfus's 'intelligence without Representation.'" *Phenomenology and the Cognitive Sciences* 1 (4): 395–401.

Armstrong, Katherine. 2011. "Covert Spatial Attention and Saccade Planning." In *Attention: Philosophical and Psychological Essays*, edited by Christopher Mole, Declan Smithies, and Wayne Wu, 78–96. Oxford: Oxford University Press.

Bishop, J. 1989. *Natural Agency: An Essay on the Causal Theory of Action*. Cambridge: Cambridge University Press.

Brand, M. 1984. *Intending and Acting*. Cambridge, MA: MIT Press.

[10] I'm very grateful to an anonymous referee for helpful comments. This material was presented to Hong Yu Wong and his Philosophy of Neuroscience group at the University of Tübingen. I am grateful to extended discussion with Hong Yu and to his group for comments and enjoyable conversation including, among others, Chiara Brozzo, Gregor Hochstetter, Alex Morgan, Krisztina Orban, Ferdinand Pöhlmann, and Katia Samoilova. Some of this material was presented at the European Society for Philosophy and Psychology meeting in Noto, Sicily in a symposium on skills and automaticity with Ellen Fridland, David Papineau, and Josh Shepherd. I've learned much from their writings and was fortunate to have their instruction and advice in Noto, especially over a fantastic Italian meal. Thanks also to Marisa Carrasco, Tim Crane, Thor Grünbaum, Bob Kentridge, Hemdat Lerman, Myrto Mylopoulos, Ian Phillips, James Stazicker, Till Vierkant, several of the "Tübingers" and others I am currently forgetting (my apologies!) for discussion and challenges during that session.

Bratman, M. 1987. *Intention, Plans and Practical Reason*. Cambridge, MA: Harvard University Press.
Burge, Tyler. 2010. *Origins of Objectivity*. OUP Oxford.
Chelazzi, L, J Duncan, E K Miller, and R Desimone. 1998. “Responses of Neurons in Inferior Temporal Cortex during Memory-Guided Visual Search.” *Journal of Neurophysiology* 80 (6): 2918–40.
Chelazzi, L, E K Miller, J Duncan, and R Desimone. 2001. “Responses of Neurons in Macaque Area V4 during Memory-Guided Visual Search.” *Cerebral Cortex (New York, N.Y.: 1991)* 11 (8): 761–72.
Chisholm, Roderick. 1966. “Freedom and Action.” In *Freedom and Determinism*, edited by Keith Lehrer, 11–44. New York: Random House.
Davidson, D. 1980a. “Actions, Reasons and Causes.” In *Essays on Actions and Events*, 3–19. Oxford: Oxford University Press.
———. 1980b. “Freedom to Act.” In *Essays of Actions and Events*, 63–81. Oxford: Oxford University Press.
Dreyfus, Hubert L. 2002. “Intelligence without Representation: Merleau-Ponty’s Critique of Mental Representation.” *Phenomenology and the Cognitive Sciences* 1: 367–83.
Fridland, Ellen. 2014. “They’ve Lost Control: Reflections on Skill.” *Synthese* 191 (12): 2729–50.
Hayhoe, Mary M., and Constantin A. Rothkopf. 2011. “Vision in the Natural World.” *Wiley Interdisciplinary Reviews: Cognitive Science* 2 (2): 158–66. doi:10.1002/wcs.113.
Hommel, B. 2000. “The Prepared Reflex: Automaticity and Control in Stimulus-Response Translation.” In *Control of Congitive Processes: Attention and Performance XVIII*, 247–73. Cambridge, MA: MIT Press.
James, William. 1890. *The Principles of Psychology, Volume 1*. Boston, MA: Henry Holt and Co.
Kentridge, R. W. 2011. “Attention Without Awareness: A Brief Review.” In *Attention: Philosophical and Psychological Essays*, edited by Christopher Mole, Declan Smithies, and Wayne Wu, 228–46. New York: Oxford University Press.
Land, M. F., and P. McLeod. 2000. “From Eye Movements to Actions: How Batsmen Hit the Ball.” *Nature Neuroscience* 3 (12): 1340–45. doi:10.1038/81887.
Mahon, Brad, and Wayne Wu. Forthcoming. “The Cognitive Penetration of the Dorsal Visual Stream?” In *Cognitive Penetration*, edited by John Zeimbekis and Athanasios Raftopoulos. Oxford: Oxford University Press.
Mann, David L., Wayne Spratford, and Bruce Abernethy. 2013. “The Head Tracks and Gaze Predicts: How the World’s Best Batters Hit a Ball.” *PloS One* 8 (3): e58289. doi:10.1371/journal.pone.0058289.
McDowell, John. 2007. “What Myth?” *Inquiry* 50 (4): 338–51.
Mele, A. 1992. *Springs of Action: Understanding Intentional Behavior*. Oxford: Oxford University Press.
Milner, A. David, and Melvyn A. Goodale. 1995. *The Visual Brain in Action*. Oxford: Oxford University Press.
Mole, Christopher. 2013. *Attention Is Cognitive Unison: An Essay in Philosophical Psychology*. OUP USA.

Neumann, O. 1987. “Beyond Capcity: A Functional View of Attention.” In *Perspectives on Perception and Action*, 361–94. Hillsdale: Lawrence Erlbaum Associates, Publishers.
Papineau, David. 2013. “In the Zone.” *Royal Institute of Philosophy Supplement* 73: 175–96.
Peacocke, Christopher. 1979. *Holistic Explanation*. Clarendon Press.
Schlosser, Markus E. 2007. “Basic Deviance Reconsidered.” *Analysis* 67 (295): 186–94.
Schneider, W., and R.M. Shiffrin. 1977. “Controlled and Automatic Human Information Processing: I. Detection, Search and Attention.” *Psychological Review* 84 (1): 1–66.
Schütz, Alexander C., Doris I. Braun, and Karl R. Gegenfurtner. 2011. “Eye Movements and Perception: A Selective Review.” *Journal of Vision* 11 (5): 9. doi:10.1167/11.5.9.
Shepherd, Joshua. 2014. “The Contours of Control.” *Philosophical Studies* 170 (3): 395–411.
Siegel, Susanna, and Nicholas Silins. forthcoming. “Attention and Perceptual Justification.” In *Festschrift for Ned Block*, edited by Adam Pautz and Daniel Stoljar. MIT Press.
Stanley, Jason, and John W. Krakauer. 2013. “Motor Skill Depends on Knowledge of Facts.” *Frontiers in Human Neuroscience* 7: 503. doi:10.3389/fnhum.2013.00503.
Strawson, Galen. 2003. “Mental Ballistics or the Involuntariness of Spontaniety.” *Proceedings of the Aristotelian Society* 103 (3): 227–57.
Wolpert, D.M., and Z. Ghahramani. 2000. “Computational Principles of Movement Neuroscience.” *Nat Neurosci* 3 Suppl: 1212–17.
Wu, Wayne. 2008. “Visual Attention, Conceptual Content, and Doing It Right.” *Mind* 117 (468): 1003–33.
———. 2011. “Confronting Many-Many Problems: Attention and Agentive Control.” *Noûs* 45 (1): 50–76.
———. 2013. “Mental Action and the Threat of Automaticity.” In *Decomposing the Will*, edited by Andy Clark, Julian Kiverstein, and Tillman Vierkant, 244–61. Oxford University Press.
———. 2014. *Attention*. Abingdon, UK: Routledge.
Yarbus, A. L. 1967. *Eye Movements and Vision*. Plenum press.