Synthese
DOI 10.1007/s11229-014-0617-9

Automatically minded

Ellen Fridland

Received: 17 April 2014 / Accepted: 24 November 2014

Abstract It is not rare in philosophy and psychology to see theorists fall into dichotomous thinking about mental phenomena. On one side of the dichotomy there are processes that I will label "unintelligent." These processes are thought to be unconscious, implicit, automatic, unintentional, involuntary, procedural, and non-cognitive. On the other side, there are "intelligent" processes that are conscious, explicit, controlled, intentional, voluntary, declarative, and cognitive. Often, if a process or behavior is characterized by one of the features from either of the above lists, the process or behavior is classified as falling under the category to which the feature belongs. For example, if a process is implicit this is usually considered sufficient for classifying it as "unintelligent" and for assuming that the remaining features that fall under the "unintelligent" grouping will apply to it as well. Accordingly, if a process or behavior is automatic, philosophers often consider it to be unintelligent. It is my goal in this paper to challenge the conceptual slip from "automatic" to "unintelligent". I will argue that there are a whole range of properties highlighted by the existing psychological literature that make automaticity a much more complex phenomenon than is usually appreciated. I will then go on to discuss two further important relationships between automatic processes and controlled processes (C-processes) that arise when we think about automatic processes in the context of skilled behavior. These interactions should add to our resistance to classifying automaticity as unintelligent or mindless. In Sect. 1, I present a few representative cases of philosophers classifying automatic processes and behaviors as mindless or unintelligent. In Sect. 2, I review trends in the psychology of automaticity in order highlight a complex set of features that are characteristic, though not definitive, of automatic processes and behaviors. In Sect. 3, I argue that at

E. Fridland (✉)
Department of Philosophy, King's College London, Room 507, Philosophy Building,
Strand, London WC2R 2LS, UK
e-mail: ellen.fridland@kcl.ac.uk

Published online: 30 January 2015

least some automatic processes are likely cognitively penetrable. In Sect. 4, I argue that the structure of skilled automatic processes is shaped diachronically by practice, training and learning. Taken together, these considerations should dislodge the temptation to equate "automatic" with "unintelligent".

Keywords Cognition · Intelligence · Automaticity · Skill · Cognitive penetrability · Selective attention

It is not rare in philosophy and psychology to see theorists fall into dichotomous thinking about mental phenomena.[1] On one side of the dichotomy there are processes that I will label "unintelligent." These processes are thought to be unconscious, implicit, automatic, unintentional, involuntary, procedural, and non-cognitive.[2] On the other side, there are "intelligent" processes that are conscious, explicit, controlled, intentional, voluntary, declarative, and cognitive. Often, if a process or behavior is characterized by one of the features from either of the above lists, the process or behavior is classified as falling under the category to which the feature belongs.[3] For example, if a process is implicit this is usually considered sufficient for classifying it as "unintelligent" and for assuming that the remaining features that fall under the "unintelligent" grouping will apply to it as well. Accordingly, if a process or behavior is automatic, philosophers often consider it to be unintelligent.

In this paper, I do not attempt to present comprehensive definitions or analyses of the concepts "intelligent" and "unintelligent." Rather, I will assume that these concepts fit somewhat neatly into the familiar mind-body dichotomy, which has dominated philosophical thinking for centuries. I will assume that "intelligent" processes are those that need to be cashed out in semantic or psychological terms—prototypically, propositional states that are conceptual, compositional, recombinatorial, generalizable, and that can enter into logical reasoning. While "unintelligent" processes can be explained according to mechanistic, causal principles, without need to appeal to concepts, semantics, or personal-level intentional states. As such, the distinction between "intelligent" and "unintelligent" is, at its root, just the distinction between the psychological and the physical/biological, the semantic and the causal, the mind and the brain.

1 For similar claims see Logan (1985), Karmiloff-Smith (1994), Tzelgov (1997a, b), Hommel (2007).

2 Philosophers and psychologists use the labels "cognitive" and "noncognitive" in different ways: for psychologists and cognitive scientists "cognitive" usually means something like "mental." In this way, perception, memory, learning, etc. are all cognitive phenomena. Philosophers, on the other hand, use "cognitive" to mean something like "intelligent" such that it makes sense to ask whether perception is cognitive or cognitively penetrable as does Fodor (1983), Pylyshyn (2000), Prinz (2006), Siegel (2010). It is in order to avoid confusion, that I use "intelligent" and "unintelligent" instead of "cognitive" and noncognitive" as my general categories of classification above.

3 This kind of move is not altogether accidental since one of the most influential theories in the philosophy of mind, Fodor's (1983) *Modularity of Mind*, is committed to precisely this clustering of features. A hallmark of modularity is that a certain class of processes exhibit a suite of characteristics are indicative of their modularity. I will not argue against modularity in this article but, in focusing on automaticity, I will show that dual-mode theories often overlook the complexity of the relationships between various often but not always co-occurring features.

Therefore, my claim that many automatic processes are not unintelligent is a claim that such processes cannot be understood in purely brute-causal, mechanistic terms. However, I should be clear that I do not thereby mean to claim that these automatic processes are intelligent. I want to resist this conclusion since it is not obvious to me that the best account of automatic processing is to be given in terms of propositional thought or personal-level intentions. Specifically, my claim is that some automatic processes are not unintelligent since they bear robust, systematic relationships to personal-level intentional contents. And this means that at least some automatic processes cannot be accounted for in non-semantic, causal terms. In this way, we can see that at least some automatic processes are not like billiard balls colliding, water running down hill, circulation or digestion. However, being sensitive to semantic content does not entail that these states are themselves conceptual, compositional, generalizable or that they can enter into logical reasoning. As such, the automatic processes that I consider do not seem to naturally belong to the "intelligent" category either. The main conclusion I will draw from these considerations, then, is that many automatic processes are neither "unintelligent" nor "intelligent" in the normal way. As such, reflecting on automaticity challenges us to think about intermediate categories of cognition that possess some but not all the properties of higher-order cognition but which are not easily accounted for in causal or mechanical terms. Accordingly, thinking about automatic but not unintelligent processes forces us to reconsider the general categories that we apply to the mind.

I should also note that my project here is largely conceptual. That is, though I appeal to empirical work, especially in Sects. 2 and 4, my main point is that there are a whole host of conceptual reasons for resisting the move from "automatic" to "unintelligent". Moreover, in spelling out the interactions between automatic and C-processes in Sects. 3 and 4, my goal will be to highlight the semantic coherence that must characterize these interactions, if they are to occur, no matter what their underlying neurobiological implementation. For present purposes, it is enough that I provide plausible descriptions of automatic-semantic interactions that challenge familiar dichotomies, which philosophers often fall into when categorizing mental states and events.

In brief, it is my goal in this paper to challenge the conceptual slip from "automatic" to "unintelligent". I will argue that there are a whole range of properties highlighted by the existing psychological literature that make automaticity a much more complex phenomenon than is usually appreciated. I will then go on to discuss two further important relationships between automatic processes and controlled processes (C-processes) that arise when we think about automatic processes in the context of skilled behavior. These interactions should add to our resistance to classifying automaticity as unintelligent or mindless. The considerations that I'll present concerning automaticity can be taken as having wide cautionary consequences for dichotomous thinking in general. In Sect. 1, I present a few representative cases of philosophers classifying automatic processes and behaviors as mindless or unintelligent. In Sect. 2, I review trends in the psychology of automaticity in order highlight a complex set of features that are characteristic, though not definitive, of automatic processes and behaviors. In Sect. 3, I argue that at least some automatic processes are likely cognitively penetrable. In Sect. 4, I argue that the structure of skilled automatic processes is shaped diachronically by practice, training and learning. Taken together, these considerations should dislodge the temptation to equate "automatic" with "unintelligent".

1 Philosophy and automaticity

In this section, in order to illustrate the philosophical trend of equating "automatic" with "unintelligent, " rather than providing a comprehensive review of the literature, I'll simply point to a few representative instances of this style of thinking.

Stanley (2011), in *Know How*, appeals to automatic triggering mechanisms in order to solve Gilbert Ryle's famous regress (Ryle 1949). Ryle's regress maintains that not all knowledge can be propositional in nature since, if it were, a second proposition would be needed in order to direct the appropriate application of the first proposition, and then a third to direct the appropriate application of the second, and so on *ad infinitum*.[4] For example, if my knowing how to play a sequence of notes on the piano were reducible to propositional knowledge then in order to know how to play those notes, I'd have to know a particular proposition governing how and when to play them. However, to know how and when to apply that proposition correctly, I'd need to know another proposition specifying how and when I should determine the proper application of the first proposition, and so on *ad infinitum.* Stanley posits that instead of an additional proposition, the reasonable intellectualist[5] can simply appeal to *automatic mechanisms* to do the triggering of propositional knowledge and thus solve the regress.[6] Following Fodor (1983), Stanley conceives of automatic mechanisms as unintelligent. If he didn't, the regress could not be solved using them since this would imply that automatic mechanisms were intelligent but not propositional—a position that is in inconsistent with intellectualism.

Similarly, Dreyfus (Dreyfus and Dreyfus 1986; Dreyfus 2002), who approaches expertise from the opposing school of thought unwittingly falls into the same dichotomies of intelligent, rational, conscious, controlled processes on one hand and unintelligent, automatic, passive, unconscious routines on the other.

On Dreyfus's account, expert skill is a case of non-deliberative, non-reflective, arational, atheoretical action. [7] It is the novice, the advanced beginner, and the competent person, i.e., those persons at the lower stages of skill development, who need to think. In contrast, the expert just does. Dreyfus writes, "The expert driver, generally without any awareness, not only feels when slowing down on an off-ramp is required, he or she knows how to perform the appropriate action without calculating and comparing alternatives. *What must be done, simply is done*" (2002, p. 372, emphasis in original). In Dreyfus's hands, the automaticity of expert action highlights its opposition to the thoughtful, calculating, deliberate, controlled, conscious qualities of rational

[4] Ryle writes, "The consideration of propositions is itself an operation the execution of which can be more or less intelligent, less or more stupid. But if, for any operation to be intelligently executed, a prior theoretical operation had first to be performed and performed intelligently, it would be a logical impossibility for anyone ever to break into the circle (1949, p. 30)."

[5] Intellectualism is the view that the cognitive or intelligent aspects of skill or know how are reducible to knowing an appropriate proposition governing that skill.

[6] Stanley writes, "[T]riggering representations is something done by an input systems rather than a central system, by a module rather than a central processor. Such triggering is something we do automatically" (2011, p. 16).

[7] As Dreyfus and Dreyfus (1986) write, "*competent performance is rational; proficient is transitional; experts act arationally"* (p. 36, emphasis in original).

processes. As Sutton et al. (2011) describe the general trend of anti-intellectualists to sap the intelligence out of automatic, expert skill:

> Although they start from a plausible rejection of the idea that action is driven by explicit rules or inner blueprints accessed by way of conscious reflective deliberation, both philosophers and scientists go too far in the other direction by treating expertise as entirely intuitive, the sole product (as Dreyfus put it) of "attractive and repulsive forces drawing appropriate activity out of an active body". In thus taking embodied activity right out of the psychological realm, these theories paradoxically reinforce dichotomies between doing and knowing, or acting and thinking, which we might have hoped to overcome (p. 92).

More recently, we can see this trend continued by Di Nucci (2013) who describes automatic behaviors, habitual actions, skilled activities, and conventional behaviors as mindless. In fact, Di Nucci goes as far as to title his recent book on automatic behaviors, *Mindlessness*. Again, highlighting the dichotomy between automatic processes and behaviors on the one hand and intelligent, mindful ones on the other. I hope that these three brief examples of automaticity in the philosophical literature are sufficient to expose the general trend of philosophers to equate "automatic" with "unintelligent."

2 Automaticity: a brief overview of the psychological literature

Automaticity is one of the most explored phenomena in all of psychology (Bargh et al. 2012). Automatic processes have been investigated in the domains of "perception, decision making, moral judgments, close-relationships, emotional processes, face perception and social judgment, motivation and goal-pursuit, conformity, behavioral contagion, embodied cognition, and the emergence of higher-level automatic processes in early-childhood" (Bargh et al. 2012). What has emerged is a view of human cognition where "most moment to moment psychological life occurs through nonconscious means" (Bargh and Chatrand 1999). That is, "most of a person's life is determined not by their conscious intentions and deliberate choices but by features of the environment that operate outside of conscious awareness and guidance" (Bargh and Chatrand 1999, p. 1).[8] The ubiquity of automatic, psychological processes makes elucidating the nature of automaticity imperative for understanding the human mind.

The consensus view that has developed in psychology is that automatic processes are much more complex than initially presumed.[9] Originally, automaticity was defined in broadly contrastive and simplistic terms. For instance, an automatic process was defined by Shiffrin and Schneider (1977) as "the activation of a sequence of nodes that (a) nearly always becomes active in response to a particular input configuration and (b) the sequence is activated automatically without the necessity of active control or attention by the subject" (p. 2). That is, Shiffrin & Schneider emphasized the fact

[8] See Logan (1985), Tzelgov (1999), Moors and Houwer (2006), Di Nucci (2013), Wu (2013a), for similar claims.

[9] In order to identify a process or behavior as potentially automatic in the first place, I think that the best we can do is to begin with the processes and behaviors that psychologists have generally studied as automatic.

that automatic processes were unintentional or obligatory, and that they do not require attention or control. Posner and Snyder (1975a, b) "chose the features unintentional, unconscious and producing no interference as necessary features" (Moors and Houwer 2006, p. 298) of automaticity. Hasher and Zacks (1979) emphasized the fact that automatic processes operated in the absence of attention. In general, as Bargh (1994) argues, automaticity was defined largely by appeal to four major features: (1) unintentional, (2) unconscious, (3) uncontrolled/uncontrollable, and (4) attention-independent or effortless.[10]

However, as individual automatic phenomena were investigated, it became clear that hardly any processes could be characterized neatly by the co-occurrence of these features (Logan 1985; Bargh 1992, 1994; Moors and Houwer 2006; Wu 2013a). As early as 1988, Schiffrin wrote that "there do not seem to be any simple defining features of automatic and attentive processes that can be applied in complete generality" (p. 765). And in 1994, Bargh writes, "It has since become increasingly clear that mental processes at the level of complexity studied by social psychologists are not exclusively automatic or exclusively controlled, but are in fact combinations of the features of each. In cognitive psychology, evidence was accumulating that no process was purely automatic by the four-criteria standard" (p. 3). As Moors and De Houwer elucidate,

> Through the years, some proponents of the capacity view came close to undermining a fixed, feature-based definition of automaticity. Schneider et al. (1984, pp. 20–21) examined 12 criteria and concluded that none was necessary or sufficient for the distinction between automatic and nonautomatic processes; however, they still pointed to control and resource demands as being the least problematic. Shiffrin (1988) evaluated several criteria for automaticity, ultimately retaining not one as generally applying to all automatic processes (2006, p. 299).

For instance, the commitment that automatic processes and behaviors do not require attention (LaBerge and Samuels 1974; Posner and Snyder 1975a, b; Shiffrin and Schneider 1977; Schneider and Shiffrin 1977; Hasher and Zacks 1979; Logan 1979) has been considered a contingent feature of automaticity since the 1980s. The most oft-cited experiments, which challenge the notion of automatic processes as essentially attention independent involve demonstrating that the Stroop task,[11] a paradigmatically automatic task, can be modulated by drawing attention away from the target stimulus (Francolini and Egeth 1980; Logan 1980; Kahneman and Henik 1981; Kahneman and Chajczyk 1983; Hoffman et al. 1983; Tzelgov et al. 1997). This surprising finding shows that at least some automatic behaviors are not, as was previously thought, attention independent.

The tight connection between being unconscious and being automatic (Posner and Snyder 1975a, b; Neumann 1984; Norman and Shallice 1986) has likewise been challenged from various angles. For one, Tzelgov (1997a, b, 1999); Tzelgov et al. (1997, 1999) has pointed out that the outputs of automatic processes are often conscious.

[10] "Effortless" refers to the subjective, qualitative feeling that accompanies a task that does not require attention.

[11] The Stroop task tests for effects of non-target features on the understanding of a word. E.g., reading the word "red" when the letters R-E-D are purple in color results in a slower reaction time.

In addition, the triggering of automatic processes and behaviors often occurs as the result of a conscious episode. This is especially evident in cases of skill and other goal-dependent activities where conscious intentions regularly initiate automatic processes. Bargh (1994), Bargh and Chatrand (1999), Bargh et al. (2012) even goes as far as classifying automatic behaviors into three basic groups, one of which is post-conscious. These post-conscious automatic processes require a conscious experience in order to trigger the automatic process.

Further, Sheets-Johnstone (2012) [12] and Tzelgov et al. (1997) have both pointed out the implausibility of agents remaining necessarily unaware of their skilled, habitual, automatic behaviors. They point out that if we combine the ubiquity of automaticity in skill and habit with the commitment that automatic behaviors remain unconscious then we would be forced to conclude that we are unaware of most of our skilled or habitual behaviors. Not only does this seem on its face plainly false but the conclusion is logically unwarranted. Plainly, it does not follow from the fact that an activity *can occur* in the absence of consciousness that the activity *can only occur* in the absence of consciousness. For instance, just because I *need not* consciously attend to my teeth brushing in order to succeed at brushing my teeth, this does not entail that I *cannot* be conscious of my teeth brushing if I choose to be. The same goes for driving, biking, swimming etc. Surely, I can be and often am conscious of performing the automatic behaviors implicated by these habitual and skillful activities.

It may be objected that consciousness interferes with skill and so is incompatible with automaticity. In fact, researchers interested in expertise often point out that consciousness interferes with skilled, automatic routines or with being, "in the zone" (Dreyfus 2007; Beilock 2010; Di Nucci 2013; Papineau 2013). They appeal to both popular and empirical evidence, which suggests that conscious attention to the means or mechanics of movements constitutive of skilled action can undermine the fast, fluid, successful performance of that skill. Papineau (2013) defines this phenomenon, which he calls "The Yips," as "what happens if you start thinking explicitly about the bodily movements required for some sporting performance" (p. 4). We should note that though it does seem that conscious attention to particular aspects of one's automatic movements can disrupt skill it does not follow from this that consciousness of automatic processes, in general, undermines successful performance. The most that we can conclude is that conscious attention to particular aspects of the mechanics of one's automatic processes is counterproductive for successfully executing some skills. However, a general consciousness that one is, e.g., down-shifting, or swinging a golf club, or reaching one's right arm backwards when swimming the back-stroke, that is, a general consciousness of one's automatic behaviors, can and often does figure in our conscious awareness without any hint of disruption to the skilled action.[13]

[12] Sheets-Johnstone writes, "when Luria speaks of the automatization of movement, it is important to point out that he is describing the way in which a single impulse is sufficient to activate a kinetic melody, and not asserting that one is unaware of writing one's name, that one is unconscious of doing so, or that one can nod off while the process continues by itself" (p. 52).

[13] See Montero (2010) for a defense of this position and a critique of the empirical evidence that is usually appealed to in order to support the view that attention and skill are incompatible.

Additionally, the notion that automatic processes must be unintentional or obligatory has likewise proven insufficient to capture the full array of automatic processes and behaviors. For instance, habits and skills are behaviors, which paradigmatically involve automatic processes but which are also, very often, intentional actions. For instance, being a gymnast and being in the presence of a balance beam is not sufficient for triggering the gymnast's automatic back-handspring-behavior. That is, the automatic routines that constitute the gymnast's skill are not involuntarily or mandatorily set off by the presence of a relevant stimulus, even if various environmental conditions are necessary for triggering the automatic behaviors. Skills are performed as the result of an intention to perform them and thus serve as clear counterexamples to the notion that automatic processes and behaviors are unintentional, mandatory or involuntary.[14] An even more bold connection between intentions and automatic processes has been forwarded by Hommel who, following Bargh (1989), has argued that "automatic processes are often, perhaps always, contingent on current intention and task goal" (2007, p. 166).

Wu (2013a) has recently tried to revive what he calls the "simple connection" between automaticity and intentionality (though, on Wu's account, intentionality is called "control"). On this view, any feature of a given process is either intentional or automatic but not both. Wu avoids the above challenge by claiming that for a feature to be intentional or controlled, that very feature has to figure in the content of an intention. As Wu explains this constraint with respect to automatic attention, "S's attention to X is automatic in respect of its feature F iff S's attending to X with feature F is not a result of an intention to do so in the F way" (Wu (forthcoming), p. 32). So, though, in general, e.g., my piano playing may be intentional, as long as my intention does not specify the particular way in which my fingers should move in, say, the second bar of the sonata, then the particular way in which I move my fingers in that section of the piece can be classified as automatic. More precisely, the automatic features are the ones not explicitly specified by my intention to play. As such, my overall piano playing is intentional though there are automatic features that constitute it. According to Wu, this entails that most actions will involve both automatic and intentional elements but this does not entail that the very same elements will be both automatic and intentional.

Though Wu's characterization of automaticity ought to be commended for its ability to differentiate between the intentional and automatic components of action, I think the simple connection fails for two reasons. 1- it does not provide us with an argument for why intentionality ought to play a singular role in differentiating between automatic and controlled features of processes. That is, Wu does not explain why it is that not being specified in the content of an intention rather than, say, being resistant to dual-task interference (Posner and Snyder 1975a, b; Logan 1979) or being autonomous (Bargh 1992), should act as the defining feature of automaticity.

Second, upon reflection, it appears that some seemingly automatic features of a process *can* result from an intention to perform them in a particular way and further, some intentional aspects of an action can themselves be automatic. As such, the simple connection does not survive scrutiny. For instance, if before playing the piano I

[14] See Logan (1985) for similar considerations.

visualize myself playing a piece, note by note, finger stroke by finger stroke, and then sit down to play with the intention to play the piece "like that" then the demonstrative content of my intention will include those specific moment-to-moment processes as intentional contents—but these aspects of my playing can still unfold in a way that is characterized by many features of automaticity (e.g., fast, parallel, effortless, resistant to dual-task interference, etc.). As such, it does not seem that being involved in one's intentional content should alone exclude some feature of a process from qualifying as automatic.

Moreover, it would seem that habitual actions are often performed as a result of automatic intentions. For example, I can intend to brush my teeth as I approach my bathroom sink in the morning but that intention is often formed in what seems to be an automatic way—it is triggered involuntarily, it does not require consciousness or explicit attention, it is robust to stressors, etc. But it is an intention all the same. After all, I don't brush my teeth accidentally. If this is correct, then it seems that intentional features of an action cannot be specified in contrast to automatic ones since some intentions can themselves be automatic. As such, it appears that though the simple connection highlights an important feature of many automatic processes, it is not itself sufficient as a general characterization of automaticity.

Lastly, Logan (1980, 1985), Bargh (1994), Tzelgov (1997a, b), and Hommel (2007) have all argued that being uncontrolled or uncontrollable is hardly a universal property of automatic processes. One important set of evidence supporting this claim comes from studies of unwilling racists who are able to overcome their automatic biases when motivated to do so (Devine 1989; Fiske 1989; Fazio et al. 1995; Dunton and Fazio 1997; Kawakami et al. 2000; Blair et al. 2001; Olson and Fazio 2004). That is, despite the fact that particular stimuli automatically trigger racist stereotypes or associations, subjects who were committed to egalitarian ideals were at least sometimes able to control and thus overcome the application of their automatic biases. Likewise, Logan (1985) has argued that if we begin by considering skilled performance rather than, for example, perceptual processes, we are very unlikely to conclude that automatic processes and behaviors lack control. In fact, the very opposite seems to be true: the more expert one is at a skill, the more automatic that skill becomes and the more controlled it is as well.

At this point, some semantic clarification is in order. Various theorists use the word "control" differently. Tzelgov (1997a, b), following Logan (1980, 1985), uses "control" to mean the "sensitivity of a system to changes in inputs (Tzelgov 1997a, b, p. 5). LaBerge and Samuels (1974) use "control" to mean the overcoming of a disturbance, similar to the way philosophers use "guidance control" (Frankfurt 1978; Fischer 1982; Fischer and Ravizza 1998). Both of these ways of defining "control" make it reasonable that automatic processes can be controlled. However, others, such as Shiffrin and Schneider (1977), Schneider and Chein (2003), and Wu (2013a) seem to use "control" to designate, by definition, the category that is to be contrasted with automaticity. This is not to say that various mental processes cannot have both automatic and controlled features, but it is to say that if a feature is automatic then it is not controlled. In this way, such theorists could agree that automatic processes are sensitive to information and responsive to disturbance but still insist that "control processes" serve as the contrast class to automatic ones. Both definitions of "control" have their uses and I will distinguish them by calling the feature that refers to guidance control, “control,” and by calling the category of states, processes, and behaviors that are meant to be contrasted with automaticity, “C-processes.”

Presently, in psychology, the simple view of automaticity, with a few notable exceptions, has almost universally been rejected.[15] Instead, various theorists have chosen various ways to isolate the notion of automaticity. Schneider and Chein (2003) list seven features that need to be explained by a theory of automaticity:

1. that extended consistent training is required in order to develop automatic processing, while controlled processes can be established in a few trials and under varied mapping situations. (p. 528)
2. automatic processing is fast and parallel, while controlled processing is slow and serial.
3. that automatic search requires little effort and can operate in high workload situations, whereas controlled processing requires substantial effort and interferes with other controlled processing tasks.
4. that automatic processing is rather robust to stressors
5. the difference in cognitive control that can be applied to automatic and controlled processes. Specifically, once a process becomes automatic, it becomes difficult to control.
6. The degree of learning is dependent on the amount and type of controlled processing, while there is little learning in pure automatic processing.
7. automatic attention response is dependent on the priority assigned to a stimulus itself, rather than on the context in which the stimulus occurs.

Notably, these seven characteristics are not presented in terms of necessary and sufficient conditions for automaticity but as features that need to be explained by any adequate theory. Others have also been seduced by feature lists; Shiffrin (1988) listed ten features of automaticity,[16] Palmeri (2002) lists thirteen.

In light of the dazzling complexity involved in the notion of automaticity, I shall pursue a conservative strategy and endorse the decomposable, gradual view of automaticity that is advocated by Moors and Houwer (2006) who follow Logan (1985). This conception of automaticity allows for a notion of automaticity to be circumscribed by a set of characteristic, but not necessary or sufficient features. These features will have their own time course for development and so will characterize various automatic processes to various degrees during different times in learning. This approach allows us to retain a concept of automaticity, which is both operationalizable and theoretically

15 Bargh (1992) has proposed that the central feature that all automatic processes and behaviors share is autonomy. “He defined an autonomous process as one that, once started (and irrespective of whether it was started intentionally or unintentionally), runs to completion with no need for conscious guidance or monitoring” (Moors and Houwer 2006, p.301). But, as I argued above in reference to Wu (2013a), and as Moors and DeHouwer make clear, it isn’t clear why any one feature, as opposed to any other feature, should have priority as being definitive of automaticity.

16 See Wu (2013a) for a report that Shiffrin has since given up on the possibility of any feature list as definitive of automaticity.

useful for guiding our thinking.[17] But it forces us to retain this concept in a qualified and contingent manner. The features listed above by Schneider and Chein (2003) and Bargh (1994) highlight that automatic processes are generally fast, parallel, efficient, effortless, resistant to stressors, relatively unaffected by cognitive load, do not require conscious guidance, are not the result of specific intentions to deploy them, etc. However, this does not mean that every automatic process will be equally characterized by each of these features at all times.

3 The cognitive penetrability of automatic processes

In this section, I will draw the reader's attention to another important connection that at least some automatic processes seem to hold to intentional states. This connection, I claim, makes classifying such automatic processes as "unintelligent" questionable. Specifically, the connection that I will explore is what has been called, in discussions of visual perception, cognitive penetrability. Though I am not in this instance concerned with visual perception, I hope that in elucidating the nature of cognitive penetrability, I will be able to apply the concept to action control and attention.

To start, I'd like to point out that the claim that automatic behaviors or processes are affected by intentional states can be take in one of two ways: (1) that intentional states trigger or initiate automatic behaviors, i.e., that there is a causal connection between an intentional state and an automatic process or (2) that automatic behaviors and processes are sensitive to the intentional content of intentional states, that is, that automatic states are cognitively penetrable. I think that both of these ways of being affected by intentional states are true of at least some automatic behaviors and processes but since the former is a weaker kind of claim, and seems to be relativity uncontroversial, I'll focus on the latter. That is, I will focus on defending the proposal that at least some automatic behaviors are sensitive to the intentional content of cognitive states and, thus, at least some automatic processes are cognitively penetrable. I will start by working through some of the difficulties associated with an adequate definition of cognitive penetrability and, in doing so, I will defend the need for semantic coherence as a criterion of cognitive penetrability. Next, I will consider a potential case of cognitive penetrability by focusing on automatic action control. Lastly, I will demonstrate that automatic selective attention is likely cognitively penetrable, as well.

We should notice that for a legitimate case of cognitive penetrability, we should not want just any kind of connection between an intentional state and psychological processing but a particular kind of connection: a connection where there is a meaningful or semantic interaction between content and processing. That is, for a good definition of cognitive penetrability, we need a systematic way to differentiate cases that fall under option (1) above from cases that fall under option (2). As it turns out, spelling out exactly how to do this proves rather difficult. Still, in the following section, I will

[17] To operationalize this concept, one would have to select the processes or features of automaticity that one was most interested in investigating. By combining various processes and features, one could investigate their overlap. This would make the concept of automaticity operationalizable according to various dimensions depending on the concerns of the researcher.

insist that it is semantic coherence or intelligibility that we need to keep in central focus, if we are to have an adequate account of cognitive penetrability.

First off, we should admit that the following should not qualify as an instance of cognitive penetrability, even though there is a clear causal connection between a cognitive state and perceptual processing: I want to find out who is at the door and I believe that if I turn my head then I will see who is at the door and, thus, find out who is there. So, I turn my head. As a result of my turning my head, the processing of my early visual system changes—whereas my visual system was in the first instance processing the features of my computer and dining room table, it is now processing the features of the door and the person who is standing at it. As a result, my visual experience changes. This kind of simple causal connection between intentional states and perceptual experience, if it were to qualify as cognitive penetrability, would make the concept ubiquitous and trivial. As such, this simple causal connection cannot be what cognitive penetrability is all about.[18]

For cognitive penetrability, the intentional content of a cognitive state should bear some sort of direct, meaningful relationship to the perceptual processing in question. That is, the phenomenal properties that a perceptual system outputs should be impacted by cognitive content in virtue of the content's semantic value. For instance, building on the previous example, if early vision were cognitively penetrable, then the way in which the visual system processes the phenomenal properties of the person standing at the door should potentially differ in a case where my beliefs or knowledge differ. So, for instance, if I believe the person at the door is a Girl Scout selling cookies, then perhaps my visual system will be more likely to produce green qualia that are rich or saturated. Compare this to a case where I believe that the person at the door is my mother, who has no special preference for wearing green. Then, presumably, my visual system will have no preference for processing greens over any other color. In such a case, early perceptual processing would be directly sensitive to the content of my belief state—that the person at the door is a Girl Scout and should be wearing green. Of course, we have to rule out cases where I selectively attend to the green color of the person's wardrobe, since, that would change the input and not the processing of early vision. For a true case of cognitive penetrability, e.g., it should be the early visual system itself favors green over another color, in virtue of my belief about the Girl Scout.

In an attempt to hone in on this precise connection between cognitive content and perceptual processing, Zenon Pylyshyn (2000) forwards the following definition of cognitive penetrability:

> If a system is cognitively penetrable then the function it computes is sensitive, in a semantically coherent way, to the organism‘s goals and beliefs, that is it can be altered in a way that bears some logical relation to what a person knows (p. 343).

This definition highlights the fact that cognitive penetrability requires not merely sensitivity to the presence of an intentional state but sensitivity to the intentional

[18] See Pylyshyn (2000) and Stokes (2013) for elucidation of this point.

content or semantic value of the intentional state. That is, for cognitive penetrability, we need a semantic or logical relationship between content and processing. Anything short of this will necessarily allow in bad cases.

Macpherson (2012) also emphasizes the importance of semantic coherence or what she calls "intelligibility" in her characterization of cognitive penetrability. Instead of defining cognitive penetrability, in order to get a handle on the phenomenon, Macpherson defines its opposite—cognitive impenetrability. She writes,

> Thus, perceptual experience is cognitively impenetrable if it is not possible for two subjects (or one subject at different times) to have two different experiences on account of a difference in their cognitive systems which makes this difference intelligible when certain facts about the case are held fixed, namely, the nature of the proximal stimulus on the sensory organ, the state of the sensory organ, and the location of attentional focus of the subject.

For Macpherson, we see that not only does there have to be a cognitive content that impacts perceptual experience when all other relevant visual factors are held fixed, but that cognitive content has to make the difference in perceptual experience intelligible. The need for emphasis on semantic coherence or intelligibility is made clear by the following example, which elucidates why a bar set any lower than semantic coherence will prove inadequate.

According to Stokes (2013), cognitive penetrability can be defined as follows:

> (CP) A perceptual experience E is cognitively penetrated if and only if (1) E is causally dependent upon some cognitive state C and (2) the causal link between E and C is internal and mental (p. 650).

However, as Stokes himself admits, this definition suffers from counterexamples that the intelligibility or semantic coherence condition, emphasized by Pylyshyn and Macpherson, adequately keeps out. Stokes writes:

> For instance, an example of Macpherson's might be adjusted and invoked here. Suppose I suffer extreme exam anxiety and I believe that I am about to take an exam. This belief causes, internally, another mental state, namely the pain that accompanies a migraine. This pain further causes, again internally, a series of visual experiences where everything appears in a reddish hue. On the face of it, this scenario satisfies CP. But one may worry that this is not an instance of the relevant cognitive influence, since the causal chain from cognitive state to perceptual experience takes a circuitous route – even if internal and mental. Macpherson's view is that this kind of case re-emphasizes the need for something like Pylyshyn's semantic criterion: since the red hue experience in no way semantically coheres with the belief about the exam, it is not the case of cognitive penetration (2013, p. 650).

I agree with Macpherson that such a case re-emphasizes the need for something like Pylyshyn's semantic coherence criterion (Macpherson 2012, p. 26). After all, it's not as though Stokes has replaced the problematic concept of "semantic coherence" with the uncontroversial or unambiguous concepts of "internal" and "mental." But, in addition, Stokes' definition, in not emphasizing semantic coherence or intelligibility, also allows in bad cases that the semantic coherence requirement keeps out. It would seem that in tallying up the virtues and vices of the preceding definitions, those that emphasize semantic coherence or intelligibility are still on top.

Finally, Wu (forthcoming) has forwarded a computationally minded definition of cognitive penetrability that isolates a particular relationship between information and processing outputs as a sufficient condition for cognitive penetrability (SCP). Wu's definition is as follows:

> (SCP) If C contains information R such that V computes over R where this explains V's outputting O rather than some On given input I, then C cognitively penetrates V.

This empirically minded definition is commendable in its specificity and testability, however, as with Stokes' (2013) definition, we are left without a clear way to differentiate between causal bad cases and semantically coherent good ones. This is especially clear because, for Wu, information is not cashed out in terms of "representational content as tied to semantic value" as it is for Pylyshyn (2000) but, rather, as Shanon-type statistical information (Wu forthcoming, p. 4). Wu wonders if in switching definitions of information he's changing the subject and though he answers, "not really," since there are still clear epistemological implications that follow from his definition, one could also fairly answer that question: "sorta". This is because though statistical information tracks semantic value through correlation, it isn't clear when the relevant system is computing over the information in virtue of its semantic value and when it is computing over the non-semantic features of the information. The non-semantic informational features surely correlate with semantic value but what we need for cognitive penetrability, as we've seen above, is for the system to responds to information in virtue of its semantic value. And since the definition that Wu presents does not provide us with a way in which to isolate those informational states within the cognitive system, which can reasonably be said to be personal-level states and those that simply carry information on a syntactic level, we are left without an adequate characterization of the cognitive side of cognitive penetrability. As such, it seems that Wu's definition of cognitive penetrability also has a difficult time differentiating between the purely causal impacts of informational states and the impacts of those states that follow from their meaning or semantic content—from what a person knows or believes. And since it is the latter that we are after, Wu's definition seems to fall short.

So, although an emphasis on semantic coherence with the intentional contents of a cognitive state may appear to be a somewhat vague requirement of cognitive penetrability, it also seems that any conception of cognitive penetrability that does not emphasize this aspect of the interaction between intentional content and psychological processing will not be in a position to decisively differentiate between the good cases of cognitive penetrability, where an explanatory, intelligible connection is present and the bad kinds of cases, where only a causal connection between the intentional state and psychological processing obtains. As such, it seems that we need to continue to incorporate the notion of semantic coherence into our understanding of cognitive penetrability. And it will be this particular criterion of cognitive penetrability that I will focus on in the following sections.

Before going on to argue that some automatic processes and behaviors are cognitively penetrable, it is worth noting that if one held that automatic behaviors and processes were sensitive to intentional states in a causal way but not in a semantic or intelligible way, one could hold that intentional states trigger or initiate automatic processes or behaviors (option 1 above) and still deny that this entails that automatic processes or behaviors are not unintelligent. The same sort of position does not appear tenable if automatic processes are cognitively penetrable since a systematic connection to semantic, conceptual contents takes a system out of the running for being explained in brute-causal terms. That is, a semantically coherent connection to intentional content would take a system or process out of the space of causes and place it squarely into the space of reasons.

3.1 Action control and cognitive penetrability

Recently, several theorists have insisted that the intentional, goal-directed, cognitive aspects of skilled performances have not been adequately appreciated in the philosophical literature. For instance, Sutton et al. write that, "expertise often requires the rapid switching of modes and styles *within* the performance context. Grooved embodied action must thus be open, under certain circumstances to the influence of explicit knowledge, specific memories and particular decisions" (2011, p.93). Sutton et al. (2011) emphasize that the dynamic, complex context-dependence, flexible nature of skilled performances require that automatic or grooved actions display a sensitivity to generally minded, cognitive processes.[19]

Likewise, in developing a general picture of how intentions impact the automatic behaviors constitutive of sporting skill, David Papineau writes,

> At any stage of an innings, a competent batsman will have assessed the situation and formed a view about how to bat—a conscious *intention* to adopt a certain strategy. As with any intention, this will then set the parameters of the basic action-control system. It will direct that system to bat aggressively, say. It will take one*raft* of conditional dispositions from the batsman's repertoire, and reconfigure that basic control system so that it embodies just those dispositions…Having been so reset, the basic action-control system will then respond accordingly, without any further intrusion of conscious thought" (2013, p. 191).

For Papineau, what's crucial is that the intention to bat in certain way, e.g., aggressively, selects a certain class of dispositions, i.e., aggressive batting dispositions, which will run automatically in response to the appropriate environmental stimuli.[20] Though Papineau does not believe that automatic actions are influenced directly by conscious thought, he does emphasize that the action-control system must be sensitive to the intentions, which guide strategy at the personal-level (2013, p. 191). One can think of

[19] See also Sutton (2007).

[20] For similar claims see Wu (2011, 2013b).

Papineau's proposal in the following manner: strategic intentions initiate the selection of a relevant action folder, which contains a set of automatic motor routines that have been developed through practice, training, drilling, etc. Once the folder is selected, the automatic motor routines within it, which Papineau describes as reflex-like,[21] run autonomously.

To note, I believe that optimal control theory[22] gives us good reason to reject the strict hybrid view of sporting skill forwarded by Papineau. Accordingly, this discussion of Papineua's account should not be seen as an endorsement by me of an account of motor control where automatic motor routines are characterized as more or less fixed reflexes while strategic intentions are cashed out as rational, cognitive, personal-level, conceptual states. Despite my disagreement with Papineau, I think that considering his view is valuable in that it will reveal that even if one holds a fairly conservative view of automatic basic actions, an adequate account of embodied skill will still require processes that are both automatic and cognitively penetrable. As such, my goal in this section is not to endorse Papineau's account of skill but, rather, to show that even on a conservative view of skill, where the motor component is cashed out in brute-causal, mechanistic terms and intentional states are characterized independently at the personal-level, we will still have to posit automatic but cognitively penetrable processes.

Now, returning to Papineaus account, to begin, I'll present a few variations on a theme to help ensure that we have the exact phenomenon that Papineau is isolating clearly in mind:

a. I am driving home from work. The road is clear, the sun is shining, the radio is on. I am done with work for the day. My intention: drive home/relaxed. In this case, my intention dictates that I drive home in a certain manner: a relaxed manner. This means that I won't drive too fast, I won't pass too many cars, I won't change lanes excessively, I'll slow down at a yellow light rather than speeding up, I'll sit back in my seat, steer with one hand, I'll sing along to the radio, etc.
b. I am driving home from work. The road is clear, the sun is shining, the radio is on, I get a call that there is an emergency at home. My intention: drive home/hurry/determined! In this case, my intention dictates that I drive home in a certain manner: hurried and determined. This intention will guide me to drive fast, to pass when I can, to change lanes as often as possible in order to pass cars driving slower than me, to speed up at a yellow light rather than slow down. I'll lean forward in my seat, I'll drive with two hands, I won't notice the music, etc.
c. I am driving home from work. The road is snowy and icy, it is dark, the radio is on. My intention: drive home/carefully/ slowly. In this case, my intention dictates that I drive home in a certain manner: slowly and carefully. This means that I will not speed, I will not pass other cars on the road, I will only change lanes when absolutely necessary, I'll be attentive to the feel of the road under the tires, I'll sit

[21] "We have seen ample reason to think that top-level batting is more like an automatic reflex than any consciously controlled sequence of movements" (Papineau 2013, p.184).

[22] For more, see Sect. 4 below. Also, see Todorov and Jordan (2002), Todorov (2004), Liu and Todorov (2007), Deidrichsen (2007).

on the edge of my seat, I'll grip the steering wheel tightly with both hands, I won't pay attention to the music, etc.

These examples fit well with Papineau's proposal for integrating intention and automatic action. In each case, I have a suite of well-rehearsed automatic motor routines that are selected by my particular intention to drive a certain way. My motor-control system must be sensitive to the intentional content of my intentions and not simply to the presence of an intention, in general. That is, my action-control system must be sensitive to my goal to drive in a relaxed, hurried, or careful manner. Otherwise, I would not be able to initiate the appropriate motor routines in response to my goals. In this case, the individual automatic motor routines are not cognitively penetrable but the action-control system as a whole, which is an automatic control system, is cognitively penetrable.

It's important to see that the action-control system is automatic in two ways: it is triggered automatically with the formation of the appropriate intentions and it runs automatically under the guidance of those intentions. First off, we can see that the action-control system is triggered automatically since after a strategic intention has been formed, there is no further monitoring, guidance, or intention needed in order to activate the action-control system. That is, the agent need not deliberately, consciously, intentionally, or effortfully initiate the working of the action-control system. The system is initiated automatically as a result of the formation of strategic intentions; it is not initiated as a result of an intention to initiate the system.

Moreover, the action-control system must run automatically as well. That is, the action-control system must have the capacity to automatically select the appropriate folder of dispositions, given the agent's goals. After all, it wouldn't be a very effective system if it required the agent to continuously monitor or guide the selection of the appropriate action folder to accomplish the task intended. I should also note that the selection of the appropriate action folder might itself be far from a simple task. For example, it stands to reason that the more skilled an agent, the more action folders she will have related to the same action. This means that not only does the action-control system have to be sensitive to the intentions of the agent but it will also need to select the most appropriate amongst various relevant and closely related action folders.[23] This does not strike me as a straightforward task.

To conclude, we can see that the action-control system is both automatically triggered and automatically run. However, this system must also have a strong, systematic relation to the semantic contents of the agent's goal states at the personal-level. Otherwise, it could not choose the right folder of routines in accordance with the agent's changing goals. Though on Papineau's account, the automatic motor routines contained within the system's folders are not themselves cognitively penetrable, the automatic mechanisms of the action-control system, which select the set of relevant automatic motor routines most certainly are.

On this account, though motor routines are integrated with intentional actions, they only need to be sensitive to semantic content indirectly. Nonetheless, the mechanisms that trigger those motor routines have been shown to be both automatic and cognitively

[23] See Fridland (2013a) for more on the problem of selecting between many fine-grained action routines.

penetrable. This finding may tempt us to generalize that individual automatic processes are only indirectly impacted by cognitive states. However, there is reason to think that such a move is unjustified.

As such, even if we may have reason to reject Papineau's view on empirical grounds, we should notice that for theoretical adequacy, Papineau's account still requires positing processes that are both automatic and intelligent. That is, without the participation of automatic but intelligent processes, the account of skill that Papineau offers could not do the work that he proposes for it to do.

3.2 Selective attention and cognitive penetrability

At least some automatic processes seem to be directly cognitively penetrable. This becomes clear when we consider selective attention.[24] This kind of attention is responsible for selecting the relevant features of an action space to which an agent must attend in order to gather the relevant information, given the agent's goals, plans, and strategies. Importantly, this kind of attention improves with training, is often automatic, and is directly sensitive to the semantic content of intentional states at the personal-level.

It is widely accepted selective attention shows significant improvement as a result of practice and training. As Pylyshyn (2003) has argued, many of the perceptual improvements observed in skilled agents can be attributed to the development of selective attention. For instance, at least one important difference between a novice and an expert hockey player can be detected in how each attends to the puck and, thus, how much relevant information each is able to extract from the same observation (2003, p. 85).[25]

Though selective attention is central to expertise, an agent need not intentionally or consciously guide her attention. That is, an agent need not explicitly or deliberately direct her attention towards the relevant features, areas, or segments of her action space. Rather, selective attention is deployed automatically, once the trained agent initiates intentional action. For example, when I decide to move through a sequence of yoga postures, I need not explicitly direct my attention to the relevant portion of my body or of the mat. This is not to say that I can never consciously direct my attention to these features of my action space but it is to say that as a result of training, I readily selectively attend to the most salient and informative features of my perceptual array without intentionally doing so. Another way of putting this is that the representational content of my intentional state does not include the intention to attend in a specific way.[26]

---

24 See Pylyshyn (2003), Wu (2011, forthcoming), and Fridland (2014a) for more on this kind of automatic, selective attention.

25 "Such anticipation is based, for example, on observing initial segments of the motion of a ball or puck or the opponent's gestures. Except for a finding of generally better attention-orientation abilities visual expertise in sports, like the expertise found in the Chase and Simon studies of chess skill, appears to be based on the nonvisual abilities related to the learned skills of identifying, predicting and therefore attending to the most relevant places" (2003, p. 85).

26 See Wu (2011, forthcoming) for more on the automaticity of selective attention.

In fact, we know that skilled agents cannot always deploy selective attention deliberately because skilled agents often do not know which features they attend to in skilled performance. The fact is that empirical evidence consistently demonstrates experts either failing to report or falsely reporting the cues to which they selectively attend during skillful action (Reed et al. 2010; Wallis 2008; Berry and Broadbent 1984; Brehmer et al. 1980; Reber and Lewis 1977).

The most significant feature of selective attention for my purposes here is that automatic selective attention is very often semantically integrated with the personal-level, intentional states of the agent. That is, even though selective attention is not deployed as a result of some specific intention to attend in this or that way, selective attention is still sensitive to the semantic content of the intentional states, which are guiding an agent's actions or decisions. In fact, we see that in order for automatic, selective attention to function effectively, it must be sensitive, not simply to the existence, but to the content of the reasons, goals, and strategies that the agent possesses at the personal-level.

A very nice example of this kind of cognitive penetrability is presented by Wayne Wu (forthcoming) in his discussion of the Yarbus (1967) experiments. Yarbus presented subjects with a painting entitled, "The Unexpected Visitor." In that painting, four people surround a table looking towards a fifth person who has just entered the room through a door that is being held open by a maid. Yarbus fitted participants with a primitive eye-tracking device and asked them, "How long has the visitor been gone?" In response to this question, participant's eyes moved in long sweeping motions from the person at the door to the people sitting down at the table and back. In contrast, when asked, "How old are the people in this picture?", participants' eyes circled on the individual faces of the people in the room.

The participant's eye saccades responded in a systematic way to the distinct questions they were being asked. Their eyes moved in a manner that would allow them to extract the relevant information from the scene in order to provide answers to the specific questions that they were asked. When it comes to telling someone's age, the face is a great place to attend. When trying to figure out how long someone has been away, the expressions of greeting between people is a better indicator. Of course, it should go without saying that participants did not deliberately direct their eyes in this or that way in order to extract the relevant information from the scene. The eye movements were automatic. As such, these experiments demonstrate that automatic selective attention, as revealed by the eye movements of the participants, is sensitive to the semantic properties of the goals of participants.[27]

From these considerations we can conclude that at least some automatic processes are cognitively penetrable. This is important to keep in mind when thinking about automaticity because it shows that at least some automatic processes are notably different from brute-causal processes like billiard balls colliding, circulation, and digestion. The

[27] As Wu has said, "Attention in the form of eye movements tracks intention, even where the intention is not explicitly to attend in that way. It won't be strange, colloquially, to say that once the person starts looking, a pattern of eye movements happens automatically which makes perfect sense given what the person is looking for. The specific pattern of the movements is a feature of the subject's attention, one that is not represented in the content of the intention…At the same time, this automatic feature is clearly influenced by the goals we have, for the pattern of movements makes sense given the intention, so it is top-down influenced (personal communication).

sensitivity to meaningful contents at the personal-level by automatic processes should encourage us to pause before categorizing automaticity as necessarily subpersonal, low-level, mechanistic, and causal. In short, the cognitive penetrability of automatic processes provides us with a reason to doubt that automatic processes and behaviors are necessarily unintelligent.

I should note that it is not my goal here to give a full account of the neurobiological mechanisms, which make the cognitive penetrability of attention possible. Rather, I am interested in highlighting that, at a conceptual level, the connections between attention and higher-order goals must be such that they qualify as instances of cognitive penetrability. For a detailed discussion of attention and the mechanisms of top-down cognitive penetrability, see Wu (forthcoming, 2014).

4 Automatic motor routines, skill and practice

In this section, I will start by presenting considerations to support the claim that the automatic motor routines constitutive of skilled actions may be diachronically shaped by intelligent processes. I will appeal to optimal control theory to provide empirical support for this suggestion. If automatic motor routines are shaped under the guidance of higher-order goals, then I take this as evidence that at least some automatic processes bear a robust diachronic connection to personal-level intentional states such that their function cannot be understood in purely brute-causal terms. Thus, these considerations should provide us with another reason for questioning whether automatic processes can easily be filed into the "unintelligent" box.

First off, we should notice that even if automatic motor routines are not under the direct cognitive control of a skilled agent at the time of performance, this still leaves open the possibility that automatic motor routines may be diachronically shaped by personal-level, cognitive states through practice and drilling. After all, it is widely accepted that focused, deliberate, effortful practice over hundreds or thousands of hours is key to developing expertise. And as Boutin et al. (2014) have recently shown, heightened conscious awareness to motor sequence execution enhances the sensorimotor learning of a skill. Additionally, it has become clear that explicit knowledge and instructions contribute to skill learning and performance in various important ways Taylor et al. (2014). As such, it would seem that personal-level states have some important role to play in expertise. The specific point that I'll try to establish in this section is that intentional states have an explanatorily significant role to play in the shaping or structuring the automatic motor routines constitutive of skilled action.

It seems clear that during the development of a skill, not only does an agent direct effort towards improving the likelihood of attaining the goal at which a skill is aimed (i.e., making a basket or landing on one's feet after a back-handspring) but deliberate effort is also directed at improving the way, manner, means or technique by which the skill and its sub-components are executed.[28] Moreover, the improvement of the way or manner in which a skill is performed can be understood as occurring at distinct levels: certainly practice allows one to develop better and more appropriate strategies

[28] See Fridland (2013b, 2014b), Fridland and Moore (2014c). Montero (forthcoming) for more on practice.

for executing a skill; that is, practice improves judgment. But practice also allows one to develop better and more effective automatic motor control (Haith and Krakauer 2013). In this way, practice affects not just what we do but also how we do it; it affects both motor planning and motor execution.

Let us focus, then, on motor execution since this is the part of skill that is most likely to automatize.[29] To begin, we should notice that the fact of a skill's automaticity is far from the only difference between the expert and the novice.[30] That is, automaticity doesn't just allow one to perform certain actions without conscious effort or attention, faster or in parallel, but automaticity in a skill also allows one to perform certain actions in a more or less effective, elegant, and refined way. And it is in the manner, in the meat of the automaticity, that diachronic effects of learning and practice, I claim, have their impact. This claim is supported by empirical work on automaticity, which shows that automatic actions are not just faster than non-automatic ones but, rather, the structure or shape of automatic processes differs from the structure or shape of nonautomatic processes (Haith and Krakauer 2013; Izawa et al. 2008; Saling and Philips 2007; Moors and Houwer 2006; Ericsson and Charness 1994; Rosenbloom and Newell 1986; Cheng 1985).

Specifically, I want to suggest that the structure and shape of automatic motor routines develops diachronically under the guidance of higher-order intentions. This fact is clearly demonstrated by various studies that support optimal control theory. Optimal control theory highlights the fact that, as motor skills develop, sensorimotor routines undergo a reduction in movement variability along task-relevant dimensions (Bernstein 1967; Cole and Abbs 1986; Scholz and Schoner 1999, Domkin et al. 2002; Todorov and Jordan 2002; Nagengast et al. 2009).

As Todorov and Jordan (2002) explain,

> trial-to-trial fluctuations in individual degrees of freedom are on average larger than fluctuations in task-relevant movement parameters—motor variability is constrained to a redundant subspace (or 'uncontrolled manifold') rather than being suppressed altogether (p. 1226).

The task-relevant reduction in kinematic details conforms to what is known as the Minimum Intervention Principle (MIP). According to MIP, agents "only correct perturbations that interfere with the achievement of task goals. If a perturbation is irrelevant to the task, for instance, if your elbow is knocked during a reaching movement without affecting your hand position then there is no need to correct for it—jut maintain the new elbow posture during the rest of the movement" (Haith and Krakauer 2013, p. 16).

When applied to high-level motor skills, we see that the automatic motor routines developed through practice do not undergo a uniform, undifferentiated, brute, reduction in variability. But, rather, as Yarrow et al. (2009) explain, "stabilization of movement is greater for those aspects of posture that contribute directly to desired out-

[29] See Christensen et al. (in progress) for a defense of this claim.

[30] In emphasizing that expert skills are automatic, I do not mean to endorse a position that excludes the possibility of experts occasionally performing the automatic components of their skill in non-automatic ways. This may occur, perhaps, when experts deliberately practice and refine particular aspects of a skill.

come" (p. 586). That is, automatic movements do not develop uniformly as a result of brute repetition but are crafted to reflect the task at hand. Moreover, and perhaps more importantly, differentiation between task-relevant and task-irrelevant dimensions of movement is not achieved by priming or triggering a fixed or pre-determined movement trajectory but, rather, results from ongoing low-level sensorimotor sensitivity to the higher-order goals.[31]

This fact is established by Liu and Todorov (2007) who demonstrate that corrections for reaching errors remain uncorrected even when there is time to correct them. That is, they demonstrate that corrections remain uncorrected not because those corrections cannot be corrected but because they are irrelevant for task success. This finding supports the notion that fine-grained sensorimotor control is flexible insofar as corrections are made in an intelligent way—not simply to conform to a pre-determined trajectory, but in order to achieve one's goal (See also Deidrichsen 2007). This much should be clear since, if a correction is unnecessary for task success, even if it was part of an original motor plan, after perturbation, it remains uncorrected. As such, we see that even fine-grained sensorimotor executions show important connections to intentional, personal-level goal-states. Importantly, these systematic connections between goals and automatic motor control are not compatible with a purely bottom-up view of automatic motor control.[32]

Returning to high-level skills, we can generalize to the fact that the deliberate repetition of certain sub-elements of a skill, e.g., the kick up to a handstand, the jump onto the board of the vault, the pop of the shoulders away from the ground in a back or front handspring, becomes automatized in ways that reflect task-relevant dimensions of the skill, the relevance of which is established under the guidance of higher-order goals. As such, it seems that through practice and repetition, an agent sculpts her automatic motor routines in a way that is responsive to personal-level features of intentional goal states. This means that the shape or structure of a skilled automatic motor routine is not formed in a brute, bottom-up fashion but in a way that bears systematic connections to intentional states. As such, optimal control theory shows us that automatic motor routines are not simply causally impacted by higher-order goals but are internally shaped in response to the semantic features of intentional states.

To end, I' d like to note that when performed all at once in competition, it may seem as though the automatic processes constituting high-level skills run on their own, perhaps only guided or selected by general intentions or strategies at the personal-level. However, upon reflection on the diachronic development of these automatic processes,

[31] "At the heart of the framework is the relationship between high-level goals, and the real-time sensorimotor control strategies most suitable for accomplishing those goals" (Todorov 2004, p. 907).

[32] This body of evidence [supporting optimal control theory] is fundamentally incompatible, with models that enforce a strict separation between trajectory planning and trajectory execution. In such serial models, the planning stage resolves the redundancy inherent in the musculoskeletal system by replacing the behavioral goal (achievable via infinitely many trajectories) with a specific 'desired trajectory'. Accurate execution of the desired trajectory guarantees achievement of the goal, and can be implemented with relatively simple trajectory-tracking algorithms. Although this approach is computationally viable (and often used in engineering), the many observations of task-constrained variability and goal-directed corrections indicate that online execution mechanisms are able to distinguish, and selectively enforce, the details that are crucial for goal achievement. This would be impossible if the behavioral goal were replaced with a specific trajectory (Todorov and Jordan 2002, p. 1226).

we see that the particular shape of an automatic motor routine, a shape that is gleaned in the dimensions of movement that are refined and those that are not, results from a semantic connection between intentional states and automatic motor processes and not from a simple bottom-up process of brute repetition. As such, it would seem that, in this case, the automatic motor routines that constitute skilled actions are difficult to class among the unintelligent phenomena.

5 Conclusion

From the previous considerations, it is clear that at least some automatic processes and behaviors are different enough from digestion or water running down hill to prevent us from categorizing them as unintelligent. That is, both the synchronic and diachronic impacts of C-processes on some automatic behaviors and processes should make us rethink the tendency we may have had to characterize them on a purely subpersonal, causal, brute level. Still, these processes and behaviors seem to share a lot of features with the "unintelligent" side of the mental dichotomy. That is, many of these processes and behaviors are deployed without explicit intention to deploy them, they are unconscious, fast, parallel, efficient, effortless, resistant to cognitive load, etc. Moreover, it isn't at all obvious that such processes are best construed as "intelligent" either. That is, it is questionable whether such processes ought to be characterized as propositional, conceptual, truth preserving, states.[33] It seems to me that the many distinctions between automatic processes and paradigmatically intelligent states should prevent us from simply trying to assimilate automatic processes to personal-level thoughts, beliefs and knowledge.

I take it that faced with this tension—that some processes are both not unintelligent but not intelligent either, that is, that some processes are not strictly personal-level, semantic, cognitive phenomena nor brute-causal, physical states, leaves us with at least two options: the first is to rethink the "unintelligent" category. In doing so, we might notice that most processes and behaviors that fall into the "unintelligent" category have many features like the capacity to adapt, learn, or respond to semantic content such that we were wrong to ever conceptualize them as mechanical, causal or unintelligent in the first place. We may conclude that nothing is brute-causal. The other option, and the one I favor, would be to distinguish those processes and behaviors in the "unintelligent" camp that are characterized by their responsiveness to C-processes from those that are not and then to establish an intermediate category of cognition that is neither "unintelligent" nor "intelligent" but somewhere in between; maybe "proto-intelligent." This move would challenges the long-standing dichotomy between mind and body in a way that potentially opens up a more veridical and explanatorily powerful way of categorizing mental states and processes.

In either case, I take it that carefully examining automatic processes and identifying several important ways in which they are cognitively integrated will help prompt us to reformulate our basic notions of cognition going forward. There is much work left to do. So let's get to it.

[33] Though, some would disagree (See Stanley 2011).

Acknowledgments I would like to thank the Naturalistic Approaches to Mind reading Group at King's College London for numerous helpful comments and suggestions. I would like to thank, especially, David Papineua, Nick Shea, Alexander Clark, Natalie Gold, Matteo Mameli, Karen Neander and Stephen Mann. I would also like to thank Jennifer Corns, David Lowenstein, John Michael, Eliot Michaelson, Micheal Brownstein and Matt Parrott for their thoughtful comments. Lastly, I would like to thank the two anonymous reviewers for their probing and insightful feedback.

References

Bargh, J. (1989). Conditional automaticity: Varieties of automatic influence in social perception and cognition. In J. S. Uleman & J. A. Bargh (Eds.), *Unintended thought*. New York: Guilford Press.

Bargh, J. (1992). The ecology of automaticity: Toward establishing the conditions needed to produce automatic processing effects. *American Journal of Psychology*, *105*(2), 181–199.

Bargh, J. (1994). The four horsemen of automaticity: Awareness, intention, efficiency, and control in social cognition. In R. S. Wyer & T. K. Srull (Eds.), *Handbook of social cognition* (Vol. 1, pp. 1–40). Hillsdale, NJ: Erlbaum.

Bargh, J., & Chatrand, T. (1999). The unbearable automaticity of being. *American Psychologist*, *53*(7), 462–479.

Bargh, J., Schwader, K., Hailey, S., Dyer, R., & Boothby, E. (2012). Automaticity in social-cognitive processes. *Trends in Cognitive Sciences*, *16*(12), 593–605.

Beilock, S. (2010). *Choke: What the secrets of the brain reveal about getting it right when you have to*. New York: Free Press.

Bernstein, N. I. (1967). *The coordination and regulation of movements*. Oxford: Pergamon.

Berry, D., & Broadbent, D. (1984). On the relationship between task performance and associated verbalizable knowledge. *Quarterly Journal of Experimental Psychology*, *36*(A), 209–231.

Blair, I. V., Ma, J. E., & Lenton, A. P. (2001). Imagining stereotypes away: The moderation of implicit stereotypes through mental imagery. *Journal of Personality and Social Psychology*, *81*(5), 828–841.

Boutin, A., Blandin, Y., Massen, C., Heuer, H., & Badets, A. (2014). Conscious awareness of action potentiates sensorimotor learning. *Cognition*, *133*, 1–9.

Brehmer, B., Hagafors, R., & Johansson, R. (1980). Cognitive skills in judgment: Subjects' ability to use information about weights, function forms, and organizing principles. *Organizational Behavior and Human Performance*, *26*, 373–385.

Cheng, P. W. (1985). Restructuring versus automaticity: Alternate accounts of skill acquisition. *Psychological Review*, *92*, 414–423.

Christensen, W., Sutton, J., & McIlwain, D. (in progress). *Cognitive control in skilled action*.

Cole, K. J., & Abbs, J. H. (1986). Coordination of three-joint digit movements for rapid finger-thumb grasp. *Journal of Neurophysiology*, *55*, 1407–1423.

Deidrichsen, J. (2007). Optimal task-dependent changes of bimanual feedback control and adaptation. *Current Biology*, *17*, 1675–1679.

Devine, P. G. (1989). Stereotypes and prejudice: Their automatic and controlled components. *Journal of Personality and Social Psychology*, *56*, 680–690.

Di Nucci, E. (2013). *Mindlessness*. Cambridge: Cambridge Scholars Publishing.

Domkin, D., Laczko, J., Jaric, S., Johansson, H., & Latash, M. L. (2002). Structure of joint variability in bimanual pointing tasks. *Experimental Brain Research*, *143*, 11–23.

Dreyfus, H., & Dreyfus, S. (1986). *Mind over machine*. Mumbai: The Free Press.

Dreyfus, H. (2007). The return of the myth of the mental. *Inquiry*, *50*(4), 352–365.

Dreyfus, H. (2002). Intelligence without representation: Merleau Ponty's critique of representation. *Phenomenology and the Cognitive Sciences*, *1*, 367–383.

Dunton, B. C., & Fazio, R. H. (1997). An individual difference measure of motivation to control prejudiced reactions. *Personality and Social Psychology Bulletin*, *23*, 316–326.

Ericsson, A. K., & Charness, N. (1994). Expert performance: Its structure and acquisition. *American Psychologist*, *49*, 725–774.

Fazio, R. H., Jackson, J., Dunton, B., & Williams, C. (1995). Variability in automatic Activation as an unobtrusive measure of racial attitudes: A bona fide pipeline? *Journal of Personality and Social Psychology*, *69*, 1013–1027.

Frankfurt, H. (1978). The problem of action. *American Philosophical Quarterly*, *15*(2), 157–162.

Fischer, J. M. (1982). Responsibility and control. *Journal of Philosophy*, *89*, 24–40.
Fischer, J. M., & Ravizza, M. (1998). *Responsibility and control: An essay on moral responsibility*. Cambridge: Cambridge University Press.
Fiske, S. T. (1989). Examining the role of intent: Toward understanding its role in stereotyping and prejudice. In S. Uleman & A. Bargh (Eds.), *Unintended thought* (pp. 253–283). New York: Guilford Press.
Fodor, J. (1983). *The modularity of mind*. Cambridge: MIT Press.
Francolini, C. M., & Egeth, H. (1980). On the automaticity of "Automatic" activation: Evidence of selective seeing. *Perception & Psychophysics*, *27*, 331–342.
Fridland, E. (2013a). Problems with intellectualism. *Philosophical Studies*, *165*(3), 879–891.
Fridland, E. (2013b). Imitation, skill learning, and conceptual thought: An embodied, developmental approach. In L. Swan (Ed.), *The origins of mind: Springer book series in biosemantics*.: Springer.
Fridland, E. (2014). They've lost control: Reflections on skill. *Synthese*, *91*(12), 2729–2750.
Fridland, E. (2014b). Skill learning and conceptual thought: Making our way through the wilderness. In B. Bashour & H. Muller (Eds.), *Philosophical naturalism and its implications*. London: Routledge.
Fridland, E., & Moore, R. (2014). Imitation reconsidered. *Philosophical Psychology*,. doi:10.1080/09515089.2014.942896.
Haith, A., & Krakauer, J. (2013). Theoretical models of motor control and motor learning. In A. Gollhofer, W. Taube, & J. B. Nielsen (Eds.), *Routledge handbook of motor control and motor learning* (pp. 1–28). London: Routledge.
Hasher, L., & Zacks, R. T. (1979). Automatic and effortful processes in memory. *Journal of Experimental Psychology: General*, *108*, 356–388.
Hoffman, J. E., Nelson, B., & Houck, M. R. (1983). The role of attentional resources in automatic detection. *Cognitive Psychology*, *51*, 379–410.
Hommel, B. (2007). Consciousness and control: Not identical Twins. *Journal of Consciousness Studies*, *14*(1–2), 155–176.
Izawa, J., Rane, T., Donchin, O., & Shadmehr, R. (2008). Motor adaptation as a process of reoptimization. *The Journal of Neuroscience*, *28*(11), 2883–2891.
Kahneman, D., & Chajczyk, D. (1983). Test of the automaticity of reading: Dilution of stroop effects by color-irrelevant stimuli. *Journal of Experimental Psychology: Human Perception and Performance*, *9*, 497–509.
Kahneman, D., & Henik, A. (1981). Perceptual organization and attention. In M. Kubovy & J. Pomerantz (Eds.), *Perceptual organization* (pp. 181–212). Hillsdale, NJ: Erlbaum.
Karmiloff-Smith, A. (1994). Précis of beyond modularity: A developmental perspective on cognitive science. *Behavioral and Brain Sciences*, *17*(4), 693–735.
Kawakami, K., Dovidio, J. F., Moll, J., Hermsen, S., & Russin, A. J. (2000). Just say no (to stereotyping): Effects of training in the negation of stereotypic associations on stereotype activation. *Journal of Personality and Social Psychology*, *78*(5), 871–888.
LaBerge, D., & Samuels, S. J. (1974). Toward a theory of automatic information processing in reading. *Cognitive Psychology*, *6*, 293–323.
Liu, D., & Todorov, E. (2007). Evidence for the flexible sensorimotor strategies predicted by optimal feedback control. *The Journal of Neuroscience*, *27*(35), 9354–68.
Logan, G. D. (1979). On the use of concurrent memory load to measure attention and automaticity. *Journal of Experimental Psychology: Human Perception and Performance*, *5*, 189–207.
Logan, G. D. (1980). Attention and automaticity in stroop and priming tasks: Theory and cata. *Cognitive Psychology*, *12*, 523–553.
Logan, G. (1985). Skill and automaticity: Relations, implications, and future directions. *Canadian Journal of Psychology*, *39*(2), 367–386.
Macpherson, F. (2012). Cognitive penetration of colour experience: Rethinking the issue in light of an indirect mechanism. *Philosophy and Phenomenological Research*, *84*(1), 24–62.
Milner, D., & Goodale, M. A. (2006). *The visual brain in action*. Oxford: Oxford University Press.
Montero, B. (2010). Does bodily awareness interfere with highly skilled movement? *Inquiry*, *53*(2), 105–122.
Moors, A., & De Houwer, J. (2006). Automaticity: A theoretical and conceptual analysis. *Psychological Bulletin*, *132*(2), 297–326.
Nagengast, A. J., Braun, D. A., & Wolpert, D. M. (2009). Optimal control predicts human performance on objects with internal degrees of freedom. *PLoS Computational Biology*, *5*(6), e1000419.

Neumann, O. (1984). Automatic processing: A review of recent findings and a plea for an old theory. In W. Prinz & F. Sanders (Eds.), *Cognition and motor processes* (pp. 255–293). Berlin: Springer.
Norman, D. A., & Shallice, T. (1986). Attention to action: Willed and automatic control of behaviour. In R. J. Davidson, G. E. Schwartz, & D. Shapiro (Eds.), *Consciousness and self-regulation* (Vol. 4, pp. 1–18). New York: Plenum Press.
Olson, M. A., & Fazio, R. H. (2004). Trait inferences as a function of automatically-activated racial attitudes and motivation to control prejudiced reactions. *Basic and Applied Social Psychology*, *26*, 1–11.
Palmeri, T. J. (2002). Automaticity. In L. Nadel, et al. (Eds.), *Encyclopedia of cognitive science* (pp. 390–401). London: Nature Publishing Group.
Papineau, D. (2013). In the Zone. *Royal Institute of Philosophy Supplement*, *73*, 175–196.
Posner, M. I., & Snyder, C. R. R. (1975a). Attention and cognitive control. In R. L. Solso (Ed.), *Information processing and cognition: The Loyola symposium* (pp. 153–175). Hillsdale: Erlbaum.
Posner, M. I., & Snyder, C. R. R. (1975b). Facilitation and inhibition in the processing of signals. In P. M. A. Rabbit & S. Dornic (Eds.), *Attention and performance V* (pp. 669–682). Waltham: Academic Press.
Prinz, J. (2006). Beyond appearances : The content of sensation and perception. In T. Gendler & J. Hawthorne (Eds.), *Perceptual experience* (pp. 434–460). Oxford: Oxford University Press.
Pylyshyn, Z. W. (2000). Is vision continuous with cognition? *Behavioral and Brain Sciences*, *22*(3), 341–365.
Pylyshyn, Z. (2003). *Seeing and visualizing*. Cambridge, MA: The MIT Press.
Reber, A., & Lewis, S. (1977). Implicit learning: An analysis of the form and structure of a body of tacit knowledge. *Cognition*, *5*, 333–361.
Reed, N., McLeod, P., & Dienes, Z. (2010). Implicit knowledge and motor skill: What people who know how to catch don't know. *Consciousness and Cognition*, *19*, 63–76.
Rosenbloom, P. S., & Newell, A. (1986). The chunking of goal hierarchies: A generalized model of practice. In R. S. Michaliski, J. G. Carbonell, & T. M. Mitchell (Eds.), *Machine learning: An artificial intelligence approach* (Vol. 2, pp. 247–288). Los Altos, CA: Morgan Kaufmann.
Ryle, G. (1949). *The concept of mind*. Chicago: The University of Chicago Press.
Saling, L., & Philips, J. G. (2007). Automatic behavior: Efficient not mindless. *Brain Research Bulletin*, *73*, 1–20.
Scholz, J. P., & Schoner, G. (1999). The uncontrolled manifold concept: Identifying control variables for a functional task. *Experimental Brain Research*, *126*, 289–306.
Schneider, W., & Chein, J. (2003). Controlled & automatic processing: Behavior, theory, and biological mechanisms. *Cognitive Science*, *27*, 525–559.
Schneider, W., & Shiffrin, R. M. (1977). Controlled and automatic human information processing. I. Detection, search, and attention. *Psychological Review*, *84*(1), 1–66.
Schneider, W., Dumais, S. T., & Shiffrin, R. (1984). Automatic and control processing and attention. In R. S. Parasuraman & D. R. Davies (Eds.), *Varieties of attention* (pp. 1–27). New York: Academic Press.
Shiffrin, R. M., & Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning, automatic attending and a general theory. *Psychological Review*, *84*(2), 127–190.
Shiffrin, R. M. (1988). Attention. In R. C. Atkinson, R. J. Hernstein, G. Lindzey, & R. D. Luce (Eds.), *Stevens' handbook of experimental psychology* (Vol. 2, pp. 739–811). Hoboken: Wiley.
Siegel, S. (2010). *The contents of visual experience*. Oxford: Oxford University Press.
Stanley, J. (2011). *Know how*. Oxford: Oxford University Press.
Sheets-Johnstone, M. (2012). Kinesthetic memory: Further critical reflections and constructive analyses. In S. C. Koch, T. Fuchs, M. Summa, & C. Müller (Eds.), *Body, metaphor, memory and movement* (pp. 43–72). Amsterdam: John Benhamins Publishing.
Stokes, D. (2013). Cognitive penetrability of perception. *Philosophy Compass*, *8*(7), 646–663.
Sutton, J. (2007). Batting, habit and memory: The embodied mind and the nature of skill. *Sport in Society*, *10*(5), 763–786.
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.
Taylor, J. A., Krakauer, J. W., & Ivry, R. B. (2014). Explicit and implicit contributions to learning in a sensorimotor adaptation task. *Journal of Neuroscience*, *34*, 3023–3032.
Todorov, E. (2004). Optimality principles in sensorimotor control. *Nature Neuroscience*, *7*(9), 907–915.

Todorov, E., & Jordan, M. I. (2002). Optimal feedback control as a theory of motor coordination. *Nature Neuroscience*, *5*(11), 1226–1235.
Tzelgov, J. (1997a). Specifying the relations between automaticity and consciousness: A theoretical note. *Consciousness and Cognition*, *6*, 441–451.
Tzelgov, J. (1997b). Automatic but conscious: That is how we act most of the time. In R. S. Wyer Jr (Ed.), *Advances in social cognition* (Vol. 10, pp. 217–230). Mahwah, NJ: Erlbaum.
Tzelgov, J. (1999). Automaticity and processing without awareness. *Psyche*, *5*(3), 18–23.
Tzelgov, J., Ganor, D., & Yehene, V. (1999). Automatic processing results in conscious representations. *Behavioral and Brain Sciences*, *22*(5), 786–787.
Tzelgov, J., Henik, A., & Porat, Z. (1997). Automaticity and consciousness: Is perceiving the word necessary for reading it? *American Journal of Psychology*, *3*, 429–448.
Wallis, C. (2008). Consciousness, context and know-how. *Synthese*, *160*, 123–153.
Wu, W. (2011). Confronting many-many problems: Attention and agentive control. *Noûs*, *45*(1), 50–76.
Wu, W. (2013a). Mental action and the threat of automaticity. In A. Clark, J. Kiverstein, & T. Vierkant (Eds.), *Decomposing the will* (pp. 244–261). Oxford: Oxford University Press.
Wu, W. (2013b). Visual spatial constancy and modularity: Does intention penetrate vision? *Philosophical Studies*, *165*(2), 647–669.
Wu, W. (2014). *Attention: New problems of philosophy series*. London: Routledge.
Wu, W. (forthcoming). Shaking up the Mind's ground floor: The cognitive penetration of visual attention.
Yarbus, A. (1967). *Eye movements and vision*. New York: Plenum Press.
Yarrow, K., Brown, P., & Krakauer, J. W. (2009). Inside the brain of an elite athlete: The neural processes that support high achievement in sports. *Nature Reviews Neuroscience*, *10*, 585–596.