



Implicit biases in visually guided action

Berit Brogaard¹

Received: 15 February 2018 / Accepted: 17 February 2020 / Published online: 7 April 2020
© Springer Nature B.V. 2020

Abstract

For almost half a century dual-stream advocates have vigorously defended the view that there are two functionally specialized cortical streams of visual processing originating in the primary visual cortex: a ventral, perception-related ‘conscious’ stream and a dorsal, action-related ‘unconscious’ stream. They furthermore maintain that the perceptual and memory systems in the ventral stream are relatively shielded from the action system in the dorsal stream. In recent years, this view has come under scrutiny. Evidence points to two overlapping action pathways: a dorso-dorsal pathway that calculates features of the object to be acted on, and a ventro-dorsal pathway that transmits stored information about skilled object use from the ventral stream to the dorso-dorsal pathway. This evidence suggests that stored information may exert significantly more influence on visually guided action than hitherto assumed. I argue that this, in turn, supports the notion of skilled automatic action that is nonetheless agential. My focus here will be on actions influenced by implicit biases (stereotypes/prejudices). Action that is biased in this way, I argue, is in an important sense intentional and agential.

Keywords Dorso-dorsal stream · Implicit Bias · Intentional action · Microaggression · Motor representation · Skilled irreflective action · Stereotype · Ventro-dorsal stream · Vision for action · Vision for perception

1 Introduction

On your evening stroll in your secure suburban neighborhood, you instinctively clutch your purse as you pass your Muslim neighbor but politely greet the Swedish immigrant from across the street a few minutes later; at the local restaurant where you work as a hostess you seat an interracial couple by the kitchen in the back of the nearly empty establishment without consciously intending to seat them in that location and without even realizing that your action is guided by an implicit bias against interracial

✉ Berit Brogaard
brogaardb@gmail.com

¹ University of Miami, Coral Gables, FL, USA

couples; when a fellow school mom tells you that she and her husband are struggling to make ends meet on one income, you ask whether she thinks she might go back to work eventually. You only realize your mistake when she snarkily replies: ‘You mean “whether I think my husband might go back to work eventually”? He’s the stay-at-home parent’.

As these cases illustrate, our implicit biases (stereotypes/prejudices) have a way of interjecting themselves into our social actions in ways that are barely noticeable to us yet are discriminatory against the targeted individuals. We call these subtle yet discriminatory actions ‘microaggressions’ (Sue 2010; Fatima 2017). To a first approximation, we can say that a microaggression is an action triggered by an implicit bias that conveys a hostile message to the targeted person because of his or her group membership. A person who acts micro-aggressively is thus engaging in discriminatory behavior.

There is no widespread consensus on what implicit biases are or on what the various measurements of implicit biases actually measure (see e.g. Mandelbaum 2016; Toribio 2018a). But it is fair to say that it is generally thought that because implicit biases are encoded on an unconscious level of the brain, discriminatory behavior influenced by them is automatic (or unconscious) and unintentional. In other words, implicit biases insert themselves in our actions by altering unconscious processes that guide action, not by modifying our intentions. However, I will argue that newer research that builds on aspects of David Milner and Melvyn Goodale’s dual-stream hypothesis lends support to a different picture (Goodale et al. 1991; Goodale and Milner 1992; Milner and Goodale 1995).

Milner and Goodale’s dual-stream hypothesis grew out of brain research suggesting that there are two anatomically segregated and functionally specialized cortical streams of visual processing originating in the primary visual cortex: a ventral, perception-related ‘conscious’ stream and a dorsal, action-related ‘unconscious’ stream. The findings also seemed to support the view the perceptual and memory systems in the ventral stream are relatively shielded from the action system in the dorsal stream.

In recent years, the dual-stream hypothesis has come under scrutiny. Evidence points to two overlapping action pathways: a dorso-dorsal pathway that calculates features of the object to be acted on, and a ventro-dorsal pathway that transmits stored information about skilled object use from the ventral stream to the dorso-dorsal pathway (Jeannerod et al. 1994; Rizzolatti and Matelli 2003; Daprati and Sirigu 2006; Buxbaum and Kalénine 2010; Grafton 2010; Binkofski and Buxbaum 2013; see also Gallese et al. 1999). The ventro-dorsal pathway provides a way for object representations in the ventral stream to continuously influence the processes underlying motor representations in the action system (Cisek 2007).

The new data, I intend to show, challenge the idea that because implicit biases are encoded on an unconscious level, their biasing effect is unintentional. As we will see, the finding that stored information can have an ongoing influence on online action suggests that implicit biases can negatively affect our actions by modulating our intentions. Action influenced by biasing intentions, I will argue, is a kind of skilled, irreflective behavior that is nonetheless agential.

2 Vision for action

The idea that the systems for perception and cognition are functionally separate from the system for action goes back to David Milner and Melvyn Goodale's research on the two visual streams.¹ The results of their studies indicate that there are two functionally specialized cortical streams of visual processing originating in striate (V1) cortex: a ventral, perception-related 'conscious' stream and a dorsal, action-related 'unconscious' stream (Goodale and Milner 1992; Goodale et al. 1991; Milner and Goodale 1995, 2008; Goodale 2014). The ventral stream runs into the temporal lobe and then connects to other temporal and frontal lobe structures that are responsible for episodic memory, working memory, reporting, decision-making, and so on. The dorsal stream runs upwards through the occipital into the parietal lobe and continues until it makes contact with the premotor cortex.

According to Milner and Goodale, the perceptual/cognitive system in the (visual) ventral stream processes information about color, size, texture and shape and relational properties of objects in allocentric (scene-based) space (e.g. *being larger than*, *being darker than*). Allocentric space is a representation of objects as they are perceived to be—independently of the perceiver's specific perspective (and the specific viewing conditions) at the time. The properties of objects in allocentric space are also known as 'allocentric properties'. In order for you to see equal-sized trees on the side of a road as having the same size, the retinal imprint of the trees must occupy different portions of your visual field (Peacocke 1992). The sizes of the occupied portions do not reflect properties of the trees independently of your specific perspective at the time, so they are not allocentric properties. Being same-sized, on the other hand, is a property the trees are represented as having independently of a particular vantage point of view, so being same-sized is an allocentric property (regardless of whether the trees are in fact same-sized). Although the ventral system can operate in the absence of visual awareness (Weiskrantz 1986, for a review see Brogaard 2011), the ventral stream is considered the correlate of conscious experience.²

By contrast to the visual system in the ventral stream, the action system in the dorsal stream normally operates in the absence of visual awareness and only computes the information needed for the organism to carry out the initiated action. Information needed for reaching to and grasping an object, for example, includes the estimated size, weight and orientation of the object; the hand aperture, body pose and move-

¹ There are auditory, tactile and emotional dual streams that parallel the dual visual streams (Romanski et al. 1999; Reed et al. 2005; Brogaard et al. 2017). However, most research has focused on vision. Here the main focus will also be on vision.

² Milner and Goodale's claim that conscious vision is associated with representations in allocentric space is not universally accepted. On the attended intermediate-level representation theory of consciousness (AIR), advocated by Jesse Prinz (2000, 2012), conscious experiences are attended intermediate-level representations. Intermediate-level representations differ from representations in allocentric space. Representations in allocentric space represent abstract viewpoint-independent features of three-dimensional objects that are generated by abstracting away from the vantage point of view and surface details. Intermediate-level representations, by contrast, represent objects and the features they instantiate from the perceiver's point of view. Because they reflect the retinal imprint (when veridical) yet capture information about Gestalt grouping, e.g., depth and orientation, they are also referred to as '2½-D sketches' in David Marr's (1982) terminology.

ments needed to reach to and grasp the object; and viewpoint-dependent properties of the object in egocentric space, including the object's location, its accessibility to the perceiver and the route that needs to be taken to get to the object (Schenk 2006).³ Egocentric space is the representation of the object relative to the perceiver, so we can call the properties of the object in egocentric space 'egocentric properties'. In a task where you need to reach your coffee mug in order to take a sip from it, the action system must determine the location and orientation of the mug relative to you; the most efficient path to the mug; the approximate size of the handle of the mug; the grib needed for the fingers to fit around the handle; the approximate weight of the mug (Brouwer et al. 2006); and the force that will need to be applied when lifting it (Brenner and Smeets 1996).


The dual-stream hypothesis is backed by studies showing that damage to structures in the dorsal stream can impair visuomotor control while leaving visual perception intact, and damage to structures in the ventral stream can impair visual perception while leaving visuomotor control intact. Goodale and Milner conducted a series of neuropsychological studies on a patient, D.F., with severe visual form agnosia, which involves damages to structures of the ventral stream (Goodale and Milner 1992; Goodale et al. 1991). D.F. could consciously see some color and texture, but she was unable to recognize shapes. Her dorsal stream was intact. She could accurately grasp objects. For example, she could post a card into a slot she could not describe. She could also adjust her finger-thumb grip size perfectly to the width of a rectangular block, even though she could not report on its width.

The dual-stream hypothesis was further tested on patients with optic ataxia, a component of Balint's syndrome. Optic ataxia is the mirror syndrome of visual agnosia. It has been shown that patients with optic ataxia, who have retained the ability to consciously identify objects, are unable to perform visually guided action that strongly depends on functional dorsal activity (Milner et al. 2001). For example, they are unable to adjust their handgrip to the size of objects in real time. They normally retain the ability to perform skilled object-dependent pantomimes as well as action guided by visual perception (Perenin and Vighetto 1988; Buxbaum et al. 2005). They also sometimes have a spared ability to adjust their handgrip when required to grasp familiar objects (Jeannerod et al. 1994).

Further evidence for the two-streams hypothesis comes from optical illusions. Research points to differential effects of optical illusions on perception and action (Aglioti et al. 1995). For example, in the Ebbinghaus Illusion, a central circle is surrounded either by large circles (right) or small circles (left). Although the central circle has the same dimensions in the two contexts, it perceptually appears smaller in the large-circle context than in the small circle context. However, in studies in which peo-

³ It should be noted that the hypothesis that properties like the size and location of an object are computed in the dorsal stream in the absence of visual awareness should not be taken to imply that these properties are not consciously available. One possibility consistent with Milner and Goodale's two-stream hypothesis is that the size and location of an object are also computed in the ventral stream, where they can be consciously accessed. On the three-stream hypothesis, which we will introduce below, information about the size and location of an object needed for the action system to generate a motor representation may be computed in the ventral stream before entering the action system.

Fig. 1 The Ebbinghaus illusion. Studies have shown that this illusion leads to a misperception of the size of the central circle but only marginally affects grasping behavior directed at the central circle



ple were asked to grasp the central circle in the two contexts, their hand aperture was only marginally influenced by the illusion (Fig. 1).

The differential effects of illusions on motor and perceptual tasks add further support for the double dissociation of the ventral and dorsal streams (see, however, Franz et al. 2000; Bruno and Franz 2009; Franz and Gegenfurtner 2008; for reviews see Bruno 2001; Smeets et al. 2002; Smeets and Brenner 2006).

Milner and Goodale originally concluded on the basis of these and other studies that our visually guided real time (non-delayed) actions are not in the direct control of what we consciously see (Goodale and Milner 1992; Goodale et al. 1991; see also Goodale 2014). The dorsal system, they reasoned, is dedicated to the rapid and accurate guidance of our movements and computes information about viewpoint-dependent properties (e.g., the object's position relative to the body) and absolute properties (e.g., absolute size) required to accurately reach to and grasp the object online, that is, as a result of a crude, fast, and automatic visuomotor transformation process. On this view, dorsal stream information thus guides programming and unfolding of real-time action needed when delayed action is counterproductive. However, while the dorsal stream mediates real-time action, under delayed movement conditions dorsal stream representations decay and action becomes mediated by the ventral stream.

The ventral system, on the other hand, is responsible for object recognition and classification (Grill-Spector et al. 2001; Kriegeskorte et al. 2008). It codes abstract allocentric (scene-based) information for storage in and retrieval from memory, and it often correlates with visual awareness. The ventral stream furthermore allows us to plan actions offline, that is, without acting immediately (e.g., by simulating them, imagining them, or calculating how to do them).

In recent years, the dogma that there is minimal interaction between the two visual streams in preparation for online action has come under scrutiny (Smeets et al. 2002, Smeets and Brenner 2006; Franz et al. 2009; Schenk and McIntosh 2010; Briscoe and Schwenkler 2015; Caiani and Ferretti 2017; Binkofski and Buxbaum 2013). More recent evidence points to two overlapping action pathways: a dorso-dorsal pathway that calculates features of the object to be acted on, and a ventro-dorsal pathway that transmits stored information about skilled object use from the ventral stream to the dorso-dorsal pathway (Jeannerod et al. 1994; Rizzolatti and Matelli 2003; Daprati and Sirigu 2006; Buxbaum and Kalénine 2010; Grafton 2010; Binkofski and Buxbaum




Fig. 2 Schematic relative location of the ventral stream (blue) and the dorso-dorsal (red) and the ventro-dorsal (green) substreams of the dorsal stream. From Binkofski and Fink (2005). (Color figure online)

2013). The ventro-dorsal pathway provides a way for object representations in the ventral stream to continuously influence the processes underlying motor representations in the action system.

3 The use and the grasp pathways

There is growing evidence that the dorsal stream is divided into two relatively segregated parieto-frontal subcircuits: a dorso-dorsal and a ventro-dorsal pathway (Jeannerod et al. 1994; Rizzolatti and Matelli 2003; Daprati and Sirigu 2006; Buxbaum and Kalénine 2010; Grafton 2010; Binkofski and Buxbaum 2013). The bilateral dorso-dorsal pathway, or the ‘Grasp’ system, involves dorsolateral fronto-parietal areas of the dorsal stream, whereas the left lateralized ventro-dorsal pathway, or ‘Use’ system, involves superior temporal and inferior parietal areas of the dorsal stream (Binkofski and Buxbaum 2013) (Fig. 2).

The dorso-dorsal Grasp system diverts from the ventral stream after V1/V2. From here it runs through dorsal V3, which is tuned to orientation and depth and dorsomedial visual areas (V6), which contain a representation of the entire visual field. From here it runs through the intraparietal sulcus (IPS) and adjacent gyri, which have a role in directing visual attention of space and organizing limb and eye movement (Corbetta and Shulman 2002). The Grasp system terminates in the dorsal premotor cortex. The ventro-dorsal Use system runs through motion areas (MT/MST) in the superior temporal cortex. From here it runs through the inferior parietal lobule (IPL),

which processes spatial, emotional and sensory information for organizing eye and limb movements in space, and terminates in the ventral premotor cortex.

The two subcircuits of the dorsal stream mediate different behavioral goals. The dorso-dorsal system is the most direct and immediate visual pathway for online control of fast and crude action. It processes pragmatic properties of currently-viewed objects needed for reaching, grasping and lifting, such as, allocentric size and weight and egocentric location, orientation, texture, object rigidity, hand shape and affordances (Jeannerod et al. 1994; Rizzolatti and Matelli 2003; Binkofski and Buxbaum 2013). On the traditional Gibsonian view, affordances are action possibilities afforded by the current state of the object to be acted on (Gibson 1977). Whether or not an object affords the possibility of a particular interaction is determined by situational and object properties and the motor capacities of the agent (e.g. strength, size, agility). For example, a handle on a coffee mug provides an affordance for you to grasp the mug only if the handle is not too small for your hands.

The ventro-dorsal use system processes sensorimotor ‘functional use’ information for skilled action for pantomime or real object use. Unlike simple behavior, skilled action requires access to information about stable object properties, stable object affordances and grip type (e.g. typical grip type for a marble) stored for the long term in implicit memory (Binkofski and Buxbaum 2013). Stable properties of familiar objects and object types include approximate size, shape, texture, weight and orientation in allocentric space. For example, it is a stable feature of open bottles containing liquids that they are oriented vertically. Stable affordances are stable action possibilities afforded by familiar objects (Sakreida et al. 2016). For example, the handles on coffee mugs provide stable affordances for grasping and holding them. The output of processing of sensorimotor information in the ventro-dorsal substream are motor representations mediating space perception, action recognition and the guidance of skilled action, such as the proper manipulation of a carving knife, a pool cue or a dental drill.

Since egocentric features and variable affordances are continually changing as an action unfolds, they are not stored in long-term memory but are the immediate ‘objects’ of rapidly decaying motor representations that typically last only for a few milliseconds. Information processing is far more laborious and protracted when it involves the ventro-dorsal Use system than when it bypasses the Use system. The Use system also maintains information over longer periods of time than the Grasp system (Buxbaum and Saffran 2002).

Research on familiarity effects on action provides support for the hypothesis that stable properties, such as familiar size, encoded in the ventral stream can have a biasing effect on grasp. McIntosh and Lashleya (2008) demonstrated that expectations about familiar size can significantly affect hand aperture and amplitude of movement needed to reach and grasp an object. The researchers asked subjects to reach to and grasp the standard large Swan Vestas and the standard small Scottish Bluebell match box in a series of baseline trials. In a series of perturbation trials, subjects were instructed to reach for a smaller replica of the Swan Vestas matchbox and a larger replica of the Scottish Bluebell match box (Fig. 3). The findings demonstrated that the expected size of the match boxes affected both the pre-shaping of the hand and the amplitude of reaches to grasp them. The researchers hypothesized that the grasp effects could arise



Fig. 3 The four matchboxes. Upper row: the standard Swan Vestas and Scottish Bluebell match boxes. Lower row: 0.8-scale replica Swan match box and 1.25-scale replica Bluebell match box. From McIntosh and Lashleya (2008)


either because the retinal size of the targets was modified by familiar size or because familiar size contributed more directly to the programming of grasp formation.

The results indicate that familiar size—which is encoded in higher cortical brain regions—influences the guidance of action. In the matchbox experiment, grasp is influenced by object representations of matchboxes as having a specific familiar size, viz. the sizes encoded via feedback from the dorsal stream to the ventral stream in the baseline trials.⁴

The claim that the visually guided action system is rooted in a dorsal and a ventral pathway within the dorsal stream is also supported by lesion studies. Optic ataxia, the condition that originally provided evidence for the dual-stream hypothesis, supports the further distinctions between the Use and Grasp systems within the dorsal pathway. Patients with optic ataxia have lesions to superior parietal lobule (SPL) and areas around the intraparietal sulcus (IPS) that impair the function of the dorso-dorsal ‘Grasp’ pathway, particularly hand–eye coordination and reach-to-grasp ability. Patients with isolated optic ataxia have intact visual fields, stereoscopic vision, voluntary eye control, proprioceptive knowledge of position of limbs, movement patterns and postures and cerebellar function (Andersen et al. 2014). The deficits to the parietal regions result in a dysfunction of the Grasp system owing to difficulties integrating sensory processing with the processing of hand aperture. As a result, they have difficulty preshaping their hand for visually guided reaching to objects in the contralesional

⁴ There has been a lot of debate about whether these kinds of influences count as instances of cognitive penetration of the action system. See e.g. Nanay (2013a, b), Mahon and Wu (2015), Burnston (2017a, b), Toribio (forthcoming). There is a parallel debate about whether the perceptual system in the ventral stream is cognitively penetrable (see e.g. Toribio 2018b for one stance in this debate). As there appear to be substantial cross-communication between the perception/cognition system in the ventral stream and the action system in the dorsal streams, even in online activities, these debates may well be interconnected. I will not take a stance on the question of cognitive penetration here. However, it is arguable that the more philosophically interesting question is that of whether the activation of implicit biases can bias how we act. Here I argue that it probably can bias how we act by modulating our intentions as opposed to modulating our motor representations or actions directly.

Fig. 4 Patient with optic ataxia, reaching beyond the pencil when asked to touch it. From Chechlacz (2018)




visual field, and they fail to correct their reaches once they realize their mistake (Fig. 4) (Andersen et al. 2014). Although patients with optic ataxia know the identity and location of the objects they are asked to manipulate in their impaired visual field, their lesions to the Grasp system prevent them from processing information needed for action guidance, especially when the visually available information is limited to contralesional peripheral vision.

Lesions specific to the ventro-dorsal ‘Use’ pathway lead to impairments to skilled action, including pantomime of object use but leave the ability to grasp objects intact (Randerath et al. 2010; Evans et al. 2016). This kind of selective impairment has been observed in patients with ideational apraxia. Patients with ideational apraxia have no problem grasping objects in front of them, but they cannot imitate the action in the absence of the object. Nor can they replicate the action if verbally requested to do so. They also have difficulties retrieving information about how to use objects, including information about stable object properties, stable affordances and grip style, especially when the task is to manipulate several objects, for example, in the case of lighting a candle, cutting a banana with a knife, or combining Lego pieces. As this form of apraxia impairs the ventro-dorsal Use system but not the dorso-dorsal Grasp system, it lends support to a double dissociation of the two dorsal substreams.

Other studies have found that patients with ideomotor apraxia—a disorder of the left inferior parietal lobe of the ventro-dorsal pathway—have deficits in imitating and recognizing skilled object manipulation and object-dependent gesturing, like hammering, cutting with scissors, using a comb, brushing teeth or eating with a fork, but they typically retain some ability to grasp novel items as well as some capacity for imitating and recognizing intransitive gesturing, like saluting, hitchhiking, waving goodbye or signaling stop (Mozaz et al. 2002; Pelgrims et al. 2005; Buxbaum et al.

Fig. 5 Schematic diagram of the VOF projections in the right hemisphere. Gray regions depict the projection zone of VOF (light gray, dorsal; dark gray, ventral). From Takemura et al. (2016). Reprinted with permission




2005). The results of these studies show that the ventro-dorsal system is crucial not only for skilled object manipulation but also for the ability to imitate and recognize skilled action, which supports the idea that information from the dorsal stream can be encoded in implicit memory in the ventral stream.

Although studies have shown that a further anatomical and functional subdivision exists within the dorsal stream, the two dorsal substreams are not anatomical and functional disconnected. As the two substreams share neural correlates, they overlap anatomically (Borghetti and Riggio 2015). For example, it has been shown that the perceptual motion area V5/MT is activated when the task is to catch a moving object rather than grasp a stationary object (Schenk et al. 2005; van Polanen and Davare 2015). There is also evidence of cross-talk between the two substreams, for example between V5/MT and the intraparietal sulcus (Pisella et al. 2006; Binkofski et al. 2007; Borra et al. 2008; Nelissen and Vanduffel 2011; Binkofski and Buxbaum 2013; Cloutman 2013). In fact, a considerable number of projections from the ventral stream reach the ventro-dorsal pathway (Takemura et al. 2016).

Using brain imaging and fiber tractography, a 3D modeling technique for visually representing nerve tracts, neuroscientist Hiromasa Takemura and colleagues identified a major white matter pathway, known as ‘the vertical occipital fasciculus’ (VOF), that allows information in the ventral stream to enter the ventro-dorsal pathway, and vice versa (Takemura et al. 2016) (Fig. 5). A substantial number of projections were identified running from areas coding for color, form and identity in the ventral stream to regions that integrate information about the object’s relative location into the action plan (e.g., your coffee mug’s location relative to your hand and your plan to grab it). As VOF seems crucial for transmitting signals between the ventral stream and the ventro-dorsal pathway, VOF can be seen as constituting an interface between the perception/cognition system and the action system. A major white matter communication pathway between the dorsal and ventral streams, VOF is also substantially involved in cognitive activities that require online coordination between the identification of the color, form and identity of an object and motor activities. For example, when we read, we coordinate our eye movements in accordance with our recognition of

Fig. 6 Delboeuf Illusion. The inner circle appears smaller when surrounded by a large circle than when surrounded by a small circle



graphemes, words and text fragments. Recognizing text, in turn, depends on seeing the illumination or color contrast between the text and the background and identifying its shape.

Grasping objects with the intent to manipulate them appropriately requires precise coordination between the two pathway (Binkofski and Buxbaum 2013; Garcea and Buxbaum 2019). Information about the current object must be integrated with retrieved information about stable affordances and stable object features (Buxbaum et al. 2003). The need for precise coordination between the two pathways has been demonstrated by studies showing that if there is a conflict between the hand shape or movements needed for grasping and those needed for object use, as in the case of grasping and manipulating a calculator or an iPhone, more laborious and protracted processing is needed.

The high degree of influence of information encoded in the ventral on action preparation in the case of skilled tool manipulation has been backed by studies of tool manipulation mimicking real-life situations. In one study, experienced endodontists were asked to cut cavities in real human teeth prepared to look like the teeth they normally see when making a cavity in a root canal (O'Shea et al. 2013). The teeth were randomly selected from teeth the endodontists would normally encounter, and the endodontists used their own tools and procedures for the cutting. The teeth the endodontists were provided with presented more or less potent versions of the Delboeuf illusion, an illusion in the same family of illusions of size as the Ebbinghaus Illusion (Fig. 6). The illusion makes an inner circle appear smaller when surrounded by a large circle than when surrounded by a small circle.

The researchers found that the illusion made the root canal for filling appear smaller to the endodontists, making them cut cavities in preparation for filling that were larger than needed. The size of cavity the endodontists made was found to be linearly related to the potency of the Delboeuf illusion.

The differential effect of the Ebbinghaus Illusion and the Delboeuf Illusion on visually guided action add further support the hypothesis that complex, skilled behavior (including the performance of gestures, the manipulation of tools and pantomimes) requires a considerable amount of additional input from the ventral stream compared to what is needed for simple grasping (Buxbaum et al. 2005).

4 Skilled action and implicit intentions

The dual action-system hypothesis provides us with an account of how information about familiar objects stored in implicit memory in the perception and cognition systems can influence motor representations in the action system. However, it also points to a greater influence of intentions on what may at first seem to be mindless and automatic subcomponents of actions. Before looking at intentions in the context of implicit biases and microaggression, I will look closer at the significance of intentions in the execution of action.

It is commonplace to distinguish between future- and present-directed intentions (Bratman 1984, 1985).⁵ Future-directed (or distal) intentions plausibly just are beliefs of the form *I will ϕ (later)* that represent the agent as completing a future goal rather than a Davidsonian ‘all-out, unconditional judgement’ of desirability (Davidson 1978: p. 99; see Mele 1992: pp. 143–144, 158). Deliberate action preparation is often indispensable for future-directed intentions to be realizable. For example, your intention to go to the Nicki Minaj concert next Saturday may not be realizable, if you don’t purchase tickets in advance.

Present-directed (or proximal) intentions need to do more than represent a near-future goal, as in ‘I will ϕ now’, ‘I am about to ϕ ’ or ‘I am ϕ -ing and will continue ϕ -ing’.⁶ They must also “tell” the action system to initiate or continue an initiated action and provide details of the action that can be expanded into precise motor representations apt to cause the intended action (Jeannerod 2006; see also Butterfill and Sinigaglia 2014; Ferretti 2016). However, the intentions we consciously entertain do not typically provide more than the bare minimum of action specifications. For example, I might think to myself ‘I shall now carve the turkey’ without consciously specifying any further details about how I will perform the action and without explicitly intending to perform any subsidiary actions, such as grabbing the carving knife. The unspecified details of the action are not simply left for the action system to decide in a mindless fashion, however. It is no coincidence that I use the carving knife and not the bread knife, and it’s no coincidence that I begin by cutting the skin between the breast and the thighs and drumsticks rather than by splitting the bird in half with a butcher’s knife. It is no coincidence because despite being unaware of forming any intentions about the details of the turkey carving, the subsidiary unreflective actions are indeed intended. I intend to grab the carving knife rather than the bread knife, and I intend to begin by cutting the skin between the breast and the thighs rather than by splitting the bird in half with a butcher’s knife.

⁵ While both future- and present-direct intentions are in some sense future-directed, present-directed intentions are intentions occurring immediately prior to the onset of the whole action or intentions occurring during (or in) action.

⁶ Others have made similar distinctions. For example, Wilfrid Sellars distinguishes between intentions for the future and volitions, where volitions are what intentions for the future become when it’s time to act (Sellars 1966: p. 110). John McDowell (2010) argues that when ‘volition’ is understood in this way, then it cannot be used to refer to intentions that arise for the first time when it’s time to act. For that reason, he prefers to distinguish the notion of intentions in action, where the latter can be volitions in Sellars’ sense or novel present-directed intentions, which he prefers to express as ‘I am (willfully) ϕ -ing’ rather than ‘I will ϕ now’. For example, ‘I am raising my hand’ as opposed to ‘I will raise my hand now’ or ‘my hand is rising’.

But how can I intend to, say, cut the skin between the breast and the thighs when there is no point at which I consciously entertain the content of the intention?⁷ The answer is that carving a turkey is an agential, skilled unreflective action—it's an action I have carried out hundreds of times.⁸ For me, it has become routine, something I don't need to consciously reflect on, something I can do while chatting with friends, watching a movie in Spanish, or developing a new idea for a paper. Turkey carving to perfection, using all the right tools—that's my shtick.

Because turkey carving is a skilled unreflective action for me, stable object properties and stable affordances pertaining to turkeys and my turkey carving routine are stored in representations ('category schemas') in implicit, long-term memory. For each imminent step in the actual turkey carving routine, I implicitly believe I will perform it, and the Use pathway in my brain then effortlessly accesses the required information in implicit long-term memory and, in collaboration with the Grasp pathway, performs the computations for the motor representations that guide my habitual action.

Together with our explicit intentions, the implicit routine-based and more fine-grained intentions form a hierarchy of intentions representing the intended action and sub-actions at different levels of abstraction (Pacherie 2011, 2018). For example, an intention with the coarse-grained content *I shall now grab a knife* is a top-level intention in an agglomeration of intentions hierarchically ordered from more determinate to more determinate (e.g., *I shall now grab a knife*, *I shall now grab a carving knife*, *I shall now grab a carving knife with my right hand*, etc.) (Rizzolatti et al. 1988; Jeannerod 2006; Blomberg and Brozzo 2017; Brozzo 2017). Intentions at the lowest possible level are also sometimes referred to as 'motor intentions', a term first coined by Merleau-Ponty (1945).

Although we typically lack conscious awareness of the majority of our intentions, they are qua intentions consciously accessible, at least in hindsight (Davidson 1978; Mele 2010). In fact, implicit intentions seem to bear the marks of conscious intentions:

- I. The agent, if sincere, will be able to acknowledge that they intended to do something, even if the intention is only implicit. Say I am asked whether I intended to grab the carving knife rather than, say, the bread knife. I'd unhesitantly say 'yes, of course'.
- II. Implicit intentions can be integrated with other mental states. A classic example of integration is pre-attentive feature integration, say the integration of a shape and a color. In a similar vein, implicit intentions can be integrated with personal values to yield self-direction (Doris 2015: pp. 25–26). Say I value a turkey well carved, and implicitly intend to grab the carving knife rather than the bread knife. If I carve the turkey well as a result, then my carving expresses self-direction.
- III. Intentions are the product of integrating a belief with a (top-level) desire. Similarly, implicit intentions are the product of integrating an implicit belief with an implicit (top-level) desire.

⁷ On the distinction between conscious and unconscious intentions, see also Anthony Marcel (2003) and Alfred Mele (2010).

⁸ Skilled unreflective actions are actions that (1) involve mastery of a given skill, such as turkey carving, swimming or driving a car; and (2) unfold without the agent explicitly thinking about what she is doing (Brownstein 2014; see also Marcel 2003; Velleman 2008; Railton 2009; Annas 2011).

- IV. Implicit intentions play a key role in rationalizing and explaining the performance of the action and its proper action parts (Davidson 1980, 1987). For instance, my intention to carve the turkey explains why I grabbed the carving knife rather than the bread knife.
- V. Implicit intentions are suitable for serving as premises or conclusions in practical reasoning (Bratman 1987; Blomberg and Brozzo 2017). For instance, I may reason as follows: *I want to carve the turkey, and I want to do it now. To carve the turkey, I must grab the carving knife. Hence, I will now grab the carving knife.* If it turns out that reasoning is inherently a conscious process, we turn to Davidson's own thoughts on this issue. Davidson (1978: pp. 85–86) considers a case in which you add spice to a dish you are cooking, wanting to improve the taste and believing that adding spice will improve the taste. According to Davidson, you need not consciously entertain this belief-desire pair for you to have intentionally improved the taste of the dish. If you intentionally improve the taste of the dish, it must be true that you would have reasoned in this way, had you been aware of the relevant belief and desire.

The evidence thus suggests that implicit intentions are intentions on equal footing with the intentions we consciously entertain. In fact, implicit intentions may even be a better reflection of our agency than our conscious intentions, because they reflect the regularities and frequency of our past choices and are not spur-of-the-moment conscious intentions whose realization would require us to act out of character (Bargh and Barndollar 1996).

As Michael Bratman has pointed out, the level of abstraction of our conscious intentions depends on skill-level (Bratman 1984: p. 401). For example, when intentionally tying your shoe, you are in all likelihood only aware of a very coarse-grained intention to tie your shoe. But a child tying her own shoes for the first time after months of parental assistance is likely guided by considerably more fine-grained conscious intentions to tie a knot, make a loop with one of the laces, warp the other lace around the loop, pull the shoelace through the hole to form another loop, and hold both loops and pull them tight.

Despite the fine grain of the intentions at the lowest level of abstraction, none of the hierarchically organized intentions are as fine-grained as the motor representations that ultimately guide action. Suppose again that I intend to carve a turkey. Below the level of conscious awareness, this top-level intention is expanded into more fine-grained intentions specifying more details of my intended action. But even intentions at the lowest possible level available to consciousness, i.e., motor intentions, are not as fine-grained as the motor representations that determine the finest details of the movements constituting the action (Haggard 2005; Pacherie 2011, 2018; Butterfill and Sinigaglia 2014; Blomberg and Brozzo 2017). They are merely semi-executive; they represent the intended action without providing all of the specifications needed for the execution of the action. In order for my action to be executed, my brain must expand intentions at the lowest level of abstraction into precise motor representations that specify the properties and affordances of the turkey and the tools I will need as well as the kinematics of muscles and joints, including the precise trajectory needed for reaching, the precise hand aperture required for grasping, lifting and cutting, the

speed of the movements to be performed and the force that will be applied to the knife at the different stages. As information pertaining to the mechanics of objects and the kinematics of movement are not accessible to consciousness, it is impossible for us to form an intention with the same content as the motor representation that ultimately guides and sustains the action. The mechanical and kinematic details of the intended action, which are decided only at the lowest level of representation, are in an important sense mindless and non-intentional.

Before we proceed, a word on the agential status of intentions is in order. It has been argued on numerous occasions that our conscious experience of intentions in action is a result of motor representations that are rapidly developing and decaying in frontal and parietal motor areas (see e.g. Haggard 2005). Indeed, there is a barrage of studies seeming to show that the motor representation that represents what we will do is already encoded in the brain at the time at which we become conscious of our decision to do it (see e.g. Libet 1985; Haggard and Libet 2001; Haggard and Eimer 1999).

These findings should not be given too much weight, however. The participants in the studies were asked to make arbitrary choices among ends that they were relatively uninvested in (e.g., move either the right or the left hand). The lack of investment may indeed result in the brain making an arbitrary choice that then gives rise to the appearance of making a conscious decision (Cisek 2007). Alternatively, the appearance the volunteers had of making a conscious decision about what to do may have been preceded by an unconscious intention (Mele 2010).

More importantly for our purposes, the findings do not show that the implicit motor intentions that serve as a template for the still more fine-grained motor representations are preceded by neural action representations that determine what we do. To be sure, the brain could in principle generate motor representations representing, say, turkey carving prior to the agent forming a proximal intention to carve the turkey. But this is because turkey carving normally is part of a longer-lasting routine that begins with an intention to, say, cook the turkey. Since routines are coded for the long term in implicit memory, the brain could in principle generate coarse-grained ‘preparatory’ motor-representations for turkey carving before the agent had formed any intentions to initiate the carving. But coarse-grained motor representations are merely anticipatory schemas, not the fine-grained motor representations that ultimately guide and sustain actions. These particular findings concerning our supposed lack of ‘free will’ thus fail to show that intentions are merely epiphenomenal.

5 Nefarious skills and biased intentions

The dual action-system hypothesis provides us with the resources for explaining how implicit biases can bias motor representations and online actions without influencing the agent’s intentions. For the purposes of this paper, I will assume an account of implicit bias that we might call ‘the memory-association model’.⁹ On the envisaged model, implicit biases are distributed social category representations (or category

⁹ For details of this type of view, see e.g. Gendler (2011), Holroyd (2012, 2015) and Sullivan-Bissett (2015).

schemas) encoded in long-term memory in the ventral stream. The encoding occurs as a result of frequent and consistent co-activation of a social category concept and concepts representing stereotype features, social roles (i.e., social cues about how a member of a social group ought to behave or what they ought to do) or social affordances (i.e., the social cues for how to interact with a member of a social group).¹⁰ For example, frequent and consistent exposure to associations between JEW and AQUILINE NOSE, or MOTHER and RESPONSIBLE FOR CHILD REARING or MUSLIM MAN and AVOID will encode the association in long-term memory in such a way that the mere presence of category cues (e.g., seeing a dark-skinned person) ensures co-activation. Once encoded, an implicit bias can make a mark on visually-guided behavior in one of two ways: by biasing motor representations without biasing intentions, or by biasing intentions directly but covertly.

Biases that impinge on our motor representations without affecting the intentional hierarchy are ‘reflex-like’ inaccessible components of our behavior. This kind of biased behavior is a genuine form of ‘zombie behavior’, mindless and unintentional. How frequent this type of mindless microaggression takes place can only be decided empirically. But I suspect that it is less common that mindful microaggression, which we will turn to shortly. Startle responses are genuinely reflexive behavior. If someone you have never seen before unexpectedly enters your living room, you may react with a startle response. What triggers the startle response is the potential threat presented by the unexpected person in your living room. It is plausible that implicit biases can trigger startle responses. Clutching your purse when you pass your Muslim neighbor but not when you pass your Scandinavian neighbor could be a kind of startle response reflecting your bias-induced irrational fear of Muslims. Or suppose a teacher at a very ‘white’ college with an implicit bias against black people is startled when she unexpectedly steps into a classroom full of black students on the first day of class. If her startle response is undergirded predominantly by fear rather than surprise, this would be a case of an implicit bias tainting behavior in a mindless fashion. Other examples of ‘reflex-like’ microaggressive behavior include fixating your eyes on a person with facial deformities, asking a dark-skinned person where they are from (primed by skin color), finishing people’s sentences if they have a speech defect, such as stuttering, or performing actions slower when primed with an elderly stereotype (McConnell and Leibold 2001; Banfield et al. 2003).

Whether implicit biases trigger ‘reflex-like’ behavior is an empirical question, but it seems plausible that implicit biases can, and often do, shape our behavior in a more mindful fashion by shaping the intentions that serve as a template for motor representations. This sort of influence is a function of habitually acting on the basis of an implicit bias, thereby establishing a pattern of automatized behavior. Suppose you are playing on the street when a black kid from your neighborhood named ‘Jacob’ tosses a baseball in your direction, thereby signaling that he wants you to toss it back to him. You harbor an implicit dislike of black kids and often start fights with black

¹⁰ On this model, implicit biases do not encompass implicit prejudices understood as affective responses. I take stereotypes to accommodate evaluative responses, like *Muslims are dangerous* or *Blacks are aggressive* (see Haslanger 2012, 2013). On an alternative model, implicit biases are clusters of co-activating representational and affective components (‘aliefs’ in Tamar Gendler’s 2008a, b sense; see also Gendler 2011, 2012; Amodio 2014; Madva and Brownstein 2018; Brogaard 2020).

kids in the neighborhood. You have never been in a fight with Jacob, however. He usually keeps to himself. So, when he tosses the baseball in your direction, you decide to toss it back to him. However, owing to your implicit dislike of black kids, you put an excessive amount of force into your throw and aim in the direction of Jacob's face. The ball nearly hits Jacob in the face; he is able to dodge it only by diving into the street.

In the envisaged scenario, your implicit bias exerts an influence on parts of the motor representation that are inaccessible to consciousness, such as information about how to adjust the handgrip, how much to rotate the shoulder, how much force to put into the throw before release, etc. These involuntary aspects of the throw are non-agential. But your tossing of the ball to Jacob with considerable force is quite unlike 'reflex-like' behavior that lack an initial intention and mindful guidance. Your action exemplifies your routine behavior of being aggressive towards black kids. Aggressing against black kids is something you regularly and intentionally with the intention of getting into a fight with them, and we can stipulate that this is also the case in this scenario. On a conscious level, you merely intended to toss the baseball back to Jacob, but this intention by itself does not suffice for your dorsal stream to generate a motor representation. The dorsal stream is not in the business of arbitrarily deciding between a toss with a force that can knock someone out and a gentle toss. It needs guidance by a finer-grained intention in order for it to compute the precise physical parameters (e.g., the exact force and trajectory) required for the execution of the action. Of course, intentions can misfire. You can intend to throw the ball gently, yet accidentally throw it hard. This can happen because you lack the skill or the ability to exercise it. If, however, you have the skill and the ability to throw the ball in a skilled way, as we have been assuming, then you throwing the ball back to Jacob with considerable force was an intended action, even if you were unaware of it at the time.

The question, of course, is whether your answers to the Anscombean questions 'What are you doing?', 'Why are you doing it?'—a hallmark of intention—will support this conclusion (Anscombe 1957). Replying with 'tossing the ball back to Jacob' and 'he tossed the ball to me' do not suggest that you intended to hurl the ball back to Jacob rather than tossing it gently. But we cannot expect you to be in a position to consciously access your covert intention without a prompt. If you were asked why you were throwing the ball so forcefully, you might reply with 'I don't know' or some confabulation along the lines of 'That's how it's done. I didn't know he doesn't know how to play' or 'He seemed like he was going to start a fight' (Sullivan-Bissett 2015). Both kinds of responses are admissions of intention. They are not admissions of intentions on the basis of implicit biases, but this is unsurprising given the relative, cognitive inaccessibility of implicit biases.

Here is another scenario exemplifying how implicit biases might influence our intentions. Suppose you are a male philosopher with an implicit bias against female philosophers. Your implicit bias routinely makes you give more attention to male philosophers than female philosophers in social group settings. At a conference you are attending, you are introduced to two graduate students, Sally and Sid, and you decide to stay and chat with them for a while. Yet you automatically direct your attention to Sid and start asking him questions about his research, mostly ignoring Sally—an archetypal example of micro-aggression against women in male-dominated

academic fields. In fact, you would characterize your action this way, if you later came to realize what you had done. So, it's not the case that you fail to know that the type of behavior you are unreflectively engaging in is discriminatory. Rather, because you are following a set routine, it doesn't occur to you that you are intentionally engaging in this type of behavior.

In the envisaged scenario, your action involves mastery of a skill, not an admirable skill but rather a nefarious skill, viz. that of inadvertently interacting mostly with male philosophers in social group settings. Despite its vicious nature, your discriminatory behavior is an exemplar of a skilled unreflective yet agential action. Your action, while not anticipated by conscious intention, is nonetheless intentional. By unreflectively following a retrieved schema for this type of interactions, you implicitly intend to interact mostly with the male graduate student. In fact, you will be able to answer the Anscombean questions, 'What are you doing?', 'Why are you doing it?' ('I am asking Sid questions about his work', 'I am interested in what Sid is working on', 'Sally wasn't really very talkative, kind of shy, maybe new to the field'). It is true, of course, that you didn't deliberately intend to discriminate against a female graduate student, but that is exactly what distinguishes discriminatory behavior you consciously intend to execute, and discriminatory behavior that is unreflective and automatized yet agential. Excusing yourself by saying 'It's true that I didn't ask Sally about her work but I didn't mean to act biased against her' is on a par with a murderer saying 'Yes, I stabbed her with a knife but I didn't mean to harm her'. In the court of law, this excuse would be exculpatory only under conditions of legal insanity, which is clearly irrelevant to our purposes here.

The effect of implicit biases on action can also be a self-fulfilling prophecy (Bargh et al. 1996; Bargh and Chartrand 1999). Suppose after chatting with Sid for a while, you turn to Sally and ask her about her work. Unbeknownst to you, however, her feeling of exclusion has made her anxious, she suddenly can't remember the details of her work and replies with disengaging phrases like 'I am not really sure' or 'I don't really remember'. As your attempts to engage with Sally are unsuccessful, you give up and continue chatting with Sid. If your colleague later asks you why you didn't talk to Sally, you might again simply reply with 'she wasn't really very talkative, kind of shy, maybe new to the field'.

The types of biased action just discussed have a level of complexity that we don't see in empirical studies of visually guided action. However, it is not hard to imagine scenarios involving bias-driven action that are more similar in complexity to the kinds of visually guided action that have been studied empirically, the skilled use of tools. Dentistry requires a barrage of tools, such as syringes for administering novocaine, a sickle probe used for finding cavities or gum disease and scraping away tartar and plaque, a suction device for removing excess saliva, dental drills for removing tooth decay, and extraction tools. Using these tools requires expert skill and precise visually guided action. But now suppose you are a dentist with an implicit bias against black people. The information about black people encoded in long-term memory represents blacks as athletic, physically and emotionally strong, thick-skinned and with a high tolerance for pain. As a result, you administer less novocaine to your black patients than to your white patients, and you are considerably less gentle with you dentist tools when it's a black person who is sitting in the dental chair. A lot of black patients in

your clinic are in a lot of pain while you work on their teeth, but they behave stoically and never complain, so you never become aware of your harsh treatment of your black patients. In the envisaged case, your implicit bias against black people modulates how you use your dental tools to fix black people's teeth, and it does so by covertly shaping the motor intentions that, via more fine-grained motor representations, will serve as a template for the execution of your actions.

6 Concluding remarks: Are we accountable for microaggression?

According to the two-stream hypothesis, the brain's visual system is constituted by two anatomically and functionally distinct pathways: a dorsal stream for the control of actions and a ventral stream for conscious identification and recognition of objects. It was originally thought that the interaction between the two streams was rather limited. However, recent research has shown significant interaction for the case of complex skilled action, such the manipulation of tools. The dorsal stream is divided into a dorso-dorsal 'Grasp' pathway and a ventro-dorsal 'Use' pathway that serves as an interface between the perception and cognition systems in the ventral stream and the action system in the dorsal stream.

The new findings suggest that skilled action, despite its automaticity, may be under mindful control to a much greater extent than previously thought. Our conscious intention to perform an action likely only specifies a very coarse-grained version of the intention that serves as a template for the motor representation that ultimately guides the action. For example, you may intentionally grab your coffee mug by its handle in order to take a sip of coffee without consciously intending to grab it this way. Knowledge of how you can manipulate different types of objects and how you normally do it is stored in implicit, long-term memory. But this knowledge enters the action system via the ventro-dorsal pathway, where it is used to expand your conscious top-level intention(s) into more fine-grained intentions that can be used to generate a motor representation suitable for guiding the action.

The new findings suggest that when an implicit bias inserts itself into our everyday actions, the action can nonetheless still be skill-based and agential, despite being irreflective. For example, suppose while chatting with two graduate students, a male and a female, you engage mostly with the male graduate student. This behavior may be the result of an implicit bias that makes you more inclined to engage intellectually with male students than female students. Your discriminatory behavior is then a skillful irreflective action guided by intentions that do not reach conscious awareness.

This hypothesis bears on the question of whether you are accountable for your biased behavior. 'Reflex-like' reactions to an activation of an implicit bias, such as an involuntary shift in gaze when passing a person in a wheelchair, are discriminatory but they do not raise questions about accountability (Doris 2015).¹¹ An intentional agential action that is discriminatory, by contrast, is clearly something the agent should

¹¹ Doris (2015) offers an appealing, compatibilist account of "accountability" (or "responsibility") in terms of "agency." To a first approximation, an agent exercises (full-blown rather than mere causal) agency when her action is an expression of *her* values. Whether all implicitly intentional, discriminatory actions will turn out to be expressions of agential values is a question I will leave for a future occasion.

be held accountable for (Machery et al. 2010; Holroyd 2012; Saul 2012, 2013; Doris 2015; Brogaard 2019). It does not reflect as badly on the agent as it does on someone who engages in discrimination fully aware of doing so, but a lack awareness of one's actions or omissions is not by itself exculpatory—it is not when the action is agential (King and Carruthers forthcoming).

In performing a skilled, agential action the agent acted negligently just when her action was avoidable, and she failed to take reasonable care to avoid it. The last clause can be rendered as follows: 'a reasonable person would have realized that it was more likely than not that the action would inflict harm on the other person and therefore would have taken measures to try to avoid it'. An action, unlike 'reflex-like' behavior, can be avoided either directly or indirectly (Mandelbaum 2016). This distinction can be cashed out as follows:

Directly avoidable One can cease to perform the action directly by intending to do so. E.g. Displaying a caricature of a Native American on a game uniform. *Future preventive measure* Stop wearing the uniform.

Indirectly avoidable One can cease to perform the action by intending to do something else that will prevent the action or make it less likely to occur. E.g., cutting off a female student before she is done speaking owing to an overestimate of the duration of the time women (but not men) hold the floor. *Future preventive measure* When you are about to cut off a female student, wait another 2–3 min.

Various combinations of methods can be implemented to counter microaggression and help minimize implicit biases, including: peer monitoring (Ex: Caro: 'Did you see Jon plug his new headware on Snapchat every day for the past week?' Bel: 'Dude! He's Jewish'.), self-monitoring (e.g., paying attention to patterns of default diffuse or focal attention or patterns of idiosyncratic stereotyping, Saul 2012; Holroyd 2015), group depolarization (e.g., joining a heterogeneous discussion forum like *Baseball Fanatics* or *Addicted to Dogs*; Brogaard 2020), outgroup socializing (e.g., befriending or 'ingrouping' outgroup members; Dasgupta and Greenwald 2001), cultural conditioning (e.g., broadcasting stereotype-negating films, songs and podcasts; displaying stereotype-canceling signs and décor; implementing a stereotype-discordant curriculum; playing counter-stereotypical games, Kawakami et al. 2007; Mendoza et al. 2010; Saul 2012), and outgroup camouflaging (e.g., anonymized job searches, voice-disguised Skype interviews, anonymized auditions for orchestra).

Acknowledgements For helpful comments or discussion of issues addressed in this paper, I am grateful to three anonymous reviewers for this journal, Bartek Chomanski, Jonathan Cohen, Andy Cui, John Doris, Dimitria E. Gatzia, Sally Haslanger, Aleks Hernandez, Bob Kentridge, Casey Landers, Azenet Lopez, Eric Mandelbaum, Chris Peacocke, Jesse Prinz, David Rosenthal, Mark Rowlands, Michael Slote, Thomas Alrik Sørensen, and audiences at the 2017 Meeting of the SPP, the 2018 Meeting of the SSPP, Duke University, NYU, Stanford University, University of Miami, University of Missouri, St. Louis, University of Oslo, University of Aalborg, University of Washington University, St. Louis, Aarhus University, and the Center for Humanities Fellowship Recipients, University of Miami 2018–2019.

References

- Aglioti, S., Goodale, M. A., & DeSouza, J. F. X. (1995). Size-contrast illusions deceive the eye but not the hand. *Current Biology*, 5, 679–685.
- Amodio, D. M. (2014). The neuroscience of prejudice and stereotyping. *Nature Reviews Neuroscience*, 15(10), 670–682.
- Andersen, R. A., Andersen, K. N., Hwang, E., & Hauschild, M. (2014). Optic ataxia: From Balint’s syndrome to the parietal reach region. *Neuron*, 81(5), 967–983. <https://doi.org/10.1016/j.neuron.2014.02.025>.
- Annas, J. (2011). *Intelligent virtue*. New York: Oxford University Press.
- Anscombe GEM. (1957/2000). *Intention*. Cambridge, MA: Harvard University Press.
- Banfield, J. F., Pendry, L. F., Mewse, A. J., & Edwards, M. G. (2003). The effects of an elderly stereotype prime on reaching and grasping actions. *Social Cognition*, 21(4), 299–319.
- Bargh, J. A., & Barndollar, K. (1996). Automaticity in action: The unconscious as repository of chronic goals and motives. In P. M. Gollwitzer & J. A. Bargh (Eds.), *The psychology of action: Linking cognition and motivation to behavior* (pp. 457–481). New York, NY: Guilford Press.
- Bargh, J. A., & Chartrand, T. L. (1999). The unbearable automaticity of being. *American Psychologist*, 54, 462–479.
- Bargh, J. A., Chen, M., & Burrows, L. (1996). Automaticity of social behavior: Direct effect of trait construct and stereotype activation on action. *Journal of Personality and Social Psychology*, 71, 230–244.
- Binkofski, F., & Buxbaum, L. J. (2013). Two action systems in the human brain. *Brain and Language*, 127, 222–229.
- Binkofski, F., & Fink, G. (2005). Apraxien (Apraxias). *Nervenarzt*, 76, 493–511.
- Binkofski, F., Reetz, K., & Blangero, A. (2007). Tactile agnosia and tactile apraxia: Cross talk between the action and perception streams in the anterior intraparietal area. *Behavioral and Brain Sciences*, 30(2), 201–202.
- Blomberg, O., & Brozzo, C. (2017). Motor intentions and non-observational knowledge of action: A standard story. *Thought*, 6, 137–146.
- Borghi, A. M., & Riggio, L. (2015). Stable and variable affordances are both automatic and flexible. *Frontiers and Human Neuroscience*, 9(351). <https://doi.org/10.3389/fnhum.2015.00351>.
- Borra, E., Belmalih, A., Calzavara, R., Gerbella, M., Murata, A., Rozzi, S., et al. (2008). Cortical connections of the macaque anterior intraparietal (AIP) area. *Cerebral Cortex*, 18(5), 1094–1111.
- Bratman, M. E. (1984). Two faces of intention. *Philosophical Review*, 93(3), 375–405.
- Bratman, M. E. (1985). “Davidson’s theory of intention,” reprinted in *Faces of Intention* (1999) (pp. 209–224). Cambridge: Cambridge University Press.
- Bratman, M. E. (1987). *Intention, plans and practical reason*. Cambridge, MA: Harvard University Press.
- Brenner, E., & Smeets, J. B. J. (1996). Size illusion influences how we lift but not how we grasp an object. *Experimental Brain Research*, 111, 473–476.
- Briscoe, R., & Schwenkler, J. (2015). Conscious vision in action. *Cognitive Science*, 39(7), 1435–1467.
- Brogaard, B. (2011). Are there unconscious perceptual processes? *Consciousness and Cognition*, 20, 449–463.
- Brogaard, B. (2019). Bias-driven attention, cognitive penetration and epistemic downgrade. In C. Limbeck & F. Stadler (Eds.), *Philosophy of perception* (pp. 199–216). Berlin: De Gruyter.
- Brogaard, B. (2020). *Hatred: Our most dangerous emotion*. Oxford: Oxford University Press.
- Brogaard, B., Marlow, K., Overgaard, M., Schwartz, B. L., Zopluoglu, C., Tomson, S., et al. (2017). Deaf hearing: Implicit discrimination of auditory content in a patient with mixed hearing loss. *Philosophical Psychology*. <https://doi.org/10.1080/09515089.2016.1268680>.
- Brouwer, A.-M., Georgiou, I., Glover, S., & Castiello, U. (2006). Adjusting reach to lift movements to sudden visible changes in target’s weight. *Experimental Brain Research*, 173, 629–636.
- Brownstein, M. (2014). Rationalizing flow: Agency in skilled unreflective action. *Philosophical Studies*, 168(2), 545–568.
- Brozzo, C. (2017). Motor intentions: How intentions and motor representations come together. *Mind and Language*, 32(2), 231–256.
- Bruno, N. (2001). When does action resist visual illusions? *Trends in Cognitive Sciences*, 5, 385–388.
- Bruno, N., & Franz, V. H. (2009). When is grasping affected by the Müller-Lyer illusion? A quantitative review. *Neuropsychologia*, 47, 1421–1433.
- Burnston, D. C. (2017a). Cognitive penetration and the cognition–perception interface. *Synthese*, 194, 3645–3668.

- Burnston, D. C. (2017b). Interface problems in the explanation of action. *Philosophical Explorations*, 20(2), 242–258.
- Buxbaum, L. J., & Saffran, E. M. (2002). Knowledge of object manipulation and object function: Dissociations in apraxic and nonapraxic subjects. *Brain and Language*, 82(2), 179–199.
- Butterfill, S. A., & Sinigaglia, C. (2014). Intention and motor representation in purposive action. *Philosophy and Phenomenological Research*, 88(1), 119–145.
- Buxbaum, L. J., & Kalénine, S. (2010). Action knowledge, visuomotor activation, and embodiment in the two action systems. *Annals of the New York Academy of Sciences*, 1191, 201–218.
- Buxbaum, L. J., Kyle, K. M., & Menon, R. (2005). On beyond mirror neurons: Internal representations subserving imitation and recognition of skilled object-related actions in humans. *Cognitive Brain Research*, 25(1), 226–239.
- Buxbaum, L. J., Sirigu, A., Schwartz, M. F., & Klatzky, R. (2003). Cognitive representations of hand posture in ideomotor apraxia. *Neuropsychologia*, 41, 1091–1113.
- Caiani, S. Z., & Ferretti, G. (2017). Semantic and pragmatic integration in vision for action. *Consciousness and Cognition*, 48, 40–54.
- Chechlac, M. (2018). Bilateral parietal dysfunctions and disconnections in simultanagnosia and Bálint syndrome. *Handbook of Clinical Neurology*, 151, 249–267.
- Cisek, P. (2007). Cortical mechanisms of action selection: The affordance competition hypothesis. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362(1485), 1585–1599.
- Cloutman, L. L. (2013). Interaction between dorsal and ventral processing streams: Where, when and how? *Brain and Language*, 127, 251–263.
- Corbetta, M., & Shulman, G. L. (2002). Control of goal-directed and stimulus-driven attention in the brain. *Nature Reviews Neuroscience*, 3(3), 201–215.
- Daprati, E., & Sirigu, A. (2006). How we interact with objects: learning from brain lesions. *Trends in Cognitive Sciences*, 10(6), 265–270.
- Dasgupta, N., & Greenwald, A. G. (2001). On the malleability of automatic attitudes: Combating automatic prejudice with images of admired and disliked individuals. *Journal of Personality and Social Psychology*, 81(5), 800–814.
- Davidson, D. (1978). “Intending,” reprinted in *Essays on Actions and Events* (1980) (pp. 83–102). Oxford: Oxford University Press.
- Davidson, D. (1980). *Essays on action and events*. Oxford: Oxford University Press.
- Davidson, D. (1987). Knowing one’s own mind. *Proceedings and Addresses of the American Philosophical Association*, 60, 441–458.
- Doris, J. (2015). *Talking to our selves: Reflection, ignorance, and agency*. Oxford: Oxford University Press.
- Evans, C., Edwards, M. G., Taylor, L. J., & Ietswaart, M. (2016). Impaired communication between the dorsal and ventral stream: Indications from apraxia. *Frontiers in Human Neuroscience*, 10, 8. <https://doi.org/10.3389/fnhum.2016.00008>.
- Fatima, S. (2017). On the edge of knowing: Microaggression and epistemic uncertainty as a woman of color. In K. Cole & H. Hassel (Eds.), *Surviving sexism in academia: Feminist strategies for leadership* (pp. 147–157). London: Routledge.
- Ferretti, G. (2016). Through the forest of motor representations. *Consciousness and Cognition*, 43, 177–196.
- Franz, V. H., & Gegenfurtner, K. R. (2008). Grasping visual illusions: Consistent data and no dissociation. *Cognitive Neuropsychology*, 25, 920–950.
- Franz, V. H., Gegenfurtner, K. R., Bühlhoff, H. H., & Fahle, M. (2000). Grasping visual illusions: No evidence for a dissociation between perception and action. *Psychological Science*, 11, 20–25.
- Franz, V. H., Hesse, C., & Kollath, S. (2009). Visual illusions, delayed grasping, and memory: No shift from dorsal to ventral control. *Neuropsychologia*, 47, 1518–1531.
- Gallese, V., Craighero, L., Fadiga, L., & Fogassi, L. (1999). Perception through action. *Psyche: An Interdisciplinary Journal of Research On Consciousness*, 5, 21.
- Garcea, F. E., & Buxbaum, L. J. (2019). Gesturing tool use and tool transport actions modulates inferior parietal functional connectivity with the dorsal and ventral object processing pathways. *Human Brain Mapping*, 40(10), 2867–2883.
- Gendler, T. S. (2008a). Alief and belief. *The Journal of Philosophy*, 105(10), 634–663.
- Gendler, T. S. (2008b). Alief in action (and reaction). *Mind and Language*, 23(5), 552–585.
- Gendler, T. S. (2011). On the epistemic costs of implicit bias. *Philosophical Studies*, 156, 33–63.
- Gendler, T. S. (2012). Between reason and reflex: Response to commentators. *Analysis*, 72(4), 799–811.

- Gibson, J. J. (1977). The theory of affordances. In R. Shaw & J. Bransford (Eds.), *Perceiving, acting, and knowing: Toward an ecological psychology* (pp. 67–82). Hillsdale, NJ: Erlbaum.
- Goodale, M. A. (2014). How (and why) the visual control of action differs from visual perception. *Proceedings. Biological Sciences*, 281(1785), 20140337.
- Goodale, M. A., & Milner, A. D. (1992). Separate visual pathways for perception and action. *Trends in Neurosciences*, 15, 20–25.
- Goodale, M. A., Milner, A. D., Jakobson, L. S., & Carey, D. P. (1991). A neurological dissociation between perceiving objects and grasping them. *Nature*, 349, 154–156.
- Grafton, S. T. (2010). The cognitive neuroscience of prehension: recent developments. *Experimental Brain Research*, 204, 475–491.
- Grill-Spector, K., Kourtzi, Z., & Kanwisher, N. (2001). The lateral occipital complex and its role in object recognition. *Vision Research*, 41, 1409–1422.
- Haggard, P. (2005). Conscious intention and motor cognition. *Trends in cognitive sciences*, 9(6), 290–295.
- Haggard, P., & Eimer, M. (1999). On the relation between brain potentials and the awareness of voluntary movements. *Experimental Brain Research*, 126, 128–133.
- Haggard, P., & Libet, B. W. (2001). Conscious intention and brain activity. *Journal of Consciousness Studies*, 8(11), 47–63.
- Haslanger, S. (2012). *Resisting reality: Social construction and social critique*. New York: Oxford University Press.
- Haslanger, S. (2013). *Social Meaning and Philosophical Method*, Presidential Address, Eastern Division of the American Philosophical Association.
- Holroyd, J. (2012). Responsibility for Implicit Bias. *Journal of Social Philosophy*, 43(3), 274–306.
- Holroyd, J. (2015). Implicit bias, awareness and imperfect cognitions. *Consciousness and Cognition*, 33, 511–523.
- Jeannerod, M. (2006). *Motor cognition: What actions tell the self*. New York: Oxford University Press.
- Jeannerod, M. J., Decety, J. F., & Michel, F. (1994). Impairment of grasping movements following a bilateral posterior parietal lesion. *Neuropsychologia*, 32, 369–380.
- Kawakami, K., Dovidio, J., & van Kamp, S. (2007). The impact of counterstereotypic training and related correction processes on the application of stereotypes. *Group Processes and Intergroup Relations*, 10(2), 139–156.
- King, M., & Carruthers, P. (forthcoming). Responsibility and consciousness. In D.K. Nelkin and D. Pereboom (Eds.), *Oxford Handbook on Moral Responsibility*.
- Kriegeskorte, N., Mur, M., Ruff, D. A., Kiani, R., Bodurka, J., Esteky, H., et al. (2008). Categorical object representations in inferior temporal cortex of man and monkey. *Neuron*, 60(6), 1126–1141.
- Libet, B. (1985). Unconscious cerebral initiative and the role of conscious will in voluntary action. *The Behavioral and Brain Sciences*, 8, 529–566.
- Machery, E., Faucher, L., & Kelly, D. (2010). On the alleged inadequacy of psychological explanations of racism. *The Monist*, 93(2), 228–255.
- Madva, A., & Brownstein, M. (2018). Stereotypes, prejudice, and the taxonomy of the implicit social mind. *Noûs*, 52(3), 611–644.
- Mahon, B. Z., & Wu, W. (2015). Cognitive penetration of the dorsal visual stream? In J. Zeimbekis & A. Raftopoulos (Eds.), *The cognitive penetrability of perception: New philosophical perspectives*. <https://doi.org/10.1093/acprof:oso/9780198738916.001.0001>.
- Mandelbaum, E. (2016). Attitude, association, and inference: On the propositional structure of implicit bias. *Noûs*, 50(3), 629–658.
- Marcel, A. (2003). The sense of agency: Awareness and ownership of action. In J. Roessler & N. Eilan (Eds.), *Agency and self-awareness* (pp. 48–93). Oxford: Clarendon Press.
- Marr, D. (1982). *Vision*. San Francisco: Freeman.
- McConnell, A., & Leibold, J. (2001). Relations among the implicit association test, discriminatory behavior, and explicit measures of racial attitudes. *Journal of Experimental Social Psychology*, 37(5), 435–442.
- McDowell, J. (2010). What is the content of an intention in action. *Ratio*, 23(4), 415–432.
- McIntosh, R. D., & Lashley, G. (2008). Matching boxes: Familiar size influences action programming. *Neuropsychologia*, 46(9), 2441–2444.
- Mele, A. R. (1992). *Springs of action*. New York: Oxford University Press.
- Mele, A. R. (2010). Conscious intentions. In J. K. Campbell, M. O'Rourke, & H. S. Silverstein (Eds.), *Action, ethics, and responsibility* (pp. 85–108). Cambridge, MA: MIT Press.

- Mendoza, S., Gollwitzer, P., & Amodio, D. (2010). Reducing the expression of implicit stereotypes: Reflexive control through implementation intentions. *Personality and Social Psychology Bulletin*, 36(4), 512–523.
- Merleau-Ponty, M. (1945/1965). *Phenomenology of perception*. London: Routledge.
- Milner, A. D., Dijkerman, H. C., Pisella, L., McIntosh, R. D., Tilkete, C., Vighetto, A., et al. (2001). Grasping the past: Delay can improve visuomotor performance. *Current Biology*, 11, 1896–1901.
- Milner, A. D., & Goodale, M. A. (1995). *The visual brain in action*. New York: Oxford University Press.
- Milner, A. D., & Goodale, M. A. (2008). Two visual systems reviewed. *Neuropsychologia*, 46, 774–785.
- Mozaz, M., Rothi, L., Anderson, J., Crucian, G. P., & Heilman, K. (2002). Postural knowledge of transitive pantomimes and intransitive gestures. *Journal of the International Neuropsychological Society*, 8, 958–962.
- Nanay, B. (2013a). *Between perception and action*. Oxford: Oxford University Press.
- Nanay, B. (2013b). Is action-guiding vision cognitively impenetrable? In N. J. Hillsdale (Ed.), *Proceedings of the 35th annual conference of the cognitive science society (CogSci 2013)* (pp. 1055–1060). New York: Lawrence Erlbaum.
- Nelissen, K., & Vanduffel, W. (2011). Grasping-related functional magnetic resonance imaging brain responses in the macaque monkey. *Journal of Neuroscience*, 31(22), 8220–8229.
- O’Shea, R. P., Chandler, N. P., & Roy, R. (2013). Dentists make larger holes in teeth than they need to if the teeth present a visual illusion of size. *PLoS ONE*, 8(10), e77343.
- Pacherie, E. (2011). Nonconceptual representations for action and the limits of intentional control. *Social Psychology*, 42(1), 67–73.
- Pacherie, E. (2018). Motor Intentionality. In A. Newen, L. De Bruin, & S. Gallagher (Eds.), *The Oxford handbook of 4E cognition, Chapter 19*. Oxford: Oxford University Press.
- Peacocke, C. (1992). *A study of concepts*. Cambridge, MA: MIT Press.
- Pelgrims, B., Andres, M., Seron, X., Duhamel, J.-R., Sirigu, A., & Olivier, E. (2005). *Role of the left supramarginal gyrus in coding hand gestures: Implication for ideomotor apraxia* (p. 2005). New York: Cognitive Neuroscience Society.
- Perenin, M. T., & Vighetto, A. (1988). Optic ataxia: A specific disruption in visuomotor mechanisms. *Brain*, 111, 643–674.
- Pisella, L., Binkofski, F., Lasek, K., Toni, I., & Rossetti, Y. (2006). No double dissociation between optic ataxia and visual agnosia: Multiple sub-streams for multiple visuo-manual integrations. *Neuropsychologia*, 44, 2734–2748.
- Prinz, J. (2000). A neurofunctional theory of visual consciousness. *Consciousness and Cognition*, 9(2), 243–259.
- Prinz, J. (2012). *The conscious brain*. Oxford: Oxford University Press.
- Railton, P. (2009). Practical competence and fluent agency. In D. Sobel & S. Wall (Eds.), *Reasons for action* (pp. 81–115). Cambridge: Cambridge University Press.
- Randerath, J., Goldenberg, G., Spijkers, W., Li, Y., & Hermsdörfer, J. (2010). Different left brain regions are essential for grasping a tool compared with its subsequent use. *Neuroimage*, 53(1), 171–180.
- Reed, C. L., Klatzky, R. L., & Halgren, E. (2005). What vs. where in touch: an fMRI study. *Neuroimage*, 25, 718–726.
- Rizzolatti, G., Camarda, R., Fogassi, L., Gentilucci, M., Luppino, G., & Matelli, M. (1988). Functional organization of inferior area 6 in the macaque monkey. II. Area F5 and the control of distal movements. *Exp Brain Res.*, 71(3), 491–507.
- Rizzolatti, G., & Matelli, M. (2003). Two different streams form the dorsal visual system: Anatomy and functions. *Experimental Brain Research*, 153, 146–157.
- Romanski, L. M., Tian, B., Fritz, J., Mishkin, M., Goldman-Rakic, P. S., & Rauschecker, J. P. (1999). Dual streams of auditory afferents target multiple domains in the primate prefrontal cortex. *Nature Neuroscience*, 2, 1131–1136.
- Sakreida, K., Effnert, I., Thill, S., Menz, M. M., Jirak, D., Eickhoff, C. R., et al. (2016). Affordance processing in segregated parieto-frontal dorsal stream sub-pathways. *Neuroscience Biobehavior Review*, 69, 89–112.
- Saul, J. (2012). Ranking exercises in philosophy and implicit bias. *Journal of Social Philosophy*, 43(3), 256–273.
- Saul, J. (2013). Implicit bias, stereotype threat, and women in philosophy. In F. Jenkins & K. Hutchison (Eds.), *Women in philosophy: What needs to change?* (pp. 39–60). Oxford: Oxford University Press.

- Schenk, T. (2006). An allocentric rather than perceptual deficit in patient D.F. *Nature Neuroscience*, *9*, 1345–1347.
- Schenk, T., & McIntosh, R. D. (2010). Do we have independent visual streams for perception and action? *Cognitive Neuroscience*, *1*, 52–62.
- Schenk, T., Ellison, A., Nichola, R., & Milner, A. D. (2005). The role of V5/MT+ in the control of catching movements: An rTMS study. *Neuropsychologia*, *43*, 189–198.
- Sellars, W. (1966). Thought and action. In K. Lehrer (Ed.), *Freedom and determinism*. New York: Random House.
- Shepherd, J. (forthcoming). Skilled action and the double life of intention. *Philosophy and Phenomenological Research*.
- Smeets, J. B. J., & Brenner, E. (2006). 10 years of illusions. *The Journal of Experimental Psychology: Human Perception and Performance*, *32*, 1501–1504.
- Smeets, J. B. J., Brenner, E., de Grave, D. D. J., & Cuijpers, R. H. (2002). Illusions in action: Consequences of inconsistent processing of spatial attributes. *Experimental Brain Research*, *147*, 135–144.
- Sue, D. W. (2010). *Microaggressions in everyday life: Race, gender, and sexual orientation*. New York: Wiley-Blackwell.
- Sullivan-Bissett, E. (2015). Implicit bias, confabulation, and epistemic innocence. *Consciousness and Cognition*, *33*, 548–560.
- Takemura, H., Rokem, A., Winawer, J., Yeatman, J. D., Wandell, B. A., & Pestilli, F. (2016). A major human white matter pathway between dorsal and ventral visual cortex. *Cerebral Cortex*, *26*(5), 2205–2214.
- Toribio, J. (forthcoming). Are visuomotor representations cognitively penetrable? Biasing action-guiding vision. *Synthese*, 1–19.
- Toribio, J. (2018a). Implicit bias: From social structure to representational format. *Theoria: An International Journal for Theory, History and Foundations of Science*, *33*(1), 41–60.
- Toribio, J. (2018b). Visual experience: Rich but impenetrable. *Synthese*, *195*(8), 3389–3406.
- van Polanena, V., & Davare, M. (2015). Interactions between dorsal and ventral streams for controlling skilled grasp. *Neuropsychologia*, *79*(Part B), 186–191.
- Velleman, J. D. (2008). The way of the wanton. In C. Mackenzie & K. Atkins (Eds.), *Practical identity and narrative agency*. New York: Routledge.
- Weiskrantz, L. (1986). *Blindsight: A case study and implications*. Oxford: Oxford University Press.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.