THE DO-ABLE SOLUTION TO THE INTERFACE PROBLEM*

[Penultimate draft; Final version forthcoming in *Philosophy and Phenomenological Research*]

Abstract. Philosophers and cognitive scientists increasingly recognize the need to appeal to motor representations over and above intentions in attempting to understand how action is planned and executed. But doing so gives rise to a puzzle, which has come to be known as "the Interface Problem": How is it that intentions and motor representations manage to interface in producing action? The question has seemed puzzling, because each state is thought to be formatted differently: Intention has propositional format, whereas the format of motor representation is motoric. My primary goal here is to defend a novel and attractive (dis)solution to the interface problem. I do so by connecting it with a rather different discussion about the format of intention, instigated by a minority of philosophers who reject the idea that intention should be construed as a propositional attitude. As I explain, the most compelling reason to accept the heterodox non-propositional conception of intention actually holds the key also to explaining away the interface problem. In so doing, the heterodox conception itself gains further credibility.

1. Introduction

Nearly everyone today bar eliminativists agrees that it is necessary to posit representational states – roughly, information bearing structures in the mind/brain – when explaining action and behaviour more broadly. With time and increasing levels of sophistication, the range of representational state-types postulated by philosophers and cognitive scientists has expanded beyond the stock-in-trade mentalia of beliefs, intentions, desires, images, precepts etc., to include also a widening array of sub-personal and sub-doxastic states. But the expansion brings in its wake questions about the possibility of *interface* between states of different types, specifically those state-types whose information is thought to be encoded in different formats. This paper is about specifically one such question or problem of interface, which has earned itself the honorific definite article, *The* interface problem. First identified by Stephen Butterfill and Corrado Sinigaglia

^{*} Comments and suggestions by Harvey Lederman, Arnon Levy, Josh Shepherd, and two anonymous referees for this journal greatly improved the manuscript. I am very grateful to all of them. Many thanks also to the constructive discussion of material from this paper by audiences at CUNY, The Hebrew University of Jerusalem, the University of

Milan and the ACU workshop on Action, Intention, and Language. Research for this paper was supported by the Israel Science Foundation (grant no. 381/23).

(2014), it is a problem about how intention, whose format is standardly assumed to be propositional like virtually all other intentional attitudes, manages to interact with motor representations, whose format is thought to be non-propositional. Butterfill & Sinigaglia's discussion has spurred philosophers to make several rather different attempts to diagnose the source of the problem and solve it. Nothing approaching a consensus or broad agreement on how best to handle the problem has so far emerged.

The present paper contributes to this expanding debate. It offers its own diagnosis of, and solution to, the interface problem by connecting it to a rather different discussion about the format of intention, instigated by a minority of philosophers who reject the idea that intention should be construed as a propositional attitude. According to the alternative non-propositional conception of intention, the content of intentions is given by act types corresponding to the verb phrases by which we regularly express their content ('S intends to V'). I identify the most compelling reason to accept this 'do-able' view. Somewhat surprisingly, the reason in question turns out to hold the key to explaining away the difficulties surrounding how intentions manage to causally interface with motor representations. The non-propositional conception thus allows us to (dis)solve the interface problem. In so doing, this unorthodox conception itself gains further credibility.

The next section reconstructs the interface problem, describing in some detail what is involved in positing motor representations and attributing to them a format distinct from that of intention (§2.1), and how this gives rise to worries about the possibility of interface between the two state-types (§2.2). That section also points out problems with some of the proposed solutions to the interface problem (§2.3), in order to clear ground for my own proposal in §3. The heart of the discussion there is then devoted to spelling out doubts over the propositional conception of intention (§3.1) and bringing them to bear on the interface problem (§3.2). As I explain, the more promising 'do-able' alternative to propositionalism about intention is independently motivated by considerations that allow us to remove the initial puzzlement over the possibility of interface

between intention and motor representation. §4 summarizes the argument and rounds off the discussion.

2. PROBLEMS AT THE INTERFACE

2.1 Intention, motor representation, format

To get a clear view of the interface problem, we need first a picture of the mental constituents the interface between which is thought to be puzzling or problematic. Intentions should be familiar enough, at least in broad outline. On the standard conception, developed most comprehensively by Michael Bratman (1987), intention is understood as a *sui generis* state, irreducible specifically to any corresponding belief-desire pair, and distinguished by its set of functional characteristics clustered around the state's role within the agent's *planning* economy. A prospective intention is regarded essentially as a plan-like person-level conscious state, while present directed intentions or intentions-in-action are seen as corresponding plan-states that are in the process of being executed. Their important role of facilitating our temporally extended agency informs the idea of intentions as fairly stable and robust states (compared to e.g., desires or wishes), that are moreover subject to various constraints of coherence and consistency designed to ensure the feasibility of the plans they encode.

Motor representations are a somewhat lesser-known mental constituent, at least outside the cognitive science of motor control and the philosophy thereof. An informal example (borrowed from Levy, 2017) will help to introduce the notion. Consider an expert pianist improvising in front of her audience. Throughout her performance, she responds to a great number of cues online, including feedback from her own instrument; changes in the ambiance of the room; subtle cues from her partner, the bass player; the expectations of the audience, etc. The pianist responds to all

 1 A conception I do not personally endorse (see Levy, 2021); but for present purposes I have no reason to quarrel with it.

these and other cues fluently, automatically adjusting how she plays when the need arises, without having to stop to reflect on the advisable course of action. Indeed, her adjustments are typically made unconsciously. Evidently, there must be some mechanism below the threshold of fully conscious intentions that explains how this is possible. Such a mechanism would by no means be proprietary to such *skilled* action as the pianist's. One's humdrum intention to take a sip from the cup on one's desk glosses over the many subtle motor adjustments that must be made for its successful execution. If the target object is to be appropriately grasped, lifted, and delivered smoothly to one's lips, various different factors must be computed, including the cup's location, directionality, size, and weight; the strength, velocity, and type of grip movement (pinch, hook, wrap, power, etc.) appropriate for lifting it, and so on.

The expert pianist's and the humdrum cup holder's ability to fine-tune the performance of their intended tasks swiftly and without conscious reflection or even awareness, turns on the *motor representations* or commands occurring downstream from the agents' explicit, conscious intentions. It is these motor representational states that encode the information required to carry out the finer adjustments. And it is the *content-sensitive* causal interactions between these states and their source intentions that the interface problem is thought to problematize. A quick way to avoid the problem altogether would therefore be to reject any robust sense of content-respecting causation taking place between intentions and motor representations. Might not the fine-tuning of motor behavior be accounted for by invoking only blind, non-representational sensorimotor dispositions, or at most sparse representations encoding only minimal kinematic and mechanical information?

The evidence suggests not. Motor fine-tuning is widely regarded as sensitive to task demands, affordances, and person-level goals, with the content of motor representations correspondingly seen as rather rich. Thus for example, according to an influential model of action monitoring, when an action is initiated, efference copies are used to create forward models of the expected sensory consequences of the planned movements, which are then compared against the actual

sensory feedback as the action unfolds; the efference copies are produced on the basis of the motor commands (Blakemore et al., 2002; Frith et al., 2000). In a similar vein, mirror neurons have been shown to fire when subjects observe actions aimed at the same goal but employing different means or effectors (hand, mouth, foot, etc.) This points to markers of motor processing that are correlated with action outcomes (Koch et al., 2010; Rochat et al., 2010; Cattaneo et al., 2010.)

The foregoing characterization might make motor representations seem like motor intentions of sorts – intentions or intention-like states with finer-grained content. If that were the case, there would be no interface problem. What gives rise to the problematic is the thought that intentions and motor representations differ not only in content but more fundamentally in *form*; the information carried by each state is thought to be encoded differently, creating a mismatch that renders the possibility of causal interaction puzzling (Butterfill & Sinigaglia, 2014). We have yet to see the grounds for suspecting that a mismatch of this sort obtains (§2.2). Before that, seeing as the operative notion of 'format' is not often explicated in the debate – and will be key for the constructive proposal of §3 – it pays to pause briefly to elucidate it in general terms.

To repeat, then, the idea of format pertains at bottom not to *what* information is encoded but to *how* it is encoded. Butterfill and Sinigaglia (2014: 125) give the illustration of a route that takes you to some destination. One and the same such route could be represented as a line plotted on a map or alternatively as a series of verbal directions. With this illustration, Butterfill & Sinigaglia contrast cartographic with propositional² or language-like format. Iconic or imagistic representation constitutes a third type. The intuitive distinction is clear enough. But how, more precisely, are formats individuated? A comprehensive treatment of this complex and important

² 'Propositional' is sometimes understood to mean something like 'conceptual', 'discursive' or simply 'representational'. None of these is the meaning intended here, where non-propositional formats are also seen as representational and potentially conceptual. Rather, 'propositional' here refers to a specific form of representation whose content corresponds to a proposition, and which is further characterized in the text.

issue is a demanding task for which there is no space here.³ But a few basic observations will prove helpful for what follows.

One fundamental difference between the three formats cited above has to do with the relation each of them exhibits between signified and signifier. Iconic representation works by employing a direct resemblance relation between the two, whereas propositional representation employs a merely conventional or causal referential relation, and cartographic representation exemplifies an isomorphism. This gives rise to a further important difference in how the content of each format composes. Distinct formats are governed by distinct combinatorial rules for how to combine basic informational constituents into larger wholes. The rules governing propositionally formatted content are very abstract and general. The basic principle of predication, which signifies the metaphysical relation of instantiation or property-possession, allows for a very high degree of expressive generality when combining different elements in propositional form: In principle, nearly any property can be predicated of nearly any object. Iconic and cartographic formats, in contrast, exhibit substantive limitations when it comes to the types of information they can represent. For example, since iconic representation works by replicating the physical appearance of objects, it cannot straightforwardly represent disjunctions (Camp, 2009). Similarly, a map cannot directly represent that someone, somewhere is happy. (At most it can represent something that entails that someone, somewhere is happy – e.g., by representing that someone at a specific location is happy.) There are no comparable constraints on the type of information that can be represented by a proposition (an important point that we shall come back to in §3.2 below.)

A further, related difference between representational formats concerns the other side of the coin from combination or composition of constituents, viz. *de*composition. Iconic and cartographic representations conform to what Green & Quilty-Dunn (2021) call "Iconicity" – the

_

³ To make a start, see the excellent contributions by Elizabeth Camp (2007, 2009, 2018).

⁴ See also Block (2023, ch. 4) and Crane (2009), who both argue that the iconic format does not allow to represent any logical operations among represented items, including disjunction, negation, and conjunction.

principle whereby every part of the representation represents some part of the scene represented by the whole representation (this is what lies behind the common labeling of iconic representations as 'analogue'). Susan Carey (2009: 458) illustrates the point with the following nice example:

A picture of a tiger is an iconic representation; the word 'tiger' is not. The head in the picture represents the head of the tiger; the tail in the picture represents the tail. The 't' in 'tiger' does not represent any part of the tiger.

A consequence of Iconicity is that every part of iconic representations (icons) is meaningful, carrying a meaning that corresponds to some part of the scene it represents. In contrast, propositional (and more broadly, discursive or language-like) representations can be segmented into meaningless constituents (e.g., the segment 'tiger is' from 'the tiger is in the cage'). As Green & Quilty-Dunn point out, icons are further distinguished by the fact that each of their parts represents multiple properties at once: The part of the icon corresponding to the tiger's head is represented as furry, and rounded, and orange, etc. all at the same time. Hence, icons lack separate vehicles corresponding to separate items (properties, individuals).

A significant upshot of all this is that specific ways of manipulating represented information – compiling, accessing, updating it etc. in certain ways – will prove easier or harder depending on the format in which it is represented; differently formatted states induce different performance profiles. To see an illustration of this, go back to the contrast between the cartographic representation of a route as a line on a map vs. the propositional representation of the same route as a list of instructions. With the former, it would be relatively easier to reverse the route or to determine how close one is to the river, while the latter should make it easier e.g. to navigate by listing all the landmarks along the way (Butterfill & Sinigaglia, 2014: 125.) Precisely such differences in performance profiles are exploited by the argument underlying the interface problem as telltale signs of differently formatted states. It is to this argument that we now turn.

2.2 What is the interface problem?

Butterfill & Sinigaglia present an ingenious argument for thinking that the formats of intention and motor representation are mismatched. After reconstructing the argument immediately below, we shall be in a position to appreciate why such a mismatch, if indeed it obtains, would be puzzling or problematic, and to assess proposed solutions. B&S' argumentative strategy works essentially by comparing performance profiles of processes that are known or assumed to involve motoric representations with performance profiles of processes that are known or assumed to involve propositional representations. Observed differences in performance are taken to indicate a corresponding difference in format, while observed similarities indicate a common format. The overall structure of the argument can be stated as follows:

- (1) Representations involved in imagining acting and actually acting have a common format.
- (2) Representations involved in imagining seeing and actually seeing have a common format.
- (3) Representations involved in imagining acting and imagining seeing have different formats.
- (4) Representations involved in imagining seeing have propositional format.
- (5) Representations involved in imagining acting have non propositional (motoric) format. [3, 4]
- (6) Intentions have propositional format.

So,

(7) Intentions and representations involved in imagining acting have different formats. [5, 6]

The 'representations involved in imagining acting' (and in actually acting) referred to above are just motor representations. The appeal to *imagining* acting and *imagining* seeing serves to circumvent the worry that performance differences between their actual counterparts are explained by "bodily or environmental factors only distantly related to the representations involved" (Butterfill & Sinigaglia, 2014: 127.) The first three premises of the argument are empirical. Briefly, in support

of (1), B&S cite studies suggesting that how long it takes to imagine⁵ moving an object is correlated with how long it takes to actually move the object (Decety, 1996; Jeannerod, 1994). Similarly, manipulating an object so as to make it harder to act on – e.g., by orienting the cup handle so that it becomes less convenient to grasp – also slows performance at imagining acting on it (Frak et al., 2001). Such similarities in performance profiles and patterns of interference are, as explained above, generally seen as indicating a common format. Premise (2) is supported by noting that whether or not one is able to imagine seeing an object all at once depends on the size and distance of the object, in the same way that actually looking it over would. Furthermore, how long it takes to imagine looking over an object depends on the object's subjective size (Currie & Ravenscroft, 1997). Moving on to premise (3), data suggest that limb amputation and hand posture interfere with imagining acting but not with imagining seeing (Nico et al., 2004; Fourkas et al., 2006). Likewise, judging the laterality of a rotated *hand* – which is thought to involve imagining that one is moving the hand – is harder when the position of the hand is bio-mechanically awkward. This contrasts with judging the laterality of a rotated letter, generally taken not to involve imagining that one is seeing the letter, which is not similarly affected. The import of this result is underscored when we note that patients suffering from Amyotrophic Lateral Sclerosis (ALS), whose motor representations are known to be impaired, do not experience a similar pattern of difficulty when judging the laterality of objects (Fiori et al., 2013).

When (3) is combined with the plausible sounding (4), we arrive at (5) – the claim that motor representations are non-propositionally or motorically formatted. The contours of the propositional format were sketched above. The specific contrast between it and the motoric format will be at the heart of the discussion surrounding the constructive proposal in §3. For now, note just that the latter is considerably less flexible than the former, in that the information it codes is subject to the bio-mechanical constraints as well as kinematic and dynamic principles of the

 $^{^5}$ The relevant notion of 'imagining' V-ing that B&S have in mind here is the more phenomenologically involved one, which goes beyond merely thinking about V-ing or thinking about imagining V-ing.

motor system, since it must ensure that impediments to successful action (for example, awkward limb positions) are avoided. No corresponding constraints shape the content of propositionally formatted representations. Having arrived at (5), and proceeding to combine it with the widely accepted assumption made in (6) – which will be challenged below in §3 – that intentions are propositional attitudes, B&S are able to derive their conclusion (7): intentions and motor representations have different formats.

To see why this result should be thought puzzling or problematic, recall that the general scheme for action explanation involves content-respecting causal connections linking intentions with motor representations. In terms of B&S' above illustration, this is tantamount to a route, part of which is drawn on a map and part given in verbal instructions. For the gap to be successfully bridged, some comparator mechanism whose job is to transpose or translate information coded in one format onto the other must evidently be in place. However, as B&S explain, "The difficulty is that nothing at all is known about this hypothetical translation between intention and motor representation, nor about how it might be achieved, nor even about how it might be investigated." (Ibid: 133)

Importantly, the mystery of translation that B&S describe should not be overstated. After all, one could always posit a primitive association relation to reconcile the mismatched formats. But as with virtually any unexplained primitive, especially at critical junctures, doing so incurs a substantial theoretical cost that is best avoided if possible. This is how I propose to understand the difficulty uncovered by B&S. Our lack of understanding as to how the contents of intentions

_

⁶ Talk of 'translation' or 'transposition' here should not be taken to suggest *equivalence* of content; the content of intentions is at least often pitched at a higher level of abstraction compared to the downstream motor representations. The point of invoking translation is to signal that the content in each differently formatted state must somehow be matched, if the idea of content-respecting causation between them is to be vindicated.

and motor representations come to match makes for an explanatory gap, which comes at a price for any theory that does not manage to convincingly bridge it.^{7,8}

2.3 Some extant solutions, and where they fall short

This section will point out flaws in some of the main proposed solutions to the interface problem. It will be relatively brief, primarily because most of the points it raises have already been noted by other writers. The primary aim here is to motivate the need for a new approach to handling the interface problem, which will be spelled out in §3 below.

Start with B&S' own proposal, which works by invoking deferred reference by intentions to motor representations. Illustrating once again in terms of their cartographic-to-propositional example, suppose the list of verbal instructions for navigating the route contains the directive: "follow that route!", where the reference of the demonstrative expression 'that route' is the segment plotted on the map. B&S believe a similar process facilitates the interface between intention and motor representation. As they put it, "Intentions involve demonstrative concepts ... which refer to actions by deferring to motor representations" (Butterfill & Sinigaglia, 2014: 134). When imagining performing some action V, a subject may come up with a way of V-ing, thinking to herself 'that's how I should do it!', where the demonstrative refers to a way of V-ing by deferring to the specific motor representations involved in V-ing. Since as already noted, much the same processes save for the muscle contractions occur both in imagining acting and in actually acting, what goes for the former also goes for the latter.

_

⁷ An important question to examine is whether the problem identified by B&S deserves to be preceded by the definite article. Don't interface problems emerge elsewhere in the mind as well? *Perceptual* representations for example are arguably iconic, whereas at least some cognitive representations are propositional (see Block 2023 for one recent extensive discussion). Do we have a better understanding of how the format gap between perception and cognition is bridged? I shall not attempt to answer this question here. Whether or not there is an interface problem at the perception/cognition border (or elsewhere), it does not take away from the need to solve the one between intentions and motor representations.

⁸ The cost of relying on primitive associations between differently formatted content would be considerably mitigated if complemented by a plausible story of how such associations come into play. One such story is sketched by Pacherie and Mylopoulos (2017), as discussed in §2.3 below.

Now as we have seen, B&S hold that any adequate solution to the problem must avoid presupposing that any (unexplained) translation mechanism is in place. But as several writers have pointed out, their solution seems implicitly to do just that. For, absent some process of translation, it is unclear how the deferred reference they postulate could ever get off the ground. Here are Myrto Mylopoulos and Elisabeth Pacherie making this point:

In the case of demonstrative deferral in intention, the agent must have an independent grasp of which motor representation is the appropriate one to select via such deferral. But this would require a way of translating between the intention and the motor representation being picked out, in order to establish which motor representation correctly corresponds [to it]. (Mylopoulos & Pacherie, 2017: 328)

In the same paper from which the above passage is taken, Mylopoulos & Pacherie propose an alternative solution of their own. They suggest that the content of intentions is partly constituted by what they dub 'executable action concepts', namely concepts of actions the agent knows how to perform (thus KICK would serve for most agents as an example of such a concept, whereas FLY would not.) It is these executable concepts that activate, hence casually interact with, motor representations. Having an executable action concept in turn depends on having an appropriate motor schema, which is more abstract and stable than motor representations. Glossing over the details of their sophisticated proposal, M&P's basic idea is that motor schemas "form a bridge between [motor] representations and the executable action concepts that are deployed in the contents of proximal intentions." (Mylopoulos & Pacherie, 2017: 330).

However, it seems that M&P's proposed solution falls prey to essentially the same flaw they (correctly) identify in B&S' proposal – viz., that of presupposing that some mechanism for translating between propositionally and motorically formatted content is in place, instead of explaining how it operates. Action concepts, being *concepts* which moreover partially constitute the content of intentions, certainly seem to be given in propositional format according to M&P's view.

Meanwhile motor schemata, being abstractions from motor representations, should be coded motorically. How then do the two manage to link up?

M&P sketch the contours of a plausible story that could help explain how "motor schemata can be learned in a bottom-up fashion" (Mylopoulos & Pacherie, 2017: 331), drawing on Bayesian learning. This does seem a promising direction for explaining how agents acquire motor schemata without positing an unexplained process of top-down translation as part of the account. Still, it seems to leave open the crucial question of how motor schemata *themselves* interface with presumably propositional executive concepts (Shepherd 2019: 290-1).

A number of other solutions have been put forward. For a final sample, consider Shepherd's recent proposal, on which intentions take hybrid – both motoric *and* propositional – content. As Shepherd himself puts it:

The solution to the interface problem is that intentions lead a double life. Intentions can take propositionally formatted contents that enable their integration with propositional thought. And intentions have motorically formatted contents that communicate in a fairly direct way with the operations of motoric-level action implementation [...]

What we need to understand ... is how intentions could provide guidance sufficient to render our common action successes non-accidental. Intentions do this by specifying outcomes that motoric-level action implementation processes take on board directly. (Shepherd, 2019: 294-5).

This is not the place for a comprehensive assessment of Shepherd's position that could do justice to the various sources of evidence he draws on. But two points may be noted briefly. First, if they can indeed be part of the content of intentions, then that would seem to make motor representations consciously aware (as Shepherd himself openly acknowledges). This will strike many, this writer included, as a non-starter. It does not seem enough to point out, as Shepherd does, that there is no evidence to suggest that motor representations should be inaccessible. If they

⁹ I am grateful to an anonymous reviewer for discussion here.

do indeed partially constitute the content of intentions, we should expect motor representations to be actually, and fairly readily *accessed*, as the content of intention typically is. But that does not seem to be the case. Agents very rarely reason about the type or velocity of the grip they should employ when intending to pick up a cup from the table, for example. (For further discussion see Shepherd, 2019: 295-6, and Mylopoulos & Pacherie, 2017.)

Second, positing hybrid formats raises a worry about the methodology of format individuation. §2.1 cited grounds for individuating representational formats in terms of the different principles governing content composition in each format, which leads in turn to differences in the range of contents each format can represent, as well as to functional differences in the manipulation of the content. As we have seen in §2.2, the functional differences in question are crucially exploited by Butterfill & Sinigaglia's argument as indicating that the format of intention and motor representation differs. Now if indeed, as Shepherd contends, intention has a distinct, hybrid format, then it cannot be assumed to follow the same compositionality principles and exhibit the same functional profiles as states with propositional or motoric format. In other words, postulating a distinct format calls for a complementary account of the associated principles of format individuation and detection, which Shepherd does not supply.

The above offers neither an exhaustive survey nor conclusive refutation of extant solutions to the interface problem. Hopefully however, absent counterarguments at least, it does indicate that there is room to search for more promising alternatives.

3. THE DO-ABLE SOLUTION

The heart of the challenge posed by the interface problem, as we have seen in the previous section, is that of finding a way to bridge the gap between the propositional and the motoric without presupposing any (unexplained) mechanism of translation. To meet this challenge, the approach set out below works by dissolving the problem rather than solving it. There is in fact no

mismatch of formats between intentions and motor representations that needs to be reconciled or aligned. This is because intention is not, despite popular opinion, a propositional attitude. Rather, its content is coded in infinitival or 'do-able' form which meshes naturally with the motoric. In introducing the do-ables view of intention in rough outline (§3.1), we shall focus on one consideration in particular that militates in its favour, and which offers a clue for handling the interface problem. The following section (§3.2) will then unpack the clue and put it to work in overcoming the problem. It will also deal with an important objection to do-ableism.

3.1 Intention as a do-able attitude

Any attempt to dislodge the propositional conception of intention has a mountain of philosophical orthodoxy to climb. The following is a typical expression of the standard view:

If we let 'A' stand for such attitude verbs as 'believe', 'desire', 'hope', 'intend', 'think', etc., then the propositional attitude statements all have the form: S As that P. (Aydede, 2010)

At first blush, the default classification of intention with the other propositional attitudes may seem surprising, given that the natural expression of intention's content is given by a verb phrase (as in 'she intends to V'), which doesn't obviously correspond to a proposition. Why construe intention as propositional, then? To their credit, Butterfill & Sinigaglia do not gloss over this question entirely as many other philosophers do. While they do not discuss the issue in detail, they do provide their reason for taking intentions to be propositionally formatted. In support of this idea, which features as premise (6) of their argument for the interface problem (see above, §2.2), Butterfill & Sinigaglia explain that they "take [the claim that intention is propositional] to be a consequence of the role of intention in practical reasoning and of the fact that one can have intentions involving quantification and identity." (2014: 130).

On the do-able conception, the content of intention is given by 'do-ables', ¹⁰ namely act types: wash the dishes, kick, climb a mountain, ... Whereas propositions are abstract entities that can be true or false, do-ables are abstract entities that can be performed successfully or botched. So understood, quantification and identity are available operations for do-able content, hence not a reason to prefer the propositional over the do-able conception. But quantification and identity are not the *only* available operations. The content of intention is commonly taken to support logical connectives such as negation ('S intends not to go to the party'), disjunction ('S intends to go to the party or to the theatre'), conjunction ('S intends to go to the party and to the theatre'), and ifthen ('S intends to go to the party only if N goes'). Furthermore, and closely related to this, intentions seem to be involved in valid inferences such as

- S intends to visit Paris and so does N.
 So,
- 2. There is someplace that both S and N intend to visit.

and

- 3. S intends to read all the books N has read.
- 4. *N* has read Journey to the End of the Night. So,
- 5. *S* intends to read Journey to the End of the Night.

It is straightforward to account for the validity of the above inferences if the sentences are analyzed in terms of two-place relations involving propositions. Can the do-able conception deliver the same result?

I consider it an open question (one that I explore in further work that aims to bolster the doable conception) *which* logical connectives are in fact supported by intention.¹¹ But that dispute

¹¹ See Levy (ms). For a quick-and-dirty illustration, consider that *S*'s alleged intention to not go to the party is arguably better understood as her intention to stay at home.

 $^{^{10}}$ The title is borrowed from Campbell (2019) – one of the few extant defences of the do-able conception of intention.

need not be settled here. For whichever connectives turn out to be supported – and consequently, whichever intention-involving inferences turn out to be valid – turns only on the analysis of intention exhibiting two features: a *relational structure*, and objects (arguments) with *uniform semantic value*. For these purposes, do-able objects will serve just as well as propositional ones.

A further, more serious obstacle to do-ableism is raised by B&S' remark about reasoning. David Lewis elaborates the point, in one of the few explicit statements of the case for propositionalism about intentional attitudes:

In attempting to systematize what we know about the causal roles of attitudes, we find it necessary to refer to the logical relations among the objects of the attitudes. Those relations will be hard to describe if the assigned objects are miscellaneous. Uniform propositional objects, on the other hand, facilitate systematic common-sense psychology. (Lewis 1979: 514)¹³

The worry Lewis raises about the potential havoc wreaked by a disuniform picture of the format of different psychological attitudes will be handled separately in §3.2. First, we need a clearer picture of what the alternative proposal endorsed here is all about. Why go in for the unorthodox do-ables view at all? That the content of intention is naturally expressed by a verb phrase rather than a declarative sentence is in itself hardly compelling grounds if the two turn out to be interchangeable; the difference will then amount to no more than "an artefact of presentation" (Williamson, 2017: 167). Are they interchangeable? One reason to suspect not adduced by some do-ableists is grounded in the fact that do-ables lack an explicit *subject*, which propositions contain. (Thus compare 'open the door' in 「S intends to open the door ¬ with '*she* open the door' in the propositionally construed 「S intends that she open the door ¬). This makes room for

_

 $^{^{12}}$ Grzankowski (2018) provides an analysis of non-propositional objectual attitudes (e.g. S loves N and S fears O) supporting these two features.

¹³ As Lewis is quick to acknowledge, the advantage he cites would accrue to *any* uniform treatment of the content of psychological attitudes, not only a propositional one — and indeed, in the same paper Lewis defends an alternative uniform conception of content in terms of *properties* rather than propositions.

incongruences of the sort broadly familiar from the literature on *de se* attitudes. Here is Jenifer Hornsby:

Suppose that Aelfric is a very powerful man with a finger in many pies. Aelfric intends that the Chair of the Board of Management speak to the Regent. But Aelfric has forgotten who the Chairs are of all the Boards. Luckily (as he thinks), Aelfric's office wall is adorned with labelled photographs of all the important people, and he can see a photograph labelled 'Chair of the Board of Management'. So Aelfric now decrees 'HE, that man, shall speak to the Regent'. Aelfric, however, fails to realize that the photograph he sees is actually a photograph of himself. So now Aelfric intends that HE should speak to the Regent, and actually HE is Aelfric, but he does not intend to speak to the Regent. (Hornsby, 2016: 6)

There is a ready reply to Hornsby's objection, however (as she acknowledges). Linguists standardly deploy as part of the syntax of infinitives the so-called [PRO] device, meant to refer to an implicit (silent) subject (see e.g. Hayes et al., 2013: 132-3). With the introduction of [PRO], 'Aelfric intends to call the Regent' becomes 'Aelfric intends [PRO] to call the Regent', which in the present case is equivalent to 'Aelfric intends that he himself call the Regent'. Consequently, Hornsby's argument for the do-ables view is inconclusive at best. However, there seem to be different, more serious grounds for doubting that do-ables are freely interchangeable with propositions in the content of intention. The doubts raised will in turn clear ground for the (dis)solution of the interface problem.

Consider the sentence

(1) S intends to swim in the icy lake.

According to the propositional conception of intention (henceforth, 'propositionalism' for short), 15 (1) is equivalent to

¹⁴ For further discussion, see Hornsby (2016: 10-13), and Stanley (2011: ch. 3)

¹⁵ It is common to use 'propositionalism' to refer to the view that the contents of intentional attitudes quite generally have propositional structure, where the alternatives to propositions most commonly discussed are things like objects, properties and kinds (see for example the articles in Grzankowski & Montague, 2018). However, I am

(2) S intends that she herself swim in the icy lake.

However, as Baier (1970) points out, (2) could be interpreted either as equivalent to (1) or alternatively as equivalent to

(3) S intends to cause herself to swim in the icy lake.

Now to insist on the difference between (1) and (3) is by no means to quibble. The intention to swim engages a specific set of powers and abilities. It is executed roughly when one moves one's limbs in a way conducive to propelling one's body in the water. In contrast, the intention to cause oneself to swim engages a very different set of powers. It is successfully executed when one manages to get oneself to swim, placing oneself in a position that will likely or even inevitably result in one's swimming – for example, by jumping into the icy waters, forcing oneself to swim so as not to drown (or alternatively by convincing someone to push one into the water, or by...).

The root of the problem for propositionalism seems to be this: To have an intention is not simply to intend that p *be* the case; it is to intend to *make it* the case that p. In order to express this idea, a propositional conception must include reference to the agent herself as part of the content of her intention. But once she is cited as part of the content, the agent may inevitably be understood (also) as an *object* of her own intention, thus giving rise to an unwelcome ambiguity in how the intention should be understood.¹⁷

It might seem like there must be resources the propositionalist could draw on to resolve the ambiguity. But as Baier (1970) and Campbell (2019) make clear, the prospects of doing so seem

here restricting myself to intentions and to do-able contents. It is a further and interesting question whether the arguments in the text apply to any other attitudes. Natural additional candidates to explore include *want* and *hope [to I/]*

¹⁶ As even a devout propositionalist as Donald Davidson seems to recognize when he points out that "[i]t is a mistake to think that when I close the door of my own free will anyone normally causes me to do it, even myself." (1980: 56)

¹⁷ Cf. Campbell (2019: 239).

dim. Briefly, in line with the idea that intentions are causally self-referential, ¹⁸ one might suggest that (2) should be read as

(2)* S intends that she herself swim as a direct result of this very intention,

where (2)* is equivalent to 1, while (3) should be read as

(3)* S intends that she herself swim as an indirect result of this very intention.

However, this will not solve the problem, because the direct/indirect distinction does not map onto the one between intending to V and intending to cause oneself to V. This can be seen by noticing that one can execute an intention to V indirectly *without* causing oneself to V – for example, intending to take a book off the top shelf by intending to climb a ladder. Nor, for similar reasons, can the distinction between (1) and (3) be captured by incorporating into the propositional content of intention a distinction between basic and non-basic actions (Campbell, 2019: 235-6).

It might seem like the appearance of a problem here is a mere artefact of how the intentions in (1)-(3) are presented. (2) appears ambiguous between (1) and (3) only because (1), and consequently also its propositional rendering (2), is underspecified. In fact, the thought goes, intentions such as (1) will typically contain information, however sketchy or incomplete, about *how* the intention is to be carried out. That information is absent from (1) and (2) as stated (as well as from their canonical expression in language), but it is there in the deep structure of the intention and helps to determine if (2) is equivalent to (1) or rather to (3).

The thought is plausible on its face but it fails to remove the problem, which does not stem from content (in)specificity. The indeterminacy identified by Baier will resurface at whatever level the content of intention is specified. To see this, consider the level most philosophers of action would regard as the most specific and fundamental possible – viz., so called "basic action", supposedly performed 'just like that' without employing any subsidiary means.¹⁹ And suppose the

¹⁸ For this idea, see Searle (1983, ch. 3), and Velleman (1989).

¹⁹ For doubts about the very idea of basic action, see Lavin (2012), and Thompson (2008, 106–15).

intention to swim in the icy lake is specified down to the basic level of moving one's arm thusand-so:

(1)** S intends to swim in the icy lake by $\{[...]$ and moving her left arm thus-and-so $\}$

Still, the ambiguity is not resolved. For according to propositionalism, S intending to move her left arm thus-and-so is equivalent to S intending that she herself move her left arm thus-and-so, which in turn can be interpreted *either* as her intending to move her left arm simpliciter or as her intending to cause herself to move her left arm – for example, by lifting it with her right arm.²⁰

If the argument of this section is on the right track, then it seems that intention should be construed as an irreducibly do-able attitude. What is probably not yet clear is how exactly all this bears on the interface problem. The next section will show how the specific difficulties discussed above with attempting to squeeze intention into the propositional mould in fact hold the key to handling the problem.²¹

²⁰ In a discussion related to Baier's, though drawing a stronger conclusion, Michael Thompson (2008: ch.8) argues that the do-able content of attitudes cannot be rendered equivalent to any proposition. Substituting 'intend' for his example of 'want', Thompson suggests that 'I intend that I walk to school' expresses "a habitual sense, which is nothing like what we had in mind" in saying 'I intend to walk to school'. Meanwhile, 'I intend that I am walking to school' "suggests that I am indifferent to my actually making it there", and 'I intend that I have walked to school', "if the content clause is *true*, seems to express satisfaction in a job well done, and if it is false, an idle wish". Finally, 'I intend that I will have walked to school' "...seems to express an indifference to my own agency" (Thompson, 2008: 128).

²¹ A question for future work that will develop the do-ables view in a more comprehensive manner than can be done here is whether intention has *exclusively* do-able content or whether alternatively, some forms of intention are propositionally formatted. Two candidate types for exclusion from the do-able template are: (a) Highly vague or abstract intentions; and (b) Intentions for *other* agents to *V*. Intentions of type (a) might be thought of as propositional, insofar as the specific ways of carrying them out might be left indeterminate pending future planning. Meanwhile, type (b) may be understood as calling for no active participation on one's own part, as in 'I intend that everyone have a good time', read as inequivalent to 'I intend *to get* everyone to have good time.' (See Campbell, 2019 for some discussion). My own inclination, which will have to await future elaboration and defense, is to treat all intentions as do-ably formatted. Briefly, (a)-type 'intentions' are arguably not proper intentions at all, since their vagueness and impracticality render them closer to idle wishes. Meanwhile, (b)-type intentions *do* strike me as calling for participation from the intending agent – otherwise, again, it is unclear to me that they should be considered genuine intentions rather than some other, less committing and involved attitude.

3.2 (Dis)solving the interface problem

Go back to the discussion of format individuation in §2.1. One of the key markers cited there for identifying the format of some representation was the combinatorial principle(s) governing how basic informational constituents combine to form larger wholes. Different principles will be more or less general and abstract, imposing more or less severe restrictions on the range of contents that each format can represent. Motorically formatted content is quite heavily restricted, as the principles governing it are subject to the constraints of the motor system. An object one intends to act on will be represented partly in terms of its affordances and the movements required to execute the action, reflecting implicit knowledge of bio-mechanical and kinematic constraints (§2.1). This rules out representing the object in all the different ways one *cannot* act on it - e.g., grabbing it with a biomechanically impossible joint angle, one's fingers bent backwards towards the back of one's hand. In contrast, *propositionally* formatted content is substantially less constrained. It is governed by the extremely thin principle of predication, which affords a high degree of expressive generality (recall the example of declarative sentences, which unlike maps and pictures can easily represent disjunctions, among various other things).

In fact, it is precisely this generality that seems to be ultimately at fault for the problem facing propositionalism about intention identified by Baier. A propositional construal of intention is so abstract and general that it abstracts away from the fundamental functional difference between V-ing and causing oneself to V, which exercise our agency in radically different ways. This renders an intention that one V problematically indeterminate, making it hard to see how it could combine downstream with task-specific motor schemata. How does a propositionally formatted intention able to link up with one set of motor representations (those involved in V-ing) rather than another (those involved in causing oneself to V), given that it is unable to disambiguate V-ing from causing

oneself to V?²² This, I submit, cuts to the core of the incongruence underlying the interface problem. An indifference to *how* the intending subject's agency gets exercised is the deep flaw plaguing the propositional conception of intention, and the reason why it should be displaced. And one major symptom of this flaw is the mystery it makes of how intention manages to link up with the correct motor representations crucial for executing it in particular – a mystery also known as the interface problem. In view of this mystery surrounding the match between intention and motor representation, we cannot afford a relaxed attitude towards the prospect of translation between the two. Nor can we posit a primitive association relation without thereby papering over a deep and difficult puzzle.

A much more promising tack would be to replace propositionalism about intention with the do-ables conception. Do-ably formatted intentions are adequately constrained, suffering no comparable ambiguity or indeterminacy. An intention to swim understood as having the do-able content row swim roan only be satisfied in one way, namely by performing a token of the act-type corresponding to its content – or, to put it more simply, by swimming. In contrast, rocause oneself to swim rodesignates a different do-able with different satisfaction conditions (see p. 18 above). Hence, the latter could not count as a way of satisfying the do-able intention to swim.

Do-ably understood, the interface of intentions with motor representations no longer presents any format gap that needs to be bridged: Motor representations can be seen as similarly formatted, more detailed specifications of what one intends to achieve and the means of achieving it. As an upshot, the grain of motor representations will predictably be finer than that of intentions, a point that is often noted in discussions of the interface problem. The sensorimotor control processes in which motor representations partake operate at a grain that is typically finer than that of the propositions and concepts that agents can grasp. This observation sits well with an entrenched

²² Note that the problem is *not* due to any mismatch in specificity or level of grain between the interfacing representational states (more on this below, p. 23). It is therefore no help to the propositionalist to insist that the content of intention can be highly specific or fine-grained. Thanks to an anonymous reviewer for discussion here.

picture of the cognitive architecture of action control, on which it is subserved by hierarchical levels of processing forming an "intentional cascade" (Pacherie, 2008: 188) from intentions down to motor representations.²³

That intentions and motor representations have different levels of grain is consistent with the do-able approach. Indeed, this fact partly explains why the do-able is considerably more flexible than the motoric, in that the former is available for diverse operations (various combinations with other intentional attitudes, deployment in reasoning, etc.) for which the latter is not available. Differential grain is likewise consistent with the rival propositional conception of intention, which I am trying to throw into doubt. Evidently, this is not where the difference between the two conceptions, nor the fault in the propositional conception, lies. On the way to understand and (dis)solve the interface problem proposed here, the issue with propositionally conceived intentions is *not* that they are more flexible or coarse-grained than motor representations; it is that they are semantically indeterminate. We have no clear idea of what they are intentions *for*. Do-ably conceived intentions do not suffer this crucial flaw.

One might worry, however, that the above does not dissolve the interface problem so much as push it back to a higher cognitive level. After all, intention regularly enters into relations with other attitudes, some of which are definitely propositional. If the former has do-able content, an interface problem with beliefs and other propositional attitudes would seem to resurface at the personal level. This objection was noted briefly in §3.1 by quoting Lewis, who warns that only a uniform picture of the objects of attitudes can facilitate a systematic understanding of folk psychology.

There are two related worries in Lewis' concern (or its vicinity) that should be disentangled. One is about the possibility of mental causation: If intentions and (e.g.) beliefs are formatted differently, how can one cause the other, absent some mechanism of inter-format translation? The

²³ For details and relevant discussion, see for instance Pacherie (2008) and Shepherd (2015).

second worry targets practical reasoning. To have a systematic picture of how and why agents make decisions and perform actions, it is crucial that we capture the logic of (good) reasoning. But how else could we do so except by appealing to truth-functional relations between propositions?

Start with the first worry. What makes it appear serious is that it seems to mirror the structure of the original interface problem: Two state-types managing to sustain causal connections despite having different formats. However, there is a crucial difference – viz., that the interface in question takes place entirely at the person level, unlike the original problem in which one of the causal relata is sub-personal. This difference does not necessarily make the former challenge any easier to handle; on the contrary, it may be harder, at least in some ways. But it does reveal it to be an altogether different challenge from the one this paper is concerned with. To explain: When we ask how differently-formatted person-level states can maintain causal connections, we seek an explanation of how persons come to be adept at simultaneously engaging with content in diverse formats. To take a concrete case, it is akin to asking: How are we able to describe (to ourselves or to others), in propositional language, the finer motoric details of some intricate sequence of movements we plan to perform? This is a question about the mental (conceptual) machinery required to traffic in differently formatted representations. A satisfactory answer would teach us what enables people to translate content from one format to the other, bringing them together in productive forms of thinking and reasoning. And the right kind of story to tell in answering this type of question is a developmental one. It is a story about how people develop genuinely new representational resources, with expressive powers that outstrip the ones they are innately endowed with.

In one way, switching to the person-level makes no difference: the existence of mismatched formats raises the puzzling question of what psychological machinery is required to transcend them and, if doing so involves a process of translation, what this process involves. This much is true whichever level of explanation we are operating at. However, the claim here is that

understanding how person-level format mismatches are transcended becomes much less puzzling if we have recourse to a particular developmental story, one that shows how persons acquire the conceptual capacities required to traffic in differently formatted content.

This is of course hardly the place to defend such a story in any detail. But very briefly, two influential contrary positions on the issue can reveal the nature of this debate and its bearing on the objection we are presently considering. Famously, then, Fodor (1975, 1981) argues for radical concept nativism. For him, discharging the developmental task referred to above amounts to little more than documenting the order in which various innate representational capacities mature and come on line. In contrast, Carey (2009) advocates an alternative account she dubs 'Quinean bootstrapping'. Highlighting the various different discontinuities in the human conceptual repertoire from infancy to adulthood, Carey explains how the bootstrapping process, starting off with the limited innate resources of 'core cognition', could give rise to novel lexical concepts with genuinely new expressive powers. Setting aside the details of this sophisticated hypothesis, what is important for us is that the process of transition Carey envisages also involves the emergence of differently formatted representations. According to her, the innate representations of core cognition (magnitude, objecthood, agency) are all iconic. Through the process of bootstrapping, richer representational structures emerge on the back of the core categories, including the familiar propositional ones from the adult repertoire. And much the same could in principle be said about do-able representations, if indeed such there are as this paper contends.

It is not my aim here to defend Carey's theory. However the dispute between Fodor and her should be settled, its resolution would shed light on how people manage to engage with representations in multiple formats. It would shed light, that is, on how format gaps are transcended at the person level. We may call the end product here a process of translation if we like; the important point is that it would not be a process about whose nature and development we are in the dark.²⁴

Moving on, the second version of Lewis' worry touches more directly on the logical relations between attitudes within episodes of practical reasoning (this is also, recall, the reason adduced by Butterfill and Sinigaglia for supposing that intention is propositional, citing 'the role of intention in practical reasoning'.) The worry seems to be that a non-propositional construal will not allow us to illuminate the role of intentions in reasoning by placing them within familiar patterns of deductive inference. One of the few philosophers (Lewis included) to make this point explicitly is John Broome (2001). He compares the following piece of theoretical reasoning,

[Belief] Chris will buy a boat

[Belief] For Chris to buy a boat, Chris must borrow money

So,

[Belief] Chris will borrow money

with the following piece of practical reasoning (examples changed slightly):

[Intention] I (Chris) will buy a boat

[Belief] For me (Chris) to buy a boat, I must borrow money

So,

[Intention] I (Chris) will borrow money

Of the latter, Broome says that

Since the syllogism is valid, the two premises cannot be true without the conclusion's being true. So you cannot rationally be set to make one premise true and take the other as true, without being set to make the conclusion true ... This is why your reasoning is correct

²⁴ Cf. Shepherd 2019, pp. 297-301.

practical reasoning. It is correct because of the very same validity as makes the theoretical reasoning

(Broome 2001: 178. Emphasis added) correct.

At root, then, the worry voiced by Lewis and Butterfill & Sinigaglia seems to be that to reject

the propositional conception of intention is to forgo logical validity as a standard of correctness

for practical reasoning. Transitions in reasoning involving intention could no longer be assessed

in terms of their conformity with, and violation of, valid argument structures. And it must be

conceded that this is indeed a cost of going in for the do-ables view. However, fortunately, the

price does not seem particularly high. This is because the project of subsuming practical validity

under logical validity is fraught with difficulties and requires significant revisions and allowances,

anyway. Two examples will illustrate the point. The first is from Broome himself, of what he dubs

'enkratic reasoning' (Broome 2013, ch. 16):

[Belief]

I ought to take a break

So,

[Intention] I shall take a break

The second example is from Elizabeth Anscombe (1963, §33) via Aristotle:

[Belief]

Vitamin X is good for all men over 60

[Belief]

Pigs' tripes are full of Vitamin X

[Belief]

I'm a man over 60

[Belief]

Here's some pigs' tripes

So,

[Intention] I'll have some

28

As Anscombe notes, "certainly no one could be tempted to think of [the conclusion] as a proposition entailed by the premises" (1963: 61). And the same goes for Broome's enkratic reasoning. Nonetheless, they are both instances of good reasoning. The upshot is that least in some episodes of practical reasoning, logical validity anyway does not function as the standard of correctness. Intention is already excluded from the jurisdiction of propositional validity in a good number of cases. Extending its exclusion further thus does not seem like a terribly serious cost of adopting the do-ableist conception. The benefits of doing so however, as I have tried to show, are considerable.

4. CONCLUSION

This paper has sought to bring together two distinct problems which at first sight do not seem to have much in common, and recommend one and the same move as a solution to both. That the content of intention is not equivalent to a proposition naturally leads to construing intention as a do-able attitude. What is probably more surprising is to realize that the same issue is ultimately responsible also for the presumed mismatch of formats between intention and motor representation that lies at the heart of the interface problem. The first step towards this conclusion was to clarify the nature of the interface problem: what exactly is involved in claiming that intentions and motor representations are formatted differently, and what is problematic about this mismatch of formats, if indeed it occurs. A key idea that emerged was that formats are individuated in large part by the specific rules governing the composition of their contents, with propositional composition being considerably more abstract and general than motoric composition.

Next came a brief survey of some extant solutions, designed to highlight that there is space for a new, and arguably more radical, approach. That approach was subsequently spelled out, its central tenet being that intention is in fact a do-able, not propositional attitude. The main consideration supporting this unorthodox idea was identified as the Baieresque observation that

no proposition can hope to capture the content of intention. That the problem raised by Baier resists a quick fix indicated that it is symptomatic of a deeper flaw in the propositional rendering of intention. Tracing this flaw to the abstract rule of predication governing propositionally formatted states held the key to presenting do-ableism as a solution to the interface problem, thereby connecting the two central concerns of the paper.

While the primary aim of the foregoing discussion has been to offer a defensible solution to the interface problem, a complementary aim was to support the view that affords the solution, viz. do-ableism, itself. In this I have only made a start, sifting through some more and less compelling reasons to accept the view, and attempting to allay worries to do with the disruption it brings to the propositional uniformity of intentional attitudes. Hopefully, the most important contribution of the paper to the project of defending do-ableism about intention consists in demonstrating that the view offers a convincing dissolution of the otherwise puzzling interface problem.

.

²⁵ For further grounds to accept the do-ableist conception of intention, see Levy (ms.)

References

Aydede, M. 2010. The language of thought hypothesis, in Zalta E. (ed.), *The Stanford Encyclopedia of Philosophy* Edward N. Zalta (ed.), URL = https://plato.stanford.edu/archives/win2017/entries/language-thought/

Anscombe, G. E. M. 1957. Intention (Harvard University Press).

Baier, A. C. 1970. Act and Intent. Journal of Philosophy 67 (19): 648-658.

Blakemore, S.-J., Wolpert, D. M., & Frith, C. D. 2002. Abnormalities in the awareness of action. *Trends in Cognitive Sciences*, 6, 237–242.

Block, N. 2023. The Border Between Seeing and Thinking (Oxford University Press).

Broome, J. 2013. Rationality through Reasoning. Wiley-Blackwell.

Broome, J. 2001. Normative Practical Reasoning. *Aristotelian Society Supplementary Volume* 75 (1): 175–193.

Bratman, M. 1987. Intention, Plans, and Practical Reason. Cambridge University Press

Butterfill, S. A., & Sinigaglia, C. 2014. Intention and motor representation in purposive action. *Philosophy and Phenomenological Research* 88: 119–145.

Camp, E. 2007. Thinking with Maps. *Philosophical Perspectives*, 21(1): *Philosophy of Mind*, ed. J. Hawthorne (Oxford: Wiley-Blackwell): 145–82.

Camp, E. 2009. A Language of Baboon Thought? In R. Lurz (ed.), *The Philosophy of Animal Minds* (Cambridge University Press): 108–27.

Camp, E. 2018. Why maps are not propositional. In Grzankowski & Montague (eds.), *Non-Propositional Intentionality* (Oxford University Press): 19-45.

Campbell, L. 2019. Propositionalism about intention: shifting the burden of proof. *Canadian Journal of Philosophy* 49 (2): 230-252.

Carey, S. 2009. The origin of concepts. Oxford University Press.

Cattaneo, L., Sandrini, M., & Schwarzbach, J. (2010). State-Dependent TMS reveals a hierarchical representation of observed acts in the temporal, parietal, and premotor cortices. Cerebral Cortex, 20(9): 2252–2258.

Crane, T. 2009. Is perception a propositional attitude? *The Philosophical Quarterly* 59: 452-469.

Currie, G. & Ravenscroft, I. (1997). Mental simulation and motor imagery. *Philosophy of Science*, 64(1): 161–180.

Davidson, D. 1980. Agency. in his Essays on Actions and Events (Clarendon Press): 43-61.

Decety, J. (1996). Do imagined and executed actions share the same neural substrate? *Cognitive Brain Research*, 3(2), 87–93.

Fiori, F., Sedda, A., Ferrè, E. R., Toraldo, A., Querzola, M., Pasotti, F., Ovadia, D., Piroddi, C., Dell'aquila, R., Lunetta, C., Corbo, M., & Bottini, G. (2013). Exploring motor and visual imagery in Amyotrophic Lateral Sclerosis. *Experimental brain research*, 226(4): 537–547.

Fodor, J. A. 1975. The Language of Thought. Harvard University Press.

Fodor, J. A. 1981. The present status of the innateness controversy. In Representations: Philosophical Essays on the Foundations of Cognitive Science. MIT Press

Fourkas, A. D., Ionta, S., & Aglioti, S. M. 2006. Influence of imagined posture and imagery modality on corticospinal excitability. *Behavioural Brain Research*, 168, 190-196.

Frak, V., Paulignan, Y., & Jeannerod, M. (2001). Orientation of the opposition axis in mentally simulated grasping. *Experimental Brain Research*, 136(1), 120–127.

Frith, C. D., Blakemore, S.-J., & Wolpert, D. M. 2000. Abnormalities in the awareness and control of action. *Philosophical Transactions of the Royal Society of London Series B – Biological Sciences*, 355, 1771–1788.

Green, E. J. & Quilty-Dunn, J. 2021. What Is an Object File? *British Journal for the Philosophy of Science* 72: 665-699.

Grzankowski, A. 2018. "A relational theory of non-propositional attitudes", in Granzkowski & Montague (eds.), *Non-propositional intentionality* (Oxford University Press): 134-151.

Hayes, B., Curtiss, S., Szabolcsi, A., Stowell, T., Stabler, E., Sportiche, D., Koopman, H., Keating, P., Munro, P., Hyams, N. & Steriade, D., 2013. *Linguistics: An introduction to linguistic theory* (Wiley-Blackwell).

Hornsby, J. 2016. "Intending, Knowing How, Infinitives." *Canadian Journal of Philosophy* 46 (1): 1–17.

Jeannerod, M. (1994). The representing brain: Neural correlates of motor intention and imagery. Behavioral and Brain Sciences, 17(2), 187–245.

Koch. G., Versace, V., Bonni, S., Lupo, F., Gerfo, E. L., Oliveri, M., & Caltagirone, C. (2010). Resonance of cortico–cortical connections of the motor system with the observation of goal directed grasping movements. *Neuropsychologia*, 48 (12), 3513–3520.

Lavin, D. 2012. Must There Be Basic Action? Noûs 47 (2): 273-301.

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

Levy, Y. 2021. Disjunctivism about Intending. American Philosophical Quarterly, 58, 161-178.

Levy, Y. ms. Taking Action First. Unpublished book manuscript.

Lewis, D. 1979. Attitudes De Dicto and De Se. The Philosophical Review 88: 513-543.

Mylopoulos, M. and Pacherie, E., 2017. Intentions and motor representations: The interface challenge. Review of Philosophy & Psychology, 8(2): 317–336.

Nico, D., Daprati, E., Rigal, F., Parsons, L., & Sirigu, A. (2004). Left and right hand recognition in upper limb amputees. *Brain*, 127(1), 120–132.

Pacherie E. 2008. The phenomenology of action: a conceptual framework. *Cognition* 107: 179–217.

Rochat, M. J., Caruana, F., Jezzini, A., Escola, L., Intskirveli, I., Grammont, F., Gallese, V., Rizzolatti, G., & Umilta`, M. A. (2010). Responses of mirror neurons in area f5 to hand and tool grasping observation. *Experimental Brain Research*, 204(4), 605–616.

Searle, J. R. 1983. Intentionality. (Cambridge University Press).

Shepherd, J. 2015. Conscious control over action. Mind & Language 30: 320-344.

Shepherd, J. 2019. Skilled Action and the Double Life of Intention. *Philosophy and Phenomenological Research* 98 (2): 286-305.

Stanley, J. 2011. Know How. (Oxford University Press).

Thompson, M. 2008. Life and Action (Harvard University Press).

Velleman, D. 1989. Practical Reflection. (Princeton University Press).

Williamson, T. 2017. Acting on Knowledge, in J. Adam Carter, E. C. Gordon, and B. Jarvis (eds.) *Knowledge-First: Approaches in Epistemology and Mind*, pp. 163–181. (Oxford University Press).