Synthese (2021) 198:4373–4389
https://doi.org/10.1007/s11229-019-02347-x

NEUROSCIENCE AND ITS PHILOSOPHY

Tracking intentionalism and the phenomenology of mental effort

Maria Doulatova[1]

Received: 10 October 2018 / Accepted: 25 July 2019 / Published online: 26 October 2019

Abstract
Most of us are familiar with the phenomenology of mental effort accompanying cognitively demanding tasks, like focusing on the next chess move or performing lengthy mental arithmetic. In this paper, I argue that phenomenology of mental effort poses a novel counterexample to tracking intentionalism, the view that phenomenal consciousness is a matter of tracking features of one's environment in a certain way. I argue that an increase in the phenomenology of mental effort does not accompany a change in any of the following candidate representational contents: (a) representation of externally presented features, e.g. brightness, contrast, and so on (b) representation of task difficulty, (c) representation of the possibility of error, (d) representation of trying to achieve some state of affairs, (e) representation of bodily changes like muscle tension, or (f) representation of change in cognitive resource availability and lost opportunity cost. While tracking intentionalism about some phenomenal experiences like pains might obtain, it does not seem to obtain for all phenomenal experiences. This puts the intentionalist into an uncomfortable position of trying to explain why some phenomenal experiences have representational content and not others. Since many believe that tracking intentionalism or something like it provides the best chance of naturalizing consciousness, these arguments deserve detailed consideration.

Keywords Tracking intentionalism · Mental effort · Attention · Qualia · Reductionist theories of consciousness

1 Introduction

Most of us are familiar with the phenomenology of mental effort accompanying cognitively demanding tasks, like focusing on the next chess move or performing lengthy mental arithmetic. In this paper, I argue that phenomenology of mental effort poses

✉ Maria Doulatova
mashawills@gmail.com; Doulatova.maria@wustl.edu

1 Washington University in St. Louis, St. Louis, USA

a novel counterexample to tracking intentionalism, the view that phenomenal consciousness is a matter of tracking features of one's environment in a certain way.

Intentionalists argue that, necessarily, any two mental states that differ in their phenomenal properties will differ in their representational properties.[1] A common way of arguing against intentionalism is to present experience pairs that differ in phenomenal character but fail to differ in representational content. In similar fashion, I argue that an increase in the phenomenology of mental effort (henceforth PME) does not accompany a change in any of the following candidate representational contents: (a) representation of externally presented features, e.g. brightness, contrast, and so on (b) representation of task difficulty, (c) representation of the possibility of error, (d) representation of trying to achieve some state of affairs, (e) representation of bodily changes like muscle tension, or (f) representation of change in cognitive resource availability and lost opportunity cost.

I focus on what is commonly taken to be the most promising reductionist intentionalist theory of phenomenal consciousness, i.e., *weak global* intentionalism conjoined with a tracking theory of intentionality, broadly defended by David Armstrong, Fred Dretske, Christopher Hill, William Lycan, and Michael Tye.[2] By offering a way to reduce *all* consciousness to tracking, tracking intentionalism offers a promising physicalist theory of consciousness.[3] Since many believe that this view offers the *best chance* of naturalizing consciousness, these arguments deserve detailed consideration (Pautz 2013; Cutter and Tye 2011).

2 What is intentionalism?

Intentionalism posits a close relationship between phenomenal and representational properties. Phenomenal properties of mental states concern *what it is like* for a subject to experience certain phenomenal properties. For example, when I look at the gray cat on my desk, my visual experience has a distinctive phenomenal character. In virtue

[1] Supervenience relation characteristic of weak intentionalism is a natural starting point because (i) any stronger relationship between the two will entail the supervenience relation (e.g., if these two types of properties turn out to be identical, this will entail that the phenomenal at least supervene on the representational), and (ii) any hopes of giving a naturalistic account of phenomenal properties requires delineating which facts about mental state form the minimal supervenience base.

[2] The supervenience claim could take one of various forms: (a) local versus global; and (b) intramodal versus intermodal. The first distinction concerns whether supervenience only holds for a certain class of mental states or for all mental states, e.g. whether it holds only for visual perceptual states or for all states including moods. The second distinction concerns whether supervenience holds only for pairs of states of the same sense modality, or for any random pair of states across modalities. Notably, as Speaks (2015) points out, these distinctions are a matter of degree. Since the modality of this phenomenology doesn't allow straightforward delineation, the intramodal/intermodal distinction won't provide any traction. For the same reason, since we are exploring a new type of phenomenology previously untouched in the debate, we are concerned with global rather than local intentionalism given that only global intentionalists promise to find a supervenience base for *any* pair of phenomenal experiences *whatsoever*.

[3] First, TI reduces phenomenal states to intentional states, which in turn are explained in terms of the tracking relation of broadly physical properties and states if affairs. The reduction of intentionality to tracking enjoys recent popularity (Mendelovici and Bourget 2014).

of what does my visual experience of the gray cat has its distinctive phenomenal character?

According to intentionalism, my experience of the gray cat has its distinctive phenomenal character in virtue of the way it *represents* the world as being (Cutter and Tye 2011). Roughly, this reductionist program operates in the following way: all phenomenal consciousness is spelled out in intentionalist terms, while intentionality is spelled out in tracking terms (Mendelovici and Bourget 2014). For example, tokens of a state S in an individual x represent that p in virtue of the fact that: under optimal conditions, x tokens S iff p, and because p (Cutter and Tye 2011).

2.1 Past alleged counterexamples to intentionalism

The main counterexample strategy for intentionalism is to refute the supervenience claim, i.e. to find a change in phenomenal properties of a subject's experience that is *not accompanied* by a change in the way the world is represented. A common pattern in intentionalist replies is to find some, however small, difference in the way the world is represented along with the phenomenal change in question. After seeing how the intentionalist handles past alleged counterexamples to intentionalism, we will be in a good position to consider whether a change in PME is also vulnerable to the same style of refutation, i.e., whether a change in PME is in fact necessarily accompanied by a change in the way putatively external objects are represented.

Consider the putative counterexample of blurred vision. If you take off your glasses or simply stop focusing your eyes, you will experience certain blurriness in your visual field. In other words, you will see objects *in a blurry way* without seeing them *as* blurry. According to Boghossian and Velleman (1989), what the experience represents may well be the same before and after taking off the glasses, but the phenomenology is different.

The intentionalist response to this example is a familiar one: there is a change in the way features of the external objects are presented. With blurry vision, the edges of objects have indistinct contours. There is less information about the putative location of the edges in the blurry condition than in the focused condition. Thus, there is indeed a representational difference as well as a phenomenal difference between the two conditions (Tye 2006).

This response exploits the *gestalt* function of attentional mechanisms (i.e., allowing the appearance of certain features to become brighter at the expense of the brightness of other features, or gaining more information about some features at the expense of information about others).[4] For example, as you focus on the computer monitor in front of you, you inevitably "zoom out" of your surroundings.

However, empirical findings suggest that the *gestalt* function of attention is not a necessary feature of attention. Another function of attentional mechanisms rarely discussed in the philosophical literature is to handle *interference* in processing (henceforth interference). As we will see in the following section, the subject does not notice

[4] Evidence for the gestalt function of attention could be found in studies showing that attention enhances the perception of low-level visual features including contrast (e.g., Carrasco et al. 2004).

the phenomenal change associated with this type of attentional control via *tracking* the way features are presented.

This is exactly what we need for a counterexample to go through. We need introspection to reveal a phenomenal difference between two mental states without an accompanying difference in the way the world is represented. In the next section, I propose to take a closer look at the *interference suppressing* role of attentional control mechanisms previously untouched by the philosophical literature. What will emerge is the following dissociation: the interference suppressing role of attention both (i) results in the phenomenology of mental effort, and yet (ii) does not change the way putatively external features are presented.

2.2 Why PME poses a problem for the intentionalist

Cognitive neuroscience has long been interested in the way the brain handles challenging and non-routine situations. These cognitively taxing situations have been most studied in association with one of three contexts: (1) tasks that require the overriding of prepotent responses (2) tasks that require selection among a set of equally permissible or under-determined responses, or (3) tasks that involve the commission of errors. A conflict in attentional mechanisms can occur when an automatic response needs to be inhibited, or no automatic response is available (e.g. working memory tasks like problem-solving) (Botvinick et al. 2004).

Importantly, it has been shown that these types of mental processes are typically associated with a subjective feeling of mental effort (Naccache et al. 2005; Dehaene et al. 1998). The main working paradigm for studying mental effort is to treat it as a result of interference in information processing. Certain processes interfere with one another for a variety of reasons. Some of these reasons are determined by the organization of our *cognitive* architecture. For example, the extent to which certain processes become automatized and are subsequently hard to override is determined by our cognitive organization.[5] Other reasons involve the organization of our *neural* architecture.

Van Veen et al. (2001) characterize the plurality of interference types in the following way: "In cognitive psychology, information processing is often thought of as occurring at a number of different levels, which might correspond to the different phases of task processing, for example (i) stimulus encoding, (ii) target detection, (iii) response selection, and (iv) response execution. Theoretically, conflicts might occur at any or all of these levels" (Van Veen et al. 2001, p. 1302).While (i), (ii), and (iv) might affect the way putatively external features are presented, (iii) does not (Milham et al. 2001). To make this more concrete, let's consider an example of PME where the *gestalt* features of attention are precluded by the conditions of the experiment.

Sàenz et al. (2003) explored the nature of information-processing interference while using the conflicting features of direction and speed. Consider the description of the relevant condition in the experiment:

[5] Botvinick et al. (2004) and Cohen et al. (1990).

> Observers were instructed to divide attention equally across two stimuli placed to the left and right of a central fixation point. In the first experiment, each stimulus was a circular patch consisting of two transparently overlapping fields of upward and downward moving dots. Subjects concurrently performed a speed discrimination task on one field of dots from each side, either moving in the same direction (up or down on both sides) or in different directions (up on one side and down on the other). Thus, without changing the visual display or the spatial distribution of attention, subjects divided attention across stimuli composed of either a common feature or opposing features.

As you might have guessed, subjects experienced greater PME when they had to divide attention across different directions of moving dots than when they had to divide attention across the same direction of moving dots. Processing the speed of the targets feels more effortful when they move in different directions because speed and direction are processed by the same population of neurons (Maunsell and Newsome 1987). To eliminate confounding variables, the authors made sure that no change in the way features are presented occurs:

> We used overlapping stimuli that were identical in all conditions so that differences in task performance could not be confounded with changes in the stimulus itself or with changes in the spatial distribution of attention (ibid, p. 633)

Interestingly, what is considered easy in the 'easy' condition is that subjects have to track either *only* the upward moving dots or *only* the downward moving dots on either sphere. In order to do that, the subjects inevitably "zoom out" of the dots moving in the direction that is not relevant to the task at hand. Hence, they can simply learn to *not see* task irrelevant features. However, in the "difficult" condition, this "gestalt" switching does not take place since subjects have to track *both* the upward moving dots on the right side, and the downward moving dots on the left side. Since subjects are *primed* by each feature on the opposing side, "zooming out" is not possible. Subjects can't simply learn to ignore a certain feature (either upward motion or downward motion), since it is exactly what they are supposed to be tracking on the other side.

Importantly for our purposes, the phenomenological effect of interference is felt prior to a decrease in discriminatory performance. Hence, a change in phenomenology proceeds independently of a change in the way features are presented (Morsella et al. 2009; Milham et al. 2001; Van Veen et al. 2001).[6]

Saenz et al.'s findings provide some reason to doubt the necessity of the supervenience claim. A change in PME is not tracked via a change in the presentation of external features. Now that we have a handle of the type of phenomenology at stake, I

[6] Although the subjective ratings of mental effort were not the focus of the study, the validity of their measuring techniques has precedent in the empirical literature (Morsella et al. 2009). Specifically, to further explore the subjective dimension of cognitive effort, the Morsella et al. introduced the following paradigm for measuring subjective effects of interference. Subjects were trained to introspect the particular feeling associated with incongruent conditions of the Stroop task. This introspection training was done to ensure that "participants were introspecting the same thing during both flanker and Stroop tasks" (ibid, p. 10). The experimenters found more subjective effects for incongruent conditions than for congruent conditions. Furthermore, these changes in phenomenology were reported prior to changes in response time. Interestingly, the phenomenology itself proved almost ineffable to the subjects (ibid, p. 16).

will attempt to *further* the intentionalist account by offering ways of spelling out PME in intentionalist terms. If these accounts come up short, intentionalists face an uphill battle for finding a supervenience base for PME.

3 Obvious intentionalist replies

Before presenting more lengthy replies, I will present a few possible replies that could be handled with relative ease. First, one might wonder if the relevant PME change in the difficult condition supervenes on a change in judgment. A change in judgment would surely qualify as a change in the way the subject is representing the world. For example, as the subject finds herself experiencing more mental effort, she could go from judging that “this task is hard on level 1” to “this task is hard on level 2.”.

Several replies are available. First, numerous experiments have shown that evaluation of task difficulty is dissociable from the experience of mental effort in *both* clinical and normal populations. If PME evolved to track task difficulty in the same way pain evolved to track potential or actual bodily damage, an absence of PME during cognitively demanding tasks should spell some type of malfunction. Consider the following characterization of TI: tokens of a state S (PME) in an individual x represent that p (task is difficult) in virtue of the fact that: *under optimal conditions*, x tokens S (PME) iff p (task is difficult), and because p. Hence, if subjects are tokening PME without representing the task to be difficult, or if subjects are not tokening PME if they’re representing the task to be difficult, there must be a malfunction in the attentional mechanisms. Indeed, it is easy to see how subjects unable to feel pain are at an evolutionary disadvantage (Tye 1995). Does an absence of PME during cognitively demanding tasks spell malfunction of the attentional mechanisms?

Naccache et al. (2005) show that a subject’s attentional and cognitive abilities could be functioning normally despite (i) awareness of task difficulty level, as well as (ii), an absence of PME. The experimenters record the subject saying, “Yes, I can see how this task is a tricky one,” but the subject still reports no feeling of PME.[7] Furthermore, optimal cognitive performance in normal subjects characterized by “flow states” involve awareness of task difficulty in the absence of PME (Nakamura and Csikszentmihályi 2014).[8]

Second, it appears that experiencing the task as hard is simply re-describing the phenomenology in question. Yes, you feel more effort, and this allows you to say that the task is hard, and then even harder, but this seems nothing more than introspecting on your phenomenology and finding words for it, i.e., “this feels a certain way, so it must be because the task is hard.” However, if the intentionalist wants to provide a

[7] Naccache et al. write that “Unexpectedly, control abilities of patient RMB evaluated in various versions of this Stroop tasks were amazingly preserved… we could see the presence of an efficient dynamic regulation of control abilities as indexed by Gratton and proportion effects” (ibid, p. 1319).

[8] Mental flow states are characterized by the experience of mastery or feeling “in the zone”. For example, if playing a musical instrument and sight reading, the subject might be aware of how smoothly the experience is going. There might be associated feelings of control, i.e. being able to adjust to other players at short notice. In contrast, phenomenology of mental effort surfaces when one is aware that things are being presented smoothly, and yet there is an added awareness that it is costing some effort (even if you’re not quite sure how to “apply” it). (Csikszentmihályi 1992). Thus, it is the complete opposite of the experience of transparency.

non-circular basis for a representational difference, it can't be a mere *re-description* of the phenomenology in question. Otherwise, there is a built-in response to *any* phenomenology presented as a counterexample. As Speaks (2015) points out, this built-in response risks trivializing intentionalism. Noticing a change in your phenomenology should not be the *basis* for delineating a representational change—it should be the other way around.

Relatedly, one might notice the amount of time it's taking to complete the task and thereby infer that it must be a difficult one. However, this appraisal could have nothing to do with the presence of PME. Nonchallenging, boring tasks tend to feel as if they're stretching on for a long time while both (i) not triggering PME, and (ii) leading you to infer that they might have been difficult.

Another straightforward reply on the behalf of the intentionalist is that an increase in PME is accompanied by a change in judgement about the possibility of erring. Since PME is often felt during task performance, task parameters easily delineate possibilities of erring. This judgement change could be characterized in the following way: "Whoa! I was even closer to messing up this time, but fortunately, I skated by!"

Nonetheless, it appears that the phenomenology of judging the possibility of erring is distinct from the phenomenology of mental effort. Consider the following scenario: you are asked to predict whether a coin will land heads or tails. Having made three successful predictions so far, you are asked to call it once again. You experience an even greater feeling of the possibility of erring. Even if you are just as likely to get it wrong as you are to get it right, you feel that this time around your luck has run out and you're sure to get it wrong. However, even if you feel that you're very close to erring, this phenomenology seems distinct from the phenomenology of mental effort normally experienced during solving a difficult math problem. This resultant phenomenology will be something like excitement or dread, but not mental effort. Hence, it is unlikely that an increase in the phenomenology of mental effort is tracking the possibility of erring.

A related proposal with some intuitive appeal is that the phenomenology of mental effort tracks *trying* to achieve a yet unrealized outcome or state of affairs.[9] In exploring this proposal, a few points are worth making. First, empirical findings suggest a double dissociation between goal pursuit and PME. Goal pursuit often happens outside consciousness (Aarts et al. 2010, 2012). This pursuit might trigger mental effort in the absence of goal representation.

On the other hand, conscious goal pursuit could happen in the absence of PME. For example, subjects experiencing mental flow are aware of pursuing and completing various goals but experience no accompanying PME (e.g. Csikszentmihályi 1992).

One way for the intentionalist to make sense of this data is to suggest that despite these cases, in the paradigm case, when a subject experiences mental effort, she represents herself as trying to achieve some goal. However, it seems that the phenomenology associated with representing oneself as trying to achieve some goal is distinct from the phenomenology of mental effort.

Going back to our previous example, I could represent myself as trying to predict whether the coin will land heads or tails. However, my *trying* is merely due to the

[9] I thank an anonymous referee for this suggestion.

uncertainty of the outcome and not the effort in bringing it about. Perhaps the intentionalist proposal could be refined to delineate the representational base of PME as "trying to achieve something difficult." However, the same type of reply is available. It seems that *trying* to predict whether the die will land on a certain number is difficult, but it does not involve any phenomenology of mental effort.

One might point out that what distinguishes these cases from instances of PME is the following: while you are a passive observer to the coin tosses, you're an active participant in exerting mental effort. That is, one might suggest that you're actively doing something to control the outcome when you're experiencing PME. Hence, PME is akin to the phenomenology of agency.

However, I would like to distinguish the phenomenology of mental effort from the phenomenology of agency. While the latter involves a feeling of *control*, the former does not.

Consider the following example of agentive phenomenology. As you apply more effort to pedaling, you're simultaneously aware of the effect this effort has on your muscles and the pace of the bike. Tracking the correspondence between effort and represented changes allows a feeling of control.[10]

On the other hand, no such correspondence exists between *mental* effort and the representation of worldly changes. Attempting to increase mental effort does not result in a reliable clarity increase. For example, a chess player could experience increasing levels of PME, but see no increase in the clarity of the problem. Try as we might, it is not always the case that we can simply "concentrate harder" when attempting a difficult mental problem. Unlike squinting, we can't simply "flex our mental muscles" in order to "see" the solution a little bit clearer.

Furthermore, even if no apparent change in the world is tracked, feelings of control in the phenomenology of agency still involve tracking some reliable effect our agential acts have on the represented state of the agent herself. Consider the following example: say you are trying to move an enormous boulder while noticing no accompanying change in the world (your efforts are futile since the boulder is too heavy). Even though no apparent changes in the world are tracked, you still experience a reliable correspondence between *trying itself* and some state of affairs. For example, you're aware that the harder you push the boulder, the greater the muscle burn, and so on. Furthermore, in experiencing agential phenomenology, you feel a sense of control in increasing your levels of effort appropriate to the task.

In contrast, in trying to solve a difficult math problem, you are unable to control "flexing your mental muscles" in increasing degrees. You have no feeling of control in trying to exert mental effort (level 1), and then mental effort (level 2). You might have no idea how you went from level 1 to level 2 and back to 1 again! That is, you might slip in and out of *flow states* without being aware of actively trying to change your levels of effort (e.g. Csikszentmihályi 1992).

[10] Christensen et al. (2015) illustrate the nature of agentive phenomenology in professional mountain biking. They argue that complex action involves a complex parametric structure (p. 344). For example, the agent is (somewhat) aware that applying pressure to the bike brakes influences the "control parameter" to change the speed ("target parameter"), which in turn could influence another parameter (e.g. curvature around the turn). So, a skilled agent can navigate the upcoming sharp curve via manipulating the amount of pressure she applies in order to manipulate the speed needed to make the turn.

In sum, it appears that the phenomenology of agency is distinct from the phenomenology of mental effort. While the former involves a feeling of control, the latter does not. Hence, it does not seem that PME tracks *trying* to achieve some state of affairs. Now that we've gotten some shorter replies out of the way, let's move on to considering more lengthy replies on the behalf on the intentionalist.

3.1 Is PME another type of *bodily* phenomenology?

Some tracking intentionalists might want to argue that the phenomenology of mental effort is just another type of bodily phenomenology, and thus could be handled in the same way. Some familiar examples of bodily phenomenology include the phenomenology of certain localized pains, muscle aches, and so on. For example, if you suddenly spill hot coffee on your leg, you experience a burning sensation on the skin surface of your leg. According to TI, your attention goes to the location of the bodily disturbance. That is, your pain appears to be on or in your body at location X.

Now, we could plausibly ask: "Could PME represent that there is a bodily sensation of type *d* at location *l*"? *Embodiment* arguments are frequently used for theories of emotion.[11] Prinz (2005) argues that bodily changes associated with emotions can, in fact, *induce* the feelings of emotions and thus could be taken to be *sufficient* for emotional experience (Ledoux 1996). Could a similar line of argument be used for PME?

Let's take a closer look at the empirical backbone of the embodiment hypothesis. One common feature associated with mental hard work is bodily tension. While working on a difficult math problem, some might notice a familiar tension in their neck. Bodily muscle tension could thus serve as a supervenience base for the PME, i.e., changes in PME require changes in the representation of patterned bodily muscle tension.

Unfortunately, this candidate does not hold up to empirical scrutiny. Studies have shown that patients paralyzed below the neck still experience mental effort while engaging in effortful motor imagery like mental rotation of three-dimensional objects (Alkadhi et al. 2005). Similarly, Cramer et al. (2005) show that subjects experience PME during mental imagery despite having lost all feeling below the neck.

Perhaps the embodiment hypothesis could be refined to survive these findings. Even if bodily muscle tension does not play a key role in PME, a feeling of tightening in the *facial* musculature could serve as a supervenience base for PME. After all, knitted brows are a familiar sight in academic settings. The facial feedback hypothesis has enjoyed some standing in the empirical literature on emotion. Several studies have shown that when subjects are induced to make a certain emotion-specific facial expression (grimacing, frowning, etc.), they report experiencing the corresponding emotion (disgust, anger, and so on).[12]

[11] Prinz (2005), Damasio (1994), and Tye (2008).

[12] Duclos and Laird (2001), Duclos et al. (1989), Flack et al. (1999), Kellerman and Laird (1982), for extensive review, see Bresler and Laird (1992), see also Niedenthal (2007) and Maringer and Niedenthal (2009).

Unfortunately, the facial feedback hypothesis has run into serious replication problems. The facial feedback hypothesis states that affective responses can be influenced by facial expressions, (e.g., smiling), even in the absence of emotional experiences. Specifically, Strack et al. (1988) had participants rate the funniness of cartoons while holding a pen in their mouths, thus inducing a "smile". Surprisingly, 'smiling' subjects rated the cartoons as significantly funnier than when they held the pen with their lips (inducing a 'pout'). However, this seminal study of the facial feedback hypothesis has not been replicated. In fact, the results of 17 independent direct replications of Study 1 from Strack et al. (1988), have failed to replicate (Wagenmakers et al. 2016).

Further trouble for the facial feedback hypothesis stems from case studies of subjects with facial paralysis. For example, Keillor et al. (2002) showed that a patient (F.P.) suffering from bilateral facial paralysis could still report normal emotional experiences despite her inability to convey emotions through facial expressions.

While the generalizability of results drawn from a few case studies pulls little weight, inference to the best explanation leans against the facial feedback hypothesis for PME. That is, our intuitions in the thought experiments, the replication troubles of the facial feedback hypothesis, as well as outcomes of individual case studies of facially paralyzed patients all suggest that the embodiment thesis is an unlikely option for the intentionalist.

4 Representing scarce resource depletion and opportunity cost?

So far I have argued that PME does not seem to be directed at any externally presented property. Perhaps tracking intentionalists could argue that PME represents a fact about the organism, e.g. cognitive or neural resource availability in relation to lost opportunity cost (Kurzban et al. 2013; Westbrook and Braver 2015).

A clarificatory note is in order. So far we have honored tracking intentionalists' commitment to phenomenal externalism. According to phenomeal externalism, phenomenal properties are not in the mind but are out there in the world, (or on the subject who is in the world). This section will branch out to a more recent development of phenomenal externalism in the tracking intentionalist literature. That is, phenomenology doesn't only track seemingly external object features like bodily damage and brightness, but also tracks general *states of affairs.*

TI's ability to account for tracking states of affairs has enjoyed some recent success (e.g., Hill 2009). Cutter and Tye (2011) argue that pain phenomenology, i.e., the experienced badness of pain, not only tracks bodily damage but also its threat to the organism's well-being:

> Our pain experiences do not just represent the presence of tissue damage, but also (roughly) represent our tissue damage as being bad for us to some degree. This view, we argue, is independently motivated by the phenomenology of pain experience, and we show how it is consistent with, and indeed predicted by, the tracking theory of intentionality (ibid, p. 91).

Analogously, intentionalists could argue that in representing cognitive resource depletion, PME tracks a state that is "bad for us to a certain degree." In this case, an increase in PME tracks the opportunity cost associated with continued resource expenditure.

Consider the following fodder for this hypothesis. Kurzban et al. (2013) argue that phenomenology of effort is associated with a cost/benefit computation. They theorize that given resource limitation; executive resource allocation should come with phenomenological tagging. If any resource is limited, its allocation carries an opportunity cost, i.e. the more resources are deployed in a particular task, the less resources could be deployed elsewhere. Hence, increased resource allocation should come with greater phenomenological reminder that perhaps these resources should be either (i) conserved, or (ii) redeployed elsewhere. Analogously, if physical energy expenditure and the related opportunity cost came with phenomenological tagging, it might increase one's chances of survival. If one could feel that his current task is using up the last of his energy resources, he would have a strategic advantage, e.g., entering confrontation with an inaccurate estimate of energy availability could prove deadly.

Perhaps intentionalists could extend this line of reasoning to PME. If mental resources are finite, their use should be phenomenologically marked in service of formulating competitive strategies. Hence, an increase in PME represents or tracks opportunity cost associated with continued resource expenditure.

Let's take a closer look at the predictions of this view. If PME tracks opportunity cost associated with continued resource expenditure, the greater the PME, the greater opportunity cost, as fewer resources remain for other tasks. Presumably, the longer this goes on, the harder it gets to recruit more resources to the task.

However, these predictions aren't exactly borne out. In fact, studies show a *decrease* instead of an *increase* in PME following ongoing resource allocation (Botvinick et al. 2001a, b; Carter and van Veen 2007). Using the Stroop paradigm, Botvinick et al. 2001a, b showed that incongruent (i.e. cognitively demanding) trials induce *more* PME when such trials are *rare* in comparison to congruent trials in a given set. If hefty resource allocation is rare, they require a total of *less* resource allocation. If they require a total of *less* resource allocation, they should introduce *less* opportunity cost (since a lot more resources remain to potentially redeploy elsewhere), and trigger *less* rather than more PME.

Consider the following analogy. You are asked to jog lightly around the track with intermittent sprinting intervals. In scenario one, you jog lightly for two laps with only two 100-m sprinting intervals at the end of each lap, totaling in 600 m of light jogging and 200 m of sprinting. In scenario two, you sprint every other 100 m, totaling in 400 m of sprinting and 400 m of jogging. Which scenario would produce greater feelings of effort: 600 m of jogging combined with 200 m of sprinting or 400 m of jogging combined with 400 m of sprinting? I can attest that the second scenario would induce greater amount of effort phenomenology than the first.

In sum, if PME is supposed to track opportunity cost associated with resource expenditure, and fewer resources are used if a smaller number of demanding tasks are being performed during a set interval of time, then subjects should be experiencing *less* rather than *more* PME during sets with *fewer* mentally demanding trials.

Bayne and Levy's (2006) account of mental effort provides yet another option for the intentionalist. Bayne and Levy categorize mental effort as a component of agentive experience. According to the authors, “the experience of mental effort involves a *representation* of the utilization and progressive fatigue of mental muscles” (ibid, p. 17). In other words, while undergoing PME, we represent our use of “mental strength” and the effect this use has on our remaining mental resources. Bayne and Levy characterize the phenomenology of mental effort in the following way:

> Anyone who has struggled with a difficult conceptual issue has experienced the effort involved in thinking a problem through. It gives rise to characteristic feelings of tiredness and a growing urge to stop (ibid, p. 13).

This thesis fits well with the opportunity cost view outlined above. After all, resource finitude is an essential component of opportunity cost. If the resources are not limited, their allocation does not involve an opportunity cost.

Before going over the empirical details behind this hypothesis, two distinctions should be made clear. First, as discussed in the previous section, likening PME to agentive phenomenology appears to be on the wrong track. While the latter involves a feeling of control, the former does not. Second, it is plausible to distinguish between *fatigue* phenomenology and *effort* phenomenology. The former involves an awareness of current resource availability, while the latter does not. In particular, when you are feeling fatigued, you are aware that you are close to not being able to carry on. On the other hand, when you experience effort, you might be unaware that you can’t carry on. For instance, say you are running and decide to suddenly increase your pace considerably. Unless you’re not used to running, you could very well be aware of applying effort, but remain ignorant that you can’t keep up this pace for a long time. On the other hand, if you are running and suddenly experience fatigue, you immediately feel that you cannot keep up for much longer. Similarly, the phenomenology of mental effort does not have to be accompanied by an awareness that you will not be able to carry on for much longer. For example, you could feel pleasantly challenged by a crossword puzzle and remain optimistic about your ability to carry on indefinitely.

Nevertheless, for the purposes of this discussion, I am willing to set this distinction aside and discuss Bayne and Levy’s thesis as a thesis about the phenomenology of mental effort. To review, Bayne and Levy argue that mental effort phenomenology is a type of agentive phenomenology that involves “a representation of the utilization and progressive fatigue of mental muscles”.

While this characterization seems like an attractive option for delineating the representational content of PME, it still falls short. Consider the following summary of its shortcomings:

(a) Unlike the individuals who are unable to feel pain, those who lack PME function perfectly well. These cases suggest that the phenomenology of mental effort did not evolve to track cognitive resource depletion (i.e., alert the organism of the utilization and progressive fatigue of mental muscles).
(b) Unlike physical energy use, frequent “mental energy” use in a given interval of time produces *less effort* than infrequent energy use.
(c) No systematic relation indicative of *tracking* could be established between cognitive resource depletion and PME.

Let us take a closer look at (a). An increase in PME is supposed to track increased mental resource expenditure because PME *evolved* to signal resource expenditure.[13] Hence, an absence of PME would spell malfunction. However, as discussed in the previous section, studies show that subjects lacking PME function well.[14] Critchley (2003) found that the patients missing PME had well-preserved general intellectual functions. They performed generally well on many demanding clinical tasks sensitive to frontal executive functions. In sum, while nothing precludes us from speculating that PME might turn out to have other (possibly derived) adaptive value, tracking cognitive resource depletion does not seem to be it.

We have already encountered (b) in discussing the opportunity cost account above. Briefly, if PME tracks mental resources expenditure, an increase in PME should signal more resource expenditure. However, studies show a decrease instead of an increase in PME following ongoing resource allocation (Botvinick et al. 2001a, b; Carter and van Veen 2007).

A further hurdle for Bayne and Levy's hypothesis involves narrowing down the notion of cognitive resources. After all, a *tracking* relation should be associated with some measure of systematicity. If the concept of cognitive resources defies systematic treatment, then Bayne and Levy's intentionalist characterization of PME is on the wrong track. The systematic nature of the tracking relation is nicely summarized by Hilbert and Klien (2014) in the following way.

> Given very general assumptions about physiology and the evolution of nervous systems,
> what is to be expected is that the internal states that *track environmental features* will have some systematic structure (ibid, p. 300)

How are we to understand the notion of "cognitive resources" and their systematic depletion? One non-metaphorical answer is that cognitive resources could be understood in neural terms. In other words, cognitive resources are neural changes underlying cognitive resource depletion.

Is there a systematic relationship between neural resource depletion and PME? In order to answer this question, we need to take a closer look at how the current scientific community is talking about concepts like "finite capacity" and "neural resources." What exactly is a "limited cognitive resource" and in what way can it become depleted?

Explaining degrading performance with increased task difficulty has typically been explained by metaphorical allusions to "finite working memory capacity" that is spread more "thinly" with increased "task load" (Baddeley 1996). However, any systematic or *structural* relationship between these placeholders is yet to be empirically demonstrated. Many researchers suggest that what creates a tax on attention control processing varies across contexts (Franconeri et al. 2012). Perhaps the limitation comes from the nature of neural processing prevalent throughout the cortex, i.e., surround

13 Recall that according to TI, representation is grounded in evolutionary histories. For example, pain experience tracks bodily damage and badness because it was selected to carry information about bodily damage and badness by reliably correlating with these features of the environment (Cutter and Tye 2011).

14 To review, Naccache et al. (2005) show that subjects' attentional and cognitive mechanisms function normally in the absence of PME.

inhibition of an activity 'peak' of some neural processes on other processes in their neighborhood.

Franconeri et al. summarize this point in the following way:

> Items interact destructively when they are close enough for their activity profiles to overlap, due to the inhibition zone that typically surrounds each activity peak. These suppressive surrounds sharpen the activity profiles of single items and resolve inter-item competition–a critical step especially when unitary actions are needed (e.g., a saccade to a single location) (ibid, p. 3).

In this two-dimensional 'map' architecture, capacity is not fixed but flexible, determined by a number of fluctuating variables, i.e., the space taken by the activity profile of certain items on the map, how these items interact with one another (e.g., whether they interact constructively or destructively is partly determined by the inhibition zone that surrounds each activity peak), the spacing on the items on the map, and so on. All these factors contribute to the competition between items at the neural level. As this competition is resolved, a more pronounced activity peak takes place. Furthermore, the way these competitive interactions resolve is not fixed but flexible. Hence, alluding to a limited capacity that somehow becomes systematically depleted with use is misleading. This organization gives the brain an adaptive edge, but it also has side effects. Perhaps PME is one such side effect, alerting the organism to no particular state of affairs.

The upshot is that cognitive neuroscience does not provide evidence of a match in structural relations between PME and neural changes in the way we have seen between pain phenomenology and somatosensory cortex activation (e.g., Price et al. 1994, illustrate the structural relationship between pain intensity and neuron firings in S1). Since tracking relations are characterized by a match in systematic relations (e.g., similarity, differences, equal intervals, proportions), it does not appear that PME tracks cognitive resource depletion.

Inzlicht and Marcora (2016) summarize the criticisms of the resource depletion model (i.e., "the central governor") in the following way:

> There is no credible evidence that mental effort actually consumes inordinate amounts of energy that are not already circulating in the brain. Recent modifications of the model make the central governor appear like an all knowing homunculus and unfalsifiable in principle, thus contributing very little to our understanding of why people tend to disengage from effortful tasks over time (ibid, p. 1)

Hence, it does not appear that Bayne and Levy have provided sufficient reason to think that the phenomenology of mental effort tracks mental resource depletion.

5 Conclusion

To sum up, let us briefly review the main argument for TI and how it fares against PME. Arguments from introspective difference state that, necessarily, if there is an introspectable difference in the two phenomenal properties of subjects, then there is a difference in the objects and properties those subjects *represent as* in their environment. However, a difference in PME does not accompany a difference in either (a) representation of the way external features appear, i.e., brighter, with more contrast and so on, (b) representation of task difficultly, (c) judgement of the possibility of error, (d) representation of trying to achieve some state of affairs, (e) representation of bodily changes like muscle tension, (f) representation of cognitive resource depletion and opportunity cost.

While local intentionalism about some phenomenal experiences like pains might obtain, it does not look like it obtains for *all* phenomenal experiences. This puts the intentionalist in an uncomfortable position of trying to explain why *some* phenomenal experiences have representational content and not others. Moreover, even if as much as one type of phenomenal experience doesn't have representational content, reductionist theories of consciousness are under threat.

Acknowledgements Many heartfelt thanks to people who have provided invaluable feedback on this paper, including Bernhard Nickel, Elizabeth Schechter, Casey O'Callaghan, Neil Levy, Alex Byrne, Wayne Wu, Ned Block, Edouard Machery, Jeff Speaks, Adam Pautz, generous conference commentators, and many others.

References

Aarts, H., Bijleveld, E., & Custers, R. (2010). Unconscious reward cues increase invested effort, but do not change speed-accuracy tradeoffs. *Cognition, 115,* 330–335.

Aarts, H., Bijleveld, E., Daunizeau, J., Veling, H., & Zedelius, C. M. (2012). Promising high monetary rewards for future task performance increases immediate task performance. *PLoS ONE, 7,* e42547.

Alkadhi, H., Brugger, P., Boedermaker, S. H., Crelier, G., Curt, A., Hepp-Reymond, M. C., et al. (2005). What disconnection tells about motor imagery: Evidence from paraplegic patients. *Cerebral Cortex, 15,* 131–140.

Baddeley, A. (1996). Exploring the central executive. *Quarterly Journal of Experimental Psychology, 49,* 5–28.

Bayne, T., & Levy, N. (2006). N. Sebanz & W. Prinz (Eds.), *The feeling of doing: Deconstructing the phenomenology of agency* in *disorders of volition* (pp. 49–68). Cambridge, MA: MIT Press.

Bicknell, K., Christensen, W., Mcllwain, D., & Sutton, J. (2015). The sense of agency and its role in strategic control for expert mountain bikers. *Psychology of Consciousness: Theory, Research and Practice, 2,* 340–353.

Boghossian, P., & Velleman, J. D. (1989). Color as a secondary quality. *Mind, 98,* 81–103.

Botvinick, M., Carter, C. S., Braver, T. S., Barch, D. M., & Cohen, J. D. (2001a). Conflict monitoring and cognitive control. *Psychological Review, 108,* 624–652.

Botvinick, M., Cohen, J. D., & Carter, C. (2004). Conflict Monitoring and anterior cingulate cortex: An update. *Trends in Cognitive Sciences, 8*(12), 539–546.

Botvinick, M. M., Braver, T. S., Barch, D. M., Carter, C. S., & Cohen, J. D. (2001b). Conflict monitoring and cognitive control. *Psychological Review, 108,* 624–652.

Bresler, C., & Laird, J. D. (1992). The process of emotional feeling: A self-perception theory. In M. Clark (Ed.), *Emotion: Review of personality and social psychology* (pp. 223–234). Newbury Park, CA: Sage.

Bush, F. M., Francis, M., Harkins, S. W., Long, S., & Price, D. D. (1994). A comparison of pain measurement characteristics of mechanical visual analogue and simple numerical rating scales. *Pain, 56,* 217–226.

Carrasco, M., Ling, S., & Read, S. (2004). Attention alters appearance. *Nature Neuroscience, 7*(3), 308–331.

Carter, C. S., & van Veen, V. (2007). Anterior cingulated cortex and conflict detection: An update of theory and data. *Cognitive Affective and Behavioral Neuroscience, 7,* 367–379.

Cavallaro, L. A., Flack, W. F., & Laird, J. D. (1999). Separate and combined effects of facial expressions and bodily postures on emotional feelings. *European Journal of Social Psychology, 29,* 203–217.

Cohen, J. D., Dunbar, K., & McClelland, J. L. (1990). On the control of automatic processes: A parallel distributed processing account of the Stroop effect. *Psychological Review, 97*(3), 332.

Cramer, S. C., Lastra, L., Lacourse, M. G., & Cohen, M. J. (2005). Brain motor system function after chronic, complete spinal cord injury. *Brain, 128,* 2941–2950.

Critchley, H. D. (2003). Human cingulate cortex and autonomic cardiovascular control: Converging neuroimaging and clinical evidence. *Brain, 216,* 2139–2152.

Csikszentmihályi, M. (1992). *Flow: The psychology of happiness*. Rider. ISBN 978-0-7126-5477-7.

Cutter, B., & Tye, M. (2011). Tracking representationalism and the painfulness of pain. *Philosophical Issues, 21,* 90–109.

Damasio, A. R. (1994). *Descartes' error: Emotion, reason and the human brain*. New York: Gossett/Putnam.

Dehaene, S., Chengeux, J. P., & Kerzberg, M. (1998). A neuronal model of a global workspace in effortful cognitive tasks. *Proceedings of the National Academy of Sciences of the United States of America, 95,* 14529–14534.

Duclos, S. E., & Laird, J. D. (2001). The deliberate control of emotional experience through control of expressions. *Cognition and Emotion, 15,* 27–56.

Duclos, S. E., Laird, J. D., Schneider, E., Sexter, M., Stern, L., & Van Lighten, O. (1989). Emotion-specific effects of facial expressions and postures on emotional experience. *Journal of Personality and Social Psychology, 57,* 100–108.

Franconeri, S. L., et al. (2012). Flexible cognitive resources: Competitive content maps for attention and memory. *Trends in Cognitive Sciences, 17*(3), 134–141.

Hilbert, D., Klien, C., (2014). "No problem". In R. Brown (Ed.), *Consciousness inside and out: Phenomenology, neuroscience, and the nature of experience, studies in brain and mind* (pp. 229–306). Dordrecht: Springer.

Hill, C. S. (2009). *Consciousness*. Nova York: Cambridge University Press.

Inzlicht, M., & Marcora, S. M. (2016). The central governor model of exercise regulation teaches us precious little about the nature of mental fatigue and self-control failure. *Frontiers in Psychology, 7,* 656.

Keillor, J., Barrett, A., Crucian, G., Kortenkamp, S., & Heilman, K. (2002). Emotional experience and perception in the absence of facial feedback. *Journal of the International Neuropsychological Society, 8*(1), 130–135.

Kellerman, J. M., & Laird, J. D. (1982). The effect of appearance on self-perceptions. *Journal of Personality, 50,* 296–351.

Kurzban, R., Duckworth, A., Kable, J. W., & Myers, J. (2013). An opportunity cost model of subjective effort and task performance. *Behavioral and Brain Sciences, 36,* 661–679.

LeDoux, J. E. (1996). *The emotional brain: The mysterious underpinnings of emotional life*. New York: Simon and Schuster.

Maringer, M., & Niedenthal, P. M. (2009). Embodied emotion considered. *Emotion Review, 1,* 122–128.

Maunsell, J. H., & Newsome, W. T. (1987). Visual processing in monkey extrastriate cortex. *Annual Review of Neuroscience, 10*(1), 363–401.

Mendelovici, A., & Bourget, D. (2014). Naturalizing intentionality: Tracking theories versus phenomenal intentionality theories. *Philosophy Compass, 9*(5), 325–337.

Milham, M. P., Banich, M. T., Webb, A., Barad, V., Cohen, N. J., & Wszalek, A. F. Kramer. (2001). The relative involvement of anterior cingulate and prefrontal cortex in attentional control depends on nature of conflict. *Cognitive Brain Research, 12,* 467–473.

Morsella, E., Wilson, L. E., Berger, C. C., Honhongva, M., Gazzaley, A., & Bargh, J. A. (2009). Subjective aspects of cognitive control at different stages of processing. *Attention, Perception, & Psychophysics, 71*(8), 1807–1824.

Naccache, L., Dehaene, S., Cohen, L., Habert, M. O., Guichart-Gomez, E., Galanaud, D., et al. (2005). Effortless control: Executive attention and conscious feeling of mental effort are dissociable. *Neuropsychologia, 43*(9), 1318–1328.

Nakamura, J., & Csikszentmihályi, M. (2014). The concept of flow. In *Handbook of positive psychology* (pp. 89–105). New York, NY: Oxford University Press.

Niedenthal, P. M. (2007). Embodying emotion. *Science, 316,* 1002–1005.

Pautz, A. (2013). The real trouble for phenomenal externalists. In T. Brown (Ed.), *Consciousness inside and out: Phenomenology, neuroscience, and the nature of experience* (pp. 237–298). Belrin: Springer.

Prinz, J. J. (2005). Are emotions feelings? *Journal of Consciousness Studies, 12*(8–10), 9–25.

Sàenz, M., Buracâs, G. T., & Boynton, G. M. (2003). Global feature-based attention for motion and color. *Vision Research, 43*(6), 629–637.

Speaks, J. (2015). *The phenomenal and the representational*. Oxford: Oxford University Press.
Strack, F., Martin, L. L., & Stepper, S. (1988). Inhibiting and facilitating conditions of the human smile: A nonobtrusive test of the facial feedback hypothesis. *Journal of Personality and Social Psychology, 54,* 768–777.
Tye, M. (1995). *Ten problems of consciousness: A representational theory of the phenomenal mind*. Cambridge, MA: MIT Press.
Tye, M. (2006). The puzzle of true blue. *Analysis, 66,* 173–178.
Tye, M. (2008). *The experience of emotion: An intentionalist theory* (unpublished). https://webspace.utexas.edu/tyem/www/emotions.pdf.
van Veen, V., Cohen, J. D., Botvinick, M. M., Stenger, V. A., & Carter, C. S. (2001). Anterior cingulate cortex, conflict monitoring, and levels of processing. *Neuroimage, 14*(6), 1302–1308.
Wagenmakers, E. J., Beek, T., Dijkhoff, L., Gronau, Q. F., Acosta, A., et al. (2016). Registered replication report: Strack, Martin, & Stepper 1988. *Perspectives on Psychological Science, 11*(6), 917–928.
Westbrook, A., & Braver, T. S. (2015). Cognitive effort: A neuroeconomic approach. *Cognitive, Affective, & Behavioral Neuroscience, 15,* 395–415.