Are basic actors brainbound agents? Narrowing down solutions to the problem of probabilistic content for predictive perceivers

George Britten-Neish¹

Accepted: 9 March 2021 / Published online: 31 March 2021
© The Author(s) 2021

Abstract

Clark (*Journal of Consciousness Studies*, 25(3–4), 71–87, 2018) worries that predictive processing (PP) accounts of perception introduce a puzzling disconnect between the content of personal-level perceptual states and their underlying subpersonal representations. According to PP, in perception, the brain encodes information about the environment in conditional probability density distributions over causes of sensory input. But it seems perceptual experience only presents us with one way the world is at a time. If perception is at bottom probabilistic, shouldn't this aspect of subpersonally represented content show up in consciousness? To address this worry, Clark argues that representations underlying personal-level content are constrained by the need to provide a single action-guiding take on the environment. However, this proposal rests a conception of the nature of agency, famously articulated by Davidson (1980a, b), that is inconsistent with a view of the mind as embodied-extended. Since Clark and other enactivist PP theorists present the extended mind as an important consequence of the predictive framework, the proposal is in tension with his complete view. I claim that this inconsistency could be resolved either by retaining the Davidsonian view of action and abandoning the extended-embodied approach, or by adopting a more processual, world-involving account of agency and perceptual experience than Clark currently endorses. To solve the puzzle he raises, Clark must become a radical enactivist or a consistent internalist.

Keywords Predictive processing · Perception · Basic action · Enactivism · Subpersonal explanation

✉ George Britten-Neish
uctyagn@ucl.ac.uk

¹ UCL (Philosophy), London, UK

1 Introduction

According to the predictive processing (PP) framework (Clark, 2015; Friston, 2010; Wiese & Metzinger, 2017) perception and perceptual experience are realised by a hierarchical generative model in the brain, which enables perceivers to recover the complex causal structure of their environment from a flux of sensory data. But why should probabilistic representational states, which are at the centre of PP approaches, give rise to experience that seems to have determinate or ‘unitary-coherent’ (Lu et al. 2016) content? This puzzling disconnect between what perception is like for us and the brain mechanisms that are supposed to realise it has been presented as a *prima facie* objection to PP accounts of perceptual experience (Block, 2018), while others have claimed that it represents a challenge to widely-held intuitions about the content of experience (Gross & Flombaum, 2017; Madary, 2016; Morrison, 2016). In contrast, Clark (2018) makes an ambitious attempt to save both the appearances and the theory, presenting the mutual dependence of action and perception as the solution to these worries. Since perception is constitutively linked to action, we should not be surprised that it delivers a univocal perspective on a world of determinate objects, ready to be acted on. Consequently, we can expect that, at the relevant level of the predictive processing hierarchy, representations are constrained to deliver content of this action-serving kind.

This proposal is in keeping with a picture of cognitive agents as interdependently constituted by interactions with their environment, which recent work by Clark and others (e.g. Bruineberg & Rietveld, 2014; Clark, 2015; Gallagher & Allen, 2018; Kirchhoff & Kiverstein, 2019) presents as an important consequence of PP. For defenders of *enactivist predictive processing*,¹ PP explains how interacting with a world of cognitive artefacts and agents, occupying ecological niches, and exploiting the problem-solving resources inherent in their own bodies, are all ways that organisms expand the envelope of their information-processing into their environments. In such cases, token mental phenomena are realised in part by processes that extend into body and world. *Radical* enactivist accounts of perception (e.g. Kirchhoff & Kiverstein, 2019) claim that some token perceptual experiences are extended in this way. Clark and other *moderate* enactivists deny this, but give bodily action a constitutive role in other mental phenomena. So, Clark’s response to the puzzle centres on how capacities for action constrain internally realised mechanisms underlying perceptual experience to deliver unitary-coherent content.

I claim that this solution is unstable, due to problematic (for enactivists) assumptions Clark makes about the metaphysics of action and perceptual experience. In the context of temporal atomism (Hurley, 1998; see also Ward, 2016a) about the content of perception at the personal level, his argument requires an analysis of bodily agency as constituted by ‘basic actions’ (Danto, 1965; Davidson, 1980a, 1980b) if it is to succeed on its own terms. This leads to trouble because the basic action view entails an internalist or ‘brainbound’ (Clark, 2008) account of the mind, according

¹ See Allen & Friston (2018) for an overview of the theoretical and empirical background of enactivist PP.

to which all mental events and processes are understood in terms of input/output relations across a fixed, anatomically-defined boundary. While some PP theorists endorse this outcome (e.g. Hohwy, 2016), it is incompatible with both moderate and radical enactivism (Ward et al. 2017). Alternatively, we can reformulate Clark's (2018) proposal, denying temporal atomism and the basic action view. But in this case, we no longer have reasons to see perceptual experience as exclusively internally realised. So, while an action-based solution to the puzzle is viable, it supports either radical enactivist or consistently internalist accounts of the relationship between the predictive mind and its environment. The compromise presupposed by Clark (2018) does not survive an analysis of what action would have to be to play the role he wants it to.

Section 1 introduces some key features of PP; Section 3 outlines how these give rise to two linked puzzles about the content of perception in the context of a view of the personal/subpersonal relation that construes subpersonal process and states as elements of vertical, mechanistic explanations of personal-level phenomena; Section 4 presents Clark's action-oriented solution to the puzzles and describes moderate enactivist PP; Section 5 describes temporal atomism about perceptual experience and shows why we should attribute this view to Clark; Section 6 argues that Clark's (2018) proposal can only solve both puzzles if it relies on a basic action account, further arguing that basic actions are a natural consequence of temporal atomism; Section 7 shows how this is incompatible with even moderate enactivism; Section 8 sketches an alternative that is friendly to enactivism by rejecting temporal atomism about perception and action. From this perspective, Clark's puzzles about unitary-determinacy are reframed as questions about how probabilistic perceptual mechanisms achieve phenomenal transparency onto a world of determinate objects and events.

2 Perception as hierarchical predictive processing

In an influential review article, Clark (2013) outlines the core problem for cognitive science that PP promises to solve:

[T]he task of the brain, when viewed from a certain distance, can seem impossible: it must discover information about the likely causes of impinging signals without any form of direct access to their source [...] How, simply on the basis of patterns of changes in its own internal states, is it to alter and adapt its responses so as to tune itself to act as a useful node (one that merits its relatively huge metabolic expense) for the origination of adaptive responses? (p. 3)

PP explains how the brain pulls off this seemingly impossible trick, allowing it "to infer the nature of the signal source (the world) from just the varying input signal itself" (ibid.). While not a systematic introduction to PP (see Wiese & Metzinger, 2017), this section highlights some features that are relevant to my argument.

Understanding perception as inferring environmental structure from sensory data is an old idea in cognitive science and raises many issues that will not be the focus of the current discussion. For example, describing the putative inferences underlying

perception as Bayesian (as in e.g. Clark, 2015; Friston, 2010; Hohwy, 2013) means committing to a particular form of computation: the perceptual system updates internal models of the world in light of new sensory evidence according to Bayes' rule. Predictive coding, in turn, is one mechanism by which an approximation to 'Bayes optimal' updating could be realised. This characterisation has been criticised by some enactivist PP theorists (e.g. Bruineberg et al. 2018; Orlandi, 2016), who argue that understanding predictive perceptual architectures this way imposes an illegitimate 'Helmholtzian' (Bruineberg et al. 2018) or intellectualist model onto information processing dynamics that are better understood through other frameworks (e.g. dynamical systems theory, ecological 'pick-up' of natural scene statistics, etc.). This paper does not comment on this debate, and presents PP within the broadly Bayesian framing that Clark (2015, 2018) does, assuming that similar questions about indeterminate (if not explicitly probabilistic) perceptual phenomenology arise for non-Bayesian versions of PP.

Bayes' rule describes the way an ideally rational agent should update her beliefs in light of new evidence, framed as a formula for modifying degrees of belief assigned to a hypothesis conditional on evidence that bears on that hypothesis. It reflects the fact that rationally constrained transitions of belief take into account both new evidence and prior beliefs. Bayesians talk of evidence being interpreted under such priors. Applied to perceptual processing, this is meant to capture how systematically ambiguous sensory data are interpreted in the light of what the perceptual system already 'knows' about environmental structures that cause particular patterns of perceptual input. Percepts, on this view, are the result of the perceptual system's inference to the best explanation for occurrent sensory states.

More concretely, perceptual processing is distributed across multiple levels, each of which contain two kinds of representing units (i.e. units that can be interpreted as being in content-bearing states; Wiese, 2017a): one models the input expected from the level below the other compares predictions received from higher levels with actual input. Transitions between representational states of the former kind of unit are guided, or entrained, by feedforward (low-to-high level) error signals generated by the second kind of unit at the level below. This is a biologically plausible data-compression strategy, and means that instead of directly extracting features from a feedforward signal the system updates its representations via a bidirectional flow of prediction and prediction error. In Bayesian terms, these transitions correspond to arriving at a posterior from a prior belief state. Such states, however, always reflect the uncertainty that the system operates under, and are thus expressed as probability distributions over a range of values. While these values are not thought to be explicitly represented in neural activity, the claim is that the brain is sensitive to this approximate statistical structure. Bayesians propose that this formalism captures the broad information processing dynamics of neural activity, even if the algorithmic and implementational details are not yet well understood (see Bastos et al. 2012; Huang & Rao, 2011; critical discussion in Aitchison & Lengyel, 2017), so that the system implements a form of Bayesian inference that is distributed across processing levels and played out in real time on perceptual input. Since updated higher-level

representations generate new feedback signals to compare with occurrent input, the system tends over time to minimize the error signal and so effectively track (or, more actively, transform; Friston et al. 2010) environmental structures that are ultimately responsible for sensory data.

Given the same sensory input but different priors, two systems could arrive at very different representational states. These differences might be expressed at any (higher) level of the processing hierarchy. *Hyperpriors* can be seen as representing system-wide expectations on what these priors ought to be. Since predictive processing systems exhibit strong top-down effects, hyperpriors have a pervasive influence on the way lower processing levels respond to sensory data. Mechanisms for system-internal precision estimation also play an important role in eliciting biologically realistic activity from predictive systems (accommodating capacities like selective attention and mental imagery; Feldman & Friston, 2010; Friston, 2010). Precision determines the relative importance of either sensory data or prior expectations for determining the eventual shape of the posterior distribution. Intuitively, this reflects the level of confidence the system has in prior models of the world or the incoming sensory evidence respectively. Input that has been assigned a low precision has a correspondingly reduced effect on representational states the system arrives at (i.e. encoded posterior distributions).

3 Clark's puzzle(s): probabilistic content, determinate experience

Hopefully, an outline of what PP says about the computational underpinnings of perceptual experience should be coming into view. Throughout the system, representational content is probabilistic and distributed across discrete processing levels. Perceptual content, then, “reflects a delicate combination of top-down model-based prediction, self-estimated sensory uncertainty, and bottom-up (incoming) sensory evidence” (Clark, 2018, p. 72).

While this might not seem to pose any puzzle about perceptual *experience*, since this potentially problematic content is ‘subpersonal’, a term which many theorists have equated with ‘unconscious’ (for a review of the use, and misuse, of the term in the philosophical literature, see Drayson, 2012), this easy response is not open to Clark. This is because he follows Dennett (1969, 1978) in understanding the personal/subpersonal distinction not as between two kinds of mental content, but rather as distinguishing different modes of explanation. Accordingly, subpersonal *vertical* explanations are contrasted with folk psychological, *horizontal* explanations (Drayson, 2012).² Mental content plays a central role in both of these, but while horizontal personal-level explanations account for present states and activities of a person (partly) in terms of antecedent content-bearing

² Of course, horizontal accounts of the ‘subpersonal mechanisms’ (e.g. mechanisms for statistical inference in the brain) in terms of their causal inputs and outputs are also needed. Sometimes, outputs (or inputs) of these mechanisms will just be representations of contents that we ascribe to persons. *Vertical* subpersonal explanation then is what makes such content attribution intelligible by identifying the representational vehicle as part of a causally-described system at a lower level of explanation. In these cases, what I’m calling isomorphic subpersonal explanation is appropriate. See Section 5 below.

states that are ascribed to that person, subpersonal explanations constitutively account for these personal-level phenomena in terms of how they are physically realised. In particular, subpersonal vertical explanation proceeds by functional decomposition, taking a personal-level capacity, state, or process as an explanatory target and breaking it down into elements that are not attributed to the person as a whole, but to sub-agencies (or ‘subpersons’ Drayson, 2012; see also Dennett, 1978; Lycan, 1987; Roth, 2015) that collectively constitute the person and her capacities. For example, representations of contours in early visual processing are ascribed to low-level visual systems, rather than the perceiving subject. Typically, multiple subpersonal mechanisms take on the roles required for joint achievement of whatever personal-level phenomenon is being analysed, and are in turn subject to further decomposition until an explanatorily fundamental level is reached and the intentional language of mental content is no longer required. So, while subpersonal explanations specify mechanisms that realise personal-level phenomena characterised in terms of their content, introducing a personal level of description does not introduce autonomous ‘personal-level content’. Instead, this content must be accounted for in terms of the content-bearing states of the subpersonal mechanisms that realise it.

This is how Clark’s puzzle arises. Since, according to PP, all representations in the brain are probabilistic and hierarchically distributed, and it is only the contents of these that are supposed to play a role in explaining how our experience presents the world to us, we are really presented with two questions about how to reconcile PP with personal-level perceptual experience:

Determinate Content: Mechanisms proposed by the predictive account encode distributions over ranges of values. If, for example, perceiving an object is explained by an occurrent representation of this kind, why should only one value show up?

Univocal Contents: The predictive system is hierarchical and distributed. Perception seems to present us with a single, internally consistent take on the world. Why should these multiple representations come together into a semantically unified whole?

Respectively, these are questions about what contents show up in perceptual experience and how these contents are integrated with each other. *Indeterminate Content* is indeterminate across a continuous range of values, while non-*Univocal Contents* take distinct, incompatible values. Contents can be univocal without being determinate,³ but typically we take the content of our experiences to be both. Within the PP framework described above, however, we might wonder why any

³ An anonymous reviewer points out that this might not hold in reverse. We can (roughly) paraphrase Clark’s worry by asking why the probability distributions underlying perceptual experience should be constrained to take a certain shape (Clark, 2018: pp. 73–75; Clark, 2015: pp. 187–188): unimodal distributions assigning a high precision to a single value. Across a single distribution, high precision (i.e. $p[x] > 0.5$) can’t be assigned to more than one value since the sum of probabilities must equal 1. Looked at this way, determinate content seems to imply univocal content. But since I claim that Clark’s (2018) solution, absent the basic action strategy, only delivers univocal content, this asymmetry does not affect the argument.

perceptual representations would be either. Multiple incompatible models leaving various details indeterminate might just be the best way to account for sensory data in the complex, dynamic world we inhabit.

Importantly, these questions are not about how a hierarchical predictive processing system *could* deliver appropriate representations for unitary-coherent percepts, since – notoriously (Bowers & Davis, 2012) – it is always possible for PP theorists to build in constraints (e.g. via formal hyperpriors, precision weighting dynamics) or appeal to additional mechanisms supporting threshold effects for the content of conscious perception. That is, they are not problems that the system has to solve, like the ‘binding problem’ in perceptual neuroscience (Treisman, 1996). Rather, they are puzzles about our explanatory procedure. To justify PP accounts of perceptual experience that guarantee *Determinate* and *Univocal* perceptual content, we need to say why a such a system *should* generate unitary-coherent experience, given that this does not simply fall out of the model.

4 Clark’s action-oriented solution

In line with his longstanding theoretical commitments (Clark, 1998, 2008) and the current ‘pragmatic turn’ (Engel et al. 2015) in cognitive science, Clark offers a single action-oriented solution to the puzzles of accounting for *Determinate* and *Univocal* perceptual content, proposing that:

[o]ur perceptual worlds display unity and coherence, and depict a single way things are [due to] the transformative role of action itself. It is the need for perception constantly to mandate action and choice that requires the system to opt for a single ‘overall take’ on how the world is now most likely to be. (Clark, 2018, p. 77)

We should expect there to be unitary-coherent content at the personal level, and so posit mechanisms that ensure this, because effective action requires such content. Linking perception and action in this way makes sense within the PP framework, since both are presented as emerging from the same fundamental drive to reduce prediction error.⁴ For Clark, this means the representations that explain experienced perceptual content are “fundamentally in the business of action control, [and] represent how the world is in ways that are entwined at multiple levels with information about how to act on the world” (2015, p. 181). This talk of the rich, multiple entwinement of perceptual representation and capacities for action is characteristic of the (moderate) enactivist paradigm.

It is a *moderate* view because the constitutive link it proposes between perception and bodily action is ‘indirect’ (Ward et al. 2011). An earlier paper, co-authored by Clark, distinguishes this kind of view, in which “what both explains and suffices for [...] visual perceptual experience is an agent’s direct unmediated knowledge concerning the ways in which she is currently poised [...] over an ‘action space’” (ibid. p. 383), from more radical sensorimotor enactivism (e.g. O’Regan & Noë, 2001; Hurley, 1998) that sees actual bodily processes of engaging with the environment as playing

⁴ More on this in Section 6 below.

this role.⁵ An ‘action space’, in this sense, is generated by internally realised capacities for action-planning (the authors refer to these as ‘second-order dispositions towards action’), which are called upon by processes underlying the direct performance of ordinary pragmatic action. It is a multimodal representation of the environment parsed for bodily engagement. In PP terms, this action space is construed as a generative model that encodes predictions about the regular causal structure of the environment along with the sensorimotor consequences of multiple patterns of bodily intervention:

[T]hese will be models capturing regularities across larger and larger temporal and spatial scales. Inevitably, the higher levels here will come to encode information concerning [properties and features of material objects, e.g. cats, tomatoes]. Just as the higher levels in a shape-recognition network respond preferentially to invariant shape properties (such as squareness or circularity), so we should expect to find higher-level networks that model driving sensory inputs (as filtered via all the intervening levels of prediction) in terms of tomatoes, cats, and so forth. The overall processing hierarchy, confronted with a scene involving a cat or tomato, will relax into a stable state in which these higher-level patterns are recognized to be present. (Clark, 2012, p. 762)

Not everything represented across the full hierarchical structure of the generative model will be relevant to the perceiving agent, and Clark endorses Koch’s (2004) characterisation of experience as an ‘executive summary’ that synthesises salient aspects of sensory, interoceptive and proprioceptive information at high levels of the processing architecture. It is at this level that the perceptual system is “constrained to select, moment-by moment, a single best-fit unitary model poised for the control of action and choice.” (Clark, 2018, p. 78). These representations are selected for system-wide ‘ignition’ (Dehaene & Changeux, 2011; see also Clark, 2019; Hohwy, 2013, ch. 10), becoming globally available to subsystems serving, thought, language and, crucially, bodily action. As such, they correspond to the familiar contents of perceptual experience (tomatoes, cats, etc.), which we reason about, name in public language, and direct our actions towards. While subpersonal computation of a range of ways the environment could be (as conditioned by the consequences and contingencies of possible actions in it), should and does take place in a probabilistic and distributed way, these representations are collapsed into a legible, pragmatic guide to actions available in the moment. This guide represents the recognisable world of objects and events in which organisms like us exist, poised on the brink of effective action.

5 Isomorphic vertical explanation and temporal atomism

Following Hurley (1998), we can understand this view as a kind of temporal atomism about perceptual experience. In a classic discussion of work in psychophysics on perceptual masking effects (reviewed in Breitmeyer & Ogmen, 2000; Breitmeyer, 2015; see also Dennett & Kinsbourne, 1992) she presents temporal atomism as

⁵ Significantly, sensorimotor enactivism of this kind is central to recent arguments for radical enactivist PP (e.g. Kirchhoff & Kiverstein, 2019).

motivated by intuitions that the question of whether a particular content is manifest in experience must always have a determinate answer.

[A]t whatever precise moment we take a snapshot of the process, consciousness must have determinate content. Consciousness admits of microtiming. The existence of determinate content cannot itself take time or involve dynamic processes such that if your temporal window gets too small you lose content. (Hurley, 1998, p. 31)

Here ‘determinate content’ doesn’t mean unitary-determinate content in the sense discussed above. One could be a temporal atomist about perceptual experience and claim that experienced contents fail to be determinate or univocal. Indeed, this would be true of the ‘Bayesian blur’ that Clark (2018) addresses his action-oriented solution towards. Temporal atomism by itself only entails determinacy about *whether* a given content is in experience at a given moment. This is determinacy about the answers to questions like ‘what is the content of consciousness at *t*?’ Answers to this question, which are guaranteed by temporal atomism to yield some independently (i.e. not merely instrumentally) specifiable content, are internally related to subpersonal questions about how such contents are realised.

Call Clark’s approach here *isomorphic vertical explanation*. For conscious perception, this is the view that perceptual experience literally shares content with (some) representations realised by its underlying subpersonal mechanisms. Since these always occupy some content-bearing state, it always makes sense to ask what the content of experience is (i.e. to ask which representations of the environment are currently poised to guide action and choice). Isomorphic vertical explanation offers one way of explaining the personal in subpersonal terms, but it is not the only possible approach.⁶ In Dennett’s early discussion of subpersonal explanation, relating the personal to the subpersonal is more a question of correlating two ways of understanding a complex process – the workings of a brain, within an organism, within an environmental, linguistic, and interpersonal context – than of directly identifying terms in the two descriptions (Dennett, 1969, ch. 4). (In)famously, Dennett (1989) ultimately presents content-ascription at the personal level as instrumental relative to our practices of explaining behaviour. On such a view, temporal atomism can be consistently denied (as in, e.g. Dennett & Kinsbourne, 1992). But once the content of experience and the content of representational states of the brain are identified, it follows that the personal-level phenomenon is temporally atomistic.

So, while isomorphic vertical explanation describes a way of specifying personal-level targets for subpersonal explanation, temporal atomism is a claim about the metaphysics of the phenomena thereby explained. Specifically, it is a claim about how episodes of experience are constituted across time. Discussion of the ontology of mental phenomena (e.g. O’Shaughnessy, 2000; Soteriou, 2013; Steward, 1997, 2013) often highlights the importance of a distinction between mental processes and states, based on their temporal dynamics. This gives us a framework for characterising entities referred to in both personal and subpersonal accounts. Processes are

⁶ For example, instrumentalism (Dennett, 1989) and ‘virtualism’ (Clowes & Chrisley, 2012) about personal-level content offer non-isomorphic vertical explanations of the content of perceptual experience.

characterised by having non-identical temporal parts; they unfold over time rather than simply obtaining for a stretch of time. In contrast, states are uniform across all times at which they obtain. Dispositional knowledge, at the personal level, and representational states, subpersonally, are paradigmatic mental states, since they are individuated by accuracy conditions that are timeless. If I explain an agent's dispositional knowledge, *P*, over a stretch of time by the tokening of a (usually unconscious) representation '*P*', in her brain over that period, I am providing an isomorphic vertical explanation of her knowledge.⁷ Similarly, if perceptual experience at any moment can be given by specifying the representation state of the relevant subpersonal mechanism, episodes of perceptual experience are constituted across time by percepts that share content with these representations. In this case, perceptual experience is a process only in the trivial sense that it involves transitions between representational states which make up its temporal parts. These are the fundamental units of experience, fixing personal-level content on a moment-by-moment basis and ensuring that Hurley's question always has a determinate answer. Answering puzzles about *Determinate* and *Univocal Contents* in the context of temporal atomism about perceptual experience means ensuring such atomic perceptual states have contents of this kind.

At first glance, this understanding of the metaphysics of perceptual experience might seem to imply a controversial further claim about perceptual phenomenology.⁸ If perception is constituted by atomic percepts, how is distinctively temporal phenomenology possible? Kelly (2005), for example, has argued that the experiences of change and continuity, as exemplified by following a melody, cannot be accounted for in terms of the flat temporal dynamics that characterise representational states. Am I taking Clark to deny this kind of temporal phenomenology? While I'm not aware of Clark addressing this issue directly, there is no reason to interpret him as endorsing a 'cinematic snapshot' (Busch & VanRullen, 2014) view of perceptual experience, in which the contents of perception are restricted to those that can be non-dynamically represented (e.g. spatial properties, tones, colours). Notice that, in the passage quoted above, Hurley refers to a 'snapshot' that *we* take of the perceptual process (i.e. 'from the outside'), specifying its content at that time. This does not mean that, for the subject, this experience is as of a non-temporally extended snapshot. Rather, the 'snapshot' specifies a personal-level target for subpersonal explanation. She also offers an analysis (pp. 27–30; see also Grush 2006) of the temptation to make this kind of mistake as a confusion between representational vehicles and represented contents. The point is that a representational state needn't share the temporal dynamics of the events it represents to count as representing (or failing to represent) them. Recent PP accounts of temporal phenomenology (Hohwy et al. 2016; Wiese, 2017b) illustrate this, expanding on Grush's (2005) trajectory estimation model. Here contents are presented as temporally ordered within a 'specious present', a moving window of subjective time, but this phenomenology is realised by discrete representational states of the

⁷ This is a simplification, since the full subpersonal story would also have to account for dynamic processes of retrieval of this represented content. Meanwhile, it's plausible that most dispositional knowledge is explained non-isomorphically, in terms of readily available inferences from tokened representations.

⁸ Thanks to an anonymous reviewer for encouraging me to be clearer on this point.

mechanisms underlying conscious experience.⁹ Arstila (2018) meanwhile proposes a ‘dynamic snapshot’ account, in which temporal phenomenology is distinguished from the experience of succession; succession and temporality are (atomically) represented by distinct mechanisms, and so both are realised without the specious present. We needn’t pin Clark down to any one of these options. The point I take from Hurley’s discussion is not about whether certain kinds of temporal phenomenology can be realised by internal representations, but rather what we can say regarding how personal-level phenomena constituted exclusively by such representational content occupy time. This is what matters for specifying the targets of subpersonal explanations of perceptual experience.

If the target for vertical explanation of perception is constituted by successive atomic percepts (whether these are subjectively manifested as specious presents, cinematic, or dynamic snapshots), this guarantees that any moment of perceptual experience will be realised by a temporally discrete representational state in the brain. The next two sections outline how, when combined with the argument presented in Clark (2018), this position becomes unstable. The only way that Clark can exclude both *Indeterminate* and non-*Univocal* content from perceptual experience involves a conception of action that leads to thoroughgoing internalism about the mental.

6 Temporal atomism and basic actions

How can bodily action generate a systematic constraint for unitary-coherent *atomic* perceptual content? A natural way of interpreting Clark’s action-oriented proposal (in line with the earlier ‘action space’ view put forward in Ward et al. 2011) is that ordinary, second-order action intentions play this role. This suggests an intuitive answer to the *Univocal Contents* worry, since effective personal-action constitutively depends on a semantically unified intention. As Clark comments, while I might be only 60% certain that I want to go to a party, “I cannot actually just go 60% to the party, but must end up either going or not going” (2018, p. 81), and it makes sense that whatever underlies my capacities for choice and action-planning should respect this very salient constraint. For example, the perceptual system of an athlete trying to catch a fast-moving ball might not be able to arrive at an unambiguous, precise estimate of its location and trajectory, but unless a single interpretation is allowed to guide action, she will have no chance of catching it. Here imposing more unitary-coherence on the data than is sensorially available looks like the best way to deliver on personal-level, whole-system goals.

But this approach looks less promising for the other half of Clark’s puzzle. Intentions to act set goals, and personal-level goals might exert a system-wide effect, but they do this via a mechanism of means-end reasoning that remains open to the *Indeterminate Content* that ‘action’ is invoked to exclude. Simplistically, if the

⁹ This tricky distinction is summed up nicely by Wiese (2017b), who points out that, on this kind of atomistic account “an experience of a succession has *synchronous* parts which can be mapped onto the different elements of the succession. The parts of the succession are not experienced as simultaneous (although they are simultaneously experienced) because the different parts of the experience do not all represent their targets in the same way.” (p. 2).

perceptual experience of a subject, S, depends on S grasping an action space specifying a possible response, ψ , then for many plausible ways of understanding S's ψ -ing, it is unclear why representations underlying the capacity to ψ won't themselves have indeterminate content. A range of different bodily actions can typically satisfy an intention to ψ ; S might ϕ in order to ψ , but equally she might χ . Here ' ϕ ' and ' χ ' stand for determinates of ψ ; either of them could be a means to achieving the end of ψ -ing. Plausibly then, sensorimotor representations underlying S's grip on an action space have to remain open to all the different ways S can ψ .¹⁰ But in this case, we do not yet have a reason to expect that one determinate interpretation fixes perceptual experience, rather than a broadly coherent range of values. To avoid this kind of agential indeterminacy, such complex, multiply-determinable actions must be excluded from the action space underlying momentary experience. Clark (2015) seem to be thinking along these lines when he comments that unitary-determinate perception makes adaptive sense since "we can only perform one action at a time, choosing to grasp the pen for writing or for throwing, but not both at once" (p. 188). Maybe simple bodily acts like grasping, rather than temporally-extended performances like writing or throwing, characterise the action space available to the perceiver on a moment-by-moment basis.

Work in the Davidsonian tradition of action theory suggests a natural way to develop this thought. The problem for personal-level intention as a constraint on perception has to do with the multiple-realizability of means to the intentionally specified ends. But as Davidson (1980a) points out, this is not a case of genuine vagueness or indeterminacy on the side of action, but rather on the side of the intention¹¹ to act:

[T]he event whose occurrence makes 'I turned on the light' true cannot be called the object, however intentional, of 'I wanted to turn on the light'. If I turned on the light, then I must have done it at a precise moment, in a particular way—every detail is fixed. But it makes no sense to demand that my *want* be directed to an action performed at any one moment or done in some unique manner. Any one of an indefinitely large number of actions would satisfy the want and can be considered equally eligible as its object. (p. 6; my emphasis)

This passage restates the problem for the contents (or intentional objects) of ordinary intentions: they seem to be indeterminate between many candidate bodily actions. But this apparent problem can be solved if we understand the relation between an intention and its worldly realisation causally. *Turning on the light* can describe an intention, or an action, or a series of bodily and worldly events. What it is for such events to constitute this intentional action is for them to be caused by a mental state with satisfaction conditions that are met when the light is turned off. In this case, they are said to be intentional under that description, even though this does

¹⁰ Cf. Ward et al. (2011) "An action space [...] is to be understood not as a fine-grained matrix of possibilities for bodily movement, but as a matrix of possibilities for pursuing and accomplishing one's intentional actions, goals and projects." (p. 383) This account looks open to a version Clark's (2018) worry: why should this matrix of possibilities ground a determinate percept?

¹¹ 'Intention' here refers to what Davidson (1980a) describes as a generalised 'pro-attitude' towards some worldly state of affairs combined with pragmatic beliefs about how to bring it about.

not uniquely pick out the bodily action that is performed. Linguistically described, ordinary actions can be analysed as a means to an end (i.e. she *turned on the light* in order to illuminate the room), or as an end achieved by some other, more fundamental action (she flicked the switch to *turn on the light*). Following this line of analysis brings us to an action that is a means to an end (an event, e.g. the light turning on) without itself being the end of some other intentional action. It is directly caused by the agent's intervention in the environment, rather than occurring further along a causal chain that she is the author of:

Not every event we attribute to an agent [i.e. ordinary actions like turning on a light] can be explained as caused by another event of which he is the agent: some acts must be primitive in the sense that they cannot be analysed in terms of their causal relations to the same agent. (Davidson, 1980b, p. 49)

These primitive, or basic,¹² actions are determinate bodily events that can be understood as intentional actions in virtue of their causal link to an agent's intentions (wants, reasons, etc.) and the role they play in satisfying them. This explains the link between potentially indeterminate ordinary agential intentions and a world of determinate objects and events. Intentional indeterminacy doesn't matter because this link is causal, not semantic. Importantly, basic actions are *not* causes of ordinary actions (though they do cause events that occur in descriptions of them). Flicking the switch and illuminating the room are one and the same action *de re*, (i.e. a particular movement of the body) though they fall under different descriptions. Since flicking the light switch is a way of illuminating the room, the intention to illuminate the room is realised by this bodily movement. Similarly, writing with or throwing a pen both depend in part on a more basic act of grasping it. Under normal conditions, where such bodily movements realise an agent's intentions and are caused by them, they are intentional actions. But since flicking the switch (moving the body so as to do this) and illuminating the room are in fact the same action, the action-constituting event can be precisely located in just these bodily movements as caused by the relevant intention.

Another motivation for basic actions is the need to ensure the causal chains in virtue of which we attribute agency terminate in the relevant agent. For the causal account to work, we have to pick out something that the agent does to cause all the further events that satisfy her intentions. In order for an agent to do anything, she must *basically* do something. Described subpersonally, this will involve a fine-grained account of internal causal chains composed of events that we do not attribute to the agent (e.g. low-level motor commands), but at the personal level the basic action is the point of causal interface between the person who acts and her environment. At the most fundamental (personal) level of analysis, then, there are only basic actions. For cases like switching off a light, this is already clear. More temporally extended ordinary actions, meanwhile, are analysed as concatenations of several basic actions, which are the parts of its whole performance: gasping the pen, pressing it to paper, forming the letter 'a', etc. (Davidson, 1980b; see also Amaya, 2017; Lavin, 2013). In every case, basic actions are expressed in determinate bodily

¹² Danto (1965) introduces the term, which has become standard in the literature (Amaya, 2017).

movements that, “at a precise moment, in a particular way” (Davidson, 1980a, p. 6), are simply performed.

This suggests a way that appealing to a constitutive link between perception and a *basic* action space could be promising for Clark (2018). Because there is only ever one way to perform a basic action, we can avoid the gap between intention and execution that rules out putative solutions to the *Determinate Content* problem that rely on ordinary, non-basic action intentions. This goes beyond the action space account (Ward et al. 2011), in that it introduces a constitutive link between first order dispositions to act (i.e. those tied to direct performance of actions) and perceptual experience. These needn’t have subjectively-accessible descriptions matching the precise action to be performed, but they do need to be appropriately poised (via intermediate ‘subpersonal’ motor processing) to bring determinate basic actions about. What underlies an agent’s unitary-determinate perceptual grip on the world, then, is not simply having perceptually-grounded intentions to act (‘I will that I flick the light switch’ or, reflexively, ‘flick the light switch now!’), but depends on how these are integrated with motor capacities to determinately realise them.

But since ‘basic actions’ don’t come up in Clark’s (2018) discussion, why should we think they enter into the complete view, which I claim is unstable? The account given in Clark (2020) of the relation between action and its subpersonal realisation according to PP helps to fill in this gap. In this and other work (Friston et al. 2010; Wiese, 2017c), agency is presented as emerging from the same machinery of prediction error minimisation as perception, with precision estimation again playing a central role:

Bringing about action [...] requires attenuating (assigning low precision to) the sensory information currently indexing the actual disposition of the body, so as to enable precise proprioceptive predictions (corresponding to some desired trajectory) to prevail [...] Intentional action thus depends upon a delicate balance that combines precise proprioceptive predictions with attenuated information concerning current bodily states (Clark, 2020, p. 4)

Actions occur when proprioceptive prediction errors elicit an active response: instead of higher-level representation states being updated (‘passive’ or ‘perceptual’ inference) the system actively transforms the incoming signal by realigning the body to suppress proprioceptive error. It is realised by a kind of ‘systematic misrepresentation’ (Wiese, 2017c) of the body’s position in space, with corresponding bodily movements actualising these misrepresentations. This is the subpersonal story, but how is it related to the intentions for action that we ascribe in personal-level psychological descriptions? The answer to this question comes from examining the contents of representational states across the processing hierarchy:

At the bottom of the stack are simple peripheral reflexes (for example, involving proprioceptive predictions that determine set points for stretch receptors that then automatically translate into movements). Towards the top lie more intuitively agency-reflecting (indeed, agency-constituting) predictions, such as the prediction that I will go to see such-and-such a movie tonight. (Clark, 2020, p.12)

Here, as in the perceptual case, Clark gives priority to high-level states that exert a global influence on the predictive system. It is these states that constitute personal-level agency, and as such “are self-fulfilling prophecies, capable of helping to bring about the very states of affairs that they describe” (ibid.). Once again, the commitment to isomorphism between the personal-level phenomena and states appealed to in subpersonal explanation conditions how we should interpret this passage. High-level predictions constitute personal level agency because they just are the content of the agent’s intentions. If high-level active representation states on one hand constitute personal-level agential intentions and, on the other, are causes of worldly states of affairs (proximally, the agent’s movements of her body), then the view of action that emerges looks very similar to the causal account just described. The step from this to the basic action view follows from the same reasoning, the need to pick out an event that determinately satisfies intentionally specified goals and can be unambiguously attributed to the whole acting subject. Proprioceptive predictions corresponding to the performance of these basic actions are how action is realised on a moment-by-moment basis. Though causally linked to (sometimes indeterminate) personal-level intentions, contents at this level should be determinate and univocal, since they specify a particular way the body should move and must be assigned high precision to bring this movement about (Friston et al. 2010). Given a constitutive link between perceptual experience and action, this provides a reason to expect percepts to represent the environment as a unitary-coherent field for direct bodily intervention: a basic action space.

Similar to temporal atomism about perceptual experience, the basic action view guarantees a determinate answer at any given moment to the question ‘what is S doing?’, since (if S is acting at all) we can always specify a basic action by reference to such proprioceptive predictions. This isn’t accidental. Unitary-determinate atomic percepts require determinate similarly atomic (i.e. basic) actions if Clark’s argument is to succeed. At first, it might look like this is in keeping with a wider embodied and action-oriented account. Shouldn’t the ongoing mutual interaction of motor and sensory processing be reflected at the personal level in a tight fit between the ongoing performance of actions and their perceptual guidance? Before returning to this issue in Section 8, I want to point out a serious tension between this solution and Clark’s commitment to extended-embodied mental states (and by extension any more radically enactivist proposals).

7 Brainbound basic actors

We gain a useful perspective on claims about embodiment and cognitive extension if we cast them as explanations of how personal-level phenomena are realised by the interaction of neural and non-neural resources at the subpersonal level. Clark motivates claims about extended cognition by appealing to a ‘common-sense functionalism’ (Clark, 2008) that characterises personal-level mental states by reference to their coarse-grained functional profile. Pursuing a vertical analysis of these states reveals a heterogeneous coalition at work, with the brain supplementing its processing resources with external artefacts, embodied tricks, and distributed cultural practices. Together

these form ‘temporary problem-solving wholes’ (ibid.), soft-assembled by the cognitive system in response to transient information-processing demands. Given an innocuous¹³ ‘parity principle’ (Clark & Chalmers, 1998) between neural and non-neural information processing, these additional resources enter into subpersonal explanations of the target mental phenomena on equal terms with neural mechanisms, and so we have *prima facie* reason to see them as constitutive parts of the cognitive agent. For Clark, internal mechanisms described by PP are a crucial part of subpersonal constitutive explanations of extended minds, in which “neural sub-assemblies form and dissolve in ways determined by changing estimations of relative uncertainty” so that they “recruit and are recruited by shifting webs of bodily and extra-bodily structure and resources.” (2015, p. 295). That is, according to Clark’s highly influential presentation of the PP framework, predictive minds are extended minds.

Taking Clark and Chalmers’ (1998) famous example, Otto’s notebook functions as an external memory in virtue of the way he incorporates it into his problem-solving routines. As a consequence, part of what constitutes his memory is now realised ‘outside the head’. More generally:

P is an embodied-extended personal-level mental state iff for subprocesses, p^1 - p^n , which constitutively explain P, some p is a process¹⁴ external to the brain.

Correspondingly, cognitive agents are embodied-extended to the extent that their personal-level states fit this definition. But this line of reasoning is only available if we identify the right personal-level explanatory target for vertical explanation.

If the personal-level phenomenon is a basic action, however, we can at every point draw a principled barrier around the biological core agent. This has been noticed by action theorists outside of the literature on active externalism and enactivism (see, e.g. Hornsby, 1980; Lavin, 2013; Steward, 2000). Lavin (2013) describes basic actions as the expression of a supposed metaphysical division between a temporally extended bodily and worldly event and the fundamentally atomic exercise of personal-level agency. Thus, given the view that:

physical action consists of a mere event and a condition of mind joined (in the right way) by the bond of causality [...] such events, including the movements of one’s body when one intentionally moves it, are thought to be constitutively independent of the subject’s rational capacities. Basic action is a necessary counter-measure, a sort of metaphysical containment wall needed to preserve the separate jurisdictions of the mind of the acting subject and what merely happens. (p. 274)

What I want to claim is that the ‘metaphysical containment wall’ between the rational subject and her body discussed by Lavin finds a parallel in the boundary between the organism or cognitive agent and its environment proposed by theorists like Hohwy (2013, 2016) (and argued *against* by e.g. Clark, 2017). In both cases, a

¹³ Of course, for many theorists (e.g. Sprevak 2009), the parity principle is not innocuous *at all* but staying within the broadly functionalist outlook Clark assumes, critics of active externalism tend to accept the parity principle (in *principle*), but point out putatively important functional differences between neural and extended processes (e.g. Adams & Aizawa, 2001; Rupert, 2004).

¹⁴ ‘Process’ here is used in the broad sense in which a series of states (e.g. representation states) and transitions between them are also processes.

privileged space of specifically mental states and processes is carved off from the merely bodily or environmental context it finds itself in.

The basic action view is an account of agency at the personal level of description. While the analysis of ordinary personal-level bodily actions (going to a party, writing a letter, illuminating the living room, etc.) as constituted by basic action(s) is not paradigmatically a ‘folk psychological’ explanation, neither is it a functional decomposition that identifies mechanistic sub-agencies. Instead, it provides a horizontal analysis of ordinary actions (and consequent events) in terms of internal states of the acting subject, the bodily movements that these produce, and their worldly consequences. This implies a form of temporal atomism about agency, according to which ordinary actions are constituted by strings of determinate bodily events with right kinds of mental causes. Taking any one of these as the target for vertical explanation, however, effectively circumscribes the available realisers of these mental causes to internal neural processes.

To see why this is the case, think about Davidson’s (1980a) light switch example. An agent intends to illuminate the room and moves her body so as to do this. The (basic) action consists in the movement of her body as caused by the intention, which sets off the chain of events that brings the goal of her (non-basic) intention about. Moment-by-moment, the agent only does one thing, and these action-constituting events come at the beginning of a worldly causal chain, located at the point of interface between agent and environment. Some basic action theorists place this point within an agent’s own body (e.g. Hornsby, 1980), while others move it ‘further out’ to the active body’s impression on the external environment.¹⁵ Deciding this issue turns partly on questions about the fineness of grain used in individuating events in the causal chains that constitute actions, as well as questions about the nature of the mental causes of bodily actions. (for an extended discussion and review, see Amaya, 2017). Whatever way we resolved these, the form of this type of analysis remains the same. Actions are composites of an *inner* mental cause and an *outer* bodily effect. As has been pointed out (by, e.g. O’Brien, 2007, ch. 8), this form is shared by causal accounts of action generally, and it remains controversial to what extent this is a problem for such accounts (Davis 2010). Even more generally, Drayson (2018) suggests that there is a tension between this kind of compositional analysis of mental states and embodied-extended claims. But it is not clear at this point that there are no principled ways to set outer, non-anatomically defined boundaries of the mind consistent with compositional accounts (see Clark, 2017; Rowlands, 2009). At least, nothing in the current argument turns on this. What is clear, however, is that if all actions are basic actions, their inner mental causes turn out to be brainbound by definition.

If we only think about cases like switching on a light, this doesn’t look like a problem. Here there is no pull towards thinking of the mental cause of an action in anything other than neural terms. But now consider temporally extended actions, in

¹⁵ There are textual reasons to read Davidson (1980a, 1980b) as taking the latter view: since agents and their doings cannot be located more precisely than their bodies, a basic action attributed to an agent cannot be located *inside* that agent. See also Steward (2000). In later work, Hornsby more fundamentally revises her commitment to the Davidsonian picture. See Section 8 below and Hornsby (2011, 2013).

which worldly, bodily, and neural events and structures mutually interact. Davidson (1980b) gives a straightforward example:

When I tie my shoelaces, there is on one hand the movement of my fingers, and on the other hand the movement of the laces. But is it possible to separate these events by calling the first, alone, my action? (p. 51)

Davidson answers ‘yes’ to his own question,¹⁶ even while noting that the bodily movements involved in shoe-tying might not be available to the agent without feedback from the shoes and laces themselves, and are not (as, indeed, with the light switch example) directly the object of the relevant intention. This is because they are still intentional under some description, and (once the subpersonal implementation details are filled in) are directly caused by a content-bearing mental state matching this description. But considerations that tend usually to support extended-embodied claims point in the other direction. When we subpersonally explain a token instance of shoe-tying, we find exactly the kind of soft-assembled coalition that is supposed to be the hallmark of extended processing. Motor representations offload information processing onto environmental structure via fine-grained sensorimotor feedback loops that integrate worldly events with proprioceptive and sensory processing. Notably, the structure of the task domain itself has been designed to facilitate this informational load sharing (discussed as ‘cognitive niche construction’ in Clark, 2015, pp. 275–279). From this perspective, the fact that (under normal circumstances), internal, neural resources are insufficient for producing the bodily movements (i.e. the chain of basic actions) begins to look like a reason to understand their *mental* causes as constituted, in part, by this culturally-situated information processing structure and affiliated processes of skilled bodily engagement. If we want to block off this move, we have to appeal to a thoroughly atomistic version of the basic action view, narrowly construing the mental causes of bodily movements as fairly proximal neural outputs.

What if we do this? Pursuant to a *moderate* enactivist program, we might think that attributing basic capacities for action and perception to a biological core agent still need not threaten extended-embodied claims about more sophisticated cognitive capacities. But notice that this view doesn’t only restrict the subpersonal correlates of intentional action. This is because intentional bodily movements are generally causally implicated in putatively world-involving mental states and processes. These ‘coupling’ (Clark, 2008) actions partially realise the process by which external resources are recruited by extended cognitive agents to constitute their personal-level mental states. Think about Clark and Chalmer’s (1998) original example. In order to consult his notebook, Otto must act on his notebook (i.e. his external memory store). Clark consistently presents coupling actions like this as a kind of externally realised mental action, like intentionally

¹⁶ The question here being *is this action just a bodily movement with the right kind of mental cause?* There is a related question about whether descriptions under which basic actions are intentional can refer to their extra-bodily effects. On my reading, Davidson thinks they can, but that the event that realises (and is caused by) the intention is the bodily movement only (Davidson, 1980b, pp. 51–52). The contents of personal-level intentions needn’t, and generally don’t, specify the determinate bodily actions that satisfy them. Thanks to Doug Lavin and an anonymous reviewer for pressing me on this point.

bringing something to mind or paying attention to the premises of an argument, rather than intrinsically non-reflexively available processes found at deeper levels of subpersonal explanation. So, while the intention with which Otto acts is to bring to mind whatever information he needs, the bodily movements by which he realises this are intentional actions (albeit under a fairly non-standard description). But if such actions fix the outer limit of Otto's mind, then his interaction with the notebook and the symbolically-represented information contained in it turns out to be external not only to his body but also to his mental processing. This information is allowed to re-enter the mind through perception, but the extended process is not constitutive of Otto's mental life in any theoretically interesting sense (cf. Adams & Aziwa, 2001). More generally, adopting a basic action view reaffirms the common-sense isomorphism between sensory input and motor output and the boundaries of the mind at the personal level.

This is a fundamental result that goes beyond simple cases like Otto and his notebook. Clark puts a particular notion of functional soft assembly at the centre of his embodied-extended account, but any claims that the bodily and worldly consequences of actions are constitutive of personal-level mental phenomena are vulnerable to the objection just outlined. If actions are basic actions, then functional analyses of personal-level phenomena are screened off from external elements, isolating all mentality *per se* from the constitutive contributions of an organism's body and environment.

8 Extended perceptual episodes and non-isomorphic vertical explanation

So, answering Clark's worry about *Determinate Content* by appealing to basic actions comes at the cost of the central claims of enactivist PP. Does this mean, if we are enactivists, that we must revise our account of perceptual phenomenology? Since, as argued in Section 6, ordinary action intentions plausibly constrain perceptual experience to have univocal content, this might seem like an attractive option: our perceptual worlds exhibit unity and a degree of clarity due to our need to form coherent plans of action, but a fringe of perceptual indeterminacy pervades experience from moment to moment (for a defence of this view, see Madary, 2016). Independently of theoretical commitments to enactivism and PP, recent empirical work on ensemble perception, perceptual gist, and other phenomena has also been presented as evidence for the indeterminacy of perceptual experience (Cohen et al. 2016; Stazicker, 2011). These claims need to be handled carefully. While this work reveals perceptual indeterminacy at certain temporal scales and under certain conditions, such phenomena diverge from a background phenomenology that Clark's (2018) question more fundamentally casts in doubt. Assuming temporal atomism about perceptual experience, widely shared intuitions that experience is unitary-determinate *at all* begin to look all the more puzzling. Does experience constantly deceive us that we see more, and more clearly, than we really do? Or do naïve intuitions about perception alienate us from the facts of our own perceptual phenomenology? There is a third option. We can avoid this ramifying puzzlement by reframing

the original question. Instead of asking how perceptual representations are constrained to have determinate content, we could ask how probabilistic perceptual processing achieves phenomenal *transparency* (Martin, 2002; see also Ward, 2016b).

Transparent processing. Perceptual experience is transparent to its objects: we seem (only) to perceive the objects and events we are in contact with, not a range of nearby possible states of affairs. How does a probabilistic perceptual system achieve this?

This reformulates Clark's puzzles to avoid the assumption of temporal atomism about perceptual experience. For this to be a credible move, we must have reasons to doubt temporal atomism about personal-level perception. This section suggests some, and sketches a more processual ontology of perception and action, which makes sense of Clark's phenomenological intuitions in the context of a more radically enactivist¹⁷ account of perceptual experience.

Something that makes temporal atomism about perceptual experience seem so plausible is the tendency to think that because both experiences and brain states have contents, a naturalising account of the former involves finding identical representational content realised at some level of the cognitive system. Since representational states are temporally atomistic, it is natural to understand the personal level phenomenon in this way too. But this is circular: temporal atomism defines the phenomenon in terms of the proposed explanandum. In the passage quoted at the start of this paper, Clark correctly points out that the only access the brain has to the external environment is mediated by sensory input; "it must discover information about the likely causes of impinging signals without any form of direct access to their source" (2013, p. 3). But as he and other enactivists also stress, the person who perceives is not (just) her brain. If we instead take our personal-level explanatory target be a temporally unfolding process, we can understand the phenomenal transparency of perceptual experience in terms of the densely interactive access that bodily engagement affords the perceiving agent to the world around her over time.

Psychophysics tells us that agents lack rich, veridical moment-by-moment access to their perceptual environment. On the atomistic view, this either means that experience systematically misrepresents this fact, or our theoretical intuitions disguise the indeterminacy of experience. But the phenomenology that gives rise to Clark's puzzle doesn't demand that perceptual experience be constituted by a series of unitary-determinate percepts; rather it requires that perception be as of a world of determinate, coherently-ordered objects and events. Explaining this needn't presuppose an isomorphic relation between perceptual experience and occurrent subpersonal representations (either illusorily unitary-determinate percepts or genuinely indeterminate ones). Indeed, the expectation for such a one-one mapping to usefully

¹⁷ How radical is this account? As we'll see, the enactivist view I sketch still appeals to subpersonal content-bearing states, but denies that these alone account for temporally-unfolding episodes of perceptual experience, since explanation of these is essentially embodied and world-involving. So this isn't radical enactivism in Hutto and Myin's (2012) sense, eschewing all representation talk. Whether Clark (or any other theorist) must take this further step will turn on independent arguments about the use of the representation construct in cognitive science. And, as far as the argument of this paper goes, Clark is just as entitled to take a step in the other direction and give up moderate enactivism.

capture the relation between experience and its underlying predictive processing is weakened when we consider that across the predictive system, representations of environmental structure are exploited in parallel by multiple subsystems at very different timescales, often encompassing feedback loops across brain, body, and environment (Allen & Friston, 2018), entailing diverse consequences for all aspects of an agent's pragmatic engagement with the world.

At the larger temporal scales on which this engagement plays out, the character of episodes of perceptual experience is still partly explained by changes in probabilistic representational states. But the details of the perceiver's bodily interaction with the real world objects of her perception that these guide and respond to also enters the picture. So conceived, perceptual experiences are characterised by ongoing stimulus control (Phillips 2017).¹⁸ This accounts for *Transparent Processing*, since it is the interaction between internal representations and actual distal objects (and not nearby possible states of affairs) on which the phenomenal character of episodes of perceptual experience depend. This transparency is realised by environmentally-extended feedback loops, and, since these unfold over time, accounting for this is essentially linked to denying temporal atomism. Taking temporally-extended perceptual episodes as the target of subpersonal explanation therefore reveals determinate perceptual experience to be a world-involving process. Conversely, if we narrow the temporal window of this process too much, we lose our grip on the transparency of perception.

A constitutive link between agential intentions and perception remains central to this story. The need for semantically unified action-planning still mandates a system-wide constraint towards univocal representation, with phenomenological consequences for the unity of perceptual experience. How action is understood as a personal-level phenomenon, however, is no longer construed according to the basic action model. Like the episodes of perceptual experience with which they are constitutively interdependent, actions are cast as temporally extended procedures that are not non-arbitrarily analysable into concatenations of atomic components. Doing so, as suggested in recent commentary (Hornsby, 2011; Lavin, 2015) on Anscombe's *Intention* (1957), would mean rejecting the Davidsonian account and recovering a view of ordinary agency as a process constituting the mind's activity *in* the world, rather than a consequence of basic actions *on* an external environment. Given the tensions between the basic action view and even moderate forms of enactivism, theorists who are sympathetic to these approaches have good reasons to understand personal-level perception and agency in this way.

Nothing about these alternatives contests the mechanistic story as it occurs in Clark (2015, 2018) or psychophysical results suggesting indeterminacy in momentary perception (Cohen et al. 2016; Stazicker, 2011). Instead, they revise how the personal-level phenomenon is characterised, and so what kind of explanatory connections can be drawn between it and underlying probabilistic subpersonal representations.

¹⁸ Phillips (2017) further suggests that pluralist stimulus-control criteria could mediate the perception/cognition distinction in a way that does not rely on differences in representation format, content, or sensitivity to top-down effects. Arguments along these lines might make sense of current controversies about the boundaries of cognition thrown up by PP (see also Drayson 2017; Green 2020), with the fairly strong environmental criteria outlined here serving to typify episodes of perceptual *experience*. Thanks to an anonymous reviewer for suggesting this framing.

Undergoing this shift in perspective means taking the temporal dynamics of whole-agent interaction with the environment as a better guide to the metaphysics of the personal-level phenomenon of perceptual experience than a model based exclusively on transitions between representational states. The claim is that these, while crucial, only make up part of a comprehensive subpersonal explanation. The intuition that the content of such subpersonal representational states is likely to be radically different from what introspection suggests about perceptual experience is what motivates Clark's original worry, and his commitment to analysing perceptual experience via isomorphic vertical explanation generates a theoretical drive to collapse this problematic indeterminacy at some level of the predictive system. Loosening up this theoretical constraint on the selection of targets for explanation offers a way of doing justice to the way perceptual experience seems to us while dissolving Clark's worry about *Indeterminate Content*. At the same time, a non-isomorphic approach to vertical explanation of perceptual experience allows us to see it as diachronically constituted (Kirchhoff, 2015; Kirchhoff & Kiverstein, 2019, ch. 6) in part by bodily and environmentally-extended processes. The same considerations that tell against Clark's attempt to combine internalism about the realisers of perceptual experience with externalist claims for other mental phenomena point the way to a more thoroughgoing enactivist view.

9 Conclusions

It goes beyond the above discussion to give an account of agency without basic actions or provide independent reasons to prefer the process-oriented methodology just outlined. Nonetheless, I hope to have shown that Clark's worries about the 'Bayesian blur' are partly based on his optional commitment to a view of personal-level perception and action that is in tension with even moderate versions of enactivist PP. While there are good reasons to think that a tight fit between perception and action grounds the unity of perceptual experiences, recognising this needn't motivate a view of experiential content as a series of unitary-determinate atomic percepts. I have suggested that non-linear and distributed processing within a predictive cognitive system that drives responses at many different temporal scales might end up supporting an analysis of perceptual experience in terms of structured environmentally and temporally-extended processes, rather than an internally-realised series of state-to-state transitions. This kind of view should be justified by the theoretical work it can be put to, of which one promising avenue is radically enactivist PP. That said, nothing in the above discussion rules out the internalist picture, but rather suggests that the compromise that Clark (2018) seeks to strike is untenable. If the puzzle he raises needs to be solved, answers are constrained to be either strictly internalist or radically enactivist.

Acknowledgments I would like to thank audiences at the London-Warwick Mind Forum, Andy Clark and his Critics at the University of Edinburgh, and the 4th Avant Conference: Fragmentation of Cognition at the University of Porto, where I had the chance to present earlier versions of this paper. Additionally, particular thanks are due to Wiktor Rorot, Lucy O'Brien, Pablo Fernando Velasco, Tristan Kreetz and Douglas Lavin for crucial comments and suggestions at various points in its production. During this time, I have been supported by funding from the London Arts and Humanities Partnership (LAHP) of the Arts and Humanities Research Council (AHRC).

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Adams, F., & Aizawa, K. (2001). The bounds of cognition. *Philosophical Psychology*, 14(1), 43–64.
- Aitchison, L., & Lengyel, M. (2017). With or without you: predictive coding and Bayesian inference in the brain. *Current Opinion in Neurobiology*, 46, 219–227.
- Allen, M., & Friston, K. J. (2018). From cognitivism to autopoiesis: towards a computational framework for the embodied mind. *Synthese*, 195(6), 2459–2482.
- Amaya, S. (2017). Basic actions reloaded. *Philosophy. Compass*, 12(9), e12435.
- Anscombe, G. E. M. (1957). *Intention*. Basil Blackwell.
- Arstila, V. (2018). Temporal experiences without the specious present. *Australasian Journal of Philosophy*, 96(2), 287–302.
- Bastos, A. M., Usrey, W. M., Adams, R. A., Mangun, G. R., Fries, P., & Friston, K. J. (2012). Canonical microcircuits for predictive coding. *Neuron*, 76(4), 695–711.
- Block, N. (2018). If perception is probabilistic, why does it not seem probabilistic? *Philosophical Transactions of the Royal Society B: Biological Sciences*, 373(1755), 20170341.
- Bowers, J. S., & Davis, C. J. (2012). Bayesian just-so stories in psychology and neuroscience. *Psychological Bulletin*, 138(3), 389.
- Breitmeyer, B. G. (2015). Psychophysical “blinding” methods reveal a functional hierarchy of unconscious visual processing. *Consciousness and Cognition*, 35, 234–250.
- Breitmeyer, B. G., & Ogmen, H. (2000). Recent models and findings in visual backward masking: A comparison, review, and update. *Perception & Psychophysics*, 62(8), 1572–1595.
- Bruineberg, J., & Rietveld, E. (2014). Self-organization, free energy minimization, and optimal grip on a field of affordances. *Frontiers in Human Neuroscience*, 8, 599.
- Bruineberg, J., Kiverstein, J., & Rietveld, E. (2018). The anticipating brain is not a scientist: the free-energy principle from an ecological-enactive perspective. *Synthese*, 195(6), 2417–2444.
- Busch, N., & VanRullen, R. (2014). Is visual perception like a continuous flow or a series of snapshots?. In V. Arstila & D. Lloyd (Eds.), *Subjective time: The philosophy, psychology, and neuroscience of temporality* (pp. 161–178). MIT Press.
- Clark, A. (1998). *Being there: Putting brain, body, and world together again*. MIT press.
- Clark, A. (2008). *Supersizing the mind: Embodiment, action, and cognitive extension*. Oxford University Press.
- Clark, A. (2012). Dreaming the whole cat: Generative models, predictive processing, and the enactivist conception of perceptual experience. *Mind*, 121(483), 753–771.
- Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, 36(3), 181–204.
- Clark, A. (2015). *Surfing uncertainty: Prediction, action, and the embodied mind*. Oxford University Press.
- Clark, A. (2017). How to knit your own Markov blanket. In T. Metzinger & W. Wiese (Eds.). *Philosophy and Predictive Processing: 1. MIND Group*
- Clark, A. (2018). Beyond the “Bayesian Blur”: Predictive Processing and the Nature of Subjective Experience. *Journal of Consciousness Studies*, 25(3–4), 71–87.
- Clark, A. (2019). Consciousness as generative entanglement. *The Journal of Philosophy*, 116(12), 645–662.

- Clark, A. (2020). Beyond desire? Agency, choice, and the predictive mind. *Australasian Journal of Philosophy*, 98(1), 1–15.
- Clark, A., & Chalmers, D. (1998). The extended mind. *Analysis*, 58(1), 7–19.
- Clowes, R. W., & Chrisley, R. (2012). Virtualist representation. *International Journal of Machine Consciousness*, 4(02), 503–522.
- Cohen, M. A., Dennett, D. C., & Kanwisher, N. (2016). What is the bandwidth of perceptual experience? *Trends in cognitive sciences*, 20(5), 324–335.
- Danto, A. C. (1965). Basic actions. *American Philosophical Quarterly*, 2(2), 141–148.
- Davidson, D. (1980a). *Actions, Reasons, and Causes*. Clarendon.
- Davidson, D. (1980b). *Agency*. Clarendon.
- Davis, W. A. (2010). The causal theory of action. In T. O'Connor & C. Sandis (Eds.), *A companion to the philosophy of action*, (pp. 32–39). Wiley-Blackwell.
- Dehaene, S., & Changeux, J. P. (2011). Experimental and theoretical approaches to conscious processing. *Neuron*, 70(2), 200–227.
- Dennett, D. C. (1969). *Content and Consciousness*. Routledge and Kegan Paul.
- Dennett, D. C. (1978). *Brainstorms*. MIT Press.
- Dennett, D. C. (1989). *The Intentional Stance*. MIT press.
- Dennett, D. C., & Kinsbourne, M. (1992). Time and the observer: The where and when of consciousness in the brain. *Behavioral and Brain Sciences*, 15(2), 183–201.
- Drayson, Z. (2012). The uses and abuses of the personal/subpersonal distinction. *Philosophical Perspectives*, 26(1), 1–18.
- Drayson, Z. (2017). Modularity and the predictive mind. In T. Metzinger & W. Wiese (Eds.), *Philosophy and Predictive Processing*: 1. MIND Group.
- Drayson, Z. (2018). Extended minds and prime mental conditions: probing the parallels. In Joseph Adam Carter, Andy Clark, Jesper Kallestrup, Spyridon Orestis Palermos & Duncan Pritchard (eds.), *Extended Epistemology*. Oxford University Press. pp. 147–161.
- Engel, A. K., Friston, K. J., & Kragic, D. (Eds.). (2015). *The pragmatic turn: Toward action-oriented views in cognitive science* (Vol. 18). MIT Press.
- Feldman, H., & Friston, K. (2010). Attention, uncertainty, and free-energy. *Frontiers in Human Neuroscience*, 4, 215.
- Friston, K. (2010). The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience*, 11(2), 127.
- Friston, K. J., Daunizeau, J., Kilner, J., & Kiebel, S. J. (2010). Action and behavior: a free-energy formulation. *Biological Cybernetics*, 102(3), 227–260.
- Gallagher, S., & Allen, M. (2018). Active inference, enactivism and the hermeneutics of social cognition. *Synthese*, 195(6), 2627–2648.
- Green, E. J. (2020). The perception-cognition border: A case for architectural division. *Philosophical Review*, 129(3), 323–393.
- Gross, S., & Flombaum, J. (2017). Does perceptual consciousness overflow cognitive access? The challenge from probabilistic, hierarchical processes. *Mind & Language*, 32(3), 358–391.
- Grush, R. (2005). Internal models and the construction of time: generalizing from estimation to estimation to address temporal features of perception, including temporal illusions. *Journal of Neural Engineering*, 2(3), S209–S218.
- Grush, R. (2006). How to, and how not to, bridge computational cognitive neuroscience and Husserlian phenomenology of time consciousness. *Synthese*, 153(3), 417–450.
- Hohwy, J. (2013). *The predictive mind*. Oxford University Press.
- Hohwy, J. (2016). The self-evidencing brain. *Noûs*, 50(2), 259–285.
- Hohwy, J., Paton, B., & Palmer, C. (2016). Distrusting the present. *Phenomenology and the Cognitive Sciences*, 15(3), 315–335.
- Hornsby, J. (1980). *Actions*. Routledge & Kegan Paul.
- Hornsby, J. (2011). Actions in Their Circumstances. In A. Ford, J. Hornsby, & F. Stoutland (Eds.), *Essays on Anscombe's intention*. (pp. 105–127). Harvard University Press.
- Hornsby, J. (2013). Basic activity. *Aristotelian Society Supplementary*, 87(1), 1–18.
- Huang, Y., & Rao, R. P. (2011). Predictive coding. *Wiley Interdisciplinary Reviews: Cognitive Science*, 2(5), 580–593.
- Hurley, S. L. (1998). *Consciousness in action*. Harvard University Press.
- Hutto, D. D., & Myin, E. (2012). *Radicalizing enactivism: Basic minds without content*. MIT press.

- Kelly, S. D. (2005). *The puzzle of temporal experience*. The philosophy and neuroscience movement, Cambridge University Press.
- Kirchhoff, M. D. (2015). Extended cognition & the causal-constitutive fallacy: In search for a diachronic and dynamical conception of constitution. *Philosophy and Phenomenological Research*, 90(2), 320–360.
- Kirchhoff, M. D., & Kiverstein, J. (2019). *Extended Consciousness and Predictive Processing: A Third Wave View*. Routledge.
- Koch, C. (2004). *The Quest for Consciousness: A Neurobiological Approach*. Roberts & Company.
- Lavin, D. (2013). Must there be basic action? *Noûs*, 47(2), 273–301.
- Lavin, D. (2015). Action as a form of temporal unity: on Anscombe's Intention. *Canadian Journal of Philosophy*, 45(5–6), 609–629.
- Lu, Y., Stafford, T., & Fox, C. (2016). Maximum saliency bias in binocular fusion. *Connection Science*, 28(3), 258–269.
- Lycan, W. G. (1987). *Consciousness*. MIT Press.
- Madary, M. (2016). *Visual phenomenology*. MIT Press.
- Martin, M. G. (2002). The transparency of experience. *Mind & Language*, 17(4), 376–425.
- Morrison, J. (2016). Perceptual Confidence. *Analytic Philosophy*, 57(1), 15–48.
- O'Brien, L. (2007). *Self-knowing agents*. Oxford University Press.
- O'Regan, J. K., & Noë, A. (2001). A sensorimotor account of vision and visual consciousness. *Behavioral and brain sciences*, 24(5), 939–973.
- Orlandi, N. (2016). Bayesian perception is ecological perception. *Philosophical Topics*, 44(2), 327–352.
- O'Shaughnessy, B. (2000). *Consciousness and the World*. Oxford University Press.
- Phillips, B. (2017). The shifting border between perception and cognition. *Noûs*, 53(2), 316–346.
- Rowlands, M. (2009). Enactivism and the extended mind. *Topoi*, 28(1), 53–62.
- Roth, M. (2015). I Am Large, I Contain Multitudes: The Personal, the Sub-personal, and the Extended. In C. Muñoz-Suárez & F. De Brigard (Eds.), *Content and Consciousness Revisited* (pp. 129–142). Springer.
- Rupert, R. D. (2004). Challenges to the hypothesis of extended cognition. *The Journal of Philosophy*, 101(8), 389–428.
- Soteriou, M. (2013). *The mind's construction: The ontology of mind and mental action*. Oxford University Press.
- Sprevak, M. (2009). Extended cognition and functionalism. *The Journal of Philosophy*, 106(9), 503–527.
- Stazicker, J. (2011). Attention, visual consciousness and indeterminacy. *Mind & Language*, 26(2), 156–184.
- Steward, H. (1997). *The ontology of mind: Events, processes, and states*. Oxford University Press.
- Steward, H. (2000). Do actions occur inside the body? *Mind & Society*, 1(2), 107–125.
- Steward, H. (2013). Processes, continuants, and individuals. *Mind*, 122(487), 781–812.
- Treisman, A. (1996). The binding problem. *Current Opinion in Neurobiology*, 6(2), 171–178.
- Ward, D. (2016a). Hurley's Transcendental Enactivism. *Journal of Consciousness Studies*, 23(5–6), 12–38.
- Ward, D. (2016b). Achieving transparency: An argument for enactivism. *Philosophy and Phenomenological Research*, 93(3), 650–680.
- Ward, D., Roberts, T., & Clark, A. (2011). Knowing what we can do: Actions, intentions, and the construction of phenomenal experience. *Synthese*, 181(3), 375–394.
- Ward, D., Silverman, D., & Villalobos, M. (2017). Introduction: The varieties of enactivism. *Topoi*, 36(3), 365–375.
- Wiese, W. (2017a). What are the contents of representations in predictive processing? *Phenomenology and the Cognitive Sciences*, 16(4), 715–736.
- Wiese, W. (2017b). Predictive Processing and the Phenomenology of Time Consciousness - A Hierarchical Extension of Rick Grush's Trajectory Estimation Model. In T. Metzinger & W. Wiese (Eds.), *Philosophy and Predictive Processing*: 26. MIND Group.
- Wiese, W. (2017). Action is enabled by systematic misrepresentations. *Erkenntnis*, 82(6), 1233–1252.
- Wiese, W., & Metzinger, T. (2017). Vanilla PP for philosophers: a primer on predictive processing. In T. Metzinger & W. Wiese (Eds.), *Philosophy and Predictive Processing*: 1. MIND Group