



Visual experience in the predictive brain is univocal, but indeterminate

Kathryn Nave¹ 

Accepted: 23 April 2021 / Published online: 18 May 2021
© The Author(s) 2021

Abstract

Among the exciting prospects raised by advocates of predictive processing [PP] is the offer of a systematic description of our neural activity suitable for drawing explanatory bridges to the structure of conscious experience (Clark, 2015). Yet the gulf to cross seems wide. For, as critics of PP have argued, our visual experience certainly doesn't *seem* probabilistic (Block, 2018; Holton, 2016).

While Clark (2018) proposes a means to make PP compatible with the experience of a determinate world, I argue that we should not rush to do so. Two notions of determinacy are conflated in the claim that perception is determinate: 'univocality' and 'full detail'. The former, as Clark argues, is only to be expected in any PP agent that (like us) models its world for the purpose of acting on it. But as Husserl argued, and as perceptual psychology has borne out, we significantly overestimate the degree of detail with which we perceive a univocal world.

This second form of indeterminacy is due not to the probabilistic nature of PP's model, but rather to its hierarchical structure, with increasingly coarse-grained representations as we move further from the sensory periphery. A PP system may, or may not, deliver a univocal hypothesis at each of these levels. An action-oriented PP system would only be expected to do so only at the level needed for successful action guidance. A naïve reporter's overestimation of the degree of determinate detail in their visual experience can thereby be accounted for with a more gradual version of the 'refrigerator light' effect: we experience determinate details just to the degree that they're needed – immediately as they're needed.

Keywords Predictive Processing · Phenomenology · Cognitive Science · Consciousness · Husserl · Visual Experience

✉ Kathryn Nave
Kathryn.nave1@gmail.com
<https://www.ed.ac.uk/profile/kate-nave>

¹ The University of Edinburgh, Edinburgh, UK

1 Introduction

Can predictive processing, as Andy Clark (2015) has suggested, allow us to “begin to bridge the daunting gap between the world of lived human experience and a cognitive scientific understanding of the inner (and outer) machinery of mind and reason?” (p. 237).

The fact that predictive processing characterises our cognitive operations in probabilistic terms seems a hurdle to these bridge-building efforts. According to the family of Bayesian Brain approaches, of which PP is a member, the brain forms a model that encapsulates hypotheses about the causal structure of its distal environment. Because the relationship between sensory evidence and the hidden environmental states that cause it is uncertain, so the brain must hedge its bets by weighting each hypothesis probabilistically. These probabilities are then updated via approximate Bayesian inference with each new wave of sensory input.

This probabilistic description seems to stand in conflict with the apparent determinacy of our visual phenomenology, a concern that Ned Block (2018) has recently used as the basis for an argument against probabilistic accounts of perception. Does the error lie (as is often the case) in the phenomenological description, in the sub-personal story PP tells about our cognitive architecture, or might we be mistaken to take them as being in conflict at all?

In section one I briefly motivate the claim that visual perception is determinate. In section two I provide a brief overview of the predictive processing framework and its apparent mismatch with this determinate visual phenomenology, then in section three I discuss and develop Andy Clark’s (2018) proposed solution: that, “Once we see perception aright, as the slave of action,” the PP framework falls into line with the visual experience of a determinate world” (p.1).

In section four I contrast this with the strategy taken by Michael Madary (2012, 2016), who draws upon the work of the phenomenologist, Edmund Husserl alongside more recent empirical evidence from inattentional blindness and gist perception, to claim that visual experience may not be ‘totally determinate’ after all. Taking the opposite tack from Clark, Madary argues that we have erred in our account of visual experience. Instead, we should recognize its pervasive *indeterminacy* as lending support to PP’s probabilistic story.

We now find ourselves with two seemingly conflicting options for resolving the apparent disconnect between probabilistic processes and subjective experience. We can explain why the former would be expected to produce a determinate percept for the purpose of action, or we can argue that the latter is, contrary to a widespread assumption, pervaded with indeterminacy.

In section five I argue that the apparent incompatibility of these responses is due to an ambiguity in the meaning of ‘indeterminacy’ at issue and distinguish between the question of ‘univocality’ and that of ‘full detail’. In the former case, the question is whether visual experience delivers a single take on the world, or an array of probabilistically-weighted, mutually exclusive options – the ‘Bayesian blur’ (Lu et al., 2016). In the latter, the question concerns the extent to which visual experience fails to specify some, in principle specifiable, aspects of the objects it presents.

In section six, I argue that the phenomenological indeterminacy identified by Husserl is of the latter kind: a matter of coarseness of grain, rather than probabilistic format. This fits naturally within a PP story, due to the hierarchical structure of the generative model, containing increasingly coarse-grained representations as we move further up from the sensory periphery. A PP system may, or may not, deliver a univocal hypothesis at each of these levels. An action-oriented PP system would be compelled to do so only at the (comparatively high) levels needed for action planning or linguistic report. The exact degree of detail specified is not fixed, however, but varies with the demands of the tasks that we engage in.

Why do we not typically notice the extent to which the details of our visual experience are undetermined? I propose a solution in the form of a more gradual modification of O'Regan & Noë's (2000) example of the refrigerator light illusion. Ordinary perception typically delivers an experience as fine-grained as needed, *just as it's needed*. We thus fail to attend to those unrequired details that remain undetermined – creating the misapprehension our experience is determinate simpliciter. Action-oriented PP can account for the non-probabilistic, univocal nature of visual experience, while also doing justice to its pervasive indeterminacy.

2 The case for determinate visual phenomenology

Ordinary perceptual experience, according to the determinativist account, does not equivocate. It may be wrong, it may mislead us, but it can usually be relied upon to have a definite opinion. Philosophers are similar. As such, we are not short for unequivocal endorsements of perceptual determinacy. To take a few:

From Richard Holton (2016)

“At the level of what we see, rather than that of what our unconscious visual systems are doing, we don't have a graded continuum of confidence in different hypotheses. Perceptions are all-or-nothing.” (p.10, my emphasis)


From John Searle (2015)

“We perceive objects and states of affairs in the real world in a way that the details are filled in... In non-pathological cases, they [visual experiences] ... present their conditions of satisfaction as totally determinate in a way that is never characteristic of verbal representations. To say “it is brown” leaves a range. But to see a brown colour does not in that way leave a range.” (p.69, my emphasis)


From Michael Rescorla (2015)

“Perception normally yields a determinate percept. For instance, one sees an object as having a determinate shape, not a spectrum of more or less probable shapes.” (P698)

And finally, Ned Block (2018) takes this determinacy as the starting point for a critique of probabilistic accounts of perception asking, “If perceptual representation is probabilistic, why does normal conscious perception not reflect the full probability functions that the probabilistic point of view endorses?” (p.1).

Fig. 1 Rubin's Vase

This concern seems well-motivated. I see a world of singular objects of a particular size and distance, not what Lu et al. (2016) term a 'Bayesian blur' of possible things at a range of possible locations. Such intuitive descriptions, they argue, can be further specified and supported by empirical demonstration using bistable stimuli, such as Rubin's Vase (Fig. 1), or the Necker cube (Fig. 2). Despite such stimuli

Fig. 2 Necker Cube

admitting of two *equally probable* but incompatible interpretations, our experience seems to only present the single coherent percept of either one or the other at a time.

An even stronger illustration is the case of binocular rivalry, in which a mirror stereoscope or red/green glasses are used to present one of a pair of conflicting images (such as a face or a house) exclusively to either the left or the right eye (Dieter & Tadin, 2011). Because both images originate from the same location, at the place where the eyes' visual fields overlap, the evidence received by the perceiver is evenly split between supporting the presence of an image of a house, or an image of a face, at this central location. Yet rather than a chimerical face-house percept, one's experience tends to fluctuate between either the univocal presentation of a house or the presentation of a face.

It might be noted that, at least in some cases, we do seem to perceive shadowy objects of ambiguous nature – for example, the experience of seeing a disconcertingly large silhouette slinking across the moorlands, shrouded in early morning mist. (Fig. 3). Still, the determinativist can respond that in such a situation your visual experience does not *itself* simultaneously represent an indeterminate array of potential creatures – hellhound, wildcat, yeti, alien, large sheep. Rather, visual experience determinately presents a dark shape of a certain size, shape and distance, from which an over-active imagination is liable to infer a number of possible interpretations. Any indeterminacy to be found here, according to this view, is at the level of post-perceptual judgement not in the content of perception itself.



Fig. 3 Possible photographic evidence of the beast of Bodmin Moor

3 Predictive Perception and Probabilistic Brains

Sensory evidence	Hidden Causes	
Luminance	Illumination	Surface reflectance
Retinal image size	Object distance	Object size

Cognitive science has recently undergone an increasingly influential probabilistic turn, which presents the brain as a Bayesian inference engine far more circumspect in its operations than this ‘all or nothing’ picture of perception might suggest (Chater et al., 2006; Clark, 2013; Friston, 2005, 2008; Hohwy., 2013). According to such approaches, the brain maintains an array of probabilistically-weighted prior assumptions about the structure of hidden causes in its environment. These are constantly updated with each wave of new evidence received from the senses, in approximate accordance with the optimal strategy specified by Bayes rule. As our sensory evidence always results from a combination of multiple hidden causes, it typically underdetermines the exact contribution of each (Helmholtz, 1867/1962). To take a couple of examples:

In such cases, the contribution of each component to the directly detected sensory evidence can only be pulled apart through progressive integration with additional cues: the use of motion parallax in to determine whether an object is small, or far away, or knowledge of lighting conditions to determine whether an object is blue, or white under blue-ish illumination. Thus the Bayesian brain initially hedges its bets and assigns probabilities across each possible combination of factors, rather than either holding back altogether or committing too enthusiastically to the definite selection of one particular interpretation.

How exactly the brain accomplishes this inferential process is an open question. Here I will be focusing on predictive processing as not only one of the best-developed of such probabilistic accounts, but also the first in which proponents have begun to make forays beyond the modelling of learning and behaviour towards provide a unifying account of our conscious mental life (Clark et al., 2019; Hohwy, 2012).¹

Predictive processing (Hohwy, 2013; Clark, 2013, 2015) describes the brain as instantiating a probabilistic, hierarchical, generative model. Generative, in that the brain’s core operation is the generation of signals to match those it will receive at the sensory interface. Hierarchical, because this predictive operation both supports the development of, and then recursively depends upon, a multi-layered model. These hierarchical layers track regularities at varying degrees of spatiotemporal grain: from comparatively stable hidden causes such as tables and teacups at the higher

¹ This phenomenological mismatch will nonetheless arise as a concern for any probabilistic framework to attempt the same. As my argument is not overly-concerned with the specific implementational details of PP, I hope it should be of use to Hierarchical Bayesian frameworks more generally.

levels, all the way down to their proximal effects in fluctuating activity patterns at the sensory surface. And, finally, probabilistic because, as described above, the mappings from an expected hidden cause to the expected unfolding patterns of sensory stimulation, and vice versa, are inherently noisy and uncertain.

According to PP, it is the rich content of this temporally-deep internal model, not the comparatively impoverished data streaming through the retina at the present instant alone, that directly determines perceptual experience. Sensory input is demoted to the role of model constraining error signals, suggesting the description of perceptual experience as a process of ‘controlled hallucination.’

4 A probabilistic puzzle

If perceptual experience is determined by the brain’s hierarchical generative model, and if the structure of this model is probabilistic, then why does this probabilistic content not make it into conscious awareness? As Clark (2013) puts it, “The world, it might be said, does not look as if it is encoded as an intertwined set of probability density distributions! It looks unitary and, on a clear day, unambiguous” (p.196).

This is particularly troubling given the aim of predictive processing to, “begin to bridge the daunting gap between the world of lived human experience and a cognitive scientific understanding of the inner (and outer) machinery of mind and reason” (Clark, 2015, p. 237). Yet, Clark (2018) proposes, this problem is merely at the surface level. Predictive processing can explain both how such a cognitive architecture yields a determinate percept, and why it does.

The *how* is relatively straightforward. We may account for the univocal switching of bistable stimuli and binocular rivalry described in Sect. 2, by postulating that the PP system’s probability distributions are constrained as Gaussians with a single peak. No matter the evidentiary situation, such a system is forced towards a single, though still probabilistically-weighted, winning hypothesis. To explain why a system would have such a constraint Hohwy (2013) argues we can understand this as a ‘hyperprior’, a higher level, more general prior that tracks the lawlike regularities of our environment – in this case, the fact that “only one object can exist in the same place at the same time” (p. 691).

As Clark (2018) points out, this is insufficient motivation for constantly forcing the selection of a single hypothesis in the face of evidentiary underdetermination. Even when two (or more) hypothesis about the state of the world are mutually incompatible, if both are equally well supported by present evidence then why not make use of the probabilistic representational schema to continue to preserve both at equal likelihood until further evidence resolves the situation? What would the predictive brain gain by overeager commitment?

Further, even when there is a clear victor among the competing hypotheses, we should still ask why it is this alone that shapes our conscious experience? Given that our predictive brain is also constantly tracking the performance of all the other near misses and runners up, the sometimes only marginally less likely ways that the world might be, then why is the performance of these competing hypotheses not also

kept alive in perceptual experience? We seem as Clark puts it, engaged in “a pointless refusal to profit from good information” (Ibid, p.77).

5 What is perception for?

5.1 Action versus accuracy

The solution to this puzzle, Clark (2018) claims, can be located by re-focusing our attention on what perception is *for*. An explanation of the contents of perception, he argues, must start from an understanding of how such contents are optimised for supporting the organism in acting to bring about its interests – from simple survival in the most basic creatures to the complex medley of goals and aversions characteristic of human perceivers. Looking at perceptual experience as action-oriented in this way, he claims, leads to very different expectations about its content from when we assume that perception is optimised for providing the most informative and accurate model of the brain’s environment, given currently available evidence.

That conscious perception is pragmatic, not perfectionist, is at least suggestive of why its contents might be artificially constrained. Yet as proponents of 4E cognitive science (Clark included) have argued, this is true of many, if not all, of our cognitive processes. So if Clark is right that probabilistic encodings are not needed for successful action, then why would they feature in our sub-personal model at all?

To explain why visual perception must be univocal, without jeopardizing the broader utility of sub-personal probabilistic encodings within 4E approaches, we must be clearer about visual perception’s specific role in the broader action-driven economy of the embodied brain. In particular, I argue, we will need to distinguish the particular capacity that is dependent upon conscious visual perception from both visual, *non-conscious* motor control on the one hand, and conscious, *non-visual* action-planning on the other.

5.2 The hypothesis of experience-based selection

When Clark speaks of a unitary hypothesis being necessary to ‘drive action’ it could sound as though he means to claim that the univocality of vision is something enforced by the motor control system’s need for action-guiding parameters. Yet he has previously, and repeatedly, criticised explanations of the contents of visual perception that begin with precisely this ‘assumption of experience-based control’ (Clark, 2001, 2007, 2009). There is a plethora of empirical evidence that visual experience is unnecessary to successfully perform visually-guided behaviours such as pointing, tracking and reaching – all of which ordinary participants are able to execute without conscious perception of their target (Bridgeman et al., 1981; Castiello et al., 1991; Goodale et al., 1986). This becomes particularly striking in

neurological disorders, such as visual agnosia and action-blindsight, where damage to specific areas involved in visual processing significantly disrupts visual experience without equivalent impairments to visual action-guidance.

The well-studied visual-agnosic D.F., for instance, lacks perceptual awareness of the size or shape of a slot in front of her eyes, yet, when instructed to do so, can post a letter through that same slot with perfect ease (Goodale & Milner, 1992, 1995). Based on such cases Goodale & Milner proposed that there are two cortical pathways by which visual information is processed: the ventral stream, responsible for ‘vision for perception’, and the dorsal stream, which handles ‘vision for action’. They argue that these pathways are capable of independent operation, and so propose that D.F.’s intact dorsal stream may underpin her preserved capacity for successful action-guidance, even while the damage to her ventral stream prevents her from being able to access this information for conscious report.

That there are (at least) two cortical pathways by which visual information is processed is now widely accepted (Gangopadhyay et al., 2010). If Milner and Goodale’s characterization of this is correct, as Clark (2001) argues, then conscious perception cannot be for online guidance in the unfolding of pre-established action-routines – unconscious vision handles this well enough on its own. Instead, Clark (2001) advocates the ‘hypothesis of experience-based selection’, later developed with Dave Ward and Tom Roberts (2011) into the ‘action space’ account of perceptual experience, which states that:

“...what counts for (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 (or, more accurately, the way she implicitly takes herself to be poised) over an ‘action space.’” (Ward et al., 2011 p.383)

The absence of this direct awareness of the range of action-routines currently available to her is reflected in D.F.’s behavioural capacities. To characterize her impairment as just ‘perceptual’, as Milner and Goodale (1995) did when using D.F.’s case to support a division into ‘vision for action’ versus ‘vision for perception’, misleadingly implies that her capacity for action is entirely unaffected by the loss of ventrally-supported perception. To divorce the role of the ventral stream from the support of action in this manner would be to disregard the fact that D.F. is further unable to indicate the width of the slot with her hands; to match the orientation of a letter to that of the slot without posting it; to take the initiative to post the letter without prompting or to scale her grasp to a just seen object after any delay.

D.F. is able to use visual input to guide the ongoing unfolding of a pre-specified action towards a target, once she is involved in this action. She is unable, however, to plan potential actions, to anticipatorily move her body to prepare for the execution of an action, or to perform an action in relation to visual information that is no longer directly available. What she lacks, Ward, Roberts, and Clark argue, is the ability to automatically integrate this visual input with her background knowledge and goals (according to PP, instantiated in her hierarchical generative model) in order to model the space of actions currently available to her.

5.3 Giving it your best guess

Experience-based selection neatly distinguishes the behavioural capacity associated with perceptual experience from the kind of online control that may proceed perfectly well without it. Still, if conscious visual perception is for the planning and selection of actions, rather than the ongoing control of them, then we now face the question: how does this differ from the kind of action planning I can engage in without either live vision or any corresponding perceptual experience? Neither, for instance, is required for me to sit here and decide that I would rather watch Netflix than go for a run this evening.

Clark (2007) seems to lose hold of such a distinction when he claims that in visual experience “consciously available information is used only for the specification of action types (‘don’t bruise the apples, they are delicate’) and targets, and is not used even to compute a rough sketch of the trajectory itself” (p. 576). While emphasizing the importance of currently transduced visual information to the constitution of an action space, Ward et al. (2011) similarly define the action space in terms of the selection of such high-level action categories. Yet it is unclear what the necessary role for the *ongoing* integration of visual information is in such a selection. The decision to ‘throw the apple away, it is mouldy,’ versus ‘eat the apple on the table, it is ripe,’ may be just as well served by being told there a mouldy apple in front of me, as by perceiving it for myself.

Beyond differing sensory requirements and phenomenological profiles associated with perceiving versus reasoning, the issue this raises for Clark’s argument about perceptual univocality is this: when it comes to non-perceptual reasoning about possible action type, there seems no impediment to the use of probabilistic information. Casino bosses, stock traders and epidemiologists would all argue it’s often rather useful. Based on prior experience I can rate the likelihood that the mouldy apple will make me ill at around 20%, and the likelihood that I’d be able to successfully throw it into the bin without having to get up as about 5%. By weighing these up, I can decide a course of action: eat it and hope for the best.²

To explain why visual perception is univocal then, we must also explain how its role in supporting action-selection differs from this kind of non-experience-based reasoning about actions. To do so, I argue, we must disagree with Clark’s (2007) position that experience-based selection is concerned only with generic action types. Instead, what distinguishes our perceptual action planning from such offline reasoning is that the former alone involves directing our background knowledge toward the moment-to-moment, goal-directed evaluation of those *spatiotemporally-specific* bodily movements that are available to me right now. The kind of actions that visual experience allows us to select between are not just abstract types like ‘eat’ vs ‘throw’, but rather the particular sort of throws that would, or would not, achieve my goal, given my current bodily position and environmental situation.

² Thanks to John Michael (personal communication) for raising this point.

This spatiotemporal-specificity explains why experience-based action selection depends upon live visual information, where offline action-planning need not. It also suggests an explanation for aspects of perception's phenomenological profile, such as its 'greater granularity' (Peacocke, 1992; Tye, 2006; Martin, 1992) and its 'perspectival connectedness' (Siegel, 2006), which have been proposed to distinguish it from the phenomenology of conceptual reasoning.

Most important here, however, is how the immediacy of this form of action selection explains the need for univocality. My choice to either throw the apple or to eat it is relatively stable, and may be actioned at any point over the next few hours. Reasoning about such generic action types may take a wait-and-see approach, accumulating more evidence until any uncertainty is eliminated. In contrast, the particular way I can throw said apple is a transient thing, shifting moment-by-moment with alterations in the position of my body or movement of the world around me. And it is because action-trajectories are always of the moment that perception must be univocal if it is to be useful. When it comes to down to the wire, we cannot probabilistically apportion an action across multiple possible worlds. We cannot both 60% throw the apple and 40% eat it. To act requires choosing one possibility over another. A probabilistic blur would leave us paralysed.

The world of perception is the probabilistic brain's constantly changing best guess, at what you can do *right now*, but this is not to say that it constantly mandates urgent activity. Among the options that my experience reveals as both currently available, and most desirable, could well be to continue remaining slumped listlessly in an armchair. Nor does this analysis of perceptual experience foreclose the possibility of our developing ways to put perceptual content to other uses, such as taking an aesthetic stance of simply enjoying the view. It means only that, even in your most tranquil mode, so long as you are visually perceiving the world around you, then you are poised for immediate action upon it. Were you not so poised – were your perceptual experience to leave you utterly without orientation towards the possibilities currently afforded to you, unsure if even 'remain unmoving' is a viable prospect, then 'tranquil' seems about the last word appropriate to the situation.³

By clarifying that enforced univocality is due not just to perception's being 'action-oriented', but specifically due to its time-sensitive role, we can avoid undermining the compatibility of probabilistic encodings within 4E accounts more generally. In combining a probabilistic mechanism with the enforced selection of univocal content for immediate action, the action-oriented predictive brain gets the best of both worlds. My perceptual experience may inform me of a single best guess regarding the way I can eat the cake right now, all whilst background cognitive processes continue to weigh up the estimated likelihood of my partner snaffling it first, in order to rate the long-term strategy of having it later.

³ Thanks to an anonymous reviewer for encouraging me to make this explicit.

6 Seeing the wood for the trees: the case for perceptual indeterminacy

6.1 Phenomenology and the horizon of indeterminacy

Predictive processing can be rendered compatible with determinativism about visual experience. So, case closed? Not quite. We may have proceeded a little too hastily in conceding the unequivocal determinacy of perceptual content. As Clark (2018) briefly notes (pointing to work from Seth (2014, 2017), Cohen and Dennett (2016), Kouider et al. (2010), Lettvin (1976) and Madary (2012) there is good reason to believe that some elements of visual experience may actually be better described as “in some sense ‘statistical’ in character” (p. 83).

Further, there is an alternative perspective on perceptual experience (one with a history that long precedes the development of the Bayesian brain hypothesis) which argues that indeterminacy is not only a common feature but a necessary characteristic of conscious visual experience. Within the phenomenological tradition, instigated by Husserl’s *Logical Investigations* in 1900 (Husserl, 1900/2001a), and particularly through its development in *Thing and Space* (1907/1997), *Ideas I* (1913/1982) & *Analyses concerning passive and active synthesis* (1920/2001b) this indeterminacy has been presented as not merely an occasional feature, but something inherent to the structure of visual perception.⁴

Recall, the determinativist position is that perceptual experience presents a single ‘unitary value’ for all properties presented in perception. So the contrasting indeterminativist position can be provisionally characterised as the claim that, for at least some visually perceived properties, what is presented is not a single value, but a delimited range of possible values.

A few preliminary clarifications. This Husserlian claim of indeterminacy is not to say, as Michael Tye does, “In the case of seeing blurrily, one’s visual experience [...] makes *no comment* on where exactly the boundaries lie” [emphasis added] (2002, p.81). When a noisy, low-resolution photograph of a distant object represents the edges of that object with indeterminacy, it still makes *some comment* on the possible locations where these edges fall, and where they definitely do not. As Husserl (1907/1997) describes, “Indeterminateness is never absolute or complete. Complete indeterminateness is nonsense; the indeterminateness is always delimited in this or that way.” (§18, p. 50/59).

Indeterminacy should also be distinguished from vagueness (Stazicker, 2011). If I claim that, “There are lots of kittens in this box,” then there are indeed some states of affairs – say zero kittens, or one kitten – that would sadly render my statement false. There are other more joyous possibilities – such as five or more kittens – under which it would clearly be true. However, there will also be situations, such as three or four kittens for which the statement’s truth or falsity is unclear. An indeterminate statement may also be vague, but it is not necessarily so.

⁴ For Husserl, indeterminacy is a feature of all intentional acts. Here I will be focusing specifically on the case of visual indeterminacy.

6.2 Scene perception and inattentional blindness

Where exactly should we locate this “horizon of indeterminacy” in the seemingly determinate world of tables, chairs, construction sites and cranes given to us in visual phenomenology? One example, pursued by Husserl’s follower, Maurice Merleau-Ponty (1996) is the figure/ground relation. Originating from the Gestalt psychologists (Kohler, 1947; Koffka, 1922), this describes the manner in which a perceptual object always appears against some broader background, which lacks the fidelity of the attended object. In this experiential structure, Koffka, describes, “we have a very general characteristic, namely, that the ground is always less “formed,” less outlined, than the figure.” (p. 556).

An illustration. Imagine yourself strolling through a coastal pine forest on a warm day, when you spot a flash of bright red against a tree trunk. Moving in for closer inspection, you identify this as belonging to the tiny, tufted plume that gives the rare red-cockaded woodpecker its name. As focus you upon the little feathery bundle, how does the bark texture of the tree trunk immediately surrounding it appear? I would suggest that, while perceptually present, the bark is far less specified than the clearly presented pattern of white speckling the little bird’s black wings. Such an experience gives a rough sense of what is meant by the horizon of indeterminacy.

Recently, empirical work has begun to flesh out the bones of the Husserlian claim of perceptual indeterminacy, revealing in the process that there is far less determinacy in visual phenomenology than intuitions about everyday experience would suggest (Madary, 2016). In the inattentional blindness paradigm (Simons & Chabris, 1999) participants are instructed to watch a video while performing a task, such as counting the number of times that a team of players dressed in white pass a basketball between them. While they do this, what would typically be considered a highly interesting development – such as a person in a gorilla suit walking through the middle of the scene – typically goes unnoticed. Even when this change occurs within the central, high-resolution region of the visual field, directly behind the ball that is being attended to, participants still regularly fail to observe it.⁵

This is an extremely surprising thing to experience. Viewers typically insist on re-watching the video before they are prepared to accept the appearance of the gorilla right under their nose. After all, while you attend to the basketball it certainly does not feel as though the goings-on around it vanish from awareness.

To explain this Cohen et al., (2016) recruit the notion of ensemble statistics, formed by collapsing measurements of an aspect of each individual background element – such as brightness, size, motion speed and direction – into a single average value for all the elements of a particular region. Such summaries can provide

⁵ This is similar to the change blindness paradigm—in which participants fail to notice significant changes to the scene that occur during a brief mask (O’Regan et al. 2000). However, Mitroff et al. (2004) propose an alternative explanation for this phenomenon in terms of failure to compare the pre and post-mask experiences, rather than a lack of detail in the experience of either. In inattentional blindness, there is no mask and the ‘unseen’ element is typically selected to be intrinsically worthy of attention, rather than by virtue of any comparison to the previous scene. As such I have focused on this to sidestep consideration of Mitroff et al.’s account.



Fig. 4 Ensemble statistics in scene perception (Oliva & Torralba, 2006)

coarse-grained information about the overall properties of our expansive visual field, without determinately specifying the exact properties of each and every element contained within it.

That the background scene is represented as ensemble statistic explains how we continue to experience it as surrounding the object of attention, yet fail to notice changes within it. In Simons & Chabris (1999) original inattention blindness test,

Fig. 5 Row of bricks (Pérez, 2008)






Fig. 6 Unnoticed expression changes (David et al., 2006)

the unnoticed gorilla is similar in its colour, size, shape, and motion to the average properties of the black-clothed players also moving about the scene. As such, the gorilla does not violate the established average properties of the scene. If it is these summaries that make up the content of our perceptual experience, then the fact that the gorilla's appearance does not make a difference to them explains why it does not reach our perceptual awareness. This interpretation is strengthened by several results showing that when a change disrupts the overall statistics of a scene, then participants tend to notice it easily (Alvarez & Oliva 2009; Brady et al., 2011; Victor & Conte, 2004).

Thanks to the pervasive repetition and redundancy in natural scenes, just a few ensemble statistics are sufficiently informative to specify a range of environment types (Figs. 4, 5, 6). This rapid perception of high-level categories is typically referred to as 'gist' perception and can precede the perception of the individual details that compose that scene (Greene & Oliva, 2009). Through statistical summarization then, our visual system can rapidly deliver a determinate perception of some high-level scene type while remaining indeterminate to the details of each of the individual elements of which it is composed.

6.3 The indeterminacy of attended objects

Background scene perception is one way the horizon of indeterminacy shows up in perception, but the picture so far is misleadingly incomplete. Cohen et al. (2016) conclude their discussion with the claim that "a handful of items are perceived with high fidelity, while the remainder of the world is represented as an ensemble statistic (or set of statistics)" (p.332). In other words, amid widespread perceptual indeterminacy, they nonetheless suggest that at least our experience of the present object of our attention is fully determinate. This is at odds with the Husserlian picture, and more importantly, I will argue, fails to accurately capture our visual phenomenology.

For Husserl, our experience of objects also contains further 'determinable indeterminacy' – the potential for further unpacking, for the further specification of detail that is initially experienced as unspecified.

"...every perception, or noematically speaking, every single aspect of the object in itself points to a continuity of possible new perceptions... It calls out to us, as it were, in these referential implications. "There is still more to see here," (1920/2001, §1, p.41)

Take this photo of a brick wall.

The description 'rows of weathered, light terracotta bricks, viewed at an angle' likely does a reasonable job of capturing the content of your initial experience. Yet there is, as Husserl describes, more to see here. You can attend to the colour graduation across an individual row, or to an individual brick and its mottled texture. Walls, bricks and other objects (just like scenes) are made up from a significant amount of repeating detail.⁶ And, as with scenes, it is not necessary for experience to exactly specify all these low-level details in order to adequately determine the object type.

If it were, you'd have immediately noticed the cigar sticking straight out from underneath the second row.⁷

Just as statistical summaries support the perception of the character of a scene from a coarser-grained description of its constituent objects, so they can also capture the overall texture of some surface without a determinate specification of the size, shape and location for each individual mark (Portilla & Simoncelli, 2000). And, as with scenes, the recognition of high-level object categories can occur more rapidly than recognition of the individual details of the object (Thorpe et al., 2001). Object-based inattentional blindness has been less comprehensively researched, but David et al. (2006) provide a striking example in which the expression of a directly-attended-to face slowly changed over time, right before a participant's eyes. Even though there was no masking while the change took place, 85% of participants fail to notice it.

This account of gist perception allows us to explain how we can directly and determinately perceive the wood, the water, the team, or the wall, without ever consciously experiencing each individual tree, wave, player, or brick (Bayne & McClelland, 2019). We can claim this without committing the category error of talking of the forest as an independent entity, over and above the more detailed elements of which it is composed. Instead, perception of a gist is something that constrains, but does not determinately specify, exactly the details that would be revealed if you were to attend more closely, or to bring the relevant items into closer view. It is, in Husserlian terms, the perception of a 'determinable indeterminacy.'

⁶ Given the similarities between object and scene perception on the account that I'm developing, and as indicated by some of these borderline cases, I expect that the distinction between the two is not as clear cut as it is often posed.

⁷ Perhaps you did see it. But the majority of people don't until the cigar is explicitly pointed out – following which a gestalt switch occurs and it becomes very difficult *not* to see it. The only way I have found to do so is by taking my glasses off and squinting intensely. This makes for interesting coffee-shop interactions.

7 Disambiguating two notions of determinacy

We now find ourselves with two, apparently contradictory, perspectives on the content of visual experience. The predicament that originally faced advocates of predictive processing seems inverted. We could reconcile claims about the determinacy of visual phenomenology with the sub-personal probabilistic machinery of PP, via Clark's (2018) action-oriented argument for determinate perception in the probabilistic brain. But those claims of determinate visual experience, which we earlier took for granted, have now been drawn into question by the conflict with both phenomenological arguments and empirical evidence for pervasive perceptual *indeterminacy*.

Drawing upon both the phenomenology of Husserl and the evidence of change-blindness, Michael Madary (2016) has recently mounted a defence of predictive processing, via a strategy that runs directly orthogonal to Clark's. Rather than seeking to make PP compatible with our initial assumptions about determinate phenomenology, he argues instead that it is this evidence for the *indeterminacy* of visual phenomenology (presented in the previous section) that reflects, and provides support for, the probabilistic representations of predictive processing.

Which strategy should the PP theorist adopt? One option would be to view the issue of visual determinacy as nothing more than another cautionary tale about the unreliability of naïve introspection. The evidence provided by inattentive blindness falls on the side of indeterminate perceptual experience, and so we should take Madary's course over Clark's. The problem with this strategy is that we also find evidence in binocular rivalry and bistable perception, which appears to support the enforcement of determinacy in visual perception. Further, even if the Husserlian perspective is right, we should still expect an error theory for why not only laypeople, but also cognitive scientists and philosophers from outside the phenomenological tradition have rejected, or neglected, this pervasive indeterminacy.

Rather than a choice between two contradictory strategies, I believe that what we are faced with here is instead a confusion, stemming from a failure to distinguish between two different meanings of 'determinacy', which has resulted in the conflation of two separate positions that may be asserted by ascribing it to perceptual experience:

- 1) **Univocality:** The content of ordinary visual experience is always presented as a single, coherent take on the way things are – and never as a probabilistically-weighted presentation of several mutually exclusive alternatives.
- 2) **Full Detail:** The content of visual experience always specifies all the details of the objects it presents

Neither Madary nor Clark distinguish these. Madary suggests that the indeterminacy of perceptual experience could be a reflection of the 'sub-personal probabilistic code' of a predictive brain, implying that rejecting 'the myth of full detail' and rejecting univocality in perception go hand-in-hand. Clark (2018)'s argument concerns *univocality*, not full detail. Nonetheless, he briefly considers Madary's arguments for (at least some) indeterminacy in perception and interprets this as, potentially, a

limited exception to perceptual univocality. Like Madary, he thus assumes perceptual indeterminacy must be a reflection of probabilistic representations.

I believe this conflation is a mistake. In a hierarchical representational schema, like PP, the claim that perceptual experience is univocal with respect to whichever details it does present entails nothing about how fine-grained these univocally presented details must be. An experience can be univocal, while still leaving some details of what it presents unspecified.

8 Gist all the way down

In addition to probabilistic representation, PP postulates a hierarchy of predictions at layers of increasingly coarse spatiotemporal grain: from the precise detail of fluctuations at the sensory periphery to comparatively detail-indifferent hypotheses about the presence of tables and chairs. At each level a single winning hypothesis may, or may not, be selected.

The content of our perceptual world is similarly hierarchically layered and irreducible to the presentation of a singular piece of content to be judged one way or the other. When you observe the woodpecker on the tree trunk you do not simply see that there is a bird on a tree. In one go, you might see the size and colour of the bird, the length of the grass, the roughness of the bark, the general leafiness and much more. So, when asked if the content of your experiences is presented univocally, you may well be inclined to respond, “Well, with regards to what?” If it is anything like mine, your experience definitively presents ‘that is a bird’ and ‘this is a forest’ – not a cityscape, rolling prairie, or ocean vista. Yet it does not univocally specify every fine-grained element of the scene. Gist perception explains how we might definitively see the grass and the foliage, without needing to experience the specific position of each and every single individual blade and leaf that constitute these.

Rather than positing a single division between gist perception versus the low-level details that are summarised, PP supports the idea of levels of increasing summarisation – a hierarchy of gists upon gists. The lower in the hierarchy, the more specific the representation. This fits nicely with how Husserl (1920/2001) characterises the indeterminacy he locates in the structure of visual experience:

“We also spoke of determinable indeterminacy. Indeterminacy is a primordial form of generality whose nature it is to be fulfilled in the coincidence of sense only by “specification.” As long as this specification itself has the character of indeterminacy... it can attain further specification, etc., in new steps.” (§8 p.45)

How then should we understand Searle’s (2015) claim that visual phenomenology determinately specifies ‘all the details’ of what it presents? In a hierarchical context, this sounds like the claim that perceptual experience demands one univocal hypothesis at *every level*, right down to the most fine-grained ‘superdeterminate’ properties (Funkhouser, 2006) which cannot be any more narrowly specified to some finer degree of granularity. This seems to be what Searle (2015) intends when he says:

“The rich intentional content [of visual perception] requires a hierarchical structure of lower perceptual features, all of which are part of the content of the seeing as.... In each case, the perception of the object as having the higher-level feature requires perception of the lower-level features. Eventually, if you carry through the steps, you reach rock bottom. You reach a set of properties which can be perceived without perceiving anything else by way of which you perceive them.” (p.112)

Searle does not tell us exactly what these basic features in which visual perception ‘bottoms out’ are, claiming this should be intuitive—though he briefly suggests something like ‘coloured shapes.’

As far as the sub-personal machinations of the predictive processing hierarchy are concerned, however, ‘coloured shapes’ are certainly not ‘rock bottom.’ Here the superdeterminate level is found in predictions of particular patterns of sensory stimulation. Unlike the high-level, coarse-grained, and abstract prediction of, say, a woodpecker in a forest, which may be unpacked into a wide variety of more detailed predictions, there is no range of more specific possible predictions that some particular pattern of sensory activation could be unpacked into.

Do patterns of sensory activation make an appearance in our perceptual phenomenology? If so, certainly not in ordinary experience. In defence of an intermediate level of conscious experience, Prinz (2017) points to the fact that patterns of activation in the early visual cortex, related to illusory contours (Ramsden et al., 2001) or rapid colour changes (Gur & Snodderly, 1997), do not correspond to consciously experienced properties.

Low-level sensory activity is unlikely to correspond to what Searle intends by talk of basic perceptual features. Yet is hard to see what else could be considered as suitably ‘basic’ in a PP hierarchy. As accounts of rapid gist perception demonstrate, the kinds of properties that may enter into experience and be consciously “*perceived without perceiving anything else by way of which you perceive them*” (Searle, 2015) can range from the particular details of a woodpecker’s speckled wings, to something like a very abstract and coarse-grained awareness of ‘forestry-ness’. There are many levels of detail available in a PP system and no basic level from which visual phenomenology is always composed. Thus, talk of perceptual determinacy in terms of ‘all the details’ being specified thus does not hold up within a PP account.

In rejecting maximal determinativism we can still accept Clark’s (2018) argument that perceptual experience is non-probabilistic, that it definitively presents a single winning hypothesis for action guidance. The indeterminacy of perceptual experience, contra Madary, does not involve the positive presentation of a probabilistically weighted blur of several mutually incompatible possibilities. Rather it corresponds to the fact that the univocally presented prediction that makes its way into experience is typically achieved at a higher level of the hierarchy, such that it may adequately be fulfilled by a wide range of possibilities at the levels below.

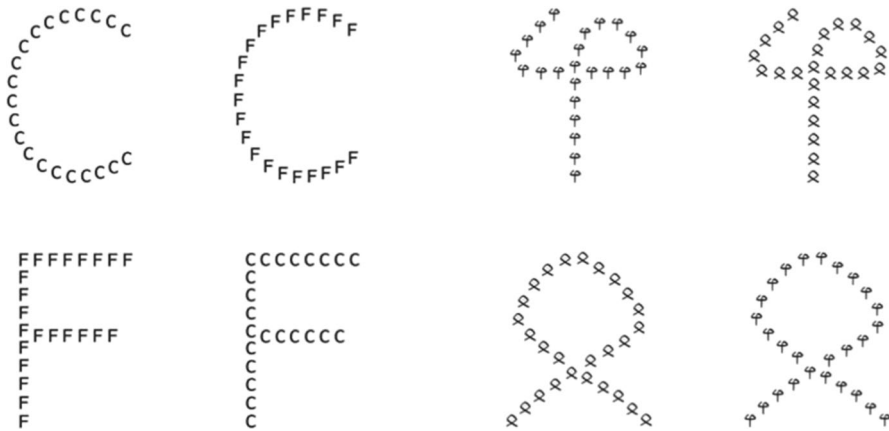


Fig. 7 Navon figures (Lachmann et al., 2014)

9 Determinate-enough perception

Two questions remain: if perceptual experience is not *maximally* determinate, then what accounts for the level of detail with which it does specify a single univocal take on the world? And why might we be misled to believe that it presents more detail than it does? The answer to the first, I will argue, is that it depends on what sort of action you are engaged in. An answer to the second question follows naturally from this.

In extending Clark's (2018) argument, I proposed that it is not just perception's need to support action which mandates its univocality, but more precisely its need to support the selection of action trajectories *specific to the immediate situation*. This does imply that perceptual content will be more detailed than the kind of content involved in thinking about potential action-types, a point reflected in their differing phenomenological profiles. But perceptual content's involving more detail than offline thought does not entail that it is always, or even ever, maximally detailed. Our actions are not only deterministic, they are also coarse-grained in their objects, operating on averages and indifferent to the fine degrees of variance that our perceptual systems are in principle capable of registering. We do not act on photons, or the thumbnail-sized orientated bars detected in the early visual cortex. We act on a landscape of graspable units, movable objects, and stable surfaces. Small variations in low-level details alter this landscape not a jot.

As such, the imperative to provide content for immediate action selection is no motivation for perception to serve up a single hypothesis about the present state of the world all the way down to these finest-grained of details. The selection of a more detailed hypothesis comes at a cost. Increased specificity brings greater uncertainty and a higher risk of getting it wrong (Otworowska et al. 2014). Further, PP's top-down account of perception describes general predictions as being formed first then unpacked in more detail, rather than being assembled piecewise from the bottom up. Gist perception for objects and scenes, discussed previously, is one more

recent demonstration of this (Alvarez & Oliva, 2009; VanRullen & Thorpe, 2001a, 2001b). But the notion can be traced back to Navon's (1977) demonstration of the global precedence effect, in which participants are shown to generally respond more quickly to the global shape of a stimulus, than to the shape of the details that constitute it (Fig. 7). There is no reason for the brain to spend the additional time and cognitive resources needed to deliver a single, consciously experienced, hypothesis about details when these make no difference to its present possibilities for action.

To say that ordinary perceptual experience delivers a determinate take only on those details that are relevant to action is not to say that this degree of detail is fixed. We constantly switch between actions operating on the world at vastly differing levels of granularity. Take your visual experience of this document. In a matter of moments, you might move from noting the typo in the word 'experience' on the line above, to scanning the number of paragraphs remaining to read, to bundling the whole page together with a stack of other papers and moving them to the side of your desk. According to this action-oriented account, as you manipulate the entire paper stack your conscious experience will not present this page with the level of detail as it did when you were engaged in reading the words printed on it. Such details are irrelevant to the action of manipulating this sheet of paper.

Indeed, when Navon figures made up from smaller letters are presented in a reading context, then the global precedence effect dissipates with the individual letter parts being recognised faster than the overall shape in which they are presented (Lachmann et al. 2014). Repeated engagement in actions that require making particularly fine-grained discriminations may even expand the possible limit of determinate detail available in visual experience altogether. Li et al. (2009) have shown that extensive playing of action video games, such as Call of Duty, leads to an improvement in a person's contrast sensitivity of between 43–58 percent.

A similar transformation likely occurs with expertise in other skills requiring the precise co-ordination of action with a finer-grained level of sensory detail than usual. Visual art is an obvious example. For the untrained, translating a perception into a convincing depiction is extremely difficult – often confusingly so. For all that one may find oneself both visually and motorically unimpaired, the results persistently suggest otherwise, and it's hard to understand exactly where things are going wrong. Artists talk of 'learning to see,' which may provoke the frustrated novice to respond that she can *see* the tree in front of her, just fine thank you very much. Indeed, said novice can perfectly well draw *a* tree, the problem is that she can't draw *this* specific tree. What the novice lacks, what the skilled drafter has developed, is the ability to coordinate her actions to visual details far more fine-grained than those required for the usual repertoire of pointing, poking, grabbing and verbally describing.

10 Fading out

We now have a solution for why one might endorse an initial impression of full detail in visual experience. The reason we do not notice a pervasive absence of determinate details throughout perceptual experience is because this experience typically serves up all the details needed to specify our current 'poise over an action

space'. When you attempt to read the words on this screen, you find these determinately enough presented in perceptual experience to do so. When you switch to attempting finer-grained actions, such as checking for dead pixels on the screen in front of you, then, at the exact moment of doing so, you find these individual pixels to be present also. It is less a matter of advances in display technology than the fact that these specific pixels were not relevant to your efforts to read the words on this page, which explains why you did not notice their lack of specification in the previous experience. Hence, we have the misleading sense that perception is determinate in all its details.

This is similar to O'Regan & Noë's (2000) explanation of change blindness under the heading of the 'refrigerator light illusion'. The suggestion here is that we do not see the gorilla, because we do not have a continuous experience of the unattended background to the object of our attention. Rather the very moment we switch our attention to this background an experience of it is formed – creating the illusion that (just as when you open the door to the fridge) the light was on all along.

Alva Noë (2004) describes this unattended world as something 'present as *accessible*', concurring with O'Regan (1992) in locating this apparently rich perceptual information as external to the mind, in the world itself as an 'outside memory.' While your eyes rove freely about the scene in a continuous process of accessing this externally-stored information, such an account seems plausible in its presentation of a rich perceptual experience as something never presented at a static moment, but rather emerging over time. Yet an account in terms of temporally-extended patterns of access is unable to do justice to the expansiveness of visual phenomenology, even when your gaze remains fixed. We still require something present in the mind itself during fixation to support this expansiveness. This, I have argued, is the univocal coarse-grained objects and scenes, present as still further *determinable* in their specific details.

Crucially, this allows us to reject the implausible dichotomy suggested by talk of visual information being either 'accessed' or not, in favour of a gradualist account that is both more phenomenologically plausible and a better fit for the varying spectrum of detail required across our action repertoire. Conscious perception is a soft vignette, not a harsh spotlight that either picks out objects with perfect fidelity, or casts them into darkness. When you read this word it is presented with a much higher degree of detail than the rest of the page, but you maintain *some* awareness of the latter in its summary statistics. An experience of surrounding type if not each individual character. Rearrangements of letters might go unnoticed, but if the entire text were to suddenly switch to 48pt Wingdings you would certainly be aware of it.

The probabilistic brain does not undermine your experience of seeing this page – and doing so determinately. Engagement in a high-stakes game of pass-the-parcel will not induce total peripheral blindness, though you should still doubt your ability to spot any subtle manipulations and underhanded trickery by surrounding players. Our subconscious probabilistic machinery delivers up a univocal, non-probabilistic world, in support of our need to select a course of action. It does not, however, offer up a single take on every possible detail of what it presents, but one that is only as detailed as the relevant opportunities for action require it to be.

Funding Supported by ERC advanced grant X-SPECT—DLV-692739.

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

- Alvarez, G. A., & Oliva, A. (2009). Spatial ensemble statistics are efficient codes that can be represented with reduced attention. *Proceedings of the National Academy of Sciences of the United States of America*, *106*, 7345–7350.
- Bayne, T., & McClelland, T. (2019). Ensemble representation and the contents of visual experience. *Philosophical Studies*, 1–21.
- 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.
- Brady, T. F., et al. (2011). A review of visual memory capacity: Beyond individual items and toward structured representations. *Journal of Vision*, *11*, 1–34.
- Bridgeman, B., Kirsch, M., & Sperling, A. (1981). Segregation of cognitive and motor aspects of visual function using induced motion. *Perception and Psychophysics*, *29*, 336–342.
- Castiello, U., Paulignan, Y., & Jeannerod, M. (1991). Temporal dissociation of motor responses and subjective awareness: A study in normal subjects. *Brain*, *114*(6), 2639–2655.
- Chater, N., Tenenbaum, J. B., & Yuille, A. (2006). Probabilistic models of cognition: Conceptual foundations.
- Clark, A. (2001). Visual experience and motor action: Are the bonds too tight? *Philosophical review*, *110*(4), 495–519.
- Clark, A. (2007). What reaching teaches: Consciousness, control, and the inner zombie. *The British Journal for the Philosophy of Science*, *58*(3), 563–594.
- Clark, A. (2009). Perception, action, and experience: Unraveling the golden braid. *Neuropsychologia*, *47*(6), 1460–1468.
- 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. (2018). Beyond the 'Bayesian blur': predictive processing and the nature of subjective experience. *Journal of Consciousness Studies*, *25*(3–4), 71–87.
- Clark, Andy, Friston, Karl and Wilkinson, Sam (2019) Bayesing Qualia: consciousness as inference, not raw datum. *Journal of Consciousness Studies*.
- Cohen, M. A., Dennett, D. C., & Kanwisher, N. (2016). What is the bandwidth of perceptual experience? *Trends in cognitive sciences*, *20*(5), 324–335.
- David, E., Laloyaux, C., Devue, C., & Cleeremans, A. (2006). Change blindness to gradual changes in facial expressions. *Psychologica Belgica*, *46*(4), 253–268.
- Dieter, K. C., & Tadin, D. (2011). Understanding attentional modulation of binocular rivalry: a framework based on biased competition. *Frontiers in Human Neuroscience*, *5*, 155. <https://doi.org/10.3389/fnhum.2011.00155>
- Friston, K. (2005). A theory of cortical responses. *Philosophical transactions of the Royal Society B: Biological sciences*, *360*(1456), 815–836.
- Friston, K. (2008). Hierarchical models in the brain. *PLoS computational biology*, *4*(11), e1000211.
- Funkhouser, E. (2006). The determinable-determinate relation. *Noûs*, 548–569.

- Gangopadhyay, N., Madary, M., & Spicer, F. (Eds.). (2010). *Perception, action, consciousness: Sensori-motor dynamics and two visual systems*. Oxford University Press.
- Greene, M. R., & Oliva, A. (2009). Recognition of natural scenes from global properties: Seeing the forest without representing the trees. *Cognitive psychology*, 58(2), 137–176.
- Goodale, M. A., & Milner, A. D. (1992). Separate Visual Pathways for Perception and Action. *TRENDS in Neurosciences*, 15(1), 20–25.
- Goodale, M. A., Péllisson, D., & Prablanc, C. (1986). Large adjustments in visually guided reaching do not depend on vision of the hand or perception of target displacement. *Nature*, 320, 748–750.
- Gur, M., & Snodderly, D. M. (1997). A dissociation between brain activity and perception: chromatically opponent cortical neurons signal chromatic flicker that is not perceived. *Vision Research*, 37(4), 377–382.
- Hohwy, J. (2012). Attention and conscious perception in the hypothesis testing brain. *Frontiers in psychology*, 3, 96.
- Hohwy, J. (2013). *The Predictive Mind*. Oxford University Press.
- Holton, R. (2016) Review of Surfing Uncertainty: Prediction, Action, and the Embodied Mind, Times Literary Supplement October 7, 10–11.
- Husserl, E. (2001a). Logical investigations. vol. 1. (J.N. Findlay, Trans.). London. Original work published 1900
- Husserl, E. (2001b). Analyses Concerning Passive and Active Synthesis. Lectures on Transcendental Logic. (A.J. Steinbock, Trans.). Dordrecht: Kluwer Academic Publishers. Original work published 1920
- Husserl, E. (1997). Thing and space: lectures of 1907 (R. Rojcewicz, Trans.). Dordrecht: Kluwer. Original work published in 1907.
- Husserl, E. (1982). Ideas pertaining to a pure phenomenology and to a phenomenological philosophy. First book: general introduction to a pure phenomenology (F. Kerstens, Trans.). Dordrecht: Kluwer. Original work published 1913.
- Köhler, W. (1947). Gestalt psychology: An introduction to new concepts in modern psychology.
- Koffka, K. (1922). Perception: an introduction to the Gestalt-Theorie. *Psychological Bulletin*, 19(10), 531.
- Kouider, S., de Gardelle, V., Sackur, J., & Dupoux, E. (2010). How rich is consciousness? *The partial awareness hypothesis*, *Trends in Cognitive Science*, 14(7), 301–307.
- Lachmann, T., Schmitt, A., Braet, W., & van Leeuwen, C. (2014). Letters in the forest: global precedence effect disappears for letters but not for non-letters under reading-like conditions. *Frontiers in psychology*, 5, 705.
- Letting, J. Y. (1976). On seeing sidelong. *The Sciences*, 16, 10–20
- Li, R., Polat, U., Makous, W., & Bavelier, D. (2009). Enhancing the contrast sensitivity function through action video game training. *Nature Neuroscience*, 12(5), 549–551.
- Lu, Stafford, and Fox. (2016). Maximum Saliency Bias in Binocular Fusion Connection Science 28:3: 258–269.
- Madary, M. (2012). How would the world look if it looked as if it were encoded as an intertwined set of probability density distributions? *Frontiers in psychology*, 3, 419.
- Madary, M. (2016). *Visual phenomenology*. MIT Press.
- Martin, M. G. F. (1992). Perception, Concepts, and Memory. *Philosophical Review*, 101, 745–763.
- Merleau-Ponty, M. (2013). *Phenomenology of perception*. Routledge.
- Milner, D., & Goodale, M. A. (1995). *The visual brain in action*. Oxford University Press.
- Mitroff, S. R., Simons, D. J., & Levin, D. T. (2004). Nothing compares 2 views: Change blindness can occur despite preserved access to the changed information. *Perception & Psychophysics*, 66(8), 1268–1281.
- Navon, D. (1977). Forest before trees: The precedence of global features in visual perception. *Cognitive psychology*, 9(3), 353–383.
- Noë, A. (2004). Perception in action.
- O'Regan, K & Noë, A. (2000) Experience is not Something We Feel but Something We Do. talk at ASSC 2000.
- O'Regan, K. K. (1992) Solving the "real" mysteries of visual perception: The world as an outside memory. *Canadian Journal of Psychology/Revue Canadienne de Psychologie*, 46 (3), 461–488.
- O'Regan, J. K., Deubel, H., Clark, J. J., & Rensink, R. A. (2000). Picture changes during blinks: Looking without seeing and seeing without looking. *Visual Cognition*, 7(1–3), 191–211.

- Oliva, A., & Torralba, A. (2006). Building the gist of a scene: The role of global image features in recognition. *Progress in brain research*, *155*, 23–36.
- Otworowska, M., Kwisthout, J., & van Rooij, I. (2014). Counter-factual mathematics of counterfactual predictive models. *Frontiers in psychology*, *5*, 801.
- Peacocke, C. (1992). Scenarios, Concepts, and Perception, in T. Crane (ed.), *The Contents of Experience: Essays on Perception*, Cambridge University Press, pp. 105–135.
- Pérez, C. (2008, 25th August). Anoché en una radio de Écija. [Blog post]. Retrieved from <https://misteriosaldescubierto.wordpress.com/2008/08/25/anoche-en-una-radio-de-ecija/>
- Portilla, J., & Simoncelli, E. P. (2000). A parametric texture model based on joint statistics of complex wavelet coefficients. *International journal of computer vision*, *40*(1), 49–70.
- Prinz, J. (2017). The Intermediate Level Theory of Consciousness. *The Blackwell Companion to Consciousness*, 257.
- Ramsden, B. M., Hung, C. P., & Roe, A. W. (2001). Real and illusory contour processing in area V1 of the primate: a cortical balancing act. *Cerebral Cortex*, *11*(7), 648–665.
- Rescorla, M. (2015). Bayesian perceptual psychology. *The Oxford handbook of the philosophy of perception*, 694–716.
- Searle, J. R. (2015). *Seeing things as they are: A theory of perception*. Oxford University Press.
- Seth, A. K. (2014). A predictive processing theory of sensorimotor contingencies: Explaining the puzzle of perceptual presence and its absence in synesthesia. *Cognitive neuroscience*, *5*(2), 97–118.
- Seth, A.K. (2017). From Unconscious Inference to the Beholder's Share: Predictive Perception and Human Experience, [Online], osf.io/preprints/psyarxiv/zvbkp.
- Siegel, S. (2006). Subject and object in the contents of visual experience. *The Philosophical Review*, *115*(3), 355–388.
- Stazicker, J. (2011). Attention, visual consciousness and indeterminacy. *Mind & Language*, *26*(2), 156–184.
- Simons, D. J., & Chabris, C. F. (1999). Gorillas in our midst: Sustained inattentive blindness for dynamic events. *Perception*, *28*(9), 1059–1074.
- Thorpe, S. J., Gegenfurtner, K. R., Fabre-Thorpe, M., & Bülthoff, H. H. (2001). Detection of animals in natural images using far peripheral vision. *European Journal of Neuroscience*, *14*, 869–876.
- Tye, M. (2002). *Consciousness, color, and content*. MIT Press.
- Tye, M. (2006). Non-Conceptual Content, Richness, and Fineness of Grain. In T. S. Gendler & J. Hawthorne (Eds.), *Perceptual Experience*. (pp. 504–530). Oxford: Oxford University Press.
- VanRullen, R., & Thorpe, S. J. (2001). The time course of visual processing: From early perception to decision-making. *Journal of Cognitive Neuroscience*, *13*, 454–461.
- VanRullen, R., & Thorpe, S. J. (2001). Is it a bird? Is it a plane? Ultra-rapid visual categorisation of natural and artificial objects. *Perception*, *30*(6), 655–668.
- Victor, J. D., & Conte, M. M. (2004). Visual working memory for image statistics. *Vis. Res.*, *44*, 541–546.
- von Helmholtz, H. (1867/1962). In J. P. C. Southall (Ed.), *Helmholtz's treatise on physiological optics*. New York: Dover.
- 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.

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