

Ensemble representation and the contents of visual experience

Tim Bayne¹  · Tom McClelland²

Published online: 23 January 2018

© The Author(s) 2018. This article is an open access publication

Abstract The on-going debate over the ‘admissible contents of perceptual experience’ concerns the range of properties that human beings are directly acquainted with in perceptual experience. Regarding vision, it is relatively uncontroversial that the following properties can figure in the contents of visual experience: colour, shape, illumination, spatial relations, motion, and texture. The controversy begins when we ask whether any properties besides these figure in visual experience. We argue that ‘ensemble properties’ should be added to the list of visually admissible properties. Ensemble properties are features that belong to a set of perceptible objects as a whole as opposed to the individuals that constitute that set. They include such features as the mean size of an array of shapes or the average emotional expression of an array of faces. Recent work in vision science has yielded compelling evidence that the visual system routinely encodes such properties. We argue that epistemological considerations provide strong reasons to think that these properties figure in visual experience. Judgements about ensemble properties are immediately warranted by our perceptual experience, and the only plausible way that a perceptual experience could confer this warrant is if it confers awareness of ensemble properties.

Tim Bayne and Tom McClelland have contributed equally to this paper, and the order of authors has been determined alphabetically.

✉ Tim Bayne
tim.bayne@gmail.com

Tom McClelland
tom.mcclelland57@gmail.com

¹ Philosophy Department, School of Philosophical, Historical and International Studies, Building 11, Monash University, Melbourne, VIC 3800, Australia

² Department of Philosophy, University of Warwick, Coventry CV4 7AL, UK

Keywords Visual perception · High-level content · Ensemble perception · Contents of consciousness · Perceptual warrant

1 Introduction

Of late philosophers of mind have been greatly exercised by a debate about the ‘admissible contents of perceptual experience’ (e.g., Bayne 2009, 2016; Brogaard 2013a; Carruthers and Veillet 2011; Fish 2013; Hawley and Macpherson 2011; Logue 2013; Lyons 2005; Masrour 2011; McClelland 2016; Nanay 2011; Price 2009; Prinz 2013; Reiland 2014; Siegel 2006, 2010). The issue, roughly put, concerns the range of properties that ‘figure in’ perceptual experience.¹ The focus of the debate has been very much on the contents of visual experience. It is relatively uncontroversial that colour, shape, illumination, spatial relations, motion, and texture can all figure in the contents of visual experience. The controversy begins when we ask whether any properties besides these are also visually experienced. If a clear case could be made for adding some class of properties to the list, this would be a valuable result. This paper makes just such a case for ‘ensemble properties’—features that belong to a set of perceptible objects as a whole as opposed to the individuals that constitute that set. Ensemble properties include such features as the *mean size of an array of shapes* or the *average emotional expression of an array of faces*. Recent research has yielded compelling evidence that the visual system routinely encodes such properties. Combining this with a number of philosophical considerations, we conclude that ensemble properties can (and often do) figure in visual experience.

2 Ensemble representation

Ensemble representation involves the representation of properties of an ensemble *as* properties of that ensemble. Alvarez offers the following useful exposition:

An ensemble representation is any representation that is computed from multiple individual measurements, either by collapsing across them or by combining them across space and/or time. For instance, any summary statistic (e.g. the mean) is an ensemble representation because it collapses across individual measurements to provide a single description of the set. (Alvarez 2011: 122)

Ariely’s (2001) investigation into discrimination of the mean size of visual objects is a classic study of ensemble representation. In two trials, subjects were presented with a set of dots for a 500 ms interval immediately followed by a test dot for a second 500 ms interval (see Fig. 1). On *member-identification trials* subjects were

¹ We focus here on the properties that are admissible with respect to normal states of human consciousness in typical environments, and leave to one side questions about what properties might be perceptually admissible *tout court*.



Fig. 1 Example trial in which subjects saw a set of 16 dots with a similarity factor of 1.4 followed by a single test spot (reproduced with permission Ariely 2001: 158)

asked whether the test dot was a member of the set of dots presented in the first interval, and on *mean-discrimination trials* subjects were asked whether the single test dot was smaller or greater than the mean dot size of the preceding set.

Ariely et al. found that subjects were unable to discriminate membership above chance, even though non-member test dots differed from member dots in size by at least 18%. However, they were able to tell whether the test dot was smaller or larger than the mean size of the presented dots. This capacity to discriminate the size of the test dot from that of the mean was found even when: the test dot's difference from the mean size was as low as 4%; the size of the set was large; and the variance of sizes in the set was great. These results led Ariely to conclude that '[o]bservers knew little or nothing about the sizes of the individual items in a set [but] encode quite precise information about the mean of a set' (2001: 160).

Ariely's study focused only on the representation of size, but subsequent studies have found evidence of ensemble coding for a range of other properties. Haberman and Whitney (2011), for instance, used a change-localization paradigm to establish ensemble representation of emotional expression. Subjects were briefly presented with an ensemble of 16 faces displaying a variety of expressions (see Fig. 2). Four of the faces in the set then changed from one extreme of emotional expression to the opposite extreme.

The experimenters found that:

Observers performed poorly when asked to locate any of the faces that changed (change blindness). However, when asked about the ensemble (which set was happier, on average), observer performance remained high. Observers were sensitive to the average expression even when they failed to localize any specific object change. (Haberman and Whitney 2011: 1)

A variety of paradigms have revealed ensemble effects for a wide range of properties, including: mean orientation (Parkes et al. 2001); mean location (Alvarez and Oliva 2008); mean gender of faces (Haberman and Whitney 2007); mean race of faces (Thornton et al. 2014); a crowd's mean direction of gaze (Sweeny and Whitney 2014); a crowd's mean direction of motion (Sweeny et al. 2012); and the mean degree of 'animacy' or 'lifelikeness' of a set of stimuli (Leib et al. 2016). For all of these properties, the data highlight the astonishing speed and reliability with which averages are computed and, most strikingly, that subjects can consciously

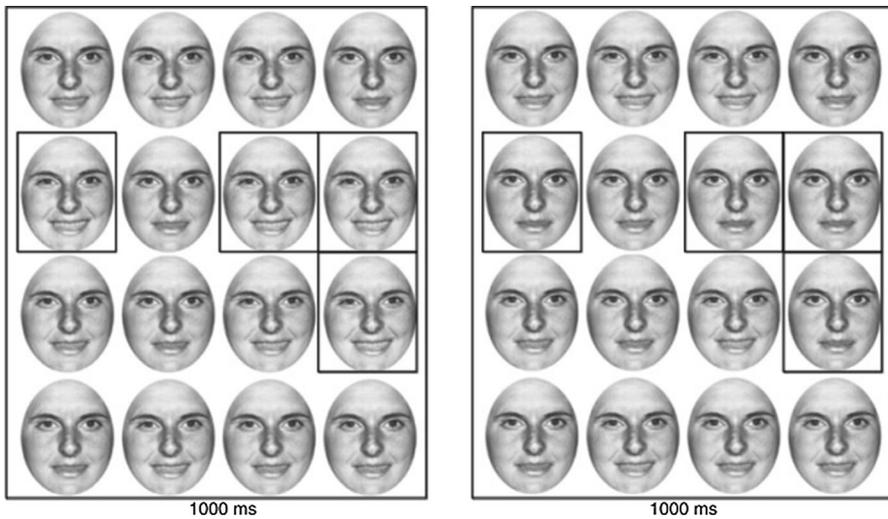


Fig. 2 Example stimuli. Sets were displayed successively for 1000 ms each, separated by a 500-ms interval. On each trial, observers had to indicate (1) which set had the happier average expression (two-interval forced choice, 50% guess rate) and (2) any one of the four items that changed between the two sets (indicated here by the black outlines, not seen by participants; 25% guess rate). (Reproduced with permission Haberman and Whitney 2011: 3)

access ‘summaries’ of a set of stimuli even when they appear to lack conscious access to information about the individual members of that set.²

Is ensemble representation genuinely *perceptual*? The papers in which this research is reported certainly give that impression, for perceptual talk is ubiquitous in this field. Consider the very title of Ariely’s (2001) ground-breaking paper: “Seeing Sets: Representation by Statistical Properties”. In this paper we use the expressions ‘ensemble coding’, ‘ensemble representation’ and ‘ensemble processing’ in order not to prejudice the very issue that is at stake in this debate, but it is noteworthy that the scientists who study this phenomenon typically refer to it as ‘ensemble *perception*’. Any non-perceptual treatment of ensemble coding would be at odds with the received view within cognitive science (see also Cohen et al. 2016). Indeed, vision scientists seem to regard the perceptual account as *uncontroversial*, for they rarely even attempt to justify their use of perceptual language in connection with ensemble representation.

Nonetheless, the perceptual treatment of ensemble representation clearly *is* controversial from the perspective of the admissible contents debate. After all, ensemble properties do not appear on any of the lists that philosophers of perception provide of those properties that are obviously perceptually admissible (see e.g. Brogaard 2013a, b; Hawley and Macpherson 2011; Logue 2013; Lyons 2005; Masrour 2011; Price 2009; Siegel 2010); indeed, to date ensemble properties have

² Although we focus on ensemble perception in vision, studies have indicated that ensemble representation can also be found in audition (e.g. Albrecht et al. 2012) and it is an open possibility that there is ensemble perception in other modalities.

not even figured in lists of properties that are *arguably* perceptual. Thus, if it could be shown that the properties represented in ensemble perception enter into the contents of perceptual experience, then this would be a valuable result. How might the perceptual treatment of ensemble representation be defended?

One approach that might be considered involves an appeal to the method of phenomenal contrast (Siegel 2006, 2010), in which one focuses on two experiential scenarios that are alleged to differ in experiential character. The idea is that the phenomenal contrast between them is best explained by supposing that one of the scenarios involves the visual representation of a certain kind of ‘high-level’ property, whereas the other does not. Employing this method, Siegel (2006) argues that the experiential changes that occur when one learns to recognize pine trees by sight is explained by the fact that the process of perceptual learning involves the acquisition of the capacity to visually represent the property of being a pine tree. Bayne (2009) deploys the method of phenomenal contrast to argue that in certain forms of visual agnosia visual experience is no longer able to represent artifactual kinds, such as being a telephone or a watch.

We doubt that the phenomenal contrast method can be successfully employed in defence of a perceptual account of ensemble representation. For one thing, previous applications of the method have been highly controversial, with detractors arguing that the phenomenal changes that result from perceptual learning or visual agnosia can be fully accounted for in terms of changes in the representation of purely ‘low-level’ features (see e.g. Hansen in press; Prinz 2013). More importantly, the very nature of ensemble representation suggests that it might not be possible to formulate a phenomenal contrast argument even if ensemble coding is perceptual. To construct a plausible phenomenal contrast argument one needs two scenarios, one of which contains only the low-level features, and one of which contains both the relevant low-level features and the target high-level feature. In other words, the method requires that the target high-level feature is dissociable from the relevant low-level features that typically accompany it. For example, Siegel’s contrast argument relies on the fact that the capacity to recognize pine trees is learnt, while Bayne’s contrast argument relies on the (alleged) fact that the capacity to recognize objects (such as telephones) can be selectively lost. However, it is not at all clear that the capacity for representing ensemble properties *is* dissociable from the capacity to represent the low-level features that underpin ensemble representation. If the capacity for ensemble representation is not learnt, and if it cannot be selectively lost, then the preconditions for a contrast argument cannot be met.

Another approach that one might pursue in defending a perceptual treatment of ensemble coding would be to look at the functional profile of such representations, and to argue that they have the functional profile associated with perception as opposed to cognition. For example, one might argue that ensemble processing is perceptual on the grounds that it is extremely fast (Chong and Treisman 2003), mandatory (Ball and Sekuler 1980; Corbett and Melcher 2014; Dubé and Sekuler 2015), cognitively impenetrable (Sweeny et al. 2012), and processed in parallel (Whitney et al. 2014).

The problem with this line of argument is that it is not clear that these features—either individually or conjunctively—are unique to perception. The functional

feature that arguably provides the strongest evidence for a perceptual view of ensemble coding is adaptation. Adaptation effects have been found for a wide range of ensemble properties, including average orientation (e.g., Gibson and Radner 1937), average direction of motion (e.g., Wataminiuk and McKee 1998), average texture density (Durgin 1995, 2008; Durgin and Huk 1997), numerosity (Burr and Ross 2008) and mean size (Corbett et al. 2012). Given that both philosophers (e.g. Block 2014; Burge 2014; Fish 2013) and psychologists (e.g. Whitney et al. 2014; Corbett et al. 2012) often treat adaptation as a defining mark of the perceptual, this work surely provides significant motivation for treating ensemble representation as perceptual. However, appeals to adaptation will fail to convince those who have doubts about whether adaptation is a defining feature of perception, and our own view is that it is very much an open question whether adaptation is a mark of perception. (One could of course stipulate that showing adaptation suffices for being perceptual, but that move threatens to turn the debate about whether ensemble representation is perceptual into a purely verbal dispute).

Appeals to the functional profile of ensemble processing—its speed, mandatory-ness, cognitive impenetrability, parallel nature, and (in particular) the fact that it seems to be subject to adaptation—provide *prima facie* support for the perceptual account, but they fall short of providing a fully comprehensive case for the perceptual view. We turn now to an argument for the perceptual account that we regard as more compelling.

3 Epistemic justification and the perception of ensemble properties

This section presents an epistemic argument for the perceptual view of ensemble processing. At the heart of this argument is the claim that the contents of perceptual experience warrant certain kinds of judgements, and that by identifying what judgements are warranted by an experience we can draw inferences about the properties it represents.

We can begin by distinguishing between two ways in which perceptual experience confers warrant: *mediately* and *immediately*. Siegel and Silins characterise mediate justification as follows:

Your experience *E* gives you mediate justification to believe that *p* just in case *E* gives you justification to believe that *p*, in a way which depends on your having justification to believe some proposition, from some source other than *E* (2015: 785)³

Seeing the position of the fuel gauge, for instance, warrants the judgement that the car is running out of petrol in a way that depends on having warranted beliefs about the relationship between the marks on the gauge and the amount of petrol in the tank (Pryor 2005). One's judgement is thus *mediately* warranted by one's perceptual

³ Although some theorists distinguish epistemic warrant from epistemic justification, we take Siegel and Silins' notion of justification to map fairly directly onto our notion of warrant.

state. Sometimes, however, perception provides one with *immediate* warrant. As Siegel and Silins put it, “*E* gives you immediate justification to believe that *p* just in case *E* gives you justification to believe that *p* that is not mediate justification to believe that *p*” (2015: 785). For example, having a visual experience as of a red object warrants the judgement that the object before one is red in a way that does not depend on the warrant of one’s prior beliefs (Pollock 1986).⁴

One might worry that in such a situation one might not *know* that the object before one is red. (Perhaps the lighting conditions generate illusions of redness; perhaps one is hallucinating.) However, this kind of possibility needn’t concern us as our target is *warrant* (or justification) rather than knowledge. A warranted belief constitutes knowledge only if certain background conditions are met. Our concern is with the *prima facie* warrant conferred by one’s perceptual experience, and not with the satisfaction of the further conditions required for knowledge. In visually experiencing an object as red one has a defeasible reason to believe that the object before one is red. If one also believes that the lighting conditions are poor or that one is hallucinating then this reason might be defeated, but in the absence of defeaters one is justified in taking oneself to be looking at a red object.

Equipped with the notion that perceptual experiences can immediately confer warrant on certain beliefs, let us consider what this means for the *contents* of perceptual experience. It is plausible that perceptual experiences confer warrant on beliefs in virtue of their content.⁵ Consider again the experience as of a red object. This perceptual experience warrants the belief that there is a red object before one in virtue of having the content *there is a red object before one*. (Or perhaps: *this particular object is red*.) Moreover, it confers this warrant *immediately*.

Thus far we have focused on the idea that knowing the contents of a subject’s perceptual experience tells us something about what judgements that subject is immediately warranted in making on the basis of that perceptual experience. The crucial thing for our purposes is that it is also possible to make an inference in the other direction: knowing what judgements a subject is immediately warranted in making on the basis of her perceptual experience can tell us something about the contents of her perceptual experience. Because of this, one can use the judgements that a subject is immediately warranted in making as a guide to the contents of her perceptual experience.

The idea that there is an intimate connection between epistemic warrant and the contents of consciousness is not a new one. It can be seen, for example, in Dretske’s epistemic test for awareness. “The rough idea...is that if you can see (and thereby know) that *x* is *F* for some value of *F*, you must be aware of *x*. You can’t know, by

⁴ Some deny that our judgements are ever justified in this way, instead claiming that warrant is only ever conferred immediately (e.g. Davidson 1986). Reflection on specific cases might be taken to rule out such a view, and a priori considerations also count against it (Pryor 2005). Worries about an infinite regress of justification, for instance, suggest that some of our beliefs must be justified immediately, and beliefs warranted purely by one’s perceptual experience are the best candidates to play such a foundational role.

⁵ Perhaps there are experiences that lack content yet still confer immediate warrant: one might regard a pain as a contentless state that immediately warrants the belief that one is in pain (Siegel and Silins 2015: 793). Nevertheless, the point remains that where a perceptual experience does have content, the content it has is relevant to what beliefs it justifies.

seeing, that x is F without being conscious of x .” (2007: 220). However, where Dretske targets perceptual awareness of objects, our concern is primarily with the perceptual awareness of properties.

A view that more closely resembles our own is Johnston’s:

... it is distinctive of what I have been calling immediate perceptual judgment that what one judges or predicates of an item is some feature of which one is also aware. So, when I taste the astringency of the calvados, I am not only aware of the calvados—a certain liquid in my mouth—but I am also aware of its astringency. (2006: 280)

Although it isn’t explicit in this passage, Johnston also emphasizes that when one perceptually experiences some property of an object, one is *warranted* in making a judgement that attributes that property to the object. This gives us the following principle:

Johnston’s Maxim: If one’s judgement that x is ϕ is warranted and immediate, then one perceptually experiences x and its ϕ -ness.

We will consider the merits of Johnston’s Maxim shortly, but three clarificatory remarks must first be made.

First, Johnston’s Maxim is neutral on whether perceptual experiences have propositional content. Our discussion thus far has been framed in terms of the judgement *that p* being immediately warranted by a perceptual experience with the content *that p* , but Johnston’s Maxim does not presuppose a propositional account of perceptual content. We leave open the possibility that perceptual content is non-propositional, and that a subject represents *that the calvados is astringent* only in judgement. Second, Johnston’s Maxim is neutral on whether or not the contents of experience are world-involving. Johnston himself has a model of perception in which sense experience ‘takes in’ worldly items, but Johnston’s Maxim does not itself entail direct realism. Third, Johnston’s Maxim is neutral on whether or not the subject *sees* the objects and properties in question. Our target is visual experience as such, and talk of seeing raises a number of issues that we wish to avoid, such as debates about the facticity of seeing.

Armed with Johnston’s Maxim, we are now in a position to return to ensemble representation. In the studies discussed, subjects make judgements about the ensemble properties of an array of objects. For convenience we will focus on studies of mean size, though our argument can be adapted for any of the ensemble properties that have been examined. Let e be some ensemble of presented circles and F be the ensemble property of having such-and-such a mean size. The subject judges *that e is F* , and we will stipulate that the case in question is one in which mean size is judged accurately. Our argument is as follows:

1. The subject’s judgement that e is F is warranted.
2. The subject’s judgement that e is F is immediate.
3. If a judgement that x is ϕ is both warranted and immediate, then one perceptually experiences x and its ϕ -ness.
4. Therefore the subject perceptually experiences e and its F -ness.

Call this ‘the argument from warrant’. This argument yields two noteworthy results. First, it entails that ensembles are perceptible objects. This fits with data suggesting that during the process of ensemble perception the visual system parses arrays of individual objects into groups, suggesting that the group as a whole is a target for property-attribution (e.g. Oriet and Brand 2013). Although one might worry that an ensemble isn’t best described as an object, we should remember that in this context talk of ‘objects’ is rather loose. (Although shadows, after-images and rainbows might not qualify as objects in other contexts, they do count as perceptual objects.) There is thus nothing untoward about the claim that ensembles are perceptual objects.

The more dramatic consequence of the argument from warrant is that ensemble *properties* are perceptible. If true, this would mean that in addition to visually experiencing size, orientation and facial expression (for example), we can also visually experience *mean* size, *mean* orientation and *mean* facial expression. Properties of this kind would be a significant addition to the list of properties that are regarded as perceptually admissible.

4 Objections and replies

We turn now to consider a number of objections to the argument from warrant. The first two objections will prove to be relatively easy to address, but the third will take us into more uncertain terrain.

4.1 Are ensemble judgments warranted?

Johnston’s Maxim applies only to cases in which perceptual judgement is warranted. Thus, if doubt can be cast on the warrant for ensemble judgements, a critic could then argue that Johnston’s Maxim is irrelevant, and thus that there are no grounds for our conclusion.

How might one cast doubt on the idea that ensemble judgments are warranted? One line of argument would be to suggest that the processes underlying ensemble representation are not reliable. After all, many have argued that the epistemic status of a judgment is constitutively connected to the reliability of the process by means of which it was formed (e.g. Goldman 1986).

However, there are two problems with this line of criticism. First, the claim that warrant is constitutively connected to reliability is controversial, and internalists of various stripes reject it (see e.g., Pollock 1986). Second, appeals to reliabilism are unlikely to provide much encouragement to the critic, for one of the striking features of ensemble perception is its reliability. Chong and Treisman (2013) found that subjects were as good at identifying which of two sets of circles had the greater mean size as they were at identifying the larger of two individual circles. Since we would certainly regard typical judgements of comparative individual size as warranted, it’s hard to deny that judgements of comparative mean size can also be warranted.

Pointing to the phenomenon of blindsight, a critic might suggest that ensemble judgements could be unwarranted even if they are underwritten by a reliable process. Dretske (2007) suggests that although blindsight patients are able to make accurate reports about some of the properties of stimuli that are presented in the relevant ‘blind’ part of their visual field, their reports will not be warranted because they are not grounded in perceptual experience. Inspired by Dretske’s comments, a critic of the argument from warrant might argue that even if ensemble judgements are reliable, they are unwarranted because they are not appropriately grounded in perceptual experience.

We agree with Dretske that the (presumably) unconscious nature of the representations that underlie blindsight leaves patients in the unusual situation of having reliably-formed-but-nonetheless-unwarranted judgements. However, it seems unlikely that ensemble judgements are relevantly analogous to the judgements of blindsight subjects. One of the considerations that counts against the warrant of blindsight judgements is that they are ‘guesses’ which subjects typically make only in forced-choice conditions. Although many ensemble studies do employ forced-choice conditions, others do not (e.g. Alvarez et al. 2009) and in these studies subjects willingly and reliably make judgements about ensemble properties. Furthermore, blindsight subjects typically report that they are just guessing, and they deny having any reason for making the judgements they make. In contrast, some ensemble studies specifically probe whether subjects take themselves to have experienced the relevant stimuli. In Bronfman et al. (2014) subjects were given a primary task of reporting a post-cued letter in an array but were also asked to report the colour diversity of a row of letters (a task that involves a form of ensemble representation). Bronfman et al. explain that ‘...participants were instructed to press an escape button in case they had no impression of the noncued letters’ colors and were discouraged from guessing the color diversity’ (2014: 1397). This strongly suggests that the (accurate) judgements subjects made regarding colour diversity were based on their visual experiences. These points are hard to reconcile with the claim that subjects are making unwarranted guesses.

For those who remain unconvinced by these points, we recommend performing some of the tasks from ensemble studies for themselves. We predict that far from guessing, you will be able to make the required ensemble judgements confidently.

4.2 Is the warrant for ensemble judgments immediate?

Johnston’s Maxim applies only to judgements that are both warranted *and* immediate. Thus, there is room for a critic to argue that although ensemble judgements are warranted their warrant is not immediate, but instead derives from background beliefs just as the warrant for beliefs about the amount of fuel in one’s petrol tank does. What can we say in response to this objection?

The first point to make is that the critic needs to tell us a plausible story about what the subject perceptually experiences that, in combination with relevant background beliefs, warrants her ensemble judgement. Since the critic denies that the subject perceptually experiences the ensemble and its mean size, she is committed to the claim that the subject’s ensemble judgement is (mediately)

warranted by her perceptual experience of *individuals* in the ensemble and their respective sizes. The question is, “Which individuals?” As we see it, there are only two broad options here.

The first option (“SAMPLE”) holds that the subject perceptually experiences only those items that belong to a small sample of the ensemble, and then uses that sample to make a (warranted) judgement about the ensemble as a whole. SAMPLE fits with the idea that perceptual experience has a limited bandwidth, allowing us to visually experience only a handful of items (say 3–4) at any given time. The problem with SAMPLE, however, is that it cannot be reconciled with the data. Although it was once suggested that sampling strategies of this kind might explain the performance of subjects in ensemble perception studies (e.g. Allik et al. 2013, 2014; Marchant et al. 2013), that view is now regarded as having been conclusively rebutted, for ensemble judgements are simply too accurate to be explained in terms of limited samples (Whitney et al. 2014; Utochkin and Tiurina 2014; Haberman and Whitney 2012; Brady and Alvarez 2015; Albrecht et al. 2012).

Of course, an advocate of SAMPLE might insist that there *is* a sampling strategy that explains the data but that it hasn’t yet been identified. However, even if such a strategy can be identified, SAMPLE faces a deeper problem. In Haberman and Whitney’s (2011) study outlined in Sect. 2, subjects were good at judging changes to the mean expression of a group of faces. Any version of SAMPLE that could explain these data would need to posit a sampling strategy that includes at least one of the four faces in the array that changed expression. The subject would then need to have a warranted belief to the effect that *if the expression of a sampled face becomes happier, then the mean expression of the ensemble will have become happier*. But Haberman and Whitney found that subjects accurately identified changes in mean emotion even on trials where they failed to identify any of the specific faces that had changed. This result is extremely hard to reconcile with the suggestion that the subject perceptually experienced at least one of the changed faces.

Our view is in a far better position to accommodate these data. Subjects are good at making accurate ensemble judgements because their perceptual experience is sensitive to the properties of all (or at least most) of the ensemble’s members—not just to the properties of a small sample of the ensemble.⁶ Subjects can detect changes to the mean emotional expression of an ensemble in virtue of their perceptual experience of the ensemble’s mean emotional expression. Thus, subjects are warranted in believing that the mean emotion of an ensemble has changed even if they don’t perceptually experience a change in the emotional expression of any *particular* face. Assuming that what’s true of ensemble processing of emotional expression is also true of ensemble processing of other kinds, we conclude that the warrant had by ensemble judgements cannot plausibly be conferred mediately by perceptual experiences of small samples of the ensemble.

This leaves the critic of premise 2 with a second option—what we will call ‘OVERFLOW’ (McClelland and Bayne 2016). According to OVERFLOW, we

⁶ We add the qualifier ‘or most’ because the data suggests that ensemble perception is often based on a sample of most, but not all, of the ensemble (Alvarez 2011). The key point is that the sample far exceeds the capacity constraint of 3–4 items to which conscious perception is thought to be subject.

phenomenally experience all (or at least most) of the objects in the ensemble, but we have *cognitive access* to only a handful of those items. On this view, even when we fail to make accurate judgements involving the properties of an ensemble member, we nevertheless *perceptually experience* that individual and its properties. Our perceptual experience then confers warrant on our ensemble judgements as follows: we perceptually experience all (or most) of the circles in the array and their size; we have an independently warranted belief that if the members of an ensemble have such-and-such individual sizes, then the ensemble as a whole has such-and-such a mean size; and we then form the mediately warranted judgement that the ensemble as a whole has such-and-such a mean size.

OVERFLOW avoids the problems faced by SAMPLE. Not only can it accommodate the reliability of ensemble perception, it can also accommodate the results of the aforementioned emotional expression study. When subjects detect a change in mean emotional expression they perceptually experience a change in the emotional expression of at least one *individual* face in the array, giving them mediate warrant for the judgement that the emotional expression of the ensemble has changed, but where the subject's experience of the change in the emotional expression of an individual face *overflows* that to which they have cognitive access, they will be unable to identify which face has changed.

Despite these advantages, OVERFLOW faces serious difficulties. One difficulty concerns the very suggestion that experience can overflow cognitive access. Although this proposal has been influentially defended (e.g., Block 2007, 2011), it remains highly controversial (e.g. Kouider et al. 2010; Cohen and Dennett 2011; Phillips 2016; Ward et al. 2016). Secondly, even if we allow that phenomenology can overflow cognitive access, it is highly plausible to suppose that perceptual experiences can confer warrant only to the extent that we have cognitive access to them. And if that is right, then appeals to cognitively inaccessible experiences cannot account for the warrant that our ensemble judgements enjoy. (By contrast, the view that we have defended has a clear story as to how warrant is conferred, for the perceptual experiences that our account appeals to *are* cognitively accessible.) We conclude that premise 2 withstands scrutiny.

4.3 Is Johnston's maxim plausible?

The third premise of our argument does much of the heavy-lifting, and is an obvious target for a critic. Here, a critic might simply reject Johnston's Maxim, acknowledging that the ensemble judgement in question is immediately warranted but denying that this brings with it any commitments regarding the contents of perceptual *experience*. On this view, although ensemble representation is *unconscious*, it is able to provide ensemble judgments with immediate warrant. We call this view 'NON-EXPERIENTIALISM'.⁷

How does NON-EXPERIENTIALISM relate to the suggestion discussed earlier that ensemble perception is analogous to blindsight? The idea is this. Like blindsight

⁷ A version of this view is adopted by Lyons (2005).

judgements, ensemble judgements are based on a reliable perceptual process that is non-phenomenal. Unlike blindsight judgements, however, ensemble judgements are based on perceptual information to which the subject has direct, cognitive, access. Subjects are able to accurately report on the apparent ensemble properties of an array, and they do not take themselves to be guessing when they do so. NON-EXPERIENTIALISM claims that having such access to perceptual representations of ensemble properties is enough to warrant ensemble judgements, despite the fact that ensemble properties do not figure in one's perceptual experience. What might we say about NON-EXPERIENTIALISM?

Our first response is to insist that Johnston's Maxim shouldn't be rejected lightly. Johnston's Maxim is premised on a credible view of how perceptual judgements acquire their warrant. It tells a plausible story about the role of phenomenal consciousness in our mental economy: namely, that it functions as an 'epistemologically enabling condition' (Dretske 2007: 222). Only by being phenomenally conscious is a perceptual representation able to confer warrant on our judgements. (This thought underwrites the view that the reliable guesses of blindsight patients are unwarranted, for they are not appropriately grounded in perceptual experience.)

In response, the critic might argue that Johnston's Maxim should be judged against its ability to accommodate cases of immediately warranted perceptual judgement, and that instances of ensemble judgement are clear counter-examples to Johnston's Maxim. We thus arrive at something of an impasse: we regard Johnston's Maxim as highly plausible, and take it to provide motivation for a perceptual account of ensemble representation; our imagined critic thinks it unlikely that ensemble properties could be perceptually represented, and takes that as a reason to reject Johnston's Maxim. How might we break this dialectical impasse?

One thing we might consider is whether there are other instances in which Johnston's Maxim might be thought to deliver the wrong verdict. If the maxim is found to be untenable in other domains, there is little point in appealing to it here.

Silins (2011) considers two potential counter-examples. The first involves a case in which one sees a red object in good conditions. One's perceptual experience immediately warrants the judgement that the object before one is red. Applying Johnston's Maxim, we get the conclusion that the perceptual experience must represent the object *as red*. However, one might think that it is more plausible to suppose that the experience represents only that the object has a determinate shade of red, such as *red₂₁*. And if that's right, then perceptual experience would provide immediate warrant for the judgement that an object instantiates a property that is not itself perceptually represented.

The second potential counter-example concerns demonstrative beliefs about one's environment. Consider a case in which one sees a beach ball in good conditions and forms the judgement *that is round*. The belief is not that *some* object is round, but that *that particular* object is round. Applying Johnston's Maxim, we get the conclusion that the perceptual experience must represent *that particular object*. However, on many accounts perceptual experience does not represent particular objects as such (see e.g. Davies 1992; McGinn 1997). Again, we have here a result that would be at odds with Johnston's Maxim (Silins 2011: 352).

Although these cases undermine Johnston's Maxim as stated, the Maxim can be easily modified to take them into account. The key here is to appeal to what Silins calls *indirect awareness*. He introduces this notion as follows:

...we can be immediately justified by experiences in believing indirect contents of experiences. A content that p of an experience is indirect if the experience has the content that p , and has the content that p at least in part in virtue of having some other content that q . (2011: 354)

Regarding the colour case, we might say that one's perceptual experience indirectly represents the object's redness in virtue of directly representing its red₂₁ness. Regarding the beach ball case, we might say that one's perceptual experience indirectly represents that a particular beach ball is round in virtue of generically representing the presence of a round object. Although Silins himself rejects all versions of Johnston's Maxim (see especially 2011: 355) we can use his notion of indirect contents to refine the maxim as follows:

Johnston's Maxim*: If one's judgement that x is ϕ is warranted and immediate, then x and its ϕ -ness figure in one's perceptual experience either directly or indirectly.⁸

We have seen that Johnston's Maxim can be modified to handle objections to it. However, matters are even worse for our imagined critic than this, for careful reflection on NON-EXPERIENTIALISM indicates how strange it really is. The critic of the perceptual account of ensemble coding holds that averages are not the kind of property that can be perceived. NON-EXPERIENTIALISM purports to protect that claim insofar as it denies that ensemble properties figure in *perceptual experience*. However, NON-EXPERIENTIALISM is itself committed to the claim that ensemble properties figure in *non-experiential perceptual representations*. (Remember, NON-EXPERIENTIALISM proposes that our ensemble judgements are grounded in perceptual representations of ensemble properties that are cognitively accessible but non-phenomenal.) So according to NON-EXPERIENTIALISM, averages *can* be perceived, albeit non-phenomenally. This combination of views seems to us to be decidedly odd. After all, what reason could there be for thinking that a certain class of properties can be visually represented but not visually experienced?

Is it possible to put pressure on Johnston's maxim without resorting to NON-EXPERIENTIALISM? Recent discussion of 'seemings' might be taken to suggest a positive answer to this question (Bergmann 2013; Brogaard 2013b; Cullison 2011; Pace 2017; Reiland 2014, 2015; Tucker 2010).⁹ Advocates of seemings hold that how the world appears to us in experience is constituted by two kinds of states: a layer of perceptual 'experience' that represents exclusively low-level properties and a layer of perceptual 'seeming' that is grounded in the contents of perceptual experience and can represent high-level properties.¹⁰ This view allows one to adopt a qualified version of Johnston's Maxim according to which if one's judgement that

⁸ Although Silins assumes a propositional view of the contents of experience, our adaptation avoids such a commitment.

⁹ We are grateful to a reviewer for prompting us to consider how the advocate of perceptual seemings might account for ensemble representation.

¹⁰ This sub-set of seeming-states is labelled in different ways by different authors. We intend 'perceptual seemings' to encompass all these different views.

x is ϕ is warranted and immediate, then x and its ϕ -ness figure (either directly or indirectly) in one's perceptual experiences *or in one's perceptual seemings*.

This qualified version of Johnston's Maxim opens up an alternative understanding of ensemble representation according to which ensemble properties figure in perceptual seemings but *not* in perceptual experience. We call this view SEEMING. It is distinct from NON-EXPERIENTIALISM insofar as seemings are phenomenal states, albeit a kind of phenomenal state that is distinct from perceptual experience.

We have three objections to SEEMING. Firstly, it is unclear what precisely perceptual seemings *are*. According to one advocate of the notion, perceptual seemings are "*sui generis phenomenal events that are passive, conceptual, and represent objects as having properties*" (Reiland 2014: 180; emphasis in original). But this fails to clearly demarcate seemings from perceptual experiences, for perceptual experiences are also passive and represent objects as having properties. It is debateable whether perceptual experience has 'conceptual content', but one certainly wouldn't want to rule out that possibility by fiat given the existence of prominent conceptualist accounts of perception (e.g. McDowell 1994; Brewer 2006). It is also unhelpful to suggest that seemings can represent high-level properties whereas perceptual experience can represent only low-level properties (Reiland 2014), for whether perceptual experience can represent high-level properties is precisely what is up for debate here. Without a robust distinction between perceptual experience and perceptual seemings, there is a real possibility that SEEMING is not genuinely distinct from our own view. After all, SEEMING is in line with our claims that: (a) we have phenomenal states that represent ensemble properties; and (b) that these phenomenal states are not judgements. Substantive disagreement between our own view and SEEMING depends on whether the advocates of SEEMING can motivate the claim that seemings are a distinct kind of state that differ from perceptual experience on the one hand and perceptually-based judgement on the other. It is not clear to us that that has been done—or even that it can be done.¹¹

Second, even if a clear conceptual distinction can be made between perceptual seemings and perceptual experiences, we are unconvinced by the arguments that have been given for positing perceptual seemings. A number of theorists have claimed that we need to posit perceptual seemings in order to account for the fact that the lines of the Müller-Lyer illusion seem to be unequal despite the fact that one knows they are equal (e.g., Brogaard 2013b; Pace 2017; Reiland 2015) But those of us who are representationalists about visual experience can account for this 'seeming' without positing a new kind of mental category, for we take this seeming to be a matter of how visual experience represents the lines in question.

¹¹ Another point of obscurity is how perceptual seemings are meant to be related to other kinds of seemings. Do perceptual seemings fall under the same umbrella as epistemic seemings, for instance, or is this association merely nominal? Yet another issue is how perceptual seemings are meant to be related to perceptual experiences. Are we to understand seemings and experiences as two layers of a single representational state or as discrete representations? Our impression is that the advocates of 'seemings' do not have a shared view of what the answers to these questions are.

Third, even if we grant the proposed bifurcation of perceptual appearance into an ‘experiential’ component and a ‘seeming’ component, it’s unclear to us that ensemble properties belong on the seeming side of this divide as opposed to the perceptual experience side of it. For one thing, it seems highly implausible to think of ensemble representations as conceptual, for surely one needn’t have the concept of mean size in order to be visually sensitive to the mean size of a set of objects. Moreover, the advocates of perceptual seemings typically hold that the contents of perceptual seemings must be grounded in the contents of perceptual experience (see especially Reiland 2015; Pace 2017). It seeming to one that x is a pine tree for example must be grounded in a perceptual experience of the low-level properties characteristic of pine trees. In the case of ensemble perception though, there is no plausible story to be told about how perceptual experience might ground ensemble seemings. If perceptual experience represents only a sample of objects we run into the problems raised earlier against *SAMPLE*, but if it somehow represents the values of all members of an array we run into the problems raised earlier against *OVERFLOW*. We conclude that neither *NON-EXPERIENTIALISM* nor *SEEMING* undermine premise 3 of the argument from warrant.

5 Concluding remarks

We have argued that ensemble properties such as mean size, mean orientation and mean emotional expression deserve a place next to colour, shape, illumination, spatial relations, motion, and texture on the list of properties that can figure in visual experience. By way of conclusion, let us consider two issues that this conclusion raises.¹²

The first issue concerns the relevance of ensemble perception for the debate about ‘high-level’ perceptual content. Does ensemble content qualify as a kind of high-level content, or is it instead merely a kind of low-level content, albeit one that has been largely overlooked?

One might argue that ensemble properties are high-level on the grounds that they are abstract in the way that canonical high-level properties are. If—as high-level theorists argue—we do indeed see tomatoes as tomatoes and trumpets as trumpets, then we do so only by being visually sensitive to a number of other properties, such as form, colour, texture and so on. Similarly, we see a set of objects as having a mean size of such-and-such in virtue of being visually sensitive to the size of its members. On the other hand, ensemble properties are unlike canonical high-level properties (and more akin to canonical low-level properties) in that the notion of a perfect *doppelgänger* struggles to get any real grip here. Something can look like a tomato without actually being a tomato, just as something can look like a trumpet without actually being a trumpet. (Even if *being a tomato* and *being a trumpet* can be represented in perceptual experience, the concepts <tomato> and <trumpet> are not observational concepts in the way that <red> and <square> arguably

¹² We thank a reviewer for prompting us to say more about these topics.



Fig. 3 A scene from Alfred Hitchcock's *Strangers on a Train*

are.) But it is doubtful whether something can look to have a certain mean size without actually having that mean size.¹³ Thus, it is natural to group ensemble properties with high-level properties in certain respects, and it is natural to group them with low-level properties in other respects. We are inclined to think that the high-level/low-level contrast is not sufficiently precise for the question of whether ensemble content is high-level or low-level to have a determinate answer.¹⁴

The second question raised by our discussion is this: if ensemble properties figure in perceptual experience, why are they not generally recognized as such? Why does it take careful empirical research—not to mention philosophical argument—to draw attention to them?

In response to this question, one might be tempted to suggest that although ensemble properties figure in the kinds of perceptual experience that occur in ensemble coding experiments, they only rarely (if at all) figure in ordinary perceptual experience. Thus, this line of thought continues, it is no wonder that their existence is typically overlooked.

This suggestion might provide a plausible solution to our puzzle if it were indeed reasonable to assume that ensemble perception is restricted to laboratory contexts. However, we view that suggestion with suspicion, for it seems highly plausible to us that ensemble perception is integral to everyday experience. Consider the figures reproduced below from the tennis scene of Alfred Hitchcock's *Strangers on a Train* (Fig. 3). Here, one immediately sees that a particular spectator's direction of gaze is very different to that of his fellow spectators.

Not only does ensemble perception facilitate the detection of outliers, it also helps us to track the behaviour of a crowd (Sweeny and Whitney 2014), and facilitates the detection of scene gist (that is, the kind of environment that one is in) (Torralba and Oliva 2003) and the perception of patterns of texture (Haberman and Whitney 2009). As Cohen et al. (2016) point out, ensemble perception promises to

¹³ One can of course misperceive ensemble properties (just as one can misperceive any property), but the point at issue here is not about misperception.

¹⁴ Some authors (e.g. Logue 2013) characterise high-level properties as any properties that aren't on the list of properties that are uncontroversially perceived. On such a characterisation ensemble properties trivially come out as high-level. However, this view would leave substantive questions about how these properties compare with canonical high-level properties entirely unanswered.

reconcile the sense of richness that accompanies everyday perceptual experience with the fact that vision has a very narrow capacity of no more than 4 or so objects.

A better response to the puzzle that we have raised appeals to the fact that we rarely form explicit ensemble judgments. Explicit thought (and talk) about the shapes, sizes, colours and locations of objects is ubiquitous, whereas explicit thought (and talk) about the mean properties of a set of objects is rare—indeed, it may be largely restricted to experimental contexts. This fact, we suggest, goes some way towards accounting for the neglect of ensemble properties. But although ensemble properties are rarely the objects of *explicit* thought (and talk), representations of them do guide our behaviour in ways that are distinctive of perceptual content, for they alert us to important features of an environment, such as the direction in which a crowd is running, or which members of a herd are significantly slower or smaller than the rest.

Acknowledgements Thanks to Jakob Hohwy, Eric Mandelbaum, David Whitney and an anonymous reviewer for their extremely helpful comments on this paper. This research was supported by an Australian Research Council Future Fellowship (FT150100266) (Bayne) and a Leverhulme Early Career Fellowship (McClelland).

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Albrecht, A. R., Scholl, B. J., & Chun, M. M. (2012). Perceptual averaging by eye and ear: Computing summary. *Attention, Perception, and Psychophysics*, *74*, 810–815.
- Allik, J., Toom, M., Raidvee, A., Averin, K., & Kreegipuu, K. (2013). An almost general theory of mean size perception. *Vision Research*, *83*, 25–39.
- Allik, J., Toom, M., Raidvee, A., Averin, K., & Kreegipuu, K. (2014). Obligatory averaging in mean size perception. *Vision Research*, *101*, 34–40.
- Alvarez, G. A. (2011). Representing multiple objects as an ensemble enhances visual cognition. *Trends in Cognitive Science*, *15*(3), 122–131.
- Alvarez, G., & Oliva, A. (2008). The representation of simple ensemble visual features outside the focus of attention. *Psychological Science*, *19*(4), 392–398.
- Alvarez, G., Oliva, A., & Treisman, A. (2009). Spatial ensemble statistics are efficient codes that can be represented with reduced. *Proceedings of the National Academy of Sciences of the United States of America*, *106*(18), 7345–7350.
- Ariely, D. (2001). Seeing sets: Representation by statistical properties. *Psychological Science*, *12*(2), 157–162.
- Ball, K., & Sekuler, R. (1980). Models of stimulus uncertainty in motion perception. *Psychological Review*, *87*, 435–469.
- Bayne, T. (2009). Perception and the reach of phenomenal content. *Philosophical Quarterly*, *59*(236), 385–404.
- Bayne, T. (2016). Gist! *Proceedings of the Aristotelian Society*, *116*(2), 107–126.
- Bergmann, M. (2013). Externalist justification and the role of seemings. *Philosophical Studies*, *166*, 163–184.
- Block, N. (2007). Consciousness, accessibility, and the mesh between psychology and neuroscience. *Behavioral and Brain Sciences*, *30*, 481–548.

- Block, N. (2011). Perceptual consciousness overflows cognitive access. *Trends in Cognitive Sciences*, 15, 567–575.
- Block, N. (2014). Seeing-as in the light of vision science. *Philosophy and Phenomenological Research*, 89(3), 560–572.
- Brady, T. F., & Alvarez, G. (2015). No evidence for a fixed object limit in working memory. *Journal of Experimental Psychology, Learning, Memory, and Cognition*, 41(3), 921–929.
- Brewer, B. (2006). Perception and content. *European Journal of Philosophy*, 14(2), 165–181.
- Brogaard, B. (2013a). Do we perceive natural kind properties? *Philosophical Studies*, 162(1), 35–42.
- Brogaard, B. (2013b). Phenomenal seemings and sensible dogmatism. In C. Tucker (Ed.), *Seemings and justification: New essays on dogmatism and phenomenal conservatism* (pp. 270–289). New York: OUP.
- Bronfman, Z. Z., Brezis, N., Jacobson, H., & Usher, M. (2014). We see more than we can report: “Cost free” color phenomenality outside focal attention. *Psychological Science*, 25(7), 1394–1403.
- Burge, T. (2014). Reply to block: Adaptation and the upper border of perception. *Philosophy and Phenomenological Research*, 89(3), 573–583.
- Burr, D., & Ross, J. (2008). A visual sense of number. *Current Biology*, 18(6), 425–428.
- Carruthers, P., & Veillet, B. (2011). The case against cognitive phenomenology. In T. Bayne & M. Montague (Eds.), *Cognitive Phenomenology*. Oxford: OUP.
- Chong, S. C., & Treisman, A. (2003). Representation of statistical properties. *Vision Research*, 43(4), 393–404.
- Chong, S. C., & Treisman, A. (2013). Representation of statistical properties. *Vision Research*, 43, 393–404.
- Cohen, J., & Dennett, D. (2011). Consciousness cannot be separated from function. *Trends in Cognitive Sciences*, 15(8), 358–364.
- Cohen, J., Dennett, D. C., & Kanwisher, J. (2016). What is the bandwidth of perceptual experience? *Trends in Cognitive Sciences*, 20(5), 324–335.
- Corbett, J. E., & Melcher, D. (2014). Characterizing ensemble statistics: Mean size is represented across multiple frames of reference. *Attention, Perception, and Psychophysics*, 76(3), 746–758.
- Corbett, J. E., Wurnitsch, N., Schwartz, A., & Whitney, D. (2012). An aftereffect of adaptation to mean size. *Visual Cognition*, 20(2), 211–231.
- Cullison, A. (2011). What are seemings? *Ratio*, 23, 260–274.
- Davidson, D. (1986). A coherence theory of truth and knowledge. In E. LePore (Ed.), *Truth and interpretation*. Oxford: Basil Blackwell.
- Davies, M. (1992). Perceptual content and local supervenience. *Proceedings of the Aristotelian Society*, 92, 21–45.
- Dretske, F. (2007). What change blindness teaches about consciousness. *Philosophical Perspectives*, 21, 215–230.
- Dubé, C., & Sekuler, R. (2015). Obligatory and adaptive averaging in visual short-term memory. *Journal of Vision*, 15(4), 1–13.
- Durgin, F. H. (1995). Texture density adaptation and the perceived numerosity and distribution of texture. *Journal of Experimental Psychology: Human Perception and Performance*, 21(1), 149.
- Durgin, F. H. (2008). Texture density adaptation and visual number revisited. *Current Biology*, 18(18), R855–R856.
- Durgin, F. H., & Huk, A. C. (1997). Texture density aftereffects in the perception of artificial and natural textures. *Vision Research*, 37(23), 3273–3282.
- Fish, W. (2013). High-level properties and visual experience. *Philosophical Studies*, 162, 43–55.
- Gibson, J. J., & Radner, M. (1937). Adaptation, after-effect and contrast in the perception of tilted lines. I. Quantitative studies. *Journal of Experimental Psychology*, 20(5), 453.
- Goldman, A. (1986). *Epistemology and cognition*. Cambridge MA: Harvard University Press.
- Haberman, J., & Whitney, D. (2007). Rapid extraction of mean emotion and gender from sets of faces. *Current Biology*, 17(17), R751–R753.
- Haberman, J., & Whitney, D. (2009). Seeing the mean: Ensemble coding for sets of faces. *Journal of Experimental Psychology: Human Perception and Performance*, 35(3), 718–734.
- Haberman, J., & Whitney, D. (2011). Efficient summary statistical representation when change localization fails. *Psychonomic Bulletin*, 18(5), 855–859.
- Haberman, J., & Whitney, D. (2012). Ensemble perception: Summarising the scene and broadening the limits of visual processing. In J. Wolfe & L. Robertson (Eds.), *From perception to consciousness: Searching with Anne Treisman*. Oxford: Oxford University Press.

- Hansen, M. Perception of high-level content and the argument from associative agnosia, *Review of Philosophy and Psychology* (in press).
- Hawley, K., & Macpherson, F. (2011). *The admissible contents of experience*. Oxford: Wiley, Blackwell.
- Johnston, M. (2006). Better than mere knowledge? The function of sensory awareness. In T. Gendler & J. Hawthorne (Eds.), *Perceptual experience*. Oxford: Oxford University Press.
- Kouider, S., de Gardelle, V., Sackur, J., & Dupoux, E. (2010). How rich is consciousness? The partial awareness hypothesis. *Trends in Cognitive Sciences*, 14(7), 301–307.
- Leib, A. Y., Kosovicheva, A., & Whitney, D. (2016). Fast ensemble representations for abstract visual impressions. *Nature Communications*. <https://doi.org/10.1038/ncomms13186>.
- Logue, H. (2013). Visual experience of natural kind properties: Is there a fact of the matter? *Philosophical Studies*, 162(1), 1–12.
- Lyons, J. (2005). Perceptual belief and nonexperiential looks. In J. Hawthorne (Ed.), *Philosophical perspectives*, 19: *Epistemology* (pp. 237–256). Malden: Blackwell.
- Marchant, A. P., Simons, D. J., & de Fockert, J. W. (2013). Ensemble representations: Effects of set size and item heterogeneity on average size perception. *Acta Psychologica*, 142, 245–250.
- Masrour, F. (2011). Is perceptual phenomenology thin? *Philosophy and Phenomenological Research*, 83(2), 366–397.
- McClelland, T. (2016). Gappiness and the case for liberalism about phenomenal properties. *Philosophical Quarterly*, 66(264), 536–558.
- McClelland, T., & Bayne, T. (2016). Ensemble coding and two conceptions of phenomenal sparsity. *Trends in Cognitive Sciences*, 20(9), 641–642.
- McDowell, J. (1994). *Mind and world*. Cambridge, MA: Harvard University Press.
- McGinn, C. (1997). *The character of mind*. Oxford: Oxford University Press.
- Nanay, B. (2011). Do we see apples as edible? *Pacific Philosophical Quarterly*, 92, 305–322.
- Oriet, C., & Brand, J. (2013). Size averaging of irrelevant stimuli cannot be prevented. *Vision Research*, 79, 8–16.
- Pace, M. (2017). Experiences, seemings, and perceptual justification. *Australasian Journal of Philosophy*, 95(2), 226–241.
- Parkes, L., Lund, J., Angelucci, A., Solomon, J. A., & Morgan, M. (2001). Compulsory averaging of crowded orientation signals in human vision. *Nature Neuroscience*, 4, 739–744.
- Phillips, I. (2016). No watershed for overflow: Recent work on the richness of consciousness. *Philosophical Psychology*, 29(2), 236–249.
- Pollock, J. (1986). *Contemporary theories of knowledge*. Lanham, MD: Rowman & Littlefield.
- Price, R. (2009). Aspect-switching and visual phenomenal character. *Philosophical Quarterly*, 59(236), 508–518.
- Prinz, J. (2013). Siegel's get rich quick scheme. *Philosophical Studies*, 163, 827–835.
- Pryor, J. (2005). Is there immediate justification? In M. Steup (Ed.), *Contemporary debates in epistemology*. Oxford: Blackwell.
- Reiland, I. (2014). On experiencing high-level properties. *American Philosophical Quarterly*, 51(3), 177–187.
- Reiland, I. (2015). Experience, seemings, and evidence. *Pacific Philosophical Quarterly*, 96, 510–534.
- Siegel, S. (2006). Which properties are represented in perception? In T. S. Gendler & J. Hawthorne (Eds.), *Perceptual experience* (pp. 481–503). Oxford: OUP.
- Siegel, S. (2010). *The contents of visual experience*. Oxford: Oxford University Press.
- Siegel, S., & Silins, N. (2015). The epistemology of perception. In M. Matthen (Ed.), *Oxford handbook of philosophy of perception* (pp. 781–811). Oxford: OUP.
- Silins, N. (2011). Seeing through the veil of perception. *Mind*, 120(478), 329–367.
- Sweeny, T. D., Haroz, S., & Whitney, D. (2012). Perceiving group behaviour. *Journal of Experimental Psychology: Human Perception and Performance*, 39, 1–10.
- Sweeny, T. D., & Whitney, D. (2014). Perceiving crowd attention: Ensemble perception of a crowd's gaze. *Psychological Science*, 25, 1–11.
- Thornton, I. M., Srismith, D., Oxner, M., & Hayward, W. G. (2014). Estimating the racial composition of groups of faces: An ensemble other-race effect. In *TSPC2014: proceedings of the trieste symposium on perception and cognition, November 27th–28th* (pp. 76–79). Trieste: EUT Edizioni Università di Trieste.
- Torralba, A., & Oliva, A. (2003). Statistics of natural image categories. *Network Computation in Neural Systems*, 14, 391–412.

- Tucker, C. (2010). Why open-minded people should endorse dogmatism. *Philosophical Perspectives*, 24(1), 529–545.
- Utochkin, I. S., & Tiurina, N. A. (2014). Parallel averaging of size is possible but range-limited. *Acta Psychologica*, 146, 7–18.
- Ward, E., Bear, A., & Scholl, B. J. (2016). Can you perceive ensembles without perceiving individuals?: The role of statistical perception in determining whether awareness overflows access. *Cognition*, 152, 78–86.
- Wataminiuk, S. N. J., & McKee, S. P. (1998). Simultaneous encoding of direction at a local and global scale. *Perception and Psychophysics*, 60, 191–200.
- Whitney, D., Haberman, J., & Sweeny, T. D. (2014). From textures to crowds: Multiple levels of summary statistical perception. In J. S. Werner & L. M. Chalupa (Eds.), *The New visual neuroscience* (pp. 695–710). Cambridge: MIT.