ANNUAL REVIEWS

Annual Review of Psychology

Ensemble Perception

David Whitney[1,2,3] and Allison Yamanashi Leib[1]

[1]Department of Psychology, University of California, Berkeley, California 94720; email: dwhitney@berkeley.edu

[2]Vision Science Program, University of California, Berkeley, California 94720

[3]Helen Wills Neuroscience Institute, University of California, Berkeley, California 94720

Annu. Rev. Psychol. 2018. 69:105–29

First published as a Review in Advance on September 11, 2017

The *Annual Review of Psychology* is online at psych.annualreviews.org

https://doi.org/10.1146/annurev-psych-010416-044232

Keywords

vision, summary statistics, consciousness, crowding, texture, scene, object recognition

Abstract

To understand visual consciousness, we must understand how the brain represents ensembles of objects at many levels of perceptual analysis. Ensemble perception refers to the visual system's ability to extract summary statistical information from groups of similar objects—often in a brief glance. It defines foundational limits on cognition, memory, and behavior. In this review, we provide an operational definition of ensemble perception and demonstrate that ensemble perception spans across multiple levels of visual analysis, incorporating both low-level visual features and high-level social information. Further, we investigate the functional usefulness of ensemble perception and its efficiency, and we consider possible physiological and cognitive mechanisms that underlie an individual's ability to make accurate and rapid assessments of crowds of objects.

Contents

1. INTRODUCTION

We process the vibrant complexity of natural scenes using the relatively limited capacity of the visual system. The fidelity with which we can perceive any complex scene at a glance is restricted by finite attentional resources (Cavanagh & Alvarez 2005, Dux & Marois 2009, Simons & Levin 1997), limits of eye movements and scanning (Kowler 2011, Wolfe 1994), and minimal visual working memory capacity (Luck & Vogel 2013). Human visual processing is further constrained by coarse peripheral resolution (Anstis 1974, Virsu & Rovamo 1979) and the fundamental limits set by visual crowding (Pelli 2008, Strasburger et al. 2011, Whitney & Levi 2011). Fortunately, although natural scenes are dense with information, this clutter is not completely random. Instead, natural scenes are filled with similar or redundant groups of objects, features, and textures. The visual system is sensitive to these similarities in both natural (e.g., stand of trees, crowd of faces) and artificial (e.g., car lot, bike rack) groups in the form of ensemble or summary statistical information (Alvarez 2011, Haberman & Whitney 2012, Whitney et al. 2014). For example, in **Figure 1**, we can extract summary statistics along many dimensions, including the average hue of the tree leaves, average facial expression of the bystanders, and average speed of the cyclists. In this review, we address the types of visual information that are represented as ensembles, what perceptual and cognitive benefits ensemble perception affords, how attention is involved in representing ensembles, and which proposed mechanisms may account for many aspects of ensemble perception.

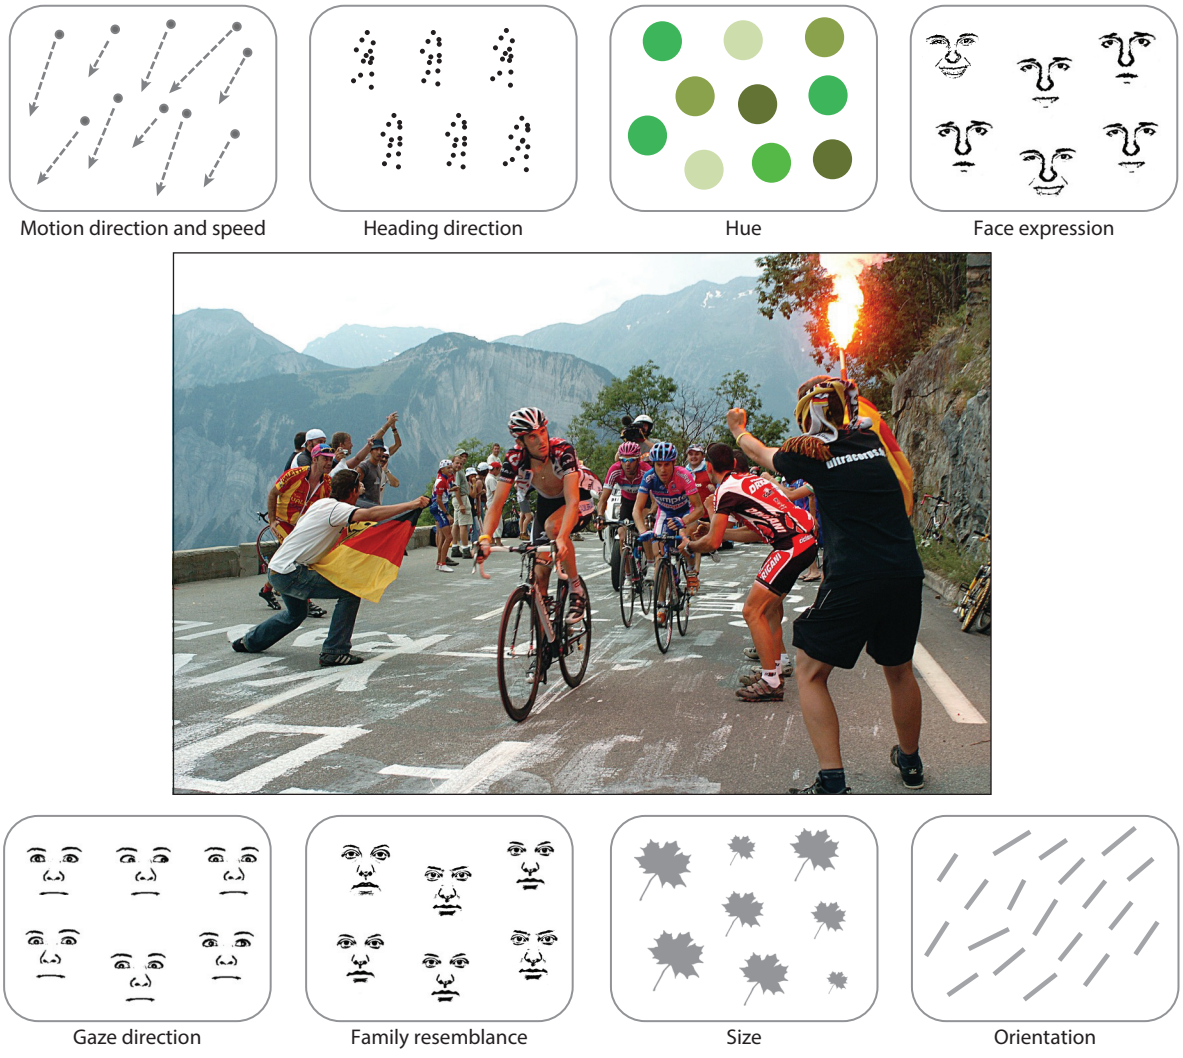

Figure 1

Multiple ensembles are present in natural scenes. Natural scenes contain numerous groups of similar and redundant stimuli. Observers perceive these groups in the form of summary statistics, such as the average orientation, size, and hue of the foliage; the average speed, motion direction, and heading of the bikers; and the average emotional expression, gaze direction, and family resemblance of the bystanders. Ensemble perception is hierarchical and occurs at many levels of visual processing. Photo credit: Max Pixel, available for reuse under the Creative Commons Zero (CC0 1.0).

2. ENSEMBLE PERCEPTS ACROSS MULTIPLE LEVELS OF VISUAL ANALYSIS

2.1. Low-Level Ensemble Perception

Ensemble perception has been reported for many low-level features, including motion, orientation, brightness, hue, and spatial position. It has been examined using a variety of psychophysical techniques (**Figure 2**). Of the low-level stimuli that have been studied, visual motion provides a canonical example of ensemble perception (Watamaniuk & McKee 1998, Watamaniuk et al. 1989).

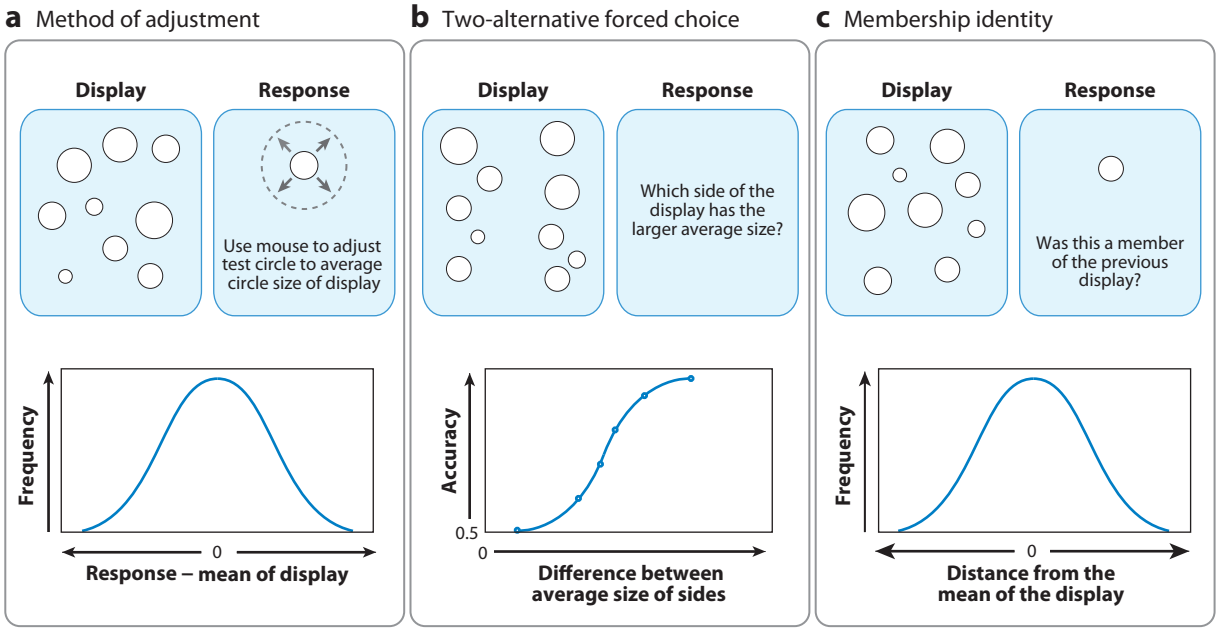

Figure 2

Methods of testing ensemble perception. (*a*) Participants match the average size of a test circle to a set of displayed circles. The spread or variance of the resulting distribution (e.g., Gaussian, Von Mises, Cauchy, etc.) reflects the matching performance. Robust sensitivity to summary statistical performance will yield a narrow distribution centered at 0. (*b*) Observers make a two-alternative forced-choice judgment about which display (left/right or which of two intervals) contains the larger average size. A psychometric function fitted to the data reveals the sensitivity to the summary statistic (e.g., discrimination threshold). (*c*) In the implicit membership identity task, observers report whether the test circle was present in the previous display. When the test circle is the average of the displayed set, participants false alarm at a high rate. The shape and location of the resulting histogram reflects sensitivity to the ensemble (average size).

Watamaniuk and colleagues (1989; Watamaniuk & McKee 1998) asked observers to judge the average motion direction of random-dot cinematograms that resembled blowing snow and found that observers accurately reported the average motion direction of the drifting dots. Similarly, observers easily estimated the average speed of random dots moving at diverse speeds with a precision comparable to their estimations of dots moving at a homogeneous speed (Watamaniuk & Duchon 1992). In addition to motion, observers can accurately discriminate, report, and reproduce the average orientation of stimuli (Dakin & Watt 1997, Miller & Sheldon 1969, Parkes et al. 2001). Further, observers can also perceive the average brightness (Bauer 2009), average hue (Demeyere et al. 2008, Webster et al. 2014), and average spatial position of a cloud of random stimuli (Alvarez & Oliva 2008, Melcher & Kowler 1999, Vishwanath & Kowler 2003). These low-level visual summary statistical representations of spatial frequency, color, and orientation may form the basis of texture recognition and discrimination (Landy 2014) as well as of scene gist impressions (Oliva & Torralba 2006).

2.2. Mid-Level Ensembles

Mid-level features and objects are also perceived as ensembles. For example, observers perceive the average size in an array of circles, even for very brief displays of less than one-tenth of a second (Ariely 2001, Chong & Treisman 2003). Ensemble size perception is somewhat controversial

(Myczek & Simons 2008), in part because, unlike perception of average motion, position, or orientation, the existence of low-level size detectors is less clear. In contrast, the size and scale of single circles could be represented as early as V1 (Schwarzkopf et al. 2011). If we consider size to be a mid-level feature, then summary statistical information about groups of circles may be associated with surface or depth perception. Consistent with this possibility, observers can perceive the average depth of a crowd of objects that have varying binocular disparity (Wardle et al. 2017), suggesting that ensemble information contributes to depth and scene perception and, perhaps, to Gestalt grouping, as well (Wagemans et al. 2012). It remains an intriguing question whether the visual system extracts ensemble information about other depth cues such as average height in field, average atmospheric perspective, and average texture gradients. Average depth information from any of these sources could facilitate recognition not only of the depth of individual objects but also of their identity and size in scenes. Ensemble percepts of average cast shadows (Koenderink et al. 2004, Sanders et al. 2010), if present, could assist in recovering global lighting in scenes and inform depth assignments of objects throughout the scene. While these are still open questions, future work on ensemble surface and mid-level visual information will be important for documenting the extent to which summary statistical representations contribute to the fundamental building blocks of surface, object, and scene perception.

2.3. High-Level Ensembles

Recent work on summary statistical perception highlights its significant role in high-level object, scene, and social perception. Most of this work has documented how ensemble perception operates on groups of face stimuli, allowing observers to rapidly access the emotional tenor or intent of a crowd. Haberman & Whitney (2007, 2009) found that observers can evaluate and discriminate the average emotional expression and gender in a crowd of faces. Observers can also accurately evaluate average facial identity (e.g., family resemblance; see Bai et al. 2015, de Fockert & Wolfenstein 2009, Neumann et al. 2013, Yamanashi Leib et al. 2012b). These studies emphasize the fact that humans can rapidly extract important social information from crowds during a brief glance, perhaps as short as 100 ms or less (Haberman & Whitney 2009, Li et al. 2016, Yamanashi Leib et al. 2016). Observers are not as sensitive to crowds of inverted or scrambled faces (Haberman & Whitney 2009, Yamanashi Leib et al. 2012b), which suggests that observers extract summary statistical information based on configural or holistic face representations. Different images of the same face and multiple viewpoints of faces can also be incorporated into a unified ensemble percept (Neumann et al. 2013, Yamanashi Leib et al. 2014), indicating that summary statistics are computed over viewpoint-invariant representations, not just two-dimensional image-level information. In addition to facial expression and identity, observers can also judge the average gaze direction and mean head rotation of the crowd (Florey et al. 2016, Sweeny & Whitney 2014), which could be useful in guiding attention and behavior. In addition to face crowds, the visual system extracts ensemble information about dynamic objects as well. Sweeny and colleagues (2013) discovered that observers viewing a crowd of point-light walkers can accurately match and discriminate their average direction. An intriguing implication of the high-level ensemble work is that complex perceptual interpretations, such as the perception of crowd panic, may be subserved by ensemble representations (see Section 5).

2.4. Multiple Ensembles

Most ensemble research focuses on the perception of one specific ensemble characteristic from a particular group of stimuli (e.g., average size or expression). However, a few studies have investigated whether participants can extract multiple ensemble characteristics from one or more groups

of stimuli. For example, Chong & Treisman (2005b) asked participants to view a group of disks containing stimuli of two different colors and found that observers were sensitive to the average size of both sets of colored disks even though their attention was divided between the two colors (Chong & Treisman 2005b). Other studies have shown that observers can successfully extract multiple ensembles from up to four groups of stimuli (Attarha & Moore 2015, Attarha et al. 2014). However, Attarha and colleagues (2014; Attarha & Moore 2015) found that undivided attention enhanced performance, suggesting that extracting multiple ensembles is not entirely capacity free (see Section 4 for the role of attention in ensemble perception).

Whereas the aforementioned studies of multiple groups of stimuli examined participants' ability to extract the same characteristic (size), other studies have investigated sensitivity to simultaneous but different ensemble characteristics. For example, Emmanouil & Treisman (2008) found that observers could perceive the average speed and size of a group of circles. Although accuracy was lower in the multiple-ensemble conditions than in the single-ensemble conditions, this result suggests that observers can perceive multiple ensemble features. Future studies should continue to explore the perceptual interactions between multiple ensembles and the capacity limits of multiple-ensemble processing.

The studies described above were restricted to visual stimuli, but it has been hypothesized that ensemble perception may be a general mechanism that operates across several sensory domains, and researchers have documented that ensemble percepts can incorporate auditory as well as visual stimuli. Listeners can perceive the average in a sequence of pure tones (Piazza et al. 2013) and can efficiently discriminate sound textures, such as those present in auditory scenes (McDermott et al. 2013). Moreover, although there is some cost when perceiving multiple ensembles within one modality (e.g., Emmanouil & Treisman 2008), there is relatively little cost in perceiving ensemble information across different modalities. For example, participants can recognize ensemble tone and visual size simultaneously, and there is little evidence of a cost associated with simultaneous displays (Albrecht et al. 2012).

2.5. Beyond Average

The most commonly measured form of ensemble representation is the perceived average of a group of items. However, the diversity or variance in a set of stimuli is also very important (**Figure 3**). For example, when walking through a crowd of people, the average emotion is informative, but equally critical is the variation of emotion present in the crowd. Haberman et al. (2015b) found that observers who viewed a crowd of up to 16 faces for 1 s successfully matched and discriminated the variance of the crowd, and subsequent tests confirmed that observers distinguished among numerous levels of variance, not merely between homogeneous and heterogeneous crowds (Haberman et al. 2015b). Variance information is useful in several respects. First, it signals the reliability of the estimated average: A homogeneous group of angry faces implies something very different than a set of faces that varies in expression. Second, ensemble variance might provide direct information about the diversity, mixture, or ambivalence of a crowd. Third, variance information might be useful to identify statistical outliers, such as deviant expressions in a crowd (Whitney et al. 2014). Ensemble variance is not available at the level of any single individual; it is an emergent property only accessible by encoding summary statistical information. Ensemble variance information is extracted for high-level (e.g., facial expression, age, and gender and racial diversity), low-level (e.g., orientation) (Dakin & Watt 1997, Morgan et al. 2008, Norman et al. 2015), and mid-level (e.g., size) features (Solomon et al. 2011). Whether third-order statistics, such as kurtosis, are extracted remains unclear. Interpreting variance in the crowd seems to be as ubiquitous a calculation as extracting summary statistical information about the average.

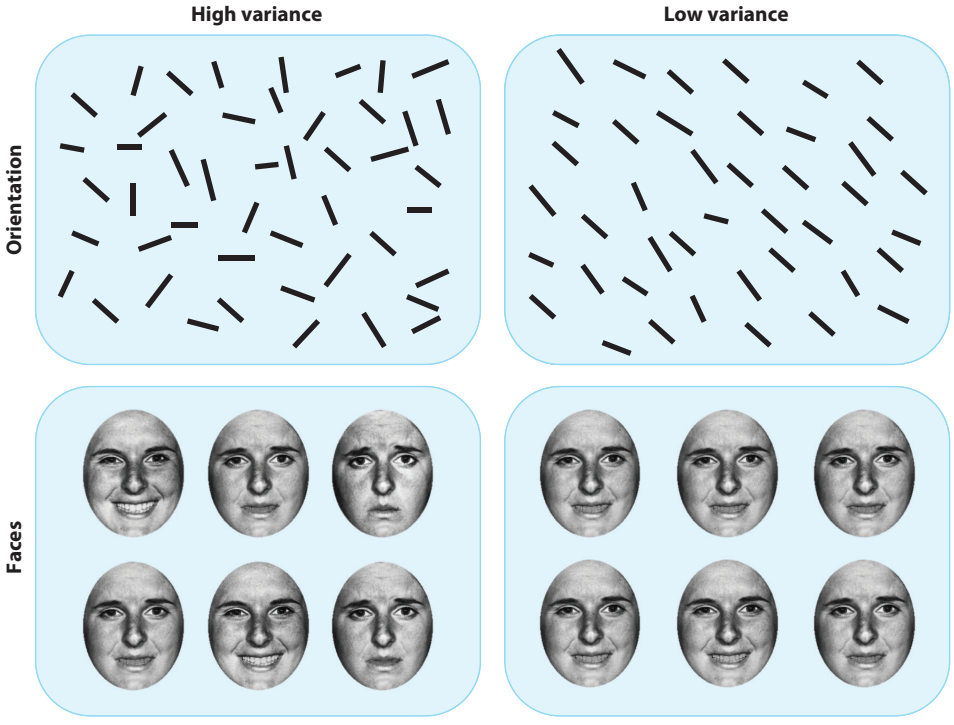

Figure 3

Extracting variance from groups of stimuli. Extracting the mean, or average, characteristic from a group of stimuli is the most commonly reported ensemble percept. However, ensemble perception may include extracting diverse statistical information, such as variance, range, or even kurtosis. In this figure, we show examples from research documenting the fact that participants can accurately evaluate the variance within groups of redundant stimuli (Haberman et al. 2015b, Solomon 2010).

2.6. Ensemble Cognition

Summary statistical representations exist at the highest levels of perceptual and cognitive processing—what we will refer to as ensemble cognition. The perceptual and cognitive evaluations we once thought required conscious scrutiny and deliberation are actually rapidly extracted as ensemble percepts. For example, subjective percepts like attractiveness can be estimated via ensemble perception (Anderson et al. 1973, Post et al. 2012, Walker & Vul 2013). In natural scenes, abstract percepts such as liveliness or animacy are also rapidly represented in the form of an ensemble (Yamanashi Leib et al. 2016). In one recent study, observers rated random crowds of animals, insects, plants, and household objects on their average lifelikeness (Yamanashi Leib et al. 2016). Observers' ratings of the crowds' lifelikeness were highly correlated with the mathematical mean of the items rated individually, even though the individual items were rated by independent observers. This indicates that observers agree about the lifelikeness of objects and crowds. Ensemble liveliness is extracted for displays as brief as 250 ms, even when observers cannot recall individual stimuli in the crowd (Yamanashi Leib et al. 2016). This suggests that ensemble information underlies observers' first impressions of the liveliness of natural scenes (Yamanashi Leib et al. 2016).

The fact that perceived liveliness can be extracted so quickly and efficiently suggests that other high-level perceptual, cognitive, and inferential processes may also rely on ensemble

representations. For example, early research hints that social labels may be evaluated through an averaging process (Leon et al. 1973), and there are intriguing findings that long-term memory consolidation may resemble ensemble representations (Richards et al. 2014). Thus, cognitive processes beyond perception may also rely on averaging mechanisms, perhaps even the same ensemble mechanisms discussed in this review.

3. OPERATIONALLY DEFINING ENSEMBLE PERCEPTION

From the discussion in the previous section, it may seem that anything—any arbitrary set of features, objects, or configurations—can be perceived as an ensemble. However, this is not the case. Ensemble perception has unique features that can help establish the foundation for an operational definition to distinguish what ensemble perception is not, identify how ensemble perception is related to other seemingly similar phenomena, and isolate the underlying neural mechanisms. A flexible operational definition of ensemble coding should include the following five concepts:

- Ensemble perception is the ability to discriminate or reproduce a statistical moment.
- Ensemble perception requires the integration of multiple items.
- Ensemble information at each level of representation can be precise relative to the processing of single objects at that level.
- Single-item recognition is not a prerequisite for ensemble coding.
- Ensemble representations can be extracted with a temporal resolution at or beyond the temporal resolution of individual object recognition.

3.1. Ensemble Perception Is the Ability to Discriminate or Reproduce a Statistical Moment

Not every group or set of things is perceived as an ensemble. We can perceive groups of random objects or interactions between features and objects that have no meaningful or consistent relationship to each other, have no underlying statistical distribution, and cannot be reported or discriminated as a set. We can also recognize Gestalt or holistic grouping cues, but these need not involve the perception of a statistical moment and, thus, are not diagnostic of ensemble processing. Gestalt grouping may interact with ensemble perception, either by constraining ensemble representations or by being generated by them. In contrast to other phenomena, sensitivity to ensemble information—to a statistical moment—depends on the variance of the underlying distribution, such that increasing variance in the dimension of interest reduces sensitivity to the summary statistic (Dakin 2001, Fouriezos et al. 2008, Haberman et al. 2015b, Im & Halberda 2013, Morgan et al. 2008, Solomon et al. 2011). However, for sets with constant variance, the shape of the underlying statistical distribution (e.g., normal, rectangular, bimodal) is less critical (Allik et al. 2013, Chong & Treisman 2003, Haberman & Whitney 2009).

3.2. Ensemble Perception Requires the Integration of Multiple Items

In terms of an operational definition for what counts as an ensemble representation for perception, the only requirement is an integration of two or more stimuli. Technically, integrating (sampling) two items is sufficient evidence for an ensemble representation, and this happens in some cases (Allik et al. 2013, Maule & Franklin 2016). However, beyond that, there is no particular quantity of items (or minimum subset of items) that is required to meet the criteria for an ensemble or summary statistical representation. In many studies, the number of features or objects integrated is greater than two (**Figure 4**; see also **Table 1**). For instance, when discriminating the average size

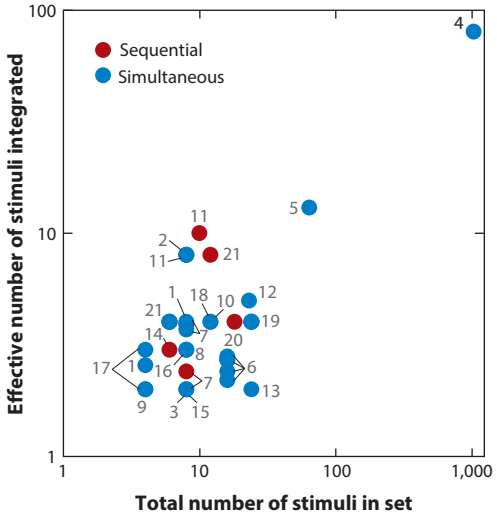

Figure 4

Efficiency of integration in ensemble perception. The number of stimuli integrated (ordinate) is plotted against the number of items in the display (abscissa) for a variety of published studies. Red circles represent experiments that displayed sequential objects in a set. Blue circles represent experiments that displayed objects simultaneously in a spatial array. The best-fitting power function for these data is $f(x) = x^{0.58}$, which suggests that, across a variety of studies, observers integrate approximately the square root of the number of displayed objects. Gray numbers denote the associated references (see **Table 1**).

or variance in a set of circles, observers performed at 60–75% efficiency, integrating at least three circles (Allik et al. 2013, Solomon et al. 2011). Observers in other studies of basic visual features incorporated information ranging from approximately 3 to 5 items per display (Im & Halberda 2013, Solomon 2010); in another study, observers incorporated information from approximately four items with replacement at approximately 5 Hz (Gorea et al. 2014). At the higher end of the range, some studies report that observers sample approximately the square root of the number of display items, even for very large set sizes (Dakin 2001). Higher-level ensemble perception studies, such as those using face crowds, biological motion, and other stimuli, indicate that observers can often integrate more than 4–8 objects (Haberman & Whitney 2010; Sweeny et al. 2013; Yamanashi Leib et al. 2014, 2016).

The estimates of efficiency or number of items integrated in ensemble representations clearly vary and are sometimes debated (Chong et al. 2008, Dakin 2001, Marchant et al. 2013, Myczek & Simons 2008, Solomon et al. 2011). This is, in part, because efficiency depends on several factors, including the stimulus type, methods used (e.g., ideal observer or equivalent noise modeling versus empirical set size manipulations), and assumptions of those methods (Solomon et al. 2011). Attention seems to influence efficiency, as well (Dakin et al. 2009), opening up the possibility of variations in estimated efficiency depending on task design, observer goals, and attentional demands (see Section 4). Further, individual differences can significantly impact how much information observers integrate (Bai et al. 2015, Haberman & Whitney 2010, Haberman et al. 2015a, Solomon 2010).

Despite the variations in methods and modeling approaches, however, there is general agreement across ensemble tasks that multiple features or objects are integrated. In fact, when a large sample of experimental estimates of efficiency are plotted together, a striking pattern emerges (**Figure 4**), suggesting that observers integrate approximately the square root of the number of

Table 1 References for Figures 4 and 6

Reference number	Reference
1	Allik et al. (2013)
2	Alvarez & Oliva (2008)
3	Chong et al. (2008)
4	Dakin (2001)
5	Dakin et al. (2005)
6	Florey et al. (2016)
7	Florey et al. (2017)
8	Gorea et al. (2014)
9	Haberman & Whitney (2009)
10	Haberman & Whitney (2010)
11	Hubert-Wallander & Boynton (2015)
12	Im & Halberda (2013)
13	Myczek & Simons (2008)
14	Piazza et al. (2013)
15	Solomon (2010)
16	Solomon et al. (2011)
17	Sweeny & Whitney (2014)
18	Sweeny et al. (2013)
19	Wolfe et al. (2015)
20	Yamanashi Leib et al. (2014)
21	Yamanashi Leib et al. (2016)

The following approach was used to estimate the approximate number integrated for each point on the graphs. If the researcher reported a range of stimuli integrated (e.g., 3–4 stimuli were integrated), we plotted the lower estimate. In some cases, the difference was minimal. In other cases the difference was substantially larger (Yamanashi Leib et al. 2014). Thus, the ordinate on the graph represents a conservative estimate, or lower bound, of effective integration across multiple ensemble perception studies. This list is not exhaustive; many studies have not estimated the efficiency or number of integrated samples, or have done so indirectly. The methods used in each study vary: Some used empirical manipulations and some used ideal observer modeling, regression approaches, or equivalent noise analysis.

objects in their summary statistical representations. Notwithstanding the individual differences and multiple factors that modulate integration efficiency, the general rule that observers effectively integrate approximately the square root of the number of stimuli thus appears reasonable under many circumstances (**Figure 4**).

Most ensemble perception research emphasizes spatial integration, but summary statistics are perceived in temporal sequences, as well. Observers successfully estimate summary statistics, including average object location, facial expression, object size, tone, and animacy, from sequentially presented objects (Albrecht & Scholl 2010, Chong & Treisman 2003, Haberman et al. 2009, Whiting & Oriet 2011, Yamanashi Leib et al. 2014). Thus, summary statistical information can be extracted flexibly over time from spatially local or global scales, in contrast to Navon figures (Navon 1977). In fact, it appears that the integration efficiency for temporally presented sets is as high as or higher than that for spatial arrays (Florey et al. 2017, Gorea et al. 2014).

Although multiple items are integrated in ensemble representations, not all items need to be weighted equally. For example, statistical outliers (deviants) are downplayed, or filtered, in the cases of color (Michael et al. 2014) and faces (Haberman & Whitney 2010). Summary statistical

perception in temporal arrays also reveals a type of weighting, in the form of primacy and recency effects, such that the first- or last-seen object can bias the estimated ensemble property (Hubert-Wallander & Boynton 2015). In addition to order- and outlier-based weighting, there is also the potential for weighting based on attention (de Fockert & Marchant 2008), eccentricity (Ji et al. 2014), and expectancy (Cheadle et al. 2014). For example, a study investigating ensemble circle size provided some evidence that attending to items biases the reported mean (de Fockert & Marchant 2008). Foveally viewed stimuli may also pull or bias estimates of the summary statistic, such as average expression (Ji et al. 2014, Wolfe et al. 2015), but foveally viewed objects are not necessary, as other studies have demonstrated integration of multiple peripheral objects without any foveal stimulation (Haberman et al. 2009, Wolfe et al. 2015).

Although **Figure 4** shows that there is little debate about the criterion that ensemble perception must involve two or more items being integrated, there is an ongoing debate about whether ensemble perception is automatic, obligatory, unconscious, parallel, or outside the focus of attention. Because these issues are not diagnostic of ensemble coding and because arguments about efficiency—the number of objects integrated into the ensemble representation—do not address these debates, we reserve discussion of these issues to Section 4, where we explore the role of attention in ensemble perception.

3.3. Ensemble Information at Each Level of Representation Can Be Precise Relative to Processing of Single Objects at That Level

Averaging cancels uncorrelated noise associated with individual items (Alvarez 2011, Galton 1907, Surowiecki 2004). As such, one might expect that sensitivity will increase with increasing sample size, and this result is sometimes found (Robitaille & Harris 2011). When individual object representations are especially noisy (e.g., brief), observers can be more sensitive to the ensemble as a whole, compared to the single item (Gorea et al. 2014, Li et al. 2016, Sweeny et al. 2013, Yamanashi Leib et al. 2014). These enhancements are not required for ensemble perception, however. Indeed, several authors have reported relatively constant sensitivity with increasing set size (Allik et al. 2013, Alvarez 2011, Ariely 2001, Chong & Treisman 2005b). The benefit of averaging across larger sample sizes may be offset by factors such as increased correlated noise and positional uncertainty, potentially yielding a pattern of results that appears as if there is constant sensitivity across set sizes. Moreover, late-stage noise may limit the apparent benefit of averaging. Nonetheless, ensemble sensitivity is generally better than would be predicted if discrimination thresholds were set by single-object discrimination. For example, there is compulsory averaging of orientation (Parkes et al. 2001), size (Allik et al. 2014), and facial expression (Fischer & Whitney 2011), and objects that are crowded and, therefore, unrecognizable nonetheless contribute to the perceived ensemble (Fischer & Whitney 2011, Ikeda et al. 2013, Parkes et al. 2001).

3.4. Single-Item Recognition Is Not a Prerequisite for Ensemble Coding

An ensemble can be perceived even if the individuals that comprise the ensemble cannot be reliably reported. Observers accurately report ensemble size and expression information during rapid serial visual presentation (RSVP) paradigms (Haberman et al. 2009, Oriet & Corbett 2008). In these tasks, the individual items are presented too quickly for observers to accurately register each one, yet the ensemble percept still incorporates them. Similarly, studies employing a change blindness paradigm found that observers extracted summary statistical information about color variance (Ward et al. 2016) or average expression from the same stimuli they were attentionally blind to (**Figure 5*b***; Haberman & Whitney 2011). Ensemble perception remains robust even

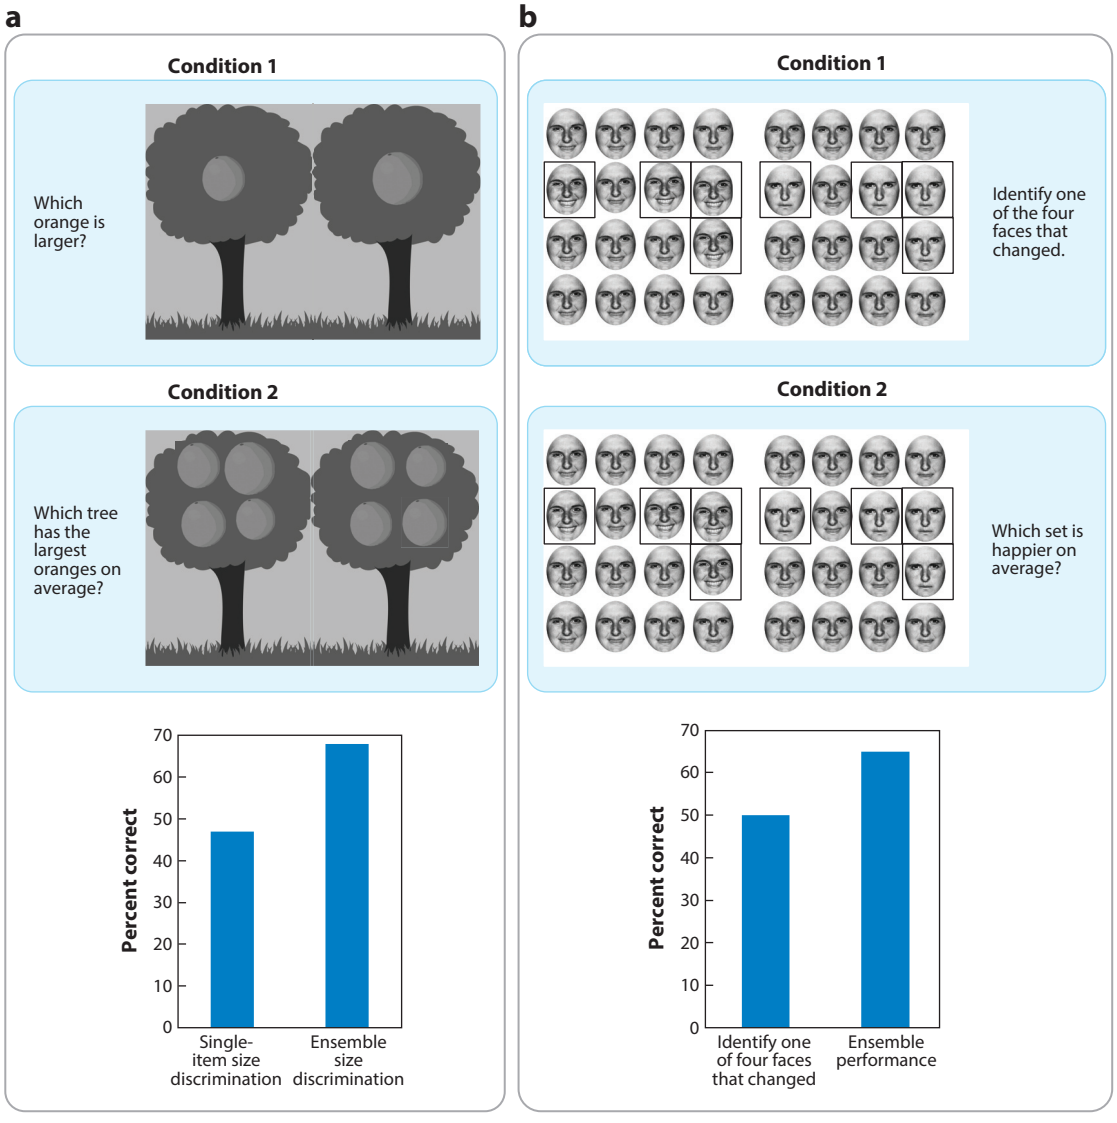

Figure 5

Ensemble perception does not require recognition of individual constituents. (*a*) Two conditions in a size discrimination task from Sweeny et al.'s (2015) study on the development of ensemble perception. In Condition 1, children chose which tree contained the largest orange. In Condition 2, children chose which tree contained the larger oranges on average. Although single-item discrimination was poor, ensemble sensitivity was robust. (*b*) Two conditions from a change localization task (Haberman & Whitney 2011). In Condition 1, participants were asked to localize any one of the four locations where the facial expression changed between the first (*left*) and second (*right*) display of faces. In Condition 2, participants were asked to report which display, on average, contained happier faces. Observers were at chance during the change localization task but nonetheless exhibited above chance ensemble perception sensitivity. Together, these and other experiments suggest that individual set members can influence ensemble percepts even when those objects go unrecognized or unnoticed.

when single-item perception is impeded by crowding, at least for orientation and faces (Fischer & Whitney 2011, Parkes et al. 2001; cf. Banno & Saiki 2012). Moreover, ensemble perception of circle size may be possible during object substitution masking, which greatly reduces the visibility of individual circles (Choo & Franconeri 2010, Jacoby et al. 2012). Further supporting this idea, studies of children, who are still developing individual object recognition, show evidence of ensemble perception (Rhodes et al. 2015, Sweeny et al. 2015). In addition, patients with simultanagnosia, unilateral visual neglect, and congenital prosopagnosia all show some evidence of ensemble perception, even when their single-item discrimination is impaired (Demeyere et al. 2008; Hochstein et al. 2015; Yamanashi Leib et al. 2012a,b). Moreover, individual differences reveal that sensitivity to ensemble information is not perfectly correlated with single-object discrimination (Haberman et al. 2015a, Sweeny et al. 2015). Taken together, a broad range of experiments suggests that ensemble coding does not depend on the recognition or memory of individual objects.

Of course, there are caveats to each example given above. For example, noise limits the conclusions that can be drawn from developing children and patients. Independent noise sources may produce a pattern where ensemble performance and single-item discrimination seem artificially disparate. Finally, although the correlation of individual differences is not perfect, ensemble discrimination and single-item discrimination are correlated to some degree (Haberman et al. 2015a). Nevertheless, the variety of ensemble perception studies using different methods and stimuli provide converging evidence that single-object representations can be lost, neglected, forgotten, unrecognized, or noisy while they are still preserved in the ensemble representation. This does not mean that ensembles are extracted outside of the focus of attention (see Section 4), but it does inform and limit the sorts of models that can be proposed to underlie ensemble perception (see Section 6).

3.5. Ensemble Representations Can Be Extracted with a Temporal Resolution at or Beyond the Temporal Resolution of Individual Object Recognition

Ensemble perception can operate at very brief durations and for fast temporal sequences (**Figure 6**). Our perception of ensemble or gist information from groups of objects can be faster than locating any particular object, such as an extreme one or one closest to the average (Haberman et al. 2009, Yamanashi Leib et al. 2014). As examples, observers perceive average facial expression (Haberman & Whitney 2009), average size (Gorea et al. 2014), average gaze (Florey et al. 2017), and average crowd animacy (Yamanashi Leib et al. 2016) even when set images are displayed as frequently as 20 per s (Haberman et al. 2009), which is beyond the limit of attentional resolution or the dwell time required to scrutinize or find individual faces in the sequence (Nothdurft 1993, Tong & Nakayama 1999, Verstraten et al. 2000). Spatial ensemble information is perceived in crowds of objects displayed for as little as 50 ms (Ariely 2001, Chong & Treisman 2003, Haberman & Whitney 2009, Li et al. 2016, Yamanashi Leib et al. 2016). In addition, ensemble information is extracted from RSVP sequences before individual objects are registered in short-term memory (Joo et al. 2009, McNair et al. 2016).

There is a distinction between the temporal resolution and integration period of ensemble coding. The temporal resolution of ensemble perception is high (and diagnostic): Stimuli in a set can be integrated even when the individuals are not recognized or recalled. However, individual stimuli are not simply blurred (Neumann et al. 2013; Yamanashi Leib et al. 2014, 2016); the individual objects must be registered before integration. Ensemble perception, therefore, resolves individual objects that cannot be uniquely resolved in memory encoding, maintenance, or retrieval.

The accuracy of ensemble estimates often improves with increasing exposure duration (Haberman et al. 2009, Li et al. 2016, Whiting & Oriet 2011), suggesting a temporal integration

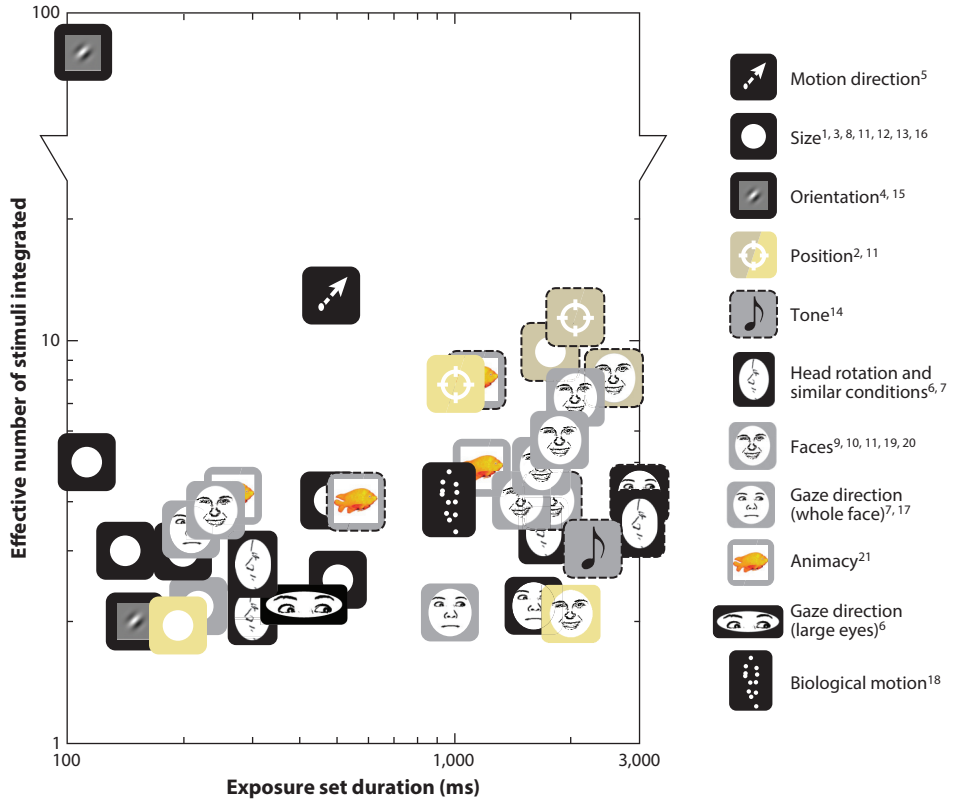

Figure 6

Ensemble perception as a function of set duration. The number of items integrated (ordinate) is plotted as a function of stimulus duration (abscissa). Data points bordered by dashed lines represent sequentially displayed set members. Published studies used different methods to estimate efficiency, indicated by different colored backgrounds. Superscript numbers in the legend denote the associated study (see **Table 1**). Overlapping data points were jittered slightly for visual clarity.

mechanism with a time constant of several hundred ms, at least for mid- and high-level stimuli (Chong & Treisman 2003, Haberman et al. 2009, Whiting & Oriet 2011). Beyond a few hundred ms, the number of integrated objects is fairly constant (**Figure 6**). Thus, there is some degree of duration invariance, at least when considering a large sample of studies that tested a range of different stimuli. Whether there are unique integration periods for different types of stimuli remains unclear [low-level features may have shorter temporal integration periods (e.g., Watamaniuk et al. 1989)], but, given the fact that ensembles are extracted more or less independently from different levels of visual analysis (**Figure 1**), it seems reasonable that there may be distinct time courses for different forms of ensemble coding.

The fact that ensemble properties like average size or expression are represented quickly, or are derived from sets of stimuli presented too fast to individuate or recall, does not mean that ensembles are necessarily unaffected by, isolated from, or calculated before attentional processes. On the contrary, attention sometimes plays an important role in ensemble perception, as will be reviewed in the following section. Therefore, attention (or lack thereof) should not be used as a

diagnostic criterion for the purposes of operationally defining what counts as summary statistical or ensemble coding.

Likewise, memory is not a diagnostic criterion for ensemble perception because ensembles can be formed on sets whose objects are not accurately encoded or recalled (Ariely 2001, Alvarez & Oliva 2008, Haberman & Whitney 2007, Haberman et al. 2009). Similarly, although ensemble representations may inform statistical learning (Fiser & Aslin 2001, Solso & McCarthy 1981), summary statistical perception occurs at first sight for novel stimuli and dimensions (Haberman & Whitney 2007, Yamanashi Leib et al. 2016) and without training or learning (Haberman & Whitney 2012).

4. WHAT IS THE ROLE OF ATTENTION IN PERCEIVING ENSEMBLES?

There remains some debate regarding whether ensemble coding requires attention, whether serial or parallel mechanisms are involved, and whether this is even a valid distinction. Some experiments suggest that directed attention is not necessary for ensemble perception. For example, using a divided attention task, Alvarez & Oliva (2008) found that observers could report the centroid or average final position of clouds of dots with comparable accuracy whether they were attended or ignored. Other studies have also demonstrated that attention directed to individual set members is not necessary to obtain ensemble estimates of color variance (Bronfman et al. 2014), circle size (Chong & Treisman 2005b), and orientation (Alvarez & Oliva 2009). Consistent with these results, ensemble motion perception (Allik 1992, Allik & Dzhafarov 1984, Watamaniuk et al. 1989) and adaptation (Harp et al. 2007) occur even when crowding makes it impossible to individuate each item (Whitney & Levi 2011). Finally, obligatory averaging reported in some domains (Fischer & Whitney 2011, Parkes et al. 2001) could suggest that ensemble percepts are extracted when and where attention cannot be deployed (Joo et al. 2009, Oriet & Brand 2013; but see also McNair et al. 2016).

Although attention may not be necessary for ensemble perception, it may strongly modulate ensemble perception. For example, attention may bias estimates of average set size (Chong & Treisman 2005a, de Fockert & Marchant 2008). Moreover, some studies have shown that dividing attention between two sets of stimuli incurs a cost in performance accuracy (Brand et al. 2012, Huang 2015). Another study showed that diverting attention reduced efficiency—the number of integrated samples—in an ensemble orientation discrimination task (Dakin et al. 2009). Consistent with this finding, limiting attentional resources may reduce the number of faces sampled to estimate average expression (McNair et al. 2016), perhaps by modulating the spatial distribution of integration. Indeed, observers who were primed with a task requiring global attention prior to performing an ensemble perception task performed significantly better than those who were primed with local attention tasks (Chong & Treisman 2005a). Finally, it is possible that attention could limit other processes like the spatial resolution of perception (e.g., crowding), working memory, decision processes, and motor control—which may make ensemble perception appear attention dependent even if it is not directly dependent on attention (Attarha & Moore 2015, de Fockert & Marchant 2008). Thus, although focused attention may not be strictly necessary for ensemble perception to occur, there is ample evidence that attention facilitates it.

Thus, the concepts of directed attention and ensemble perception are not at odds (Allik et al. 2013), and the black-and-white dichotomy between serial and parallel ensemble processing is not especially useful. Taken together, these findings reveal that ensemble perception can be valuable in situations that allow directed attention to a crowd and similarly useful in situations where attention is limited. Instead of exploring an oversimplified dichotomy, future work should attempt

to characterize the interactions between attention and ensemble perception at multiple levels of visual processing to identify the mechanism through which attention facilitates summary statistical representations.

For example, it is worth considering attention's interaction with the hierarchical nature of ensemble perception. Attention may operate at the level of the ensemble even if it does not operate at the level of the individual components within the ensemble. There is substantial evidence for gist representations without explicit knowledge of the individual components, including ensemble orientation discrimination (Parkes et al. 2001), ensemble face recognition (Haberman & Whitney 2007), ensemble change detection (Haberman & Whitney 2011), and ensemble size discrimination (Allik et al. 2013, Chong & Treisman 2005b, Choo & Franconeri 2010, Oriet & Brand 2013). In each of these cases, task demands required observers to attend to the ensemble characteristic as a whole, so attention to the relevant ensemble dimension may have been necessary even if awareness of the individual member was not (see Section 3.4). A promising recent approach to addressing this issue is to measure aftereffects following adaptation to summary statistical information. In such an experiment, observers' attention (and awareness) can be controlled during the adaptation period. Prior adaptation studies have investigated adaptation to an ensemble statistic, such as average size or motion direction (Anstis et al. 1998, Corbett et al. 2012, Harp et al. 2007), without rigorously controlling for attention. One recent study that did control for attention found adaptation to the average facial expression in a rapid sequence of faces (Ying & Xu 2017). Future studies can use this and related adaptation paradigms to isolate the particular role of attention in ensemble representations.

5. USEFULNESS OF ENSEMBLE REPRESENTATIONS

Ensemble representations might be the basis of some of our fastest and richest perceptual experiences (Intraub 1981, Potter 1975, Thorpe et al. 1996), which do not rely on explicitly or consciously representing all of the individual members of the scene (see Section 3.4). For example, we can perceive the average liveliness of a scene in the briefest of glances (Yamanashi Leib et al. 2016). A face that is crowded such that it is unrecognizable nonetheless influences the ensemble expression perceived in the crowd (Fischer & Whitney 2011), and a changing object that goes unnoticed or unrecognized can still alter the perceived ensemble property of the scene as a whole (Haberman & Whitney 2011, Ward et al. 2016). One interpretation of these and similar findings is that the individual objects are phenomenally available to consciousness but unreportable (Block 2011, McClelland & Bayne 2016, Shea & Bayne 2010). Alternatively, the visual system might encode summary statistical information in crowds of objects by unconsciously processing individual object identities (Chaney et al. 2014, Cohen et al. 2016). In either case, the resulting percept is richer than would be expected when faced with the limits of visual short-term memory, cognition, language, or attention (Cohen et al. 2016). Therefore, much of what counts as our rich visual experience may take the form of ensemble representations (Block 2011, Cohen et al. 2016, McClelland & Bayne 2016).

Ensemble representations may be especially useful at the highest levels of perceptual processing because they carry emergent and social information—unique characteristics of crowds, environments, and social interactions that can only be specified at the level of the group. For example, observers are sensitive to the ambivalence, mixture of emotion, or racial diversity of a crowd, but these cannot be conveyed at the level of individual faces (Haberman et al. 2015b). Other emergent ensemble percepts may include the overall threat of a crowd, its gaze direction (Mareschal et al. 2016, Sweeny & Whitney 2014), and its heading direction (Sweeny et al. 2013). For example, the perception of crowd panic is probably based on summary statistical information involving a

calculation of heading direction (Sweeny et al. 2013) as well as variance in direction and speed. Visuosocial summary statistical information not only is important for recognition and awareness but also serves as an important cue to guide action: Crucial behavioral decisions in crowd navigation (where to walk next, speed of walking, etc.) may be driven by summary statistical information. Finally, ensemble information may be constructive, amplifying the perception of summary statistical dimensions (Price et al. 2014). For example, the perceived attractiveness of faces in a crowd may be exaggerated (van Osch et al. 2015, Walker & Vul 2013).

Ensemble representations may be a critical component of visual working and long-term memory. Recent research suggests that the average circle size in scenes biases subsequent estimates of individual object size in memory tasks (Brady & Alvarez 2011), and that recalled locations of individual objects in a cluster are pulled toward the ensemble centroid location (Lew & Vul 2015). Thus, individual objects in memory are not treated simply as independent entities but as part of a hierarchy that includes information about individual details and ensembles (Brady & Alvarez 2011). This is advantageous because statistical structure or ensemble information affords more information than would be available from only encoding independent individual objects. For example, grouping proximate sets of circles by ensemble characteristics increases the capacity of visual working memory (VWM) (Im & Chong 2014), and ensemble representations might also facilitate statistical learning and category boundary formation (Oriet & Hozempa 2016). Insofar as memory consolidation results in summary statistical-like representations (Richards et al. 2014), an intriguing possibility is that ensemble coding might also facilitate long-term memory for outliers, just as summary statistical perception improves deviance detection and pop-out in visual search (Whitney et al. 2014). More broadly, we can better model and understand the mechanisms of memory if we incorporate the important role of ensemble representations into theories of memory encoding and consolidation (Brady & Tenenbaum 2013).

6. POSSIBLE PHYSIOLOGICAL AND COGNITIVE MECHANISMS OF ENSEMBLE PERCEPTION

The relationship between ensemble perception and other visual phenomena has raised the possibility that they result from shared mechanisms. One of the most common associations we can examine for insight on these mechanisms is that between ensemble coding and visual crowding. Visual crowding is the deleterious effect of clutter on object recognition and awareness in the peripheral visual field (Pelli 2008, Whitney & Levi 2011). Traditional models of visual perception and crowding argue that the visual system lacks the bandwidth to encode detailed information outside of the fovea, so peripheral high-fidelity visual information is irreversibly lost. Instead, what might emerge from crowded scenes is summary statistical information. In that sense, crowding and ensemble percepts might be thought of as two sides of the same coin (Parkes et al. 2001) or caused by similar pooling processes (Balas et al. 2009, Freeman & Simoncelli 2011). However, ensemble perception can occur with or without crowding (Bulakowski et al. 2011, Dakin et al. 2009), so ensemble representations do not require crowding.

Models of crowding also face challenges in describing ensemble perception. Pooling models, which often use variations of texture synthesis algorithms (Portilla & Simoncelli 2000), may help account for some aspects of low-level texture perception in crowding (Balas et al. 2009, Freeman & Simoncelli 2011; but see also Wallis et al. 2016). However, these models do not operate at the level of object representations and, thus, cannot explain object-level crowding (Farzin et al. 2009, Ikeda et al. 2013, Kimchi & Pirkner 2015, Louie et al. 2007, Manassi et al. 2012) or high-level ensemble perception (Whitney et al. 2014). Nor can these models explain how crowded and unrecognized faces can prime subsequent valence judgments (Faivre et al. 2012, Kouider et al. 2011) or how these

faces can influence ensemble expression perception (Fischer & Whitney 2011). Current pooling and texture synthesis models (Balas et al. 2009, Freeman & Simoncelli 2011, Pelli 2008, Rosenholtz et al. 2012) would not generate these effects. Likewise, scene gist (Oliva & Torralba 2006) and texture models stop short of describing object-level processing and cannot explain higher-level ensemble percepts, such as average identity, expression, and animacy or the perceived variance of these properties (de Fockert & Wolfenstein 2009, Haberman & Whitney 2007, Haberman et al. 2015a, Neumann et al. 2013, Yamanashi Leib et al. 2016).

The research discussed above challenges any model of ensemble coding to explain how object information can be unrecognizable due to crowding but retained for subsequent ensemble perception (Faivre et al. 2012, Fischer & Whitney 2011, Ikeda et al. 2013, Parkes et al. 2001). One way in which this problem could be reconciled is if the visual system maintains and passes high-fidelity representations through each level of visual analysis but these representations cannot be selected with sufficient resolution to recognize particular objects (Chaney et al. 2014). According to this hierarchical sparse selection model, when observers attempt to select a feature or object in a crowd, they may not be able to resolve that feature or object, although they can extract an ensemble of that feature or object dimension. Thus, we see that crowding occurs, and ensembles can be extracted, at multiple levels of visual processing (Whitney & Levi 2011). Modified versions of the pooling models mentioned above could be implemented hierarchically at multiple levels of visual processing along these lines to explain the multiple levels of crowding and ensemble perception. An updated reverse hierarchy visual model (Hochstein et al. 2015) may also be invoked to help explain how high-fidelity individual object information contributes to ensemble percepts even when crowding severely impedes recognition of that object.

The neural mechanism(s) of ensemble perception remain unknown, but it is unlikely that there is a single unified mechanism. Psychophysical evidence suggests that ensemble representations are distributed, or available at multiple levels of visual processing. For example, the individual differences in ensemble perception for low- and high-level objects do not correlate well, suggesting that there may be independent mechanisms for different types of ensembles (Haberman et al. 2015a, Sweeny et al. 2015). This interpretation also potentially explains the lack of patients who exhibit a unique deficit in ensemble processing. Many patients do not have impaired ensemble discrimination or exhibit less difficulty than would be predicted based on their single-item impairments alone (Demeyere et al. 2008; Hochstein et al. 2015; Karaminis et al. 2017; Yamanashi Leib et al. 2012a,b, 2014; cf. Rhodes et al. 2015). There is also physiological and neuroimaging evidence for multiple stages of ensemble representation, including motion, color, and textures, in occipital visual areas (e.g., Okazawa et al. 2015), MT+ (Born & Bradley 2005) and the anterior-medial ventral visual cortex (Cant & Xu 2012). The neural loci and mechanisms involved in coding crowds of faces, biological motion, and animacy remain to be explored. Because ensemble representations can be easily calculated from population codes that occur at nearly every level of visual processing, a distinct possibility is that ensemble representations will be found at virtually any stage examined (Chaney et al. 2014).

7. CONCLUSIONS

Ensemble perception is ubiquitous. It occurs at multiple levels of visual analysis, ranging from low-level orientation processing to high-level social impressions (e.g., the emotional tenor or the liveliness of a crowd). In this review, we have proposed an operational definition of perceptual ensemble coding, which includes five factors: the perception of a statistical moment in a crowd, the integration of multiple stimuli (approximately the square root of the number of stimuli in the scene), precise representation of the ensemble property, lack of a requirement for sensitivity

to particular individual set members, and high temporal resolution. This operational definition helps distinguish ensemble perception from other phenomena. Ensemble perception provides an efficient way to access group-level information in a quick glance without the need to scrutinize or recall individual objects. Even in circumstances when single-item analysis may be possible, ensemble perception of the integrated group provides emergent, functionally useful information that cannot be attained from any single group member. The accumulating evidence reviewed above suggests that summary statistical perception is a significant contributing factor to visual perception and may generate much of what contributes to a rich conscious experience during rapid, first-glance assessments of visual scenes.

FUTURE ISSUES

1. How stable are the individual differences in ensemble perception? How do group differences such as culture, gender, and (typical and atypical) cognitive development influence ensemble perception?
2. How does ensemble perception interact with visual search functions, such as outlier detection or pop-out? For example, are individual differences in visual search performance predicted by or correlated with the fidelity of summary statistical representations?
3. What is the capacity limit for multiple parallel ensembles? Do multiple-ensemble representations interact with each other, and how do multiple ensembles inform fast scene recognition?
4. How is ensemble information integrated across modalities within scenes? For example, is the perception of ensemble biological motion in the whole scene (e.g., an orchestra) independent from or integrated with ensemble auditory information from the parts of the scene (e.g., the collection of instruments in the orchestra)?
5. Can ensemble perception be trained and improved? What cognitive, social, and developmental benefits are conferred by enhanced ensemble perception? Does it enhance emotional intelligence? Does it improve interactions with the world, such as driving? Does it improve scene recognition or the fidelity of memory and richness of conscious experience?

DISCLOSURE STATEMENT

The authors are not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

ACKNOWLEDGMENTS

Thanks to Jason Haberman, Mauro Manassi, and Tim Sweeny for helpful discussions. Thanks to Zixuan Wang for editorial assistance. This work was supported by National Science Foundation grant 074689 and National Institutes of Health grant EY018216.

LITERATURE CITED

Albrecht AR, Scholl BJ. 2010. Perceptually averaging in a continuous visual world: extracting statistical summary representations over time. *Psychol. Sci.* 21(4):560–67

Albrecht AR, Scholl BJ, Chun MM. 2012. Perceptual averaging by eye and ear: computing summary statistics from multimodal stimuli. *Atten. Percept. Psychophys.* 74(5):810–15
Allik J. 1992. Competing motion paths in sequence of random dot patterns. *Vis. Res.* 32(1):157–65
Allik J, Dzhafarov EN. 1984. Motion direction identification in random cinematograms: a general model. *J. Exp. Psychol. Hum. Percept. Perform.* 10(3):378–93
Allik J, Toom M, Raidvee A, Averin K, Kreegipuu K. 2013. An almost general theory of mean size perception. *Vis. Res.* 83:25–39
Allik J, Toom M, Raidvee A, Averin K, Kreegipuu K. 2014. Obligatory averaging in mean size perception. *Vis. Res.* 101:34–40
Alvarez GA. 2011. Representing multiple objects as an ensemble enhances visual cognition. *Trends Cogn. Sci.* 15(3):122–31
Alvarez GA, Oliva A. 2008. The representation of simple ensemble visual features outside the focus of attention. *Psychol. Sci.* 19(4):392–98
Alvarez GA, Oliva A. 2009. Spatial ensemble statistics are efficient codes that can be represented with reduced attention. *PNAS* 106(18):7345–50
Anderson NH, Lindner R, Lopes LL. 1973. Integration theory applied to judgments of group attractiveness. *J. Personal. Soc. Psychol.* 26(3):400–8
Anstis S. 1974. A chart demonstrating variations in acuity with retinal position. *Vis. Res.* 14:589–92
Anstis S, Verstraten FAJ, Mather G. 1998. The motion aftereffect. *Trends Cogn. Sci.* 2(3):111–17
Ariely D. 2001. Seeing sets: representation by statistical properties. *Psychol. Sci.* 12(2):157–62
Attarha M, Moore CM. 2015. The capacity limitations of orientation summary statistics. *Atten. Percept. Psychophys.* 7(4):1116–31
Attarha M, Moore CM, Vecera SP. 2014. Summary statistics of size: fixed processing capacity for multiple ensembles but unlimited processing capacity for single ensembles. *J. Exp. Psychol. Hum. Percept. Perform.* 40(4):1440–49
Bai Y, Yamanashi Leib A, Puri AM, Whitney D, Peng K. 2015. Gender differences in crowd perception. *Front. Psychol.* 6:1300
Balas B, Nakano L, Rosenholtz R. 2009. A summary-statistic representation in peripheral vision explains visual crowding. *J. Vis.* 9(12):13
Banno H, Saiki J. 2012. Calculation of the mean circle size does not circumvent the bottleneck of crowding. *J. Vis.* 12(11):13
Bauer B. 2009. Does Stevens's power law for brightness extend to perceptual brightness averaging? *Psychol. Res.* 59:171–86
Block N. 2011. Perceptual consciousness overflows cognitive access. *Trends Cogn. Sci.* 15(12):567–75
Born RT, Bradley DC. 2005. Structure and function of visual area MT. *Annu. Rev. Neurosci.* 28:157–89
Brady TF, Alvarez GA. 2011. Hierarchical encoding in visual working memory: Ensemble statistics bias memory for individual items. *Psychol. Sci.* 22(3):384–92
Brady TF, Tenenbaum JB. 2013. A probabilistic model of visual working memory: incorporating higher order regularities into working memory capacity estimates. *Psychol. Rev.* 120(1):85–109
Brand J, Oriet C, Tottenham LS. 2012. Size and emotion averaging: costs of dividing attention after all. *Can. J. Exp. Psychol.* 66(1):63–69
Bronfman ZZ, Brezis N, Jacobson H, Usher M. 2014. We see more than we can report: "cost free" color phenomenality outside focal attention. *Psychol. Sci.* 25(7):1394–403
Bulakowski PF, Post RB, Whitney D. 2011. Dissociating crowding from ensemble percepts. *Atten. Percept. Psychophys.* 73(4):1003–9
Cant JS, Xu Y. 2012. Object ensemble processing in human anterior-medial ventral visual cortex. *J. Neurosci.* 32(22):7685–700
Cavanagh P, Alvarez GA. 2005. Tracking multiple targets with multifocal attention. *Trends Cogn. Sci.* 9(7):349–54
Chaney W, Fischer J, Whitney D. 2014. The hierarchical sparse selection model of visual crowding. *Front. Integr. Neurosci.* 8:73
Cheadle S, Wyart V, Tsetsos K, Myers N, de Gardelle V, et al. 2014. Adaptive gain control during human perceptual choice. *Neuron* 81(6):1429–41

Chong SC, Joo SJ, Emmanouil TA, Treisman A. 2008. Statistical processing: not so implausible after all. *Percept. Psychophys.* 70(7):1327–34
Chong SC, Treisman A. 2003. Representation of statistical properties. *Vis. Res.* 43(4):393–404
Chong SC, Treisman A. 2005a. Attentional spread in the statistical processing of visual displays. *Percept. Psychophys.* 67(1):1–13
Chong SC, Treisman A. 2005b. Statistical processing: computing the average size in perceptual groups. *Vis. Res.* 45(7):891–900
Choo H, Franconeri SL. 2010. Objects with reduced visibility still contribute to size averaging. *Atten. Percept. Psychophys.* 72(1):86–99
Cohen MA, Dennett DC, Kanwisher N. 2016. What is the bandwidth of perceptual experience? *Trends Cogn. Sci.* 20(5):324–35
Corbett JE, Wurnitsch N, Schwartz A, Whitney D. 2012. An aftereffect of adaptation to mean size. *Vis. Cogn.* 20(2):211–31
Dakin SC. 2001. Information limit on the spatial integration of local orientation signals. *J. Opt. Soc. Am. A* 18(5):1016–26
Dakin SC, Bex PJ, Cass JR, Watt RJ. 2009. Dissociable effects of attention and crowding on orientation averaging. *J. Vis.* 9(11):28
Dakin SC, Mareschal I, Bex PJ. 2005. Local and global limitations on direction integration assessed using equivalent noise analysis. *Vis. Res.* 45(24):3027–49
Dakin SC, Watt RJ. 1997. The computation of orientation statistics from visual texture. *Vis. Res.* 37(22):3181–92
de Fockert JW, Marchant AP. 2008. Attention modulates set representation by statistical properties. *Percept. Psychophys.* 70(5):789–94
de Fockert JW, Wolfenstein C. 2009. Rapid extraction of mean identity from sets of faces. *Q. J. Exp. Psychol.* 62(9):1716–22
Demeyere N, Rzeskiewicz A, Humphreys KA, Humphreys GW. 2008. Automatic statistical processing of visual properties in simultanagnosia. *Neuropsychologia* 46(11):2861–64
Dux PE, Marois R. 2009. How humans search for targets through time: a review of data and theory from the attentional blink. *Atten. Percept. Psychophys.* 71(8):1683–700
Emmanouil TA, Treisman A. 2008. Dividing attention across feature dimensions in statistical processing of perceptual groups. *Percept. Psychophys.* 70(6):946–54
Faivre N, Berthet V, Kouider S. 2012. Nonconscious influences from emotional faces: a comparison of visual crowding, masking, and continuous flash suppression. *Front. Psychol.* 3:129
Farzin F, Rivera SM, Whitney D. 2009. Holistic crowding of Mooney faces. *J. Vis.* 9(6):18
Fischer J, Whitney D. 2011. Object-level visual information gets through the bottleneck of crowding. *J. Neurophysiol.* 106(3):1389–98
Fiser J, Aslin RN. 2001. Unsupervised statistical learning of higher-order spatial structures from visual scenes. *Psychol. Sci.* 12(6):499–504
Florey J, Clifford CWG, Dakin S, Mareschal I. 2016. Spatial limitations in averaging social cues. *Sci. Rep.* 6:32210
Florey J, Dakin SC, Mareschal I. 2017. Comparing averaging limits for social cues over space and time. *J. Vis.* 17(9):17
Fouriezos G, Rubenfeld S, Capstick G. 2008. Visual statistical decisions. *Percept. Psychophys.* 70(3):456–64
Freeman J, Simoncelli EP. 2011. Metamers of the ventral stream. *Nat. Neurosci.* 14(9):1195–201
Galton F. 1907. Vox populi (the wisdom of crowds). *Nature* 75:450–51
Gorea A, Belkoura S, Solomon JA. 2014. Summary statistics for size over space and time. *J. Vis.* 14(9):22
Haberman J, Brady TF, Alvarez GA. 2015a. Individual differences in ensemble perception reveal multiple, independent levels of ensemble representation. *J. Exp. Psychol. Gen.* 144(2):432–46
Haberman J, Harp T, Whitney D. 2009. Averaging facial expression over time. *J. Vis.* 9(11):1
Haberman J, Lee P, Whitney D. 2015b. Mixed emotions: sensitivity to facial variance in a crowd of faces. *J. Vis.* 15(4):16
Haberman J, Whitney D. 2007. Rapid extraction of mean emotion and gender from sets of faces. *Curr. Biol.* 17(7):751–53

Haberman J, Whitney D. 2009. Seeing the mean: ensemble coding for sets of faces. *J. Exp. Psychol. Hum. Percept. Perform.* 35(3):718–34
Haberman J, Whitney D. 2010. The visual system discounts emotional deviants when extracting average expression. *Atten. Percept. Psychophys.* 72(7):1825–38
Haberman J, Whitney D. 2011. Efficient summary statistical representation when change localization fails. *Psychon. Bull. Rev.* 18(5):855–59
Haberman J, Whitney D. 2012. Ensemble perception: summarizing the scene and broadening the limits of visual processing. In *From Perception to Consciousness: Searching with Anne Treisman*, ed. J Wolfe, L Robinson, pp. 1–21. Oxford, UK: Oxford Univ. Press
Harp TD, Bressler DW, Whitney D. 2007. Position shifts following crowded second-order motion adaptation reveal processing of local and global motion without awareness. *J. Vis.* 7(2):15
Hochstein S, Pavlovskaya M, Bonneh YS, Soroker N. 2015. Global statistics are not neglected. *J. Vis.* 15(4):7
Huang L. 2015. Statistical properties demand as much attention as object features. *PLOS ONE* 10(8):e0131191
Hubert-Wallander B, Boynton GM. 2015. Not all summary statistics are made equal: evidence from extracting summaries across time. *J. Vis.* 15(4):5
Ikeda H, Watanabe K, Cavanagh P. 2013. Crowding of biological motion stimuli. *J. Vis.* 13(4):20
Im HY, Chong SC. 2014. Mean size as a unit of visual working memory. *Perception* 43(7):663–76
Im HY, Halberda J. 2013. The effects of sampling and internal noise on the representation of ensemble average size. *Atten. Percept. Psychophys.* 75(2):278–86
Intraub H. 1981. Rapid conceptual identification of sequentially presented pictures. *J. Exp. Psychol. Hum. Percept. Perform.* 7(3):604–10
Jacoby O, Kamke MR, Mattingley JB. 2012. Is the whole really more than the sum of its parts? Estimates of average size and orientation are susceptible to object substitution masking. *J. Exp. Psychol. Hum. Percept. Perform.* 39(1):233–44
Ji L, Chen W, Fu X. 2014. Different roles of foveal and extrafoveal vision in ensemble representation for facial expressions. In *Engineering Psychology and Cognitive Ergonomics*, ed. D Harris, pp. 164–73. Berlin: Springer
Joo SJ, Shin K, Chong SC, Blake R. 2009. On the nature of the stimulus information necessary for estimating mean size of visual arrays. *J. Vis.* 9(9):7
Karaminis T, Neil L, Manning C, Turi M, Fiorentini C, et al. 2017. Investigating ensemble perception of emotions in autistic and typical children and adolescents. *Dev. Cogn. Neurosci.* 24:51–62
Kimchi R, Pirkner Y. 2015. Multiple level crowding: crowding at the object parts level and at the object configural level. *Perception* 44(11):1275–92
Koenderink JJ, Van Doorn AJ, Pont SC. 2004. Light direction from shad(ow)ed random Gaussian surfaces. *Perception* 33(12):1405–20
Kouider S, Berthet V, Faivre N. 2011. Preference is biased by crowded facial expressions. *Psychol. Sci.* 22(2):184–89
Kowler E. 2011. Eye movements: the past 25 years. *Vis. Res.* 51(13):1457–83
Landy MS. 2014. Texture analysis and perception. In *The New Visual Neurosciences*, ed. JS Werner, LM Chalupa, pp. 639–52. Cambridge, MA: MIT Press
Leon M, Oden GC, Anderson NH. 1973. Functional measurement of social values. *J. Personal. Soc. Psychol.* 27(3):301–10
Lew TF, Vul E. 2015. Ensemble clustering in visual working memory biases location memories and reduces the Weber noise of relative positions. *J. Vis.* 15(4):10
Li H, Ji L, Tong K, Ren N, Chen W, et al. 2016. Processing of individual items during ensemble coding of facial expressions. *Front. Psychol.* 7:1332
Louie EG, Bressler DW, Whitney D. 2007. Holistic crowding: selective interference between configural representations of faces in crowded scenes. *J. Vis.* 7(2):24
Luck SJ, Vogel EK. 2013. Visual working memory capacity: from psychophysics and neurobiology to individual differences. *Trends Cogn. Sci.* 17(8):391–400
Manassi M, Sayim B, Herzog MH. 2012. Grouping, pooling, and when bigger is better in visual crowding. *J. Vis.* 12(10):13
Marchant AP, Simons DJ, de Fockert JW. 2013. Ensemble representations: effects of set size and item heterogeneity on average size perception. *Acta Psychol.* 142(2):245–50

Mareschal I, Otsuka Y, Clifford MD. 2016. "Are you looking at me?" How children's gaze judgments improve with age. *Dev. Psychol.* 52:695–703
Maule J, Franklin A. 2016. Accurate rapid averaging of multihue ensembles is due to a limited capacity subsampling mechanism. *J. Opt. Soc. Am. A.* 33:A22–29
McClelland T, Bayne T. 2016. Ensemble coding and two conceptions of perceptual sparsity. *Trends Cogn. Sci.* 20(9):641–42
McDermott J, Schemitsch M, Simoncelli EP. 2013. Summary statistics in auditory perception. *Nat. Neurosci.* 16(4):493–98
McNair NA, Goodbourn PT, Shone LT, Harris IM. 2016. Summary statistics in the attentional blink. *Atten. Percept. Psychophys.* 79:100–16
Melcher D, Kowler E. 1999. Shapes, surfaces and saccades. *Vis. Res.* 39(17):2929–46
Michael E, de Gardelle V, Summerfield C. 2014. Priming by the variability of visual information. *PNAS* 111(21):7873–78
Miller AL, Sheldon R. 1969. Magnitude estimation of average length and average inclination. *J. Exp. Psychol.* 81(1):16–21
Morgan MJ, Chubb C, Solomon JA. 2008. A "dipper" function for texture discrimination based on orientation variance. *J. Vis.* 8(11):9
Myczek K, Simons DJ. 2008. Better than average: alternatives to statistical summary representations for rapid judgments of average size. *Percept. Psychophys.* 70(5):772–88
Navon D. 1977. Forest before trees: the precedence of global features in visual perception. *Cogn. Psychol.* 9:353–83
Neumann MF, Schweinberger SR, Burton AM. 2013. Viewers extract mean and individual identity from sets of famous faces. *Cognition* 128(1):56–63
Norman LJ, Heywood CA, Kentridge RW. 2015. Direct encoding of orientation variance in the visual system. *J. Vis.* 15:3
Nothdurft HC. 1993. Faces and facial expressions do not pop out. *Perception* 22(11):1287–98
Okazawa G, Tajima S, Komatsu H. 2015. Image statistics underlying natural texture selectivity of neurons in macaque V4. *PNAS* 112(4):E351–60
Oliva A, Torralba A. 2006. Building the gist of a scene: the role of global image features in recognition. *Prog. Brain Res.* 155:23–36
Oriet C, Brand J. 2013. Size averaging of irrelevant stimuli cannot be prevented. *Vis. Res.* 79:8–16
Oriet C, Corbett JE. 2008. Evidence for rapid extraction of average size in RSVP displays of circles. *J. Vis.* 8(6):13
Oriet C, Hozempa K. 2016. Incidental statistical summary representation over time. *J. Vis.* 16(3):3
Parkes L, Lund J, Angelucci A. 2001. Compulsory averaging of crowded orientation signals in human vision. *Nat. Neurosci.* 4(7):739–44
Pelli DG. 2008. Crowding: a cortical constraint on object recognition. *Curr. Opin. Neurobiol.* 18(4):445–51
Piazza EA, Sweeny TD, Wessel D, Silver MA, Whitney D. 2013. Humans use summary statistics to perceive auditory sequences. *Psychol. Sci.* 24(8):1389–97
Portilla J, Simoncelli EP. 2000. A parametric texture model based on joint statistics of complex wavelet coefficients. *Int. J. Comput. Vis.* 40(1):49–71
Post RB, Haberman J, Iwaki L, Whitney D. 2012. The frozen face effect: why static photographs may not do you justice. *Front. Psychol.* 3:22
Potter M. 1975. Meaning in visual search. *Science* 187(4180):965–66
Price PC, Kimura NM, Smith AR, Marshall LD. 2014. Sample size bias in judgments of perceptual averages. *J. Exp. Psychol. Learn.* 40(5):1321–31
Rhodes G, Neumann MF, Ewing L, Palermo R. 2015. Reduced set averaging of face identity in children and adolescents with autism. *Q. J. Exp. Psychol.* 68(7):1391–403
Richards BA, Xia F, Santoro A, Husse J, Woodin MA, et al. 2014. Patterns across multiple memories are identified over time. *Nat. Neurosci.* 17(7):981–86
Robitaille N, Harris IM. 2011. When more is less: extraction of summary statistics benefits from larger sets. *J. Vis.* 11(12):18

Rosenholtz R, Huang J, Ehinger KA. 2012. Rethinking the role of top-down attention in vision: effects attributable to a lossy representation in peripheral vision. *Front. Psychol.* 3:13

Sanders K, Haberman J, Whitney D. 2010. Mean representation beyond a shadow of a doubt: summary statistical representation of shadows and lighting direction. *J. Vis.* 8(6):334

Schwarzkopf DS, Song C, Rees G. 2011. The surface area of human V1 predicts the subjective experience of object size. *Nat. Neurosci.* 14(1):28–30

Shea N, Bayne T. 2010. The vegetative state and the science of consciousness. *Br. J. Philos. Sci.* 61(3):459–84

Simons DJ, Levin DT. 1997. Change blindness. *Trends Cogn. Sci.* 1(7):261–67

Solomon JA. 2010. Visual discrimination of orientation statistics in crowded and uncrowded arrays. *J. Vis.* 10(14):19

Solomon JA, Morgan MJ, Chubb C. 2011. Efficiencies for the statistics of size discrimination. *J. Vis.* 11(12):13

Solso RL, McCarthy JE. 1981. Prototype formation of faces: a case of pseudo-memory. *Br. J. Psychol.* 72(4):499–503

Strasburger H, Rentschler I, Jüttner M. 2011. Peripheral vision and pattern recognition: a review. *J. Vis.* 11(5):13

Surowiecki J. 2004. *The Wisdom of Crowds*. New York: Random House

Sweeny TD, Haroz S, Whitney D. 2013. Perceiving group behavior: sensitive ensemble coding mechanisms for biological motion of human crowds. *J. Exp. Psychol. Hum. Percept. Perform.* 39(2):329–37

Sweeny TD, Whitney D. 2014. Perceiving crowd attention: ensemble perception of a crowd's gaze. *Psychol. Sci.* 25(10):1903–13

Sweeny TD, Wurnitsch N, Gopnik A, Whitney D. 2015. Ensemble perception of size in 4–5-year-old children. *Dev. Sci.* 18(4):556–68

Thorpe S, Fize D, Marlot C. 1996. Speed of processing in the human visual system. *Nature* 381(6582):520–22

Tong F, Nakayama K. 1999. Robust representation for faces: evidence from visual search. *J. Exp. Psychol. Hum. Percept. Perform.* 25(4):1016–35

van Osch Y, Blanken I, Meijs MHJ, van Wolferen J. 2015. A group's physical attractiveness is greater than the average attractiveness of its members: the group attractiveness effect. *Personal. Soc. Psychol. Bull.* 41(4):559–74

Verstraten FAJ, Cavanagh P, Labianca A. 2000. Limits of attentive tracking reveal temporal properties of attention. *Vis. Res.* 40(26):3651–64

Virsu V, Rovamo J. 1979. Visual resolution, contrast sensitivity, and the cortical magnification factor. *Exp. Brain Res.* 37(3):475–94

Vishwanath D, Kowler E. 2003. Localization of shapes: eye movements and perception compared. *Vis. Res.* 43(15):1637–53

Wagemans J, Elder JH, Kubovy M, Palmer SE, Peterson MA, et al. 2012. A century of Gestalt psychology in visual perception: I. Perceptual grouping and figure-ground organization. *Psychol. Bull.* 138(6):1172–217

Walker D, Vul E. 2013. Hierarchical encoding makes individuals in a group seem more attractive. *Psychol. Sci.* 25(1):230–35

Wallis TSA, Bethge M, Wichmann FA. 2016. Testing models of peripheral encoding using metamerism in an oddity paradigm. *J. Vis.* 16(2):4

Ward EJ, Bear A, Scholl BJ. 2016. Can you perceive ensembles without perceiving individuals? The role of statistical perception in determining whether awareness overflows access. *Cognition* 152:78–86

Wardle SG, Bex PJ, Alais D. 2017. Stereoacuity in the periphery is limited by internal noise. *J. Vis.* 12(6):12

Watamaniuk SN, Duchon A. 1992. The human visual system averages speed information. *Vis. Res.* 32(5):931–41

Watamaniuk SN, McKee SP. 1998. Simultaneous encoding of direction at a local and global scale. *Percept. Psychophys.* 60(2):191–200

Watamaniuk SN, Sekuler R, Williams DW. 1989. Direction perception in complex dynamic displays: the integration of direction information. *Vis. Res.* 29(1):47–59

Webster J, Kay P, Webster MA. 2014. Perceiving the average hue of color arrays. *J. Opt. Soc. Am. A* 31(4):A283–92

Whiting BF, Oriet C. 2011. Rapid averaging? Not so fast! *Psychon. Bull. Rev.* 18(3):484–89

Whitney D, Haberman J, Sweeny T. 2014. From textures to crowds: multiple levels of summary statistical perception. In *The New Visual Neuroscience*, ed. JS Werner, LM Chalupa, pp. 695–710. Cambridge, MA: MIT Press

Whitney D, Levi DM. 2011. Visual crowding: a fundamental limit on conscious perception and object recognition. *Trends Cogn. Sci.* 15(4):160–68

Wolfe BA, Kosovicheva AA, Yamanashi Leib A, Wood K, Whitney D. 2015. Foveal input is not required for perception of crowd facial expression. *J. Vis.* 15(4):11

Wolfe JM. 1994. Guided Search 2.0: a revised model of visual search. *Psychon. Bull. Rev.* 1(2):202–38

Yamanashi Leib A, Fischer J, Liu Y, Qiu S, Robertson L, et al. 2014. Ensemble crowd perception: a viewpoint-invariant mechanism to represent average crowd identity. *J. Vis.* 14(8):26

Yamanashi Leib A, Kosovicheva A, Whitney D. 2016. Fast ensemble representations for abstract visual impressions. *Nat. Commun.* 7:13186

Yamanashi Leib A, Landau AN, Baek Y, Chong SC, Robertson L. 2012a. Extracting the mean size across the visual field in patients with mild, chronic unilateral neglect. *Front. Hum. Neurosci.* 6:267

Yamanashi Leib A, Puri AM, Fischer J, Bentin S, Whitney D, Robertson L. 2012b. Crowd perception in prosopagnosia. *Neuropsychologia* 50(7):1698–707

Ying H, Xu H. 2017. Adaptation reveals that facial expression averaging occurs during rapid serial presentation. *J. Vis.* 17(1):15