A quantitative model of ensemble perception as summed activation in feature space

Maria M. Robinson and Timothy F. Brady

Abstract

Ensemble perception is a process by which we extract and consolidate redundancies in the environment, getting the ‘gist’ of a set of objects. This ability to summarize complex scenes is thought to underlie our ability to construct robust memory representations, guide attention and classify features or objects when many objects are present. Despite the relevance of ensemble perception to everyday cognition, there are few computational models that provide a formal, process-level account of ensemble perception across a range of experimental conditions and stimuli. In the current work, we take a conceptually novel approach towards developing such a model. According to the proposed theory of how ensemble representations are computed, items evoke distributed patterns of familiarity over feature space (e.g., a color evokes a pattern of familiarity over all color channels), and ensemble representations reflect the global sum of these signals across all individual items. We leverage this set of minimal assumptions to formally connect a model of memory for individual items to a model for ensembles. Our approach involves using parameter estimates of performance on a visual memory task for individual items, to predict individual differences in performance on an ensemble continuous report task. That is, we account for the full distribution of errors on the ensemble task with zero free parameters. We evaluate our ensemble model against a set of newly developed alternative models in five ensemble tasks. Our top-down modeling approach is the first to formally unify models of memory for individual items and ensembles, and opens a novel venue for building and comparing models of distinct memory processes and representations. We discuss our model and results in the context of current theories and models of ensemble processing, gist memory and neural population coding.

Human perception and cognition are grounded in a capacity limited system (e.g., Baddale, 1992; Miller, 1956; Pashler, 1984; Simon, 1990). For this reason, a basic question across research areas in the behavioral and social sciences is how people effectively represent and navigate an environment that should far exceed their processing capabilities (e.g., Kahneman, 2003). A widely accepted answer to this question is that perceptual and cognitive systems take advantage of redundancies in the environment, for example, by making use of a condensed summary or gist rather than representing all of the details of the world (Ariely, 2001; Brady & Oliva, 2008; Goldstein, et al., 2010). In the visual domain, the ability of people to extract summaries by making use of such redundancies is typically referred to as *ensemble perception* (Whitney & Yamanashi Leib, 2018). In a typical laboratory ensemble task, participants are shown a set of stimuli that share properties in a specific feature dimension (e.g., color, orientation, motion direction, face emotion), and are instructed to report on their average along that dimension. For instance, as depicted in Figure 1, participants might be shown a set of colored objects and be instructed to report on their average colors, or be shown a set of shapes and instructed to report on the average shape of the set. People are remarkably accurate at this task, and often notably better at reporting the average of the set than they are at reporting on any one individual item (e.g., Alvarez & Oliva, 2008). Extensive empirical and theoretical work suggests that ensemble processing of this type partially underlies our ability to create more robust representations of simple scenes (Alvarez & Oliva, 2008; Brady, Shafer-Skelton, Alvarez, 2017), categorize objects (Utochkin, 2015) and guide our attention (Balas, Nakano & Rosenholtz, 2009). Such tasks are also claimed to lead to critical insights about the nature of not only cognitive limits but also conscious perception: for example, preserved ensemble information in the relative absence of information about individuals has been interpreted as showing that perceptual consciousness overflows conscious access (e.g., Block, 2011; Block, 2014; c.f. Cohen, Dennett & Kanwisher, 2016).

Fig. 1 | Laboratory ensemble tasks. Examples of laboratory ensemble tasks in which participants are typically asked to report on the average along a stimulus dimension, such as color or shape, using a continuous reproduction task.

Given the fundamental role of ensemble perception in everyday behavior and theories of perception, there is immense value in developing process-based models that can elucidate the mechanisms of ensemble extraction. So far, mainstream task-general theories of ensemble processing have been primarily grounded in verbal descriptions (for review see: Whitney & Yamanashi Leib, 2018). A known limitation of verbal theories is that they often lack the precision of formal computational or mathematical models, which may be requisite for delineating key hypothetical constructs and adjudicating between competing theoretical accounts (e.g., Grahek, Schaller, & Tackett, 2021; Guest & Martin, 2021; Navarro, 2021; Oberauer & Lewandowsky, 2019). The goal of the current work is to attempt to fill this gap by developing and testing a unifying computational model of ensemble perception. According to the proposed model, perceived stimuli evoke distributed patterns of activity over feature values (e.g., different colors), and ensemble representations reflect the global sum of these patterns of activations. Critically, this modeling framework allows us to, for the first time, formally link a model of memory for individual items with a model for ensembles. More precisely, it is the first model that allows us to predict the entire distribution of errors on a range of ensemble tasks

from tasks that involve processing of individual items. As such, our modeling approach meets a high benchmark for a computational model and theory because it permits *generalizing* across different tasks (e.g., Busemeyer & Wang, 2000; Lee, 2011; Yarkoni, 2021).

Critically, we also use this model to address core questions in the ensemble and gist-memory literature, including how ensemble representations are built up from representations of individual items, as well as how ensemble representations can be more robust than representations of individual items. The proposed model is also the first to provide a high-precision account of human performance by capturing entire distributions of errors in continuous report tasks. Finally, the current modeling framework postulates probabilistic representations making it broadly consistent with contemporary population coding models of perceptual and cognitive processing (e.g., Rust, 2014).

The structure of the paper proceeds as follows. We begin by reviewing relevant theoretical work on ensemble perception, placing a special focus on dichotomies that highlight core questions researchers may want a model of ensemble processing to answer. Next, we describe the theoretical framework on which our ensemble model is based, and how it opens a venue towards linking a model of memory for individual items with a model of ensembles. This is followed by a description of the current ensemble model, as well as a set of new alternative quantitative models, which we developed to address each of the reviewed theoretical dichotomies. We then describe five empirical studies in which we systematically evaluate these models; we assess their ability to account for ensemble performance for different stimulus spaces, and for a range of ensemble task demands. Finally, we discuss how our model of ensemble processing connects to prominent cognitive and neural models of perception and memory.

Models of ensemble perception. While largely non-quantitative, existing theories of ensemble perception have been foundational for stimulating broad conceptualizations of ensemble perception and memory (for reviews see: Alvarez, 2011; Whitney & Yamanashi Leib, 2018). For this reason, we use these theories here to outline three relevant dichotomies, which highlight core desiderata for a quantitative process-model of ensemble extraction.

The first relevant dichotomy is between views that posit that processing of ensembles does versus does not involve operating over representations of individual items (Alvarez, 2011). This is a core question because it speaks to how ensemble representations are computed, and the

extent to which they are possible to compute when item information is unavailable — which is critical for questions about the role of ensembles in consciousness (e.g., Ward, Bear & Scholl, 2016; Cohen et al. 2016). One class of views posit that ensemble processing involves pooling over already-computed representations of individual items (e.g., Alvarez & Oliva, 2008; Oriet, Giesinger, & Stewart, 2020; Haberman & Whitney, 2011; Marchant, Simons, & Fockert, 2013). In contrast, other views suggest that ensemble processing involves automatically extracting an average without first representing each individual on its own (Ariely, 2001; Ward, Bear, & Scholl, 2016), or dividing the total amount of activation by the number of perceived items independent of any representation of individuals (Setic, Sveggar, & Domijan, 2007). Much of the work that seeks to address this critical question, being non-quantitative, makes the connection between individual and ensemble representation unclear. For example, Alvarez and Oliva (2008) report that representations of ensembles are present even when performance on individual items are at chance, implying distinct representations for both, but also notes that pooling noisy information about individual items can predict the ensemble data (see also Haberman & Whitney, 2011). Thus, this remains an area of continued uncertainty in the literature on ensemble perception.

The second relevant dichotomy is between views that postulate that noise accrues before versus after ensemble representations are extracted from the display. In general, this distinction has been made in the context of early versus late theories of noise (e.g., Baek & Chong, 2020), where early noise refers to noise accrued at the perceptual processing stage, whereas late noise refers to noise accrued at a post-perceptual (e.g., memory or decision) stage of processing. A more extreme version of this question has been the focus of work on whether people subsample just a few items to base ensemble judgments on (e.g., Myczek & Simons, 2008; Simons & Myczek, 2008) or instead rely on representations computed using all items (e.g., Chong et al. 2008). In the current work, we do not focus on the extreme views of subsampling vs. not, but instead focus on two distinct aspects of when noise accumulates: the role of perceptual vs. post-perceptual noise, and the role of presummarization vs. postsummarization noise (see Solomon, 2021 for latter distinction). We make the distinction between these kinds of noise because it is conceptually possible — and in fact is an assumption of the model we propose here — that representations of individual items are corrupted by perceptual noise (consistent with early, perceptual noise accrual), but that post-perceptual noise accrues after, rather than before ensembles are computed (consistent with ‘late’ postsummarization noise accrual). This is what would be expected if ensemble computation is a

perceptual process, one that operates over item representations that have already been corrupted by perceptual but not post-perceptual noise. By contrast, subsampling accounts (where only a few items are used to create ensemble representations) are more aligned with the view that ensembles are computed after post-perceptual noise, e.g., a subset of items are simply stored in working memory and the ensemble is computed only when people are tested after a retention delay (e.g., Marchant, Simons, & Fockert, 2013). Broadly, these dichotomies between different variants of early and late theories of noise provide insight into the time-course of ensemble extraction, and whether ensemble computation is more akin to a kind of perception or to an operation we can perform on items in working memory.

The last relevant dichotomy is between views that posit that ensemble representations are probabilistic versus point estimates. Relevant in this context is work by Chetverikov, Campana, and Kristjansson (2016) who examined the content of ‘ensemble’ representations in a visual search task. These authors reported evidence that people are sensitive to the entire underlying distribution of features in the external environment (uniform versus Gaussian), rather than simply an estimate of the average or variance of those features. This claim is consistent with the idea that people store entire probability distributions over visual features, at least in the kind of implicit tasks used in that work (Hansmann-Roth et al. 2021), which we note, may not generalize to the type of explicit judgment ensemble tasks used here. This probabilistic representation view contrasts with an alternative view according to which people simply represent a point estimate of the ensemble (e.g., an average in feature space of each individual item’s features; Alvarez & Oliva, 2008; Alvarez, 2011). This dichotomy speaks to the information-bearing, representational content of ensembles, and the richness of the ‘summaries’ that are computed in such tasks.

To preview, the current model posits that ensemble perception does involve pooling over individual representations. Our model also postulates that representations of individual items are corrupted by noise at an early perceptual stage, but that post-perceptual noise accrues *after* ensembles are computed (i.e., ensemble perception is distinct from simply holding items in working memory and then summarizing them after the fact). Finally, the model posits a probabilistic representation, with each item and the ensemble representation being represented as a distribution of activity over the entire feature space. In the next section, we describe and motivate the quantitative framework that serves as the conceptual and mathematical basis for the current ensemble model.

A quantitative framework: the Target Confusability Competition model of memory. We take a top-down strategy for developing a computational model of ensemble processing (for related approaches see: van Rooji, & Baggio, 2021). More precisely, we leverage an existing quantitative framework of memory (Schurgin, Wixted, Brady, 2020) and carve out a set of plausible constraints on the algorithms that may underlie ensemble perception and memory for them. These constraints are formalized within a set of computational models. The strength of our top-down approach is that our models are based on a priori, theoretically grounded predictions, rather than designed to maximize the fit to data from a particular task. Furthermore, these predictions have the potential to capture ensemble processes across a broad range of ensemble tasks, rather than accounting for data within only a specific ensemble task. Accordingly, we evaluate and compare these models across a range of experimental tasks to formally assess whether these computations are invariant across different experimental demands. Finally, as we discuss next, our theoretical and technical approach leads to a novel unifying theory for memory of individual items and ensembles (e.g., for conceptually related modeling approaches in other domains see: Anderson, 1981; Palestro, Bahg, Sederberg, Lu, Steyvers & Turner, 2018; Tuner, Forstmann, Love, Palmeri, & Van Maanen, 2017).

We conceived the current ensemble model using a quantitative framework that has recently been applied to visual memory for individual items, the Target Confusability Competition (TCC) model of memory (Schurgin et al., 2020). In a set of more than a dozen experiments, Schurgin et al., (2020) found that TCC outperformed mainstream models of visual memory both in terms of fit to data, and ability to generalize and predict data across distinct visual working memory and visual long-term memory tasks. The TCC model combines two fundamental ideas shared by a broad range of cognitive computational models, which is that memory-based decisions are made under uncertainty (Thurstone, 1927; Swets, 1986) and that information in the world is processed based on its psychophysical, rather than physical similarity structure (Luce & Galanter, 1963; Shepard, 1987; Stevens, 1936). Figure 2 depicts these two ideas within the TCC framework.

The first premise of the TCC model is that memory representations are intrinsically probabilistic and vary in strength, a core principle of Signal Detection Theory (Swets, 1986; Wickens, 2000; Wixted, 2020) and, broadly, Bayesian models of cognition (Brady, Schacter, & Alvarez, 2018; Chater, Tenenbaum, Yuille, 2006; Hemmer & Steyvers, 2009; McCarley & Benjamin, 2013). For instance, a remembered item (e.g., a purple square) is assumed to be neither completely forgotten nor completely remembered. Instead, there is a probability distribution over how well

the item is remembered, such that sometimes it is remembered with high fidelity and elicits a strong familiarity signal, and other times it is remembered with lower fidelity and elicits a weaker familiarity signal. Since memory representations are intrinsically noisy, even items that are not stored in memory also elicit some sense of familiarity. Thus, rather than being all-or-none, memory is assumed to be a graded phenomenon (e.g., MacMillan & Creelman, 1996). This point is formalized and depicted graphically by the two Gaussian distributions on the left of Figure 2, which denote the distribution of familiarity signals for old (purple) and new (yellow/orange) items, with higher up being more familiar. In line with standard signal detection theory, the standardized distance between the mean of the two distributions is captured with a single free parameter d' , which quantifies the average strength of the memory representation across trials (e.g., how much more, on average, the 'seen' item is familiar than the 'unseen' item).

The second premise of the TCC model is that familiarity spreads across feature space according to the stimulus' psychophysical properties (for a few foundational models of memory that share this general assumption see: Hintzman, 1986; Howard & Kahana, 2002; Murdock, 1982; Reder et al. 2000; Shiffrin & Steyvers, 1997). More precisely, the familiarity of a given stimulus is a function of the *psychophysical similarity* between this stimulus and the contents of memory. This assumption entails that stimuli in the environment will elicit a stronger familiarity signal if they are more psychophysically 'similar' to contents in memory. For instance, if the remembered item is a purple square, the color purple will elicit a very strong familiarity signal, as will colors that are nearly perceptually indistinguishable from purple (e.g., other purples). Colors that are somewhat similar to purple (e.g., magenta) will also elicit a familiarity signal, which will be stronger than those elicited by relatively dissimilar colors (e.g., green). In line with previous work, Schurgin et al., (2020) found that the latent psychophysical similarity function and the corresponding distribution of memory signals is approximately exponential in form, outside the zone where perceptual confusions make it somewhat flatter (right side of Figure 2; see also Nosofsky, 1987; Shepard, 1987; Sims, 2018; Tenenbaum, 1999).

Fig. 2 | The TCC framework merges principles of memory uncertainty and the ~exponential generalization gradient. Left: Signal detection theory. The two Gaussian distributions represent the distribution of familiarity signals for old (purple color) and new (yellow/orange) items, with 'higher up' being more familiar. Purple, having been seen, on average has higher familiarity, but on a given trial people judge just one sample from this distribution, so yellow/orange may feel more familiar sometimes. In line with signal detection theory, the distance between these distributions (d') quantifies memory fidelity. Right: Similarity. The psychophysical similarity function shows how average familiarity scales as a function of psychophysical similarity to the remembered item (e.g., the purple at the center of the distribution has familiarity d' , and the orange/yellow at the edges has 0).

Critically, the TCC framework formally integrates these two ideas to capture performance in both visual working- and long-term memory tasks, for both low-level (e.g., color) and high-dimensional (e.g., faces) visual stimuli, with a single psychophysical similarity function for each stimulus and a single free parameter to explain performance in each task (d'). Schurgin et al., (2020) compared this TCC model to a comprehensive set of mainstream visual memory models, and found that the TCC model outperformed these models in terms of model fit, capturing the "long tail" of errors that typically arise in difficult continuous reproduction tasks without any need for a separate 'guessing process' (e.g, Zhang & Luck, 2008). More importantly, unlike other extant models, the TCC model permitted principled generalization across different memory tasks, such as a zero-free-parameter prediction from a two-alternative forced choice to a delayed estimation visual-working memory task. Notably, the ability of the model to generalize across different tasks meets a high benchmark for model performance because it indicates that a given model is not overfitting a sample of data, but is capturing

cognitive invariants, that is, latent cognitive processes that are stable across different task structures (Anderson, 2004; Busmeyer & Wang, 2000; Lee, 2011; Newell, 1994; Yarkoni, 2021).

To summarize, TCC is a model that combines two fundamental premises about memory processes in a way that permits generalization across memory tasks with a single free parameter. The generalizability and parsimony of TCC, as well as its basis on probabilistic models of cognition and psychophysical scaling, make it a uniquely powerful framework for building cognitive architectures. For this reason, we leverage TCC to derive a model of ensembles, which we refer to as the Perceptual Summation model. We also derive a set of alternative models that make different processing assumptions, which we compare against the Perceptual Summation model. In the following section we describe the TCC model for individual items in detail, as well as how we extend this model of memory for individual items to a model of memory for ensembles.

Linking models of memory for individual items with memory for ensembles. In this section we set up the basis for formally linking a model of memory for individual items with a model for ensembles. We begin with a brief description of how the TCC framework explains performance in the standard delayed estimation visual-working memory task in which participants are instructed to remember individual items and report them after short delays, which are similar to ensemble tasks except they probe memory for individual items instead of the entire set. Figure 3 (left) shows a schematic of a typical trial in a visual working memory task; in this scenario, there are three to-be-remembered items and the task requires memorizing colored circles and their spatial locations. The TCC model postulates that each item elicits some location dependent pattern of activity. In particular, each item causes an increase in activation (and familiarity) for its respective color, but also for similar colors, in a location-dependent manner. The pattern of activity elicited by each item is then corrupted by noise, and the signal-to-noise ratio depends on many variables (e.g., memory load, encoding time, delay length) that affect how well the items are initially encoded and how much noise accumulates during the delay. At the end of the trial, the probed-item's location is queried, and the participant reports the color channel that generates the maximum familiarity signal. Formally, the most straightforward version of the TCC model for individual items is given by the following equation:

$$r_{i, VWM} = \operatorname{argmax} \left(f(x)_i d' + \sigma_{\text{Noise}} \right), \quad (1)$$

where $r_{i,VWM}$ is the predicted response on the continuous report visual working memory task on the probed item i , $f(x)$ is the measured similarity of each color x of item i to the color shown at the probed location, d' is a free parameter that quantifies the average strength of the memory representations, σ_{Noise} is a fixed amount of post-perceptual noise (noise is set to one standard deviation of a Gaussian distribution, consistent with a standard signal detection model), and $argmax$ denotes the decision rule that memory reports are based on the feature that generates the maximum familiarity signal. While the framework is conceptually compatible with any variation of the d' between items, the standard TCC model assumes equal weighting of items in memory, meaning that on average, each item is remembered with the same fidelity and, thus, the same d' is used for each item. However, this is not a necessary assumption for the framework: for example, we describe an unequal weighting (recency) model specifically designed for a sequential presentation paradigm below.

Fig. 3 | TCC links models for individual items and ensembles. The TCC model for individual items (left panel) postulates that each item in the memory array elicits a distributed pattern of activation over feature values (upper left), which is corrupted by noise (center left). When tested, people report on the feature value that generates the maximum memory signal based on the probed item (lower left). The TCC Perceptual Summation model for ensembles (right panel) postulates that each item in the memory array elicits a distributed pattern of activation over feature values, which are pooled at an early encoding stage of processing (upper right). This distributed pattern of activation is corrupted by noise (middle right) and when queried on the mean color, people report on the feature value that elicits the maximum memory signal in this ensemble representation (lower right).

We developed the Perceptual Summation ensemble model a priori from this TCC model for individual items. Like the model for individual items, the Perceptual Summation model postulates that each item in the memory array elicits some pattern of activity over feature values (Figure 3, right). However, in contrast to the model for individual items, the Perceptual Summation model postulates that the ensemble representation is extracted during encoding, before post-perceptual noise accrues over representations of individual items¹. Finally, the model postulates that post-perceptual (memory) noise accumulates over the ensemble representation, and, when probed on the average, the participant reports on the color channel that generates the maximum familiarity signal. The equation for the Perceptual Summation model is given below.

$$r_{ENS} = \operatorname{argmax} \left(\left(\sum_{i=1}^N f(x)_i d' \right) + \sigma_{Noise} \right) \quad (2)$$

Note that Equations 1 and 2 are nearly identical, with the exception that self-reports on the visual working memory task are determined by levels of activation elicited by a single probed item (i), whereas self-reports on the ensemble task (r_{ENS}) are determined by the summed levels of activation of all N items.

This creates a formal connection between the model for individual items and for ensembles. Through the lens of these models, the patterns of activation elicited by each item in the memory array is what is pooled via summation in the early perceptual stage of ensemble extraction. As described, this pattern of activation is quantified with an independently estimated psychophysical similarity function and a single free parameter d' . Conceptually, the psychophysical similarity function quantifies the predicted pattern of activation over feature values (e.g., colors) which is evoked by each item in the ensemble array. d' rescales this pattern of activation based on the signal-to-noise ratio of each memory representation for a given stimulus space, display configuration, memory load, encoding time and retention interval.

Thus, a core premise of this linking framework is that the signal-to-noise ratio for individual items is *the same* across tasks that probe individual items and ensembles. This entails that, with all task parameters (except for the instruction manipulation) held constant, the bottleneck for

¹ Like the model for individual items, we assume that each item accrues perceptual noise before the ensemble is computed (that is, that familiarity spreads according to the psychophysical similarity function for each item). This perceptual noise is assumed to be fixed, and is built into the psychophysical similarity function as explained in the Technical Details section and Schurkin et al., 2020.

encoding each item in the memory array, as well as the amount of post-perceptual noise accrued over the extracted memory representations (of a single item or the ensemble) is the same in both visual ensemble tasks and visual working memory tasks for individual items. Critically, according to the Perceptual Summation model, in ensemble tasks these patterns of activation of individual items are pooled via summation *prior to* post-perceptual (memory) noise accrual. That is, if $d' = 1$ for each individual item, and there are 5 very similar items, d' will equal nearly 5 for the ensemble task (depending on the item's similarity). For this reason, the signal-to-noise ratio of the post-summation ensemble representation will tend to be larger than it is for individual items (see Figure 3A), particularly when there is overlap in feature values (i.e., redundancies) between items in the ensemble array. Note that this contrasts with an alternative view, embodied in the post-perceptual summation alternative model below, where items are simply each held in working memory (thus each acquiring noise on their own) and summed after post-perceptual noise accrual, at the time of test.

The perceptual summation models' pooling mechanism and its effect on the signal-to-noise ratio, can be seen as a cognitive-level approximation of the types of pooling processes described in neural population coding models. Very generally, some evidence suggests that increased population size may also increase the amount of information embedded in populations of neurons (e.g., Shamir, 2014), although the relationship between population size and readout accuracy is extremely complex and an active topic of investigation (e.g., Averbeck, Latham, & Pouget, 2006; Bartolo, Saunders, Mitz, & Averbeck, 2020; Kohn, Coen-Cagli, Kanitscheider, & Pouget, 2016; Shamir, 2014), this framework provides one neurally plausible instantiation for the computations postulated in the Perceptual Summation model. A similar parallel exists between the TCC model of individual items, and neural population models of visual working memory (e.g., Bays, 2015; see Bays, 2019).

These simple processing assumptions permit us to conceptually link memory for individual items and ensembles. They also explain a key question discussed earlier, that is, how ensemble representations can be 'built-up' from representations of individual items, while still being relatively more robust than representations of individual items (e.g., Alvarez & Oliva, 2008; Haberman & Whitney, 2011). In our modeling and analysis, we link these two models formally by estimating d' from a visual working memory task for individual items and by using this d' estimate, as well as the psychophysical similarity functions of each item in the ensemble array, to compute the predicted summed pattern of activation of the ensemble. To preview, this approach permits us to, for the first time, predict entire distributions of memory errors in

continuous report ensemble tasks with *zero* free parameters. We then assess and compare the predictive accuracy of the Perceptual Summation model (proposed above) and six alternative models, which are described next.

Alternative models. Currently there are no computational models of ensemble processing that fully capture distributions of errors in a continuous self-report task and that can account for data across a range of ensemble manipulations. In order to bolster the interpretability of tests of our a priori derived Perceptual Summation model with model comparison, we derived a set of alternative contending models that rest on different processing assumptions than the Perceptual Summation model, which is based on probabilistic representations and ensembles computed after perceptual noise but before post-perceptual noise. We consider six alternative ensemble models, giving seven models in total. The first two of these alternative models are nested within the TCC framework, but rest on different processing assumptions. The remaining four alternative models are not nested in the TCC framework. These models differ systematically in whether they see representations as probabilistic vs. point representations vs. point representations plus uncertainty; whether they see ensemble extraction as corrupted by perceptual noise before vs. after summarization; and whether they see ensemble extraction as occurring before post-perceptual noise (e.g., as a perceptual process) or after post-perceptual noise (e.g., after memory noise, as though performed on working memory representations after the delay).

Fig. 4 | Competing ensemble models. The top row shows all models that fall within the framework of TCC models, and the bottom row shows models grounded in alternative frameworks. Within the TCC framework, the models differ in whether they allow perceptual noise, in the form of the spreading of familiarity according to the similarity function, to take effect before (perceptual and post-perceptual summation) or after (automatic averaging) the ensemble summary step, and whether they see the representations as accruing memory noise only after they have been summarized into an ensemble (perceptual summation; automatic averaging) or while still represented as individual items (post-perceptual summation). The non-TCC models differ in the extent to which they represent uncertainty (noisy point estimate; precision point estimate) vs do not (noise-free point estimate; average of point estimates) and the extent to which they assume an averaging of individual (average of point estimates) vs a direct extraction of the ensemble mean (other models). See main text for full model descriptions.

Figure 4 shows the full set of instantiated ensemble models. The first alternative model we refer to as the Post-perceptual Summation model. According to this model, people compute ensembles after they perceive, encode and remember individual items, e.g., a summary is computed across items only after each individual item has already been affected by the psychophysical similarity function *and* accrued memory noise. This model entails that people maintain location dependent representations of each item in memory until they are probed on their memory for the ensemble. Thus, according to this model, ensemble representations are computed at a relatively late, rather than early processing stage. This prediction contrasts to the prediction of the Perceptual Summation model, according to which noise accrues at a relatively late processing stage. The equation for the Post-perceptual Summation is given below.

$$r_{ENS} = \operatorname{argmax} \left(\sum_{i=1}^N (f(x)_i d' + \sigma_{Noise}) \right) \quad (3)$$

Note that the terms in Equations 2 and 3 are identical; the only difference between these equations is in whether summation occurs over individual items that have accrued post-perceptual noise σ_{Noise} (Post-perceptual Summation model), or before representations of individual items have accrued post-perceptual noise (Perceptual Summation model).

We refer to the second alternative model as the Automatic Summation model. According to this model, a representation of the average is automatically extracted, without processing of, or pooling over representations of individual items. Rather than assuming this representation has perfect resolution, we include post-perceptual noise in the estimate of this representation in the same way we would if it was an individual item. Thus, this model is equivalent to assuming that the ‘average’ is directly perceptually available to people in the same way that an item that is perceptually present is. As shown in Figure 3C, this model postulates that people extract a single probability distribution over the mean feature, which is also corrupted by noise.

$$r_{ENS} = \operatorname{argmax} \left(f(x)_{Mean} d' + \sigma_{Noise} \right) \quad (4)$$

Note that Equation 4 is identical to Equation 1; this is because the Automatic Average model postulates that a representation of a single item with the mean feature value is extracted automatically, entailing that the ensemble representation has the same signal-to-noise ratio as the representation of a single individual item.

The next four contending models are not nested within the TCC framework, but serve as conceptual foils to this set of models. The Noise-Free Point Estimate model (Figure 3D) postulates that people automatically extract a *noise-free* point estimate of the mean feature. In other words, the Noise-Free Point Estimate model simply predicts that self-reports on the ensemble task are the true mean ($\bar{X}$). Although this model assumes that ensemble representations are noise-free point estimates, the model also assumes a small amount of motor error when making ensemble responses (we describe how this is implemented in the model in the Modeling and Analysis section). This model is unlikely to perform well since it cannot capture the full distribution of errors in a delayed estimation task, however, we include it because it serves as a useful reference point against which to compare more plausible models.

The fifth model is the Noisy Point Estimate model, according to which people automatically extract a point estimate of the mean feature with an uncertainty interval around this value. For this model, we make the simplifying assumption that people represent a uniform uncertainty interval around the true mean value (Rahnev et al., 2021).

Neither of the first two Point estimate models formally link memory for individual items to memory for ensembles. Therefore, we consider two additional point estimate models, which link processing across the two tasks and also make more tenable assumptions about processing on ensemble tasks. The first of these models is the Average Item Point Estimate model, according to which people compute ensembles by averaging over point estimates of individual item representations. This is effectively the model proposed by Alvarez and Oliva (2008) (see also Alvarez, 2011). The second of these models is the Precision Ensemble model, according to which people maintain a point estimate of the mean, which has a Gaussian, rather than uniform uncertainty interval around it. Rather than using a fixed interval for all subjects as we did for the Noisy Point Estimate model, we calibrate this interval for each individual by obtaining a standard deviation estimate from their working memory data using the popular standard mixture model (inspired by Zhang and Luck, 2008; details discussed in the *Technical details of model fitting and prediction* section).

To summarize, neither the Noise-free Point Estimate nor the Item Average Point Estimate model postulate probabilistic representations. The Noisy Estimate and Precision ensemble model postulates partially probabilistic representations because there is uncertainty interval around the mean, but not a full probability distribution over feature values (see Rahnev et al., 2021 for a description of partially probabilistic models). Note that we do not presume that these

point-estimate models capture all possible ways in which point-estimate or partially probabilistic models could account for the data. However, in the absence of other quantitative models in the literature that can apply to such ensemble perception tasks, we created models that spanned a wide range of plausible assumptions (e.g., probabilistic vs. discrete representations; ensemble processing as occurring during perception vs. after memory delay; perceptual noise before vs. after the ensemble representation is extracted, etc). We revisit this point in the General Discussion.

Unequal weighting of items in memory. Thus far, each of the models we described posit that, on average, each item is weighted equally when computing an ensemble. This is tenable under conditions that do not lead to disproportionate prioritization of a specific item, or subset of items. In the TCC framework for visual working memory for individual items, the idea that, on average (across trials), d' is equal between items is corroborated by the finding that this variant of the TCC model can predict data across memory tasks (Schurgin et al. 2020), as well as newer findings from studies that explicitly examine whether there is variability between items (e.g., Williams, Brady, Stoermer. 2022). It is also validated in a more general signal detection framework by near-equal-variance receiver operating characteristics in visual working memory tasks for simple features like color (e.g., Williams, Robinson, Schurgin, Wixted, Brady, 2022; Robinson et al. 2020).

However, it is known that some conditions do elicit unequal weighting of items in memory. For instance, items that are shown more recently tend to be remembered better than items shown less recently and such recency effects have been shown to affect ensemble representations as well (Tong, Dube & Sekuler, 2019). Thus, to evaluate the generalizability of our modeling results we extended our modeling and analysis to conditions in which items receive unequal prioritization in memory. In addition to generalizing our theory of ensembles to a wider range of tasks, our extension of the modeling to conditions in which items receive unequal weighting, provides a stronger testbed for the view that the ensemble representation reflects a sum of local patterns of activation elicited by each item in the array, since this summation account becomes more distinct from other possible accounts when items vary in strength. Specifically, as we discuss in the General Discussion, this generalization of the model demonstrates that we are able to use the model to predict both inter- and intra-individual variations in ensemble processing.

To this end, we focused on modeling a sequential presentation ensemble paradigm because it has been shown to elicit robust recency effects. One way to generalize the TCC-based models to this situation is to simply obtain separate d' estimates for probes of each item in the sequence (e.g., d' for when the 3rd item is probed, d' for when the second item is probed, etc) and use these estimates to compute ensemble predictions. We report these results below.

However, we can also use a model of how memory changes as a function of the sequential presentation, rather than simply measuring d' for each position in the sequence. This should allow the model to be more robust and require fewer parameters and measurements to predict ensemble performance. To extend our model in this way, we used prior modeling work by Tong, Dube and Sekuler (2019) and developed a Recency TCC model for individual items and a Recency Perceptual Summation model for ensembles. This Recency TCC model quantifies recency weights using as an exponential function (without base e) over the serial position of each stimulus in the sequence. The recency model for individual items is given by the following equation:

$$r_{i,VWM} = \operatorname{argmax} \left(f(x)_i d' rate^j + \sigma_{Noise} \right), \quad (5)$$

where all terms are identical to those given in Equation 1, with the exception of a second parameter $rate$, which has the item position of item j in the exponent (where $j=1$ is the most recent item in the sequence). The rate parameter is a free parameter bounded between 0 and 1 that captures the effects of memory decay on memory, with smaller values (of the parameter) indicating stronger decay effects and, therefore, relatively higher weighting of more recent items in the ensemble (and relatively better performance for them in the visual working memory task). The equations for the Perceptual and Post-perceptual Summation ensemble models are extended in a similar way, as shown below (Equations 6 and 7, respectively):

$$r_{ENS} = \operatorname{argmax} \left(\left(\sum_{i=1}^N f(x)_i d' rate^j \right) + \sigma_{Noise} \right), \quad (6)$$

$$r_{ENS} = \operatorname{argmax} \left(\left(\sum_{i=1}^N f(x)_i d' rate^j + \sigma_{Noise} \right) \right). \quad (7)$$

Note that the Automatic Averaging model is unchanged because it postulates that people automatically extract a single representation of the mean without building up this representation from individual items. We treat this model as being conceptually equivalent to a prototype model that entails an equally weighted average (e.g., Smith & Minda, 2000).

Accounting for performance across stimuli spaces and presentation schemes. We ran five experiments to evaluate the predictive accuracy of the Perceptual Summation model against the six alternative models. The goal of each experiment was to assess the generalizability of our modeling results across different ensemble tasks. As previewed, each experiment had the same structure, meaning that participants completed one block of a visual working memory task and one block of an ensemble task (presented in random order across participants). This allowed us to measure d' in the visual working memory task, and to use this d' estimate to predict performance in the ensemble task by applying the modeling framework described above. In Experiments 1 and 2 we examined people's memory for color, and manipulated set size as well as the range of colors (e.g., how distinct they were in color space) in the ensemble task (Figure 5 top panel). In Experiments 3 and 4 we evaluated the generalizability of these results for higher-level features by evaluating people's memory for shape (Figure 5 bottom panel). Like in the color experiments, we also manipulated set size and varied the range of the shapes in the shape space in the ensemble array (Experiment 3 and 4, respectively). Finally, Experiment 5 was designed to evaluate the models in a sequential presentation task, under conditions in which memory resources were unequally, rather equally distributed across items (Figure 6).

Fig. 5 | Color and shape memory tasks. Example trial sequence used in the visual working memory (VWM) and ensemble tasks of Experiments 1-4. In the visual working memory tasks (left), participants saw a set of colors (top) or shapes (bottom) and then after a delay a single location was probed, and participants had to indicate which color or shape was in that position. In the ensemble tasks (right), participants saw a set of colors (top) or shapes (bottom) and then after a delay were probed on the mean color or mean shape (e.g., a summary of the entire set) rather than on a single individual item.

Block 1 (or 2) VWM task

First item

Last item

Self-report

What color was this object?

Block 2 (or 1) ensemble task

First item

Last item

Time

What is the mean color
of all six objects?

Fig. 6 | Sequential memory tasks. Example trial sequence used in the visual working memory and ensemble task in Experiment 5. On each trial of the visual working memory task (top), participants saw a number of colored real world objects presented one at a time and then a single object appeared at test, in grayscale, and participants had to indicate what color that particular item had been. In the ensemble task (bottom), participants saw a sequence of colors on a single real-world object, and at test had to indicate the average color of this object.

Results

The goal of our analysis was to evaluate the ability of our Perceptual Summation ensemble model to generalize performance from the visual-working memory task to the ensemble task. We formally compare the predictive accuracy of this model with other models by comparing the predicted negative log likelihood (PNLL) between it and the contending models at the level of individual subjects. This is *predictive* because we do not *fit* the ensemble data: instead, we fit the model to the visual working memory data alone and use parameters obtained from fits to the visual working memory data to predict (with no free parameters) the ensemble task data. Because PNLL is a negative log likelihood, lower scores reflect less deviance and better model predictions (that is, superior model performance). Measures of predictive accuracy such as PNLL naturally account for model complexity because they capture a model's capacity to generalize to new data samples and/or new task structures (for technical and conceptual discussion of related points see: e.g., Busemeyer & Wang, 2000; Yarkoni, 2021). More precisely, models that are overly complex and overfit a sample of data will do poorly at

prediction in new (unseen) samples of data. Conversely, models that capture substantive variance in a sample of data (i.e., latent cognitive processes rather than error variance) will do well at generalizing to new samples of data and task structures and yield relatively higher predictive accuracy (and lower PNLL scores) than competing models.

Ensemble memory for color with different set sizes. In Experiment 1, we asked participants to perform both visual working memory and ensemble tasks using color as the stimuli, and manipulated set size, which we expected to impact memory performance for individual items. This allowed us to test whether ensemble performance changed as predicted by each model when memory performance for individual items changed. Overall, we found the best performing model in predicting the ensemble data was the Perceptual Summation model. In particular, we found that PNLL was significantly smaller for the Perceptual Summation ($\bar{X} = 1,056$; $SEM = 9.6$) than the Post-perceptual Summation ($\bar{X} = 1,095$; $SEM = 7.9$) model ($t(49) = 5.5$, $p < 1.0e-05$); significantly smaller for the Perceptual Summation than the Automatic Averaging ($\bar{X} = 1,156$; $SEM = 7.8$) model ($t(49) = 12.2$, $p < 1.0e-15$); significantly smaller for the Perceptual Summation than the Noise-free Point estimate ($\bar{X} = 1,322$; $SEM = 0.2$) model ($t(49) = 28.3$, $p < 1.0e-31$); significantly smaller for the Perceptual Summation than the Noisy Point Estimate ($\bar{X} = 1,225$; $SEM = 2.3$) model ($t(49) = 22.1$, $p < 1.0e-26$); significantly smaller for the Perceptual Summation than the Average Item Point Estimate ($\bar{X} = 1,252$; $SEM = 2.2$) model ($t(49) = 23.5$, $p < 1.0e-27$); and significantly smaller for the Perceptual Summation than the Precision Ensemble ($\bar{X} = 1,244$; $SEM = 4.1$) model ($t(49) = 21.9$, $p < 1.0e-26$). The fit of the TCC model and the prediction of the ensemble Perceptual Summation model to aggregate data are shown in Figure 7A. The top panel of Figure 8 shows violin plots of the difference in PNLL between the Perceptual Summation and competing models at the level of individual participants.

Ensemble memory for color with different ranges. In Experiment 2, participants performed a color task but instead of varying set size, we manipulated the range of the colors in the ensemble task (e.g., how distinct they were from each other). We again found the Perceptual Summation model provided the best fit for performance in the ensemble conditions. In particular, we found that PNLL was significantly smaller for the Perceptual Summation ($\bar{X} = 708$; $SEM = 4.8$) than the Post-perceptual Summation ($\bar{X} = 753$; $SEM = 4.2$) model ($t(49) = 15.9$, $p < 1.0e-20$); significantly smaller for the Perceptual Summation than the Automatic Averaging ($\bar{X} = 790$; $SEM = 3.7$) model ($t(49) = 26.6$, $p < 1.0e-30$); significantly smaller for the

Perceptual Summation than the Noise-free Point ($\bar{X} = 882$; $SEM = 0.1$) estimate model ($t(49) = 37.2$, $p < 1.0e-36$); and significantly smaller for the Perceptual Summation than the Noisy Point estimate ($\bar{X} = 818$; $SEM = 1.36$) model ($t(49) = 30.1$, $p < 1.0e-32$); significantly smaller for the Perceptual Summation than the Average Item Point Estimate ($\bar{X}=835$; $SEM = 1.1$) model ($t(49) = 31.4$, $p < 1.0e-33$); and significantly smaller for the Perceptual Summation the Precision Ensemble ($\bar{X}=833$; $SEM = 2.2$) model ($t(49) = 29.4$, $p < 1.0e-32$). The fit of the TCC model and the prediction of the ensemble Perceptual Summation model to aggregate data are shown in Figure 7B. The bottom panel of Figure 8 shows violin plots of the difference in PNLL between the Perceptual Summation and competing models at the level of individual participants.

Fig. 7 | The Perceptual Summation model predicts ensemble memory for color. Graphical representation of the best performing model's (Perceptual Summation) fit and prediction of data in Experiments 1-2. In these experiments participants had to remember colors of simultaneously presented circles, where either set size (1) or range of colors (2) was manipulated in the ensemble task. The top row shows the fits of the TCC model for individual items to aggregate data from the visual working memory task in which we manipulated set size. d' estimates from the visual working memory task were substituted into the Perceptual Summation model to predict the ensemble data (second row). The third row shows the fit of the TCC model for individual items to aggregate data from a visual working memory task in which participants had to remember six items. The d' estimate from this task was used to predict the ensemble data in each of the range conditions where participants had to remember the ensemble color for six items. The core set of results are the Perceptual Summation model's capacity to predict the ensemble data. We visually show the fits of the TCC model for individual items to the visual working memory data to demonstrate that it provides a reasonable fit to the data (for extended model comparison between this model and other contending models for individual items based on fit and predictive accuracy see: Schurgin et al., 2020).

Fig. 8 | Comparison in predictive accuracy between Perceptual Summation model and competing models of ensemble memory for color. The difference in predicted negative log likelihood scores between each of the six alternative competing models ($PNLL_{Alt}$) and the main Perceptual Summation model ($PNLL_{PerSum}$) for Experiments 1 (top panel) and 2 (bottom panel). Lower values of PNLL indicate higher predictive accuracy, therefore, PNLL difference scores higher (or lower) than zero indicate support for the Perceptual Summation (or a competing) model. In both experiments, the vast majority of participants are better predicted by the Perceptual Summation model than any of the alternatives.

Ensemble memory for shapes with different set sizes. In Experiment 3, we asked participants to perform both visual working memory and ensemble tasks using shape as the stimuli, and manipulated set size, which we expected to impact memory performance for

individual items. This allowed us to test whether ensemble performance changed as predicted by each model when memory performance for individual items changed. Overall, we found the best performing model was the Perceptual Summation model. In particular, we found that PNLL was significantly smaller for the Perceptual Summation ($\bar{X}$ = 1099; SEM = 12.8) than the Post-perceptual Summation ($\bar{X}$ = 1,123; SEM = 11.4) model ($t(49) = 4.5, p < 1.0e-4$); significantly smaller for the Perceptual Summation than the Automatic Averaging ($\bar{X}$ = 1179; SEM = 9.5) model ($t(49) = 11.9, p < 1.0e-15$); significantly smaller for the Perceptual Summation than the Noise-free Point ($\bar{X}$ = 1,322; SEM = 0.2) estimate model ($t(49) = 17.9, p < 1.0e-22$); and significantly smaller for the Perceptual Summation than the Noisy Point ($\bar{X}$ = 1,239; SEM = 3.9) estimate model ($t(49) = 15.4, p < 1.0e-19$); significantly smaller for the Perceptual Summation than the Average Item Point Estimate ($\bar{X}$ =1,259; SEM = 3.3) model ($t(49) = 16.4, p < 1.0e-20$); and significantly smaller for the Perceptual Summation than the Precision Ensemble ($\bar{X}$ = 1,260; SEM = 4.1) model ($t(49) = 14.9, p < 1.0e-19$). The fit of the TCC model and the prediction of the ensemble Perceptual Summation model to aggregate data are shown in Figure 9A. The top panel of Figure 10 shows violin plots of the difference in PNLL between the Perceptual Summation and competing models at the level of individual participants.

Ensemble memory for shapes with different ranges. In Experiment 4, participants performed a shape task but instead of varying set size, we manipulated the range of the shape across the shape wheel in the ensemble task (e.g., how distinct they were from each other). We again found the Perceptual Summation model provided the best fit for performance across these conditions. In particular, we found that PNLL was significantly smaller for the Perceptual Summation ($\bar{X}$ = 766; SEM = 8.8) than the Post-perceptual Summation ($\bar{X}$ = 783; SEM = 7.2) model ($t(49) = 3.6, p < 1.0e-3$); significantly smaller for the Perceptual Summation than the Automatic Averaging ($\bar{X}$ = 814; SEM = 5.9) model ($t(49) = 8.5, p < 1.0e-10$); significantly smaller for the Perceptual Summation than the Noise-free Point ($\bar{X}$ = 882; SEM = 0.1) estimate model ($t(49) = 13.4, p < 1.0e-17$); and significantly smaller for the Perceptual Summation than the Noisy Point ($\bar{X}$ = 836; SEM = 2.4) estimate model ($t(49) = 10.7, p < 1.0e-13$); significantly smaller for the Perceptual Summation than the Average Item Point Estimate ($\bar{X}$ =845; SEM = 2.1) model ($t(49) = 11.3, p < 1.0e-14$); and significantly smaller for the Perceptual Summation than the Precision Ensemble ($\bar{X}$ =850; SEM = 2.5) model ($t(49) = 11.3, p < 1.0e-14$). The fit of the TCC model and the prediction of the ensemble Perceptual Summation

model to aggregate data are shown in Figure 9B. The bottom panel of Figure 10 shows violin plots of the difference in PNLL between the Perceptual Summation and competing models at the level of individual participants.

Fig. 9 | The Perceptual Summation model predicts ensemble memory for shape. Graphical representation of the best performing model's fit and prediction of data in Experiments 3-4. In these experiments participants had to remember simultaneously presented shapes, where either set size (1) or range of shapes (2) was manipulated in the ensemble task. The top row shows the fits of the TCC model for individual items to aggregate data from the visual working memory task in which we manipulated set size. d' estimates from the visual working memory task were substituted into the Perceptual Summation model to predict the ensemble data (second panel). The third panel shows the fit of the TCC model for individual items to aggregate data from a visual working memory task in which participants had to remember six items. The d' estimate from this task was used to predict the ensemble data in each of the range conditions where participants had to remember the ensemble shape for six items.

Predictive accuracy of ensemble models in shape memory task

Fig 10 | Comparison in predictive accuracy between Perceptual Summation model and competing models of ensemble memory for shape. The difference in predicted negative log likelihood scores between each of the six alternative competing models ($PNLL_{Alt}$) and the main Sequential Perceptual Summation model ($PNLL_{PerSum}$) for Experiments 3 (top panel) and 4 (bottom panel). Lower values of $PNLL$ indicate higher predictive accuracy, therefore, $PNLL$ difference scores higher (or lower) than zero indicate support for the Perceptual Summation (or a competing) model. In both experiments, the vast majority of participants are better predicted by the Perceptual Summation model than any of the alternatives.

Ensemble memory for sequentially presented stimuli. In Experiment 5, we present stimuli sequentially, introducing variation in the strength of the items. We found that again, the Perceptual Summation model performed best. In particular, we found that $PNLL$ was significantly smaller for the Sequential Perceptual Summation ($\bar{X} = 467$; $SEM = 5.3$) than the

Sequential Post-perceptual Summation ($\bar{X} = 491$; $SEM = 4.6$) model ($t(49) = 9.7$, $p < 1.0e-12$); significantly smaller for the Sequential Perceptual Summation than the Automatic Averaging ($\bar{X} = 483$; $SEM = 4.4$) model ($t(49) = 4.7$, $p < 1.0e-4$); significantly smaller for the Sequential Perceptual Summation than the Noise-free Point ($\bar{X} = 564$; $SEM = 0.1$) estimate model ($t(49) = 18.5$, $p < 1.0e-22$); and significantly smaller for the Perceptual Summation than the Noisy Point ($\bar{X} = 531$; $SEM = 1.49$) estimate model ($t(49) = 14.9$, $p < 1.0e-19$); significantly smaller for the Perceptual Summation than the Average Item Point Estimate ($\bar{X}=534$; $SEM = 1.3$) model ($t(49) = 15.6$, $p < 1.0e-19$); and significantly smaller for the Perceptual Summation than the Precision Ensemble ($\bar{X}=532$; $SEM = 2.3$) model ($t(49) = 15.2$, $p < 1.0e-19$). We also found that the Recency Perceptual Summation model outperformed an equal weighting Perceptual Summation ($\bar{X} = 487$; $SEM = 7.7$) model, where the same d' was used for each item in the sequence ($t(49) = 4.0$, $p < 1.0e-3$). Finally, as expected, the Recency Perceptual Summation model performed comparably in predictive accuracy to a model where we estimated a separate d' for each item in the sequence ($\bar{X} = 466$; $SEM = 5.3$; $t(49) = 1.1$, $p = .27$).

Fig. 11 | The Perceptual Summation model predicts ensemble memory for sequentially presented stimuli. Summary of results from Experiments 5, in which participants had to remember colors of sequentially presented real-world objects. The top row shows the fits of the Sequential TCC model to individual data and the second row shows the TCC Sequential Perceptual Summation model's prediction of the ensemble data in two conditions. In the clockwise (counterclockwise) condition the most recently shown items were from the clockwise (counterclockwise) direction from the mean color, producing a clockwise (counterclockwise) bias.

Predictive accuracy of ensemble models in sequential colored object task

Fig. 12 | Comparison in predictive accuracy between Sequential Perceptual Summation model and competing models of ensemble memory for sequentially presented stimuli. The difference in predicted negative log likelihood scores between each of the eight alternative competing models ($PNLL_{Alt}$) and the main Sequential Perceptual Summation model ($PNLL_{PerSum}$). Lower values of PNLL indicate higher predictive accuracy, therefore, PNLL difference scores higher (or lower) than zero indicate support for the Sequential Perceptual Summation (or a competing) model. The vast majority of participants are better predicted by the Sequential Perceptual Summation model than any of the alternatives. Note that the baseline here is the Sequential Perceptual Summation model that relies on fitting a decay rate. The independent d' Perceptual Summation model, the last model above, is the same model but without this parametric assumption about how d' changes across the items in the working memory task. This independent model is instead one in which we used separate d' estimates to quantify familiarity of items as a function of serial position, rather than a single d' and rate parameter. This model is marked with an * because it is also a version of the Sequential Perceptual Summation model and so shows comparable predictive accuracy to the main Sequential Perceptual Summation model we use, as expected.

Together, across five experiments, these results provide converging support for the Perceptual Summation ensemble model, indicating that this very simple generalization of the TCC model to ensemble tasks can capture effects across a wide range of stimulus spaces and for conditions that elicit equal and unequal weighting of items in memory, based solely on measuring memory performance for individual items and generalizing this to ensemble task performance with zero free parameters.

Predictions for other phenomena in ensemble literature. The goal of our modeling work was to develop and test a principled, generalizable computational model of ensemble processing that could account for a wide range of phenomena across different task demands. For this reason, we focused our empirical work on a range of mainstream ensemble tasks and manipulations (e.g., set size, range, sequential vs. simultaneous) that yield relatively robust effects on ensemble reports. In this section we report simulations from the Perceptual

Summation model to demonstrate how the model could account for a wider range of phenomena, with a specific focus on ensemble tasks that either yield inconsistent results (e.g., ensemble tasks with outliers), or are not typically used in the ensemble literature, but may elucidate theoretically interesting boundary conditions of the model (e.g., ensemble tasks with randomly sampled stimuli).

The first set of simulations focuses on the potential effects of outliers on ensemble processing (e.g., a single item very far from the other items in feature space). As shown in the left panel of Figure 13A, the Perceptual Summation model straightforwardly predicts discounting of outliers, which has often but not always been found in the literature (e.g., Haberman & Whitney, 2010; Li et al., 2017). In these simulations we assumed that the outlier stimulus is weighted in the same way, i.e., has the same signal-to-noise ratio, as other stimuli in the ensemble array. However, in spite of this assumption of equal representations, the Perceptual Summation model predicts that the outlier is ‘discounted’ simply because, after pooling via summation, the distributed patterns of activation elicited by other items in the ensemble array collectively outweigh the levels of activation elicited by the outlier stimulus, and so the color with maximum activation is almost entirely unaffected by the outlier. Thus, the model most naturally predicts outliers should be discounted in ensemble tasks even when they are processed in the same way as other items in the memory array.

However, in principle, the model can also be accommodated to capture increased weighting of the outlier as shown in the right panel of Figure 13A. Through the lens of the TCC framework, such increased weighting of the outlier would entail a higher signal-to-noise ratio (larger d' estimate) for the outlier relative to other stimuli in the ensemble array. There are likely situations where this occurs, such as at sufficiently long encoding times, or when the task demands are set up in such a way as to encourage increased attentional processing to salient items that do not fit the display (e.g., as modeled by Brady & Tenenbaum, 2013). The general propensity of models like the one proposed here to discount outliers may thus sometimes be counteracted by people’s tendency to encode outliers more strongly (such as in conditions where a stimulus ‘pops-out’, e.g., Treisman & Gelade, 1980), perhaps explaining why there are mixed results in the literature about whether outliers are discounted (e.g., Cant & Xu, 2020; Khayat & Hochstein, 2018; Hochstein, Pavlovskaya, Bonne, & Soroker, 2018).

To quantitatively test this set of predictions about outlier discounting, researchers would need to obtain a separate d' estimate for the relatively homogenous items and the outlier stimulus in a

visual working memory task. This is not straightforward to do, however, because participants may use high-level strategies in such a task, relying on the other items for items in the homogenous set (e.g., Nassar et al. 2018). In addition, we note that the model is not designed to capture higher-level, decision stage processes that people may use to incorporate the outlier into the ensemble if it is over-weighted relative to other items. For instance, it is possible that if people become aware there is an outlier, they do not simply read out the patterns of activation elicited by the ensemble array, but use additional compensatory strategies to calibrate their responses so they approximate the average more accurately (e.g., explicitly shift their responses toward the average of the homogenous set of items). Proposing a comprehensive theory of outlier processing is outside the scope of our work, but our simulations point to reasons why outlier tasks may vary but may generally result in some outlier discounting.

Another phenomenon which we did not focus on in the current empirical work is that ensemble representations may sometimes become more precise with an increase in set size (e.g., Allik et al., 2013; Baek & Chong, 2020; Haberman & Whitney, 2010; Lee et al., 2016; Parkes, Lund, Angelucci, Solomon, & Morgan, 2001), and sometimes not (e.g., Ji & Pourtois, 2018). As shown in Figure 13B, the Perceptual Summation model predicts such increased precision effects under some experimental conditions but not others. Specifically, under conditions where the range of feature values is the same across different set sizes, the Perceptual Summation model predicts a relative increase in precision for higher set size (compare leftmost panel with rightmost panel in Figure 13B). However, under conditions where the step size between feature values is fixed and the range varies — as in our experiments — the Perceptual Summation model predicts a small drop in precision for higher set sizes even for the same values of d' (compare the leftmost panel to the center panel in Figure 13B). We believe this may account for some of the differences in the literature with regard to whether set size improves or does not improve ensemble performance.

Thus, by the nature of the summation process and representation of items and ensembles as distributions over feature space, the model predicts differential effects of set size on the precision of ensemble representations, depending on the distribution of feature values of to-be-remembered stimuli. Notably, our empirical results and simulations thus provide a parsimonious account for a pattern of results that was previously used to motivate a specific subsampling account of ensemble processing. Specifically, Marchant, Simons and Fockert (2013) found that ensemble performance decreased when the number of items and their heterogeneity increased. To account for these inconsistent effects of set size on ensemble

processing, the authors proposed that people selectively subsample items in the ensemble array, store them in visual working memory, and use these working memory representations of individual items to compute an average at probe (similar to the prediction of the Post-perceptual Summation model). Our simulations demonstrate that these differential effects of set size on ensemble processing can be predicted in a straightforward fashion by the signal detection based Perceptual Summation model without any appeal to subsampling.

Finally, we consider how the different distributions of feature values affect ensemble reports for a fixed set size. In particular, we examine what the model predicts when feature values are drawn randomly or in clusters from a single feature continuum. We did not perform such studies in our empirical work because a basic assumption in ensemble processing in experimental and ecological conditions is that ensemble representations are computed when there are redundancies across items in the external environment. Random sampling of feature values does not ensure such overlap, and makes it conceptually unclear how people would extract an ensemble representation under such conditions.

Nevertheless, we demonstrate that the model would make conceptually plausible predictions for the distribution of activation patterns elicited by the ensemble array for distributions of feature values that are not drawn from just a single (e.g., uniform) cluster. This could potentially be used as the basis for ‘chunking’ (or the formation of ‘clusters’) in visual working memory displays (e.g., Son et al. 2020), providing a unifying principle for such chunking strategies and for ensemble perception.

In particular, as shown in Figure 13C, if there are multiple (six) objects in the ensemble array, the predicted patterns of activation of the ensemble representation will depend on how feature values are distributed. If the distribution is close to uniform (leftmost panel of Figure 13C), the predicted patterns of activation will also approach uniformity, *not* giving rise to well-represented ‘chunks.’ If there is some clustering in the ensemble array (center panel of Figure 13C), the predicted patterns of activation of the ensemble will show slight clustering, particularly for the items that are similar to each other. Finally, the rightmost panel of Figure 13C shows that if stimuli values are constrained to be sampled in clusters, these stimuli would elicit bumps of activation analogous to ‘chunks.’ While it is not clear precisely how to investigate whether participants use such a strategy to form chunks in more heterogeneous displays, the connection between the ensemble model and a potential model of chunking is a useful avenue for future work.

Perceptual Summation model predictions

A) Outlier discounting

B) Differential effects of set size depending on range

C) Chunking: Random and clustered stimuli distributions

Fig. 13 | Predictions of other phenomena from the Perceptual Summation model. (A) Predicted distribution of errors in ensemble tasks where people have to remember eight items and an outlier is present; in particular, where the distribution of feature values is -53, -38, -23, -8, 7, 22, and 112 around the true average (in both panels). In the leftmost panel all stimuli are weighted equally and, therefore, have the same d' , and the outlier — at position 112 — does not significantly skew the resulting response distribution. In the rightmost panel, the outlier is weighted more (d' of 2) than the other stimuli (d' of 0.82, from Schurgin et al. set size 8). The d' value for the outlier was chosen arbitrarily, since there are currently no theories of how much the outlier may be overweighted relative to other items in the array. In this circumstance, the outlier starts to play a role in affecting ensemble performance. (B) Predicted distribution of errors for different set sizes and different stimulus ranges (e.g., different item similarities). The d' used for these simulations were taken from real aggregate data from Schurgin et al., 2020 (d' for six and eight items is 1.1 and 0.82, respectively). The leftmost panel shows the predicted distribution of errors for six items where all items are equally spaced in feature space, and their range is 60°. The center panel shows the predicted distribution of errors for eight items where all items are equally spaced and they have the same step-size (12deg) as the six item condition, but a wider range. The rightmost panel shows the predicted distribution of errors for eight items where all items are equally spaced and they have a smaller

step-size than the six item condition, but the same range. As can be seen, the model predicts different effects of set size depending on how the range of the items is modulated as an effect of set size. (C) Potential generalization of the model to account for chunking. The predicted distribution of errors on the ensemble task for randomly drawn and clustered distribution of stimuli in the ensemble array is shown. The leftmost and center panel show predicted errors for two different simulations in which stimuli feature values are randomly sampled. In the leftmost panel, the distribution of stimuli is -129, -88, -40, 5, 57, and 83 deg from the mean value. In the center panel, the distribution of stimuli is -150, -114, -83, 7, 4, 117 deg. Both distributions were obtained by randomly generating integers using the MATLAB built in function randi on the range -180 to 180. The rightmost distribution shows the predicted distribution of errors on the ensemble task if there is clustering in the ensemble array; here the distribution of stimuli values was set to -120, -108, -96, 96, 108 and 120. The model not only provides a model of ensemble perception with a single cluster of items, but seems to provide a promising avenue for investigating how people separate displays into “chunks” as well.

General Discussion

The contribution of our work is to develop and test a novel, theoretically grounded model of ensemble memory — what we term the Perceptual Summation model — which captures performance across a range of ensemble tasks. The model postulates that ensemble representations reflect a sum of activations in feature space elicited by individual items in the memory array, and that this pooling occurs at a relatively early, encoding stage in processing. Critically, this model of ensemble processing is the first to formally link a model of memory for individual items with a model of memory for ensembles. It is also the first ensemble model that can capture entire distributions of memory errors, and is therefore, the only extant model that makes high-precision predictions of performance in continuous report ensemble tasks. In particular, it provides zero free parameter predictions of the full distribution of errors in ensemble tasks, if visual working memory performance in the same displays is measured first. Finally, it is the first general process-model of ensembles, developed a priori based on an existing theory of memory for individual items to make predictions for any ensemble task. In particular, our modeling and simulations reveal that it can make predictions for ensemble tasks that use different stimulus spaces (e.g., color or shape), tasks that use manipulations of range and memory load, and tasks that induce equal or unequal weighting of items in memory.

Testing alternative processing assumptions. We compared our Perceptual Summation model of ensembles to a suite of contending models that make alternative processing assumptions. These model comparisons permitted us to adjudicate between competing hypotheses regarding how ensembles are extracted.

First, we consider the comparison between the Perceptual Summation model and Automatic Averaging model, which provides insight into whether people do or do not use representations of individual items to compute ensembles. We note that both variants of the Perceptual Summation and Automatic Averaging model are nested within the TCC framework and, therefore, both formally capture the effects of psychophysical similarity and noise accrual on memory processes. The Perceptual Summation model entails that people pool via summation over representations of individual items, whereas the Automatic Averaging model posits that people automatically extract an average from the ensemble array without processing individual items. The Automatic Averaging model aligns with prior proposals (e.g., Ariely, 2001) that people extract an average without maintaining representations of individual items; it serves as a logical foil to the Perceptual and Post-perceptual Summation models, which both predict that ensembles are constructed from representations of individual items. Across all studies, we found that the Perceptual Summation model outperformed the Automatic Averaging model, suggesting that people use representations of individual items to extract ensembles, rather than automatically extracting an average.

The second critical comparison is between the Perceptual and Post-perceptual Summation models, which provides insight into the time-course of ensemble extraction. Both of these models are nested within the TCC framework and both postulate that people use individual item representations to compute an average. The two models differ in the predictions they make about the stage of ensemble extraction. The Perceptual-Summation model entails that people pool over individual item representations relatively early, that is, at the perceptual/encoding stage of ensemble extraction. In contrast, the Post-perceptual Summation model entails that people pool at a later processing stage, that is, after individual items are encoded and consolidated in working memory (e.g., Marchant, Simons, & Fockert, 2013). Thus, unlike the Perceptual Summation model the Post-perceptual Summation models entails that post-perceptual noise accrues before rather than after ensembles are extracted. Again, we found that the Perceptual Summation model outperformed the Post-perceptual Summation model, indicating that pooling occurs at a relatively early processing stage; that is, that ensemble perception in these settings is a perceptual process, rather than simply a complex decision process people can perform with their memory representations after the delay.

We also compared the Perceptual Summation model with four point estimate models, which provides insight into how to characterize ensemble representations. Unlike the Automatic Averaging and Summation models the point estimate models are not nested in the TCC

framework; some of them do not link memory for individual items to memory for ensembles, and critically, they do not entail fully probabilistic representations. More precisely, the Noise-free Point Estimate model entails that people make ensemble judgments based on a noise-free point estimate of the mean. The Noisy Point Estimate model makes the more tenable assumption that people use a point estimate of the mean and an uncertainty interval around the mean. The Average Item Point Estimate model assumes that people maintain a noisy point estimate of each item in the array and average across these to compute the ensemble. Finally, the Precision Ensemble model predicts that people maintain a point estimate of the average with Gaussian uncertainty interval. We found that the Perceptual Summation model outperformed these point estimates model. Relevant in this context, is active debate regarding whether perceptual and memory representations are probabilistic in the way characterized here. For instance, Rahnev et al. (2021) discusses how alternative (e.g, partial-distribution) representations can mimic behavioral outcomes in, e.g., cue combination, tasks traditionally assumed to provide a testbed for probabilistic models of cognition. We find via model comparison that this fully probabilistic model outperforms a point estimate and partial distribution model according to which ensemble representation reflects a point estimate and an uncertainty interval. Furthermore, we find that presuming such probabilistic representations permits us to make high-precision predictions of people's performance, that is, to capture the entire distribution of memory errors in a continuous report task. To summarize, the current ensemble model is consistent with the view that memory decisions in visual working memory ensemble tasks reflect a readout of a probability distribution over feature values. Accordingly, our results are also broadly consistent with neural population coding models of memory, according to which memory representations are grounded in distributed neural patterns of activation across feature values (e.g., Eckstein, 2017; Ma, 2012). Consistent with this speculation, Bays (2019) showed analytically that the TCC theory for individual items can be recast as a neural population coding model of visual working memory where the similarity function replaces the otherwise relatively unconstrained neural population tuning functions.

We conclude this section by noting that, like in all model comparisons, our inferences are qualified by the set of models we consider. That is, we certainly do not promote the view that in principle there is no alternative point estimate model that could provide a better account of the data than the Perceptual Summation model. Our goal was to develop a broad range of alternative models within and outside of the TCC framework, with varying assumptions, and implement them as fairly and generously to each theoretical position as possible. However, we

anticipate that future modeling work may provide a new suite of alternative models. As we discuss next, we believe it is critical that such work focuses on developing models that can account for performance across a range of ensemble tasks and make high precision predictions of performance.

Connections to other quantitative models of ensembles. We took a top-down approach towards model development, meaning that we used an existing framework of memory for individual items to develop, a priori, a memory for ensembles. We contrast this top-down approach to a *task-driven* model development in which models are developed to fit data to a specific task. To our knowledge, the majority of quantitative models of ensemble processes fall within the scope of such task-specific models. For instance, Solomon's (2010) ensemble model was developed to fit data for orientation ensemble tasks, the distributed attention model (Baek and Chong, 2020) was developed to fit data from an ensemble task for size judgments, and Tong, Dube, and Sekuler (2019) ensemble model was developed to fit data for sequential ensemble tasks. Although there is certainly potential for extending these models, it remains unclear how well they would generalize to a wider range of ensemble tasks and stimuli spaces. Moreover, in their current form, none of these models make high-precision predictions of performance by capturing the full distribution of memory errors in a continuous report task, nor generalize from individual item memory to ensemble memory.

Absent a quantitative comparison, we note informally that there is conceptual overlap between these models and the Perceptual Summation model. For instance, like the Baek and Chong (2020) and Solomon (2010) models, the Perceptual Summation model is grounded in signal detection theory, and postulates that ensemble representations can be corrupted at different stages of processing (here, a perceptual pre-summation stage and memory post-summation stage). In addition, we leveraged the Tong et al., (2019) model to develop the Recency Perceptual Summation model and, as such, our model is also aligned with the view that ensemble memory representations in sequential presentation tasks are subject to exponential decay. Furthermore, like the Tong et al., (2019) our ensemble model is broadly consistent with prototype models because it posits an abstraction (via summation) over representations of individual items. However, our model also postulates that the decision variable in the ensemble task reflects a sum of similarity signals from all individual item representations and, therefore, also inherits some assumptions from exemplar models (e.g., Nosofsky, 1988; Kahana & Sekuler, 2002). Through this lens, our model fits with more nuanced views (e.g., Lee & Vanpaemel, 2008; Ross & Makin, 1999; Tong et al., 2018) in which prototype and exemplar

representations can be seen as endpoints on a continuum, as opposed to mutually exclusive classes of representations.

Ensemble processing and gist memory. The Perceptual Summation model for ensembles has direct relevance for theorizing about gist memory, above and beyond tasks traditionally classified as ensemble tasks. Gist memory is broadly defined as memory for ‘generalities’ across multiple items, as opposed to memory for individual items and, as such, both gist and ensemble processes involve abstracting regularities from multiple items (e.g., Alvarez, 2011). The close link between ensemble and gist representations is also evidenced by the fact researchers use ensemble extraction tasks with sequential (e.g., Potter, 1976) and simultaneous (e.g., Zeng, Tompary, Schapiro, & Thompson-Schill, 2021) presentation to study gist memory. Our model of ensembles cannot speak to how memory for gist and individual items interact during short or long-term memory retrieval. However, it provides a candidate explanation for how gist memory representations are computed. As we discussed, our model proposes that the bottleneck during encoding of individual items is the same across visual working memory tasks for individual items, and the extraction of a pooled (ensemble or gist) representation. Furthermore, the model proposes that post-perceptual (memory) noise accrues in the same way regardless of whether people are instructed to remember a single item or ensemble. Critically, however, we postulate that pooling of representations occurs at a relatively early processing stage, prior to post-perceptual noise accrual. Thus, the Perceptual Summation provides a parsimonious and precise account of how gist representations may arise from representations of individual items, while still being more robust than representations of individual items. Our computational model also provides a relatively unambiguous account regarding how processing of individual items differs from processing of gist (e.g., Rosenbaum & Bowman, 2021). That is, instead of using theoretically underspecified constructs (as critique by e.g., Hommel et al, 2019), such as focused versus diffuse attention (e.g., Baek & Chong, 2020; Greene & Naveh-Benjamin, 2021), or pre-attentive and attentive modes of processing (e.g., Chen, Zhuang, Wang, Yanju, & Abrams, 2021; Zepp, Dube, & Melcher, 2021), we explicitly describe how different computations over the same representations can give rise to distinct types of memories. We believe such an approach has great promise for building precise and testable models of gist memory, hierarchical representations and reconstructive memory processes in the visual domain.

Limitations and future directions. One limitation of the current modeling approach is that we use a single parameter d' to measure a potentially diverse set of processes. In line with

standard Signal Detection models, d' in our modeling framework quantifies the signal-to-noise ratio of each memory representation. It is likely that many different processes at encoding, maintenance and retrieval are built into this measure. For instance, it is unlikely that memory representations or processes are themselves one-dimensional. Importantly, signal detection based accounts are fully compatible with this view (Wixted, 2010). Our measure of memory with d' simply captures how people combine their multidimensional memory representations into a single decision variable which they use to make memory judgments when their memory is probed, which is necessarily one-dimensional for a single decision task (for related points see: Gershman, 2021).

We also note that every computational cognitive model makes simplifying assumptions and the appropriate level of specificity depends on the modeling goals (e.g., Lewandowsky & Farrel, 2017). Our modeling goal was to develop a model of ensemble perception that speaks to core issues (e.g., whether ensemble representations are probabilistic vs point estimates; whether memory related noise accrues before or after summation) and to formally connect a model of visual working memory for individual items and ensembles. As a starting point towards this linking analysis, we used a quantitative measure of the signal-to-noise ratio in the visual working memory task for individual items — simply quantifying memory as ‘strength’ for each item, rather than specifying its entire representational format — and examined whether it can be substituted into a process-level ensemble model to predict individual differences in performance on the ensemble task. Thus, in light of our goal of predicting performance across different memory tasks, more parsimonious models provide a useful springboard for generalization.

One concern may be that d' is not only quantifying the signal-to-noise ratio due to memory related processes, but also other non-memory related processes, such as individual differences in task-engagement. Accordingly, the finding that d' predicts individual differences in performance across tasks, may reflect, in part, its ability to capture individual differences in third variables such as how conscientiously people approach the task. We certainly do not deny that part of the variance captured by d' may reflect individual differences in other global factors such as task engagement, which undoubtedly have downstream effects and explain inter-individual differences in memory performance in *all* memory tasks. However, our results from intra-individual analyses (within-subject conditions) — for example, how both item memory and ensemble memory change as a function of range or set size — are not possible to explain by appealing solely to overall differences in motivation across participants. In addition, data from the sequential paradigm (Experiment 5) provide strong evidence against the view that the

predictive value of d' reflects solely on such across-subject factors, and not individual differences in memory-related processes. This is because in Experiment 5 we show strong support for the view that our models and generalization approach also captures intra-individual variance in performance (as captured by changes in d' as a function of serial position), that is, memory changes as a function of the serial position of the stimulus within the same person. Such within-task predictions would be impossible to make if d' only captured non-memory related processes.

We also point out that in the current work we instantiate models according to which ensemble summaries are computed after the retention delay (e.g., the Post-perceptual summation model), however, we do not instantiate models in which participants 'sub-sample' items from the array (Myczek & Simons, 2008) by using information from the subsampled items but no information from other items. This subsampling view entails that participants can functionally ignore the remaining subset of items at an early encoding stage, a view compatible with item-based capacity visual working memory models (e.g., Cowan, 2001; Zhang & Luck, 2008), which postulates that people can store no information about some subset of items (an assumption coined 'complete information-loss' in the broader modeling literature, e.g., Kellen & Klauer, 2015). However, such a threshold-based view (that is, complete loss of all information about subsets of items) is incompatible with new working memory data which provide strong support for the view that working memory representations vary continuously in strength and are rarely or never lost completely (e.g., Robinson et al. 2020; Williams, Robinson, Schurgin, Wixted, Brady, 2022). Relatedly, this view is also incompatible with most modern accounts of visual working memory (e.g., van den Berg et al. 2012; Bays, 2014; Schneegans et al. 2020; Schurgin et al. 2020) according to which all perceived items give rise to noisy representations in visual working memory rather than being represented in an all-or-none fashion – that is, being represented with complete fidelity or not being represented at all. If future work reveals that, unlike visual working memory processes, ensemble perception does involve complete information loss about some subset of items, such models could be tested in the current framework (with a mixture of d' values for different items, with some set to zero and others set to be greater than zero).

A final limitation is that we did not model all possible phenomena in the ensemble literature. However, our goal was to formally establish a link between two different processing models, and therefore, we focused our modeling and analysis on a set of mainstream ensemble tasks where the patterns of effects are robust (e.g., changes with range, set size, stimuli, and simultaneous vs. sequential). We report model simulations to demonstrate that our model can, in principle,

capture both outlier discounting or increased weighting of outliers and that which effect should occur in a given task should be predictable by measuring memory for individual items. We also report simulations to show how the model accounts for differential effects of set size on the fidelity of ensemble representations, and effects of various (e.g., random) distributions of stimuli in the ensemble array on ensemble reports. Together, our aim is to lay a novel theoretical and methodological framework for future modeling research of ensemble and gist memory.

Methods

In each experiment, we collected data until our final sample size was $n=50$. Participants were from the University of California, San Diego community and participated in exchange for course credit. All participants were at least 18-years old, provided informed consent and reported normal or corrected-to-normal vision. The study was completed online through the University's SONA system and approved by the Institutional Review Board. We did not analyze data of participants who failed to complete the study. We also excluded data from participants if their d' estimates in any of the visual working memory task conditions (and for the last item in the sequential presentation task) were below 1.5 standard deviations of the group mean.

Experiment 1: Memory for color with manipulation of set size Figure 5 shows an example schematic of trials in the visual-working memory and ensembles task for all five experiments. Participants completed a block of a visual-working memory task and an ensemble task (order of blocks was randomized across participants). At the beginning of every trial in both the visual-working memory and ensemble tasks, participants were shown a written prompt with the current trial number and the total number of trials in that block (1,000ms). After the prompt offset, participants were shown a fixation cross in the center of the screen and (six) placeholders (1,000ms). Next, participants were briefly presented with the memory array (350ms). We manipulated memory load (randomly across trials) in the visual working memory and ensemble tasks, thus, participants were instructed to remember six (50% of trials) or eight items in both tasks. The color of each circle was randomly sampled from the CIELAB color space of Schurgin et al. (2020) with the constraint that each color had to be at least 30° away from other colors in the array. In the ensemble block, the memory array also consisted either of six (50% of trials) or eight colored circles. The step-size between colors in the ensemble task was fixed to 15° for both set sizes. There were 150 trials in the visual working memory block and 150 trials in the ensemble block (75 trials per memory load condition in each of the tasks).

The memory array in both blocks was followed by a retention interval (900ms) and the memory probe. In the visual-working memory task, participants were shown a black outline around one of the placeholders, which cued them to report on the color of the circle shown in that spatial location. In the ensemble task, participants were instructed to report on the average color. In both tasks, participants reported on the color using a color wheel.

Experiment 2: Memory for color with manipulation of range in the ensemble task The procedure of Experiment 2 was identical to the procedure in Experiment 1 with the following exceptions. First, in both the visual-working and ensemble memory block, the memory array always consisted of six colored circles. Second, in the ensemble task, the stepsize between colors in each condition was constrained to be 10° (60° range condition), 15° (90° range condition), or 20° (120° range condition). There were 75 trials of the visual-working memory task, and 225 trials of the ensemble task (75 trials in each of the three range conditions).

Experiment 3: Memory for shape with manipulation of set size Experiment 3 was identical to Experiment 1, with the exception that people were shown shapes instead of colors and we changed encoding time to 1,000ms and the retention interval to 800ms because the shape task is more difficult than the color task. Shape stimuli were taken from Li, Liang, Lee & Barense (2020).

Experiment 4: Memory for shape with manipulation of range in the ensemble task Experiment 4 was also identical to Experiment 2, with the exception that people were shown shapes instead of colors and we changed encoding time to 1,000ms and the retention interval to 800ms.

Experiment 5: Memory for sequentially presented colors The goal of Experiment 5 was to model data from a sequential instead of simultaneous presentation paradigm. Therefore, in this experiment participants were instructed to remember colors of colored pictures of real-world objects (objects taken from: Brady, Schacter, & Alvarez, 2018). We used pictures of real-world objects instead of uniform stimuli (e.g., circles) because this allowed us to easily probe an item's serial position in the sequential visual-working task for individual items by showing participants a grayscale photo of one of the objects in the sequence and probing them on that object's color.

As before, all participants completed a block of the visual working memory and ensemble task. In both tasks, participants self-advanced each trial by mouse-clicking on a fixation cross in the center of the screen. The mouse-click was followed by a brief delay (1,000ms), after which they were shown a sequence of six objects, each presented one at-a-time in the center of the computer screen. In both tasks, each object was presented for 600ms and followed by a 450ms inter-stimulus interval. In the visual-working memory task, on each trial, each object in the sequence was unique, and the color of each object was constrained to be at least 30° away from colors of the remaining object. In the ensemble task, on each trial, each object in the

sequence was the same (though different objects were presented across trials), and the stepsize between colors was 20° . In order to measure effects of recency in the ensemble task, we adapted a manipulation from Tong et al., (2019). Specifically, on half of the trials the first (or last) three objects in the sequence had colors that were counterclockwise to the mean color, whereas the last (or first) three objects had colors that were clockwise from the mean color in color space.

In both tasks, the last object in the sequence was followed by a 900ms delay. Within the delay period participants were shown a dynamic visual mask, which was displayed for 100ms, 100ms after the last object offset. The mask was used to reduce potential effects of iconic memory on recency effects in the sequential presentation design. After the retention interval, participants were probed on their memory with a continuous report. In the visual-working memory task, participants were shown a grayscale version of one of the six objects in the sequence and instructed to adjust its color to its color on that trial. In the ensemble task, participants were shown a grayscale version of the object from that trial and instructed to adjust it to have average color on that trial. There were 120 trials in the visual working memory task and each object in the sequence was probed equiprobably (on 20 trials) across the experimental block. There were 96 trials in the ensemble task, with 48 trials in the counterclockwise and clockwise conditions.

Technical details of model fitting and prediction

Models were fit separately to each participant's visual working memory data. The best-fitting parameter estimates from these fits were used to predict the same person's data on the ensemble task. In Experiments 1-4, we fit the standard TCC model for single items to the visual working memory data. The formula for this model is given in Equation 1. After obtaining a d' estimate from fitting models to the visual-working memory data, we substituted this parameter into Equations 2-4 of the Perceptual and Post-perceptual Summation and Automatic Averaging model, to predict the ensemble data. In Experiment 5, we fit the sequential version of the visual-working memory model (Equation 5), and substituted both the d' and rate parameters into Equations 6 and 7 to predict the ensemble data using the Recency variants of the Perceptual and Post-perceptual Summation models for ensembles, respectively. As noted, the Automatic Averaging model postulates that people extract a single representation of the mean without building it up from representations of individual items, therefore, for this model we only used a single d' estimate to make predictions.

To generate predictions from the Noise-free point estimate model we calculated an equally weighted average value of the ensemble (like in the Automatic Averaging model). We then added a small amount of jitter to this estimate to simulate small effects of motor noise. To simulate motor noise, we used the built-in `randn` function in MATLAB, which generates random samples from a standard Gaussian distribution. To generate predictions from the Noisy Point estimate model we drew random samples of data from a uniform distribution, which had a range of 60° and was centered on the true value of the average i.e., samples were drawn from -30 and 30° around the mean value. The assumption behind this model is that people represent a uniform uncertainty interval around the true mean value, thus, the number of samples was based on the number of trials in each ensemble condition. To generate predictions for the Average Item Point Estimate we used people's working memory data to sample n point estimates (measured in degrees of error) for each item in the array, and then averaged across these. For instance, if people had to remember six items on the visual working memory and ensemble tasks, we drew (with replacement) six samples of their self-report data (converted to error in degrees) on the visual working memory task, and averaged across these. This was repeated for each trial, to generate the predicted distribution of errors on the ensemble task. Finally, the Precision Ensemble model was implemented by fitting a Standard Mixture model (e.g., Zhang & Luck, 2008) to each individual's visual working memory data and using the standard deviation estimate from this model to compute an uncertainty interval around the true average. The latter was implemented by sampling random samples of data (based on the number of trials on the ensemble task), from a normal distribution with mean zero and the standard deviation set to the standard deviation (inverse of precision) estimate from the mixture model.

Models were fit to visual-working memory data in MATLAB using Maximum Likelihood Estimation (MLE) by minimizing the negative log likelihood. Minimization was implemented with the `fmincon` algorithm in the Optimization Toolbox. The predictive accuracy of each model was measured using the predictive negative log likelihood (PNLL). As explained, we substituted best-fitting parameters from the visual-working memory task into equations for each of the ensemble models to predict data on the ensemble task, and calculated the PNLL using data from the ensemble task conditions.

Finally, the psychophysical similarity functions in all of our experiments were estimated in prior work using standard methods (e.g., triad judgments) for psychophysical scaling similarity data. See Schurgin et al. (2020) for full description.

References

- Allik, J., Toom, M., Raidvee, A., Averin, K., & Kreegipuu, K. (2013). An almost general theory of mean size perception. *Vision Research*, 83, 25–39.
<https://doi.org/10.1016/j.visres.2013.02.018>
- Alvarez, G. A. (2011). Representing multiple objects as an ensemble enhances visual cognition. *Trends in Cognitive Sciences*, 15(3), 122–131. <https://doi.org/10.1016/j.tics.2011.01.003>
- Alvarez, G. A., & Oliva, A. (2008). The Representation of Simple Ensemble Visual Features Outside the Focus of Attention. *Psychological Science*, 19(4), 392–398.
<https://doi.org/10.1111/j.1467-9280.2008.02098.x>
- Anderson, J. R. (Ed.) (1981). *Cognitive skills and their acquisition*. Hillsdale, NJ: Erlbaum.
- Ariely, D. (2001). Seeing Sets: Representation by Statistical Properties. *Psychological Science*, 12(2), 157–162. <https://doi.org/10.1111/1467-9280.00327>
- Averbeck, B. B., Latham, P. E., & Pouget, A. (2006). Neural correlations, population coding and computation. *Nature Reviews Neuroscience*, 7(5), 358–366.
<https://doi.org/10.1038/nrn1888>
- Baddeley, A. (1992). Working memory. *Science*, 255(5044), 556–559.
- Baek, J., & Chong, S. C. (2020). Ensemble perception and focused attention: Two different modes of visual processing to cope with limited capacity. *Psychonomic Bulletin & Review*, 27(4), 602–606. <https://doi.org/10.3758/s13423-020-01718-7>
- Balas, B., Nakano, L., & Rosenholtz, R. (2009). A summary-statistic representation in peripheral vision explains visual crowding. *Journal of Vision*, 9(12), 13.
<https://doi.org/10.1167/9.12.13>
- Bartolo, R., Saunders, R. C., Mitz, A. R., & Averbeck, B. B. (2020). Information-Limiting Correlations in Large Neural Populations. *The Journal of Neuroscience*, 40(8), 1668–1678. <https://doi.org/10.1523/jneurosci.2072-19.2019>
- Bays, P. M. (2014). Noise in neural populations accounts for errors in working memory. *Journal of Neuroscience*, 34(10), 3632–3645. doi:10.1523/JNEUROSCI.3204-13.2014

- Bays, P. M. (2015). Spikes not slots: noise in neural populations limits working memory. *Trends in cognitive sciences*, 19(8), 431-438.
- Bays, P. M. (2019). Correspondence between population coding and psychophysical scaling models of working memory. *bioRxiv*, 699884.
- Block, N. (2011). Perceptual consciousness overflows cognitive access. *Trends in cognitive sciences*, 15(12), 567-575.
- Block, N. (2014). Rich conscious perception outside focal attention. *Trends in cognitive sciences*, 18(9), 445-447.
- Brady, T. F., & Oliva, A. (2008). Statistical Learning Using Real-World Scenes. *Psychological Science*, 19(7), 678–685. <https://doi.org/10.1111/j.1467-9280.2008.02142.x>
- Brady, T. F., Schacter, D. L., & Alvarez, G. (2018). The adaptive nature of false memories is revealed by gist-based distortion of true memories. *PsyArxiv*, <https://doi.org/10.31234/osf.io/zeg95>
- Brady, T. F., Shafer-Skelton, A., & Alvarez, G. A. (2017). Global ensemble texture representations are critical to rapid scene perception. *Journal of Experimental Psychology: Human Perception and Performance*, 43(6), 1160–1176. <https://doi.org/10.1037/xhp0000399>
- Brady, T. F., & Tenenbaum, J. B. (2013). A probabilistic model of visual working memory: Incorporating higher order regularities into working memory capacity estimates. *Psychological Review*, 120(1), 85–109. <https://doi.org/10.1037/a0030779>
- Bussemeyer, J. R., & Wang, Y. M. (2000). Model Comparisons and Model Selections Based on Generalization Criterion Methodology. *Journal of Mathematical Psychology*, 44(1), 171–189. <https://doi.org/10.1006/jmps.1999.1282>
- Cant, J. S., & Xu, Y. (2020). One bad apple spoils the whole bushel: The neural basis of outlier processing. *NeuroImage*, 211, 116629. <https://doi.org/10.1016/j.neuroimage.2020.116629>

- Chen, Z., Zhuang, R., Wang, X., Ren, Y., & Abrams, R. A. (2021). Ensemble perception without attention depends upon attentional control settings. *Attention, Perception, & Psychophysics*, 83, 1240-1250.
- Chater, N., Tenenbaum, J. B., & Yuille, A. (2006). Probabilistic models of cognition: where next? *Trends in Cognitive Sciences*, 10(7), 292–293. <https://doi.org/10.1016/j.tics.2006.05.008>
- Chetverikov, A., Campana, G., & Kristjánsson, R. (2016). Building ensemble representations: How the shape of preceding distractor distributions affects visual search. *Cognition*, 153, 196–210. <https://doi.org/10.1016/j.cognition.2016.04.018>
- Chong, S. C., Joo, S. J., Emmanouil, T. A., & Treisman, A. (2008). Statistical processing: Not so implausible after all. *Perception & Psychophysics*, 70(7), 1327-1334. <https://doi.org/10.3758/PP.70.7.1327>
- Cohen, M. A., Dennett, D. C., & Kanwisher, N. (2016). What is the Bandwidth of Perceptual Experience? *Trends in Cognitive Sciences*, 20(5), 324–335. <https://doi.org/10.1016/j.tics.2016.03.006>
- Cowan, N. (2001). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and brain sciences*, 24(1), 87-114.
- Eckstein, M. P. (2017). Probabilistic computations for attention, eye movements, and search. *Annual review of vision science*, 3, 319-342.
- Gershman, S. J. (2021). The rational analysis of memory. In *Oxford handbook of human memory*. Oxford University Press Oxford, UK.
- Goldstein, M. H., Waterfall, H. R., Lotem, A., Halpern, J. Y., Schwade, J. A., Onnis, L., & Edelman, S. (2010). General cognitive principles for learning structure in time and space. *Trends in Cognitive Sciences*, 14 (6):249-258.
- Grahek, I., Schaller, M., & Tackett, J. L. (2021). Anatomy of a psychological theory: Integrating construct-validation and computational-modeling methods to advance theorizing. *Perspectives on Psychological Science*, 1745691620966794.
- Greene, N. R., & Naveh-Benjamin, M. (2021). The effects of divided attention at encoding on specific and gist-based associative episodic memory. *Memory & Cognition*, 1-18.

- Guest, O., & Martin, A. E. (2021). How computational modeling can force theory building in psychological science. *Perspectives on Psychological Science*, 16(4), 789–802. <https://doi.org/10.1177/1745691620970585>
- Haberman, J., & Whitney, D. (2011). Efficient summary statistical representation when change localization fails. *Psychonomic Bulletin & Review*, 18(5), 855–859. <https://doi.org/10.3758/s13423-011-0125-6>
- Hansmann-Roth, S., Thorsteinsdóttir, S., Geng, J., & Kristjánsson, R. (2021). Temporal integration of feature probability distributions in visual working memory. *Journal of Vision*, 21(9), 1969. <https://doi.org/10.1167/jov.21.9.1969>
- Hemmer, P., & Steyvers, M. (2009). A Bayesian Account of Reconstructive Memory. *Topics in Cognitive Science*, 1(1), 189–202. <https://doi.org/10.1111/j.1756-8765.2008.01010.x>
- Hintzman, D. L. (1986). “Schema abstraction” in a multiple-trace memory model. *Psychological Review*, 93(4), 411–428. <https://doi.org/10.1037/0033-295x.93.4.411>
- Hochstein, S., Pavlovskaya, M., Bonne, Y. S., & Soroker, N. (2018). Comparing set summary statistics and outlier pop out in vision. *Journal of Vision*, 18(13), 12. <https://doi.org/10.1167/18.13.12>
- Hommel, B., Chapman, C. S., Cisek, P., Neyedli, H. F., Song, J. H., & Welsh, T. N. (2019). No one knows what attention is. *Attention, Perception, & Psychophysics*, 81(7), 2288–2303.
- Howard, M. W., & Kahana, M. J. (2002). A Distributed Representation of Temporal Context. *Journal of Mathematical Psychology*, 46(3), 269–299. <https://doi.org/10.1006/jmps.2001.1388>
- Ji, L., & Pourtois, G. (2018). Capacity limitations to extract the mean emotion from multiple facial expressions depend on emotion variance. *Journal of Vision*, 18(10), 610. <https://doi.org/10.1167/18.10.610>
- Kahana, M. J., & Sekuler, R. (2002). Recognizing spatial patterns: a noisy exemplar approach. *Vision Research*, 42(18), 2177–2192. [https://doi.org/10.1016/s0042-6989\(02\)00118-9](https://doi.org/10.1016/s0042-6989(02)00118-9)
- Kahneman, D. (2003). A Psychological Perspective on Economics. *American Economic Review*, 93(2), 162–168. <https://doi.org/10.1257/000282803321946985>

- Khayat, N., & Hochstein, S. (2018). Perceiving set mean and range: Automaticity and precision. *Journal of Vision*, 18(9), 23. <https://doi.org/10.1167/18.9.23>
- Kohn, A., Coen-Cagli, R., Kanitscheider, I., & Pouget, A. (2016). Correlations and Neuronal Population Information. *Annual Review of Neuroscience*, 39(1), 237–256. <https://doi.org/10.1146/annurev-neuro-070815-013851>
- Lee, M. D. (2011). How cognitive modeling can benefit from hierarchical Bayesian models. *Journal of Mathematical Psychology*, 55, 1–7.
- Lee, H., Baek, J., & Chong, S. C. (2016). Perceived magnitude of visual displays: Area, numerosity, and mean size. *Journal of Vision*, 16(3), 1-11. <https://doi.org/10.1167/16.3.12>
- Lee, M. D., & Vanpaemel, W. (2008). Exemplars, Prototypes, Similarities, and Rules in Category Representation: An Example of Hierarchical Bayesian Analysis. *Cognitive Science*, 32(8), 1403–1424. <https://doi.org/10.1080/03640210802073697>
- Li, V., Herce Castañón, S., Solomon, J. A., Vandormael, H., & Summerfield, C. (2017). Robust averaging protects decisions from noise in neural computations. *PLoS Computational Biology*, 13.
- Luce, R. D., & Galanter, E. (1963). Psychophysical scaling. *Handbook of mathematical psychology*, 1, 245-307.
- Ma, W. J. (2012). Organizing probabilistic models of perception. *Trends in cognitive sciences*, 16(10), 511-518.
- Macmillan, N. A., & Creelman, C. D. (1996). Triangles in ROC space: History and theory of “nonparametric” measures of sensitivity and response bias. *Psychonomic Bulletin & Review*, 3(2), 164–170. <https://doi.org/10.3758/bf03212415>
- 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(2), 245–250. <https://doi.org/10.1016/j.actpsy.2012.11.002>
- McCarley, J. S., & Benjamin, A. S. (2013). Bayesian and signal detection models. In *The Oxford Handbook of Cognitive Engineering*.

- Miller, G. A. (1956). The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological review*, 63(2), 81.
- Murdock, B. B. (1982). A theory for the storage and retrieval of item and associative information. *Psychological Review*, 89(6), 609–626. <https://doi.org/10.1037/0033-295x.89.6.609>
- Nassar, M. R., Helmers, J. C., & Frank, M. J. (2018). Chunking as a rational strategy for lossy data compression in visual working memory. *Psychological review*, 125(4), 486.
- Navarro, D. J. (2021). If mathematical psychology did not exist we might need to invent it: A comment on theory building in psychology. *Perspectives on Psychological Science*, 1745691620974769.
- Newell, A. (1994). *Unified theories of cognition*. Harvard University Press.
- Nosofsky, R. M. (1987). Attention and learning processes in the identification and categorization of integral stimuli. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 13(1), 87–108. <https://doi.org/10.1037/0278-7393.13.1.87>
- Oberauer, K., & Lewandowsky, S. (2019). Addressing the theory crisis in psychology. *Psychonomic Bulletin & Review*, 26(5), 1596–1618. <https://doi.org/10.3758/s13423-019-01645-2>
- Oriet, C., Giesinger, C., & Stewart, K. M. (2020). Can change detection succeed when change localization fails? *Journal of Experimental Psychology: Human Perception and Performance*, 46(10), 1127–1147. <https://doi.org/10.1037/xhp0000834>
- Palestro, J. J., Bahg, G., Sederberg, P. B., Lu, Z. L., Steyvers, M., & Turner, B. M. (2018). A tutorial on joint models of neural and behavioral measures of cognition. *Journal of Mathematical Psychology*, 84, 20–48. <https://doi.org/10.1016/j.jmp.2018.03.003>
- Parkes, L., Lund, J., Angelucci, A., Solomon, J. A., & Morgan, M. (2001). Compulsory averaging of crowded orientation signals in human vision. *Nature Neuroscience*, 4(7), 739–744. <https://doi.org/10.1038/89532>
- Pashler, H. (1984). Processing stages in overlapping tasks: Evidence for a central bottleneck. *Journal of Experimental Psychology: Human Perception and Performance*, 10(3), 358–377. <https://doi.org/10.1037/0096-1523.10.3.358>

- Potter, M. C. (1976). Short-term conceptual memory for pictures. *Journal of experimental psychology: human learning and memory*, 2(5), 509.
- Rahnev, D., Block, N., Denison, R. N., & Jehee, J. (2021). Is perception probabilistic? Clarifying the definitions. *PsyArXiv*
- Reder, L. M., Nhouyvanisvong, A., Schunn, C. D., Ayers, M. S., Angstadt, P., & Hiraki, K. (2000). A mechanistic account of the mirror effect for word frequency: A computational model of remember-know judgments in a continuous recognition paradigm. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 26(2), 294.
- Robinson, M. M., Benjamin, A. S., & Irwin, D. E. (2020). Is there a K in capacity? Assessing the structure of visual short-term memory. *Cognitive Psychology*, 121, 101305.
<https://doi.org/10.1016/j.cogpsych.2020.101305>
- van Rooij, I., & Baggio, G. (2021). Theory before the test: How to build high-verisimilitude explanatory theories in psychological science. *Perspectives on Psychological Science*, 16(4), 682-697.
- Rosenbaum, D., & Bowman, H. (2021). Extraction of gist without encoding of individual items in RSVP of numerical sequences.
- Ross, B. H., & Makin, V. S. (1999). Prototype versus exemplar models in cognition. *The nature of cognition*, 205-241.
- Rust, N. C. (2014). 30 Population-Based Representations. *The cognitive neurosciences*, 337.
- Schneegans, S., Taylor, R., & Bays, P. M. (2020). Stochastic sampling provides a unifying account of visual working memory limits. *Proceedings of the National Academy of Sciences*, 117(34), 20959-20968.
- Schurgin, M. W., Wixted, J. T., & Brady, T. F. (2020). Psychophysical scaling reveals a unified theory of visual memory strength. *Nature Human Behaviour*, 4(11), 1156–1172.
<https://doi.org/10.1038/s41562-020-00938-0>
- Šetić, M., ŠVegar, D., & Domijan, D. (2007). Modelling the statistical processing of visual information. *Neurocomputing*, 70(10–12), 1808–1812.
<https://doi.org/10.1016/j.neucom.2006.10.069>

- Shamir, M. (2014). Emerging principles of population coding: in search for the neural code. *Current opinion in neurobiology*, 25, 140-148.
- Shepard, R. N. (1987). Toward a universal law of generalization for psychological science. *Science*, 237(4820), 1317-1323.
- Shiffrin, R. M., & Steyvers, M. (1997). A model for recognition memory: REM—retrieving effectively from memory. *Psychonomic Bulletin & Review*, 4(2), 145–166.
<https://doi.org/10.3758/bf03209391>
- Simon, H. A. (1990). Invariants of human behavior. *Annual review of psychology*, 41(1), 1-20.
- Sims, C. R. (2018). Efficient coding explains the universal law of generalization in human perception. *Science*, 360(6389), 652-656.
- Smith, J. D., & Minda, J. P. (1998). Prototypes in the mist: The early epochs of category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 24(6), 1411–1436. <https://doi.org/10.1037/0278-7393.24.6.1411>
- Solomon, J. A. (2010). Visual discrimination of orientation statistics in crowded and uncrowded arrays. *Journal of Vision*, 10(14), 1-16. <https://doi.org/10.1167/10.14.19>
- Solomon, J. A. (2021). Five dichotomies in the psychophysics of ensemble perception. *Attention, Perception, & Psychophysics*, 83(3), 904-910.
- Son, G., Oh, B. I., Kang, M. S., & Chong, S. C. (2020). Similarity-based clusters are representational units of visual working memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 46.
- Stevens, S. S. (1936). A scale for the measurement of a psychological magnitude: loudness. *Psychological Review*, 43(5), 405–416.
- Swets, J. A. (1986). Form of empirical ROCs in discrimination and diagnostic tasks: Implications for theory and measurement of performance. *Psychological Bulletin*, 99, 181–198.
- Tenenbaum, J. B. (1999). Bayesian modeling of human concept learning. *Advances in neural information processing systems*, 59-68.

- Tong, K., Dubé, C., & Sekuler, R. (2019). What makes a prototype a prototype? Averaging visual features in a sequence. *Attention, Perception, & Psychophysics*, 81(6), 1962-1978.
- Thurstone, L. L. (1927). A law of comparative judgment. *Psychological Review*, 34(4), 273–286. <https://doi.org/10.1037/h0070288>
- Treisman, A. M., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, 12(1), 97–136. [https://doi.org/10.1016/0010-0285\(80\)90005-5](https://doi.org/10.1016/0010-0285(80)90005-5)
- Turner, B. M., Forstmann, B. U., Love, B. C., Palmeri, T. J., & van Maanen, L. (2017). Approaches to analysis in model-based cognitive neuroscience. *Journal of Mathematical Psychology*, 76, 65–79. <https://doi.org/10.1016/j.jmp.2016.01.001>
- Utochkin, I. (2015). Ensemble summary statistics as a basis for visual categorization. *Journal of Vision*, 15(12), 891. <https://doi.org/10.1167/15.12.891>
- Van den Berg, R., Shin, H., Chou, W. C., George, R., & Ma, W. J. (2012). Variability in encoding precision accounts for visual short-term memory limitations. *Proceedings of the National Academy of Sciences*, 109(22), 8780-8785.
- Ward, E. J., 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. <https://doi.org/10.1016/j.cognition.2016.01.010>
- Whitney, D., & Yamanashi Leib, A. (2018). Ensemble Perception. *Annual Review of Psychology*, 69(1), 105–129. <https://doi.org/10.1146/annurev-psych-010416-044232>
- Wickens, T. D. (2001). *Elementary signal detection theory*. New York, NY, US: Oxford University Press.
- Williams, J. R., Robinson, M. M., Schurgin, M.W., Wixted, J.T., and Brady, T.F. (2022). You can't "count" how many items people remember in working memory: The importance of signal detection-based measures for understanding change detection performance. *PsyArXiv*
- Williams, J., Brady, T.F., and Störmer, V.S. (in press). Guidance of attention by working memory is a matter of representational fidelity. *Journal of Experimental Psychology: Human Perception and Performance*.

Wixted, J. T. (2020). The forgotten history of signal detection theory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 46(2), 201–233.
<https://doi.org/10.1037/xlm0000732>

Yarkoni, T. (2021). The generalizability crisis. *Behavioral and Brain Sciences*, 1-37.

Zeng, T., Tompary, A., Schapiro, A. C., & Thompson-Schill, S. L. (2021). Tracking the relation between gist and item memory over the course of long-term memory consolidation. *ELife*, 10. <https://doi.org/10.7554/elife.65588>

Zepp, J., Dubé, C., & Melcher, D. (2021). A direct comparison of central tendency recall and temporal integration in the successive field iconic memory task. *Attention, Perception, & Psychophysics*, 83(3), 1337-1356.

Zhang, W., & Luck, S. J. (2008). Discrete fixed-resolution representations in visual working memory. *Nature*, 453(7192), 233–235. <https://doi.org/10.1038/nature06860>