Neuromethods
DOI 10.1007/7657_2019_20

Feature Distribution Learning (FDL): A New Method for Studying Visual Ensembles Perception with Priming of Attention Shifts

Andrey Chetverikov, Sabrina Hansmann-Roth, Ömer Dağlar Tanrıkulu, and Árni Kristjánsson

Abstract

We discuss how priming of attention shifts has in recent studies proved to be a useful method for studying internal representations of visual ensembles. Attentional priming is very powerful in particular when role reversals between targets and distractors occur. Such role reversals can be used to assess how expected or unexpected a particular target is. This new method for studying representations of visual ensembles has revealed that observer's representations are far more detailed than previous studies of ensemble perception have suggested where the emphasis has been on summary statistics, i.e., mean and variance. Observers can represent surprisingly complex distribution shapes such as whether a representation is bimodal or not. We discuss the details of how this feature distribution learning (FDL) method has been used to assess internal representations of visual ensembles. We also speculate that the method can prove to be an important *implicit* way of assessing how observers represent regularities in their environments.

Keywords Perceptual representations, Visual ensembles, Visual search, Priming, Feature distribution learning (FDL)

1 Intro and Background

Priming of attention shifts has been extensively investigated over the last 25 years, mainly with various forms of visual search tasks. A key finding in this literature comes from the studies of Maljkovic and Nakayama [1]. They used a paradigm introduced by Bravo and Nakayama [2] where observers searched for an oddly colored diamond among distractors of another color and had to judge whether there was a notch on its left or right side. Bravo and Nakayama had observed that when targets maintained their color between trials (e.g., the target was always the red diamond among green distractor diamonds), search was overall faster than when the target and distractor identity reversed unpredictably (from a red target diamond among green diamond distractors to a green target diamond among red distractor diamonds). There was therefore a benefit to target and distractor consistency from one trial to the next.

Maljkovic and Nakayama [1] replicated this result, finding additionally that search became faster the more often the same target color repeated. The consistency benefit therefore reflected this "priming" effect, and importantly they found that the priming effects were only minimally affected by top-down strategies (such as whether observers knew the upcoming target color or not; but see [3, 4]). Maljkovic and Nakayama also tested the cumulative effects of priming over several trials, finding that response times decreased by 20 to 25%, with no corresponding increases in error rates (see, e.g., [5]).

This basic finding on priming during feature search has been replicated many times (see [6], for review). Here, our aim is firstly to discuss basic considerations for priming paradigms and secondly to introduce a paradigm that utilizes priming to assess how human observers represent the environment.

1.1 What Can Prime?

It is important to note that not only target characteristics can prime from one attentional allocation to the next but also the identity of the distractors [1, 7–9]. If the same distractors appear from one trial to the next, search will be speeded, irrespective of whether the target identity is unchanged [1, 7]. These distractor priming effects therefore make their independent contribution to performance [8, 10, 11] and can be just as strong as the target priming effects, although the two can interact. In a paradigm where the target and distractors do not vary independently, it is impossible to disentangle the two, and this needs to be taken into account when results from priming studies are interpreted. These two separate sources of priming effects can combine so that when an odd-one-out target contains the colors of the distractor on the preceding trials, search is slowed even more, reflecting so-called role-reversal effects. Importantly, the strength of these role-reversal effects can be used to answer other questions regarding visual perception, as we discuss below.

The priming effects have typically been thought to reflect facilitation of individual features [12]. Other findings show that such priming can occur from the repetition of more complex characteristics, such as feature combinations or objects identities [13–16]. Separate *features* of the stimuli that observers search for can cause their own priming effect, and so can whole targets, depending on the circumstances [15]. For example, color, spatial frequency, and orientation can cause independent priming effects depending on task relevance [17].

A critical feature of these priming effects is that they are so strong that they sometimes seem to be able to account for effects that have typically been attributed to explicit top-down attentional guidance in the literature [14, 18, 19]. As an example, Kristjánsson, Wang, and Nakayama showed that search times were similar when target identity was always the same and when priming effects were maximal, although target identity was not known from one trial to the next, showing that large portions of effects attributed to top-down guidance were accounted for by priming. Belopolsky et al. [18] showed that so-called contingent-capture effects [20], thought to be caused by top-down guidance, could, to a large extent, be explained by priming. Theeuwes and van der Burg [19] assessed interference effects from irrelevant distractor stimuli, finding that observers could not use top-down set (from verbal or symbolic cues) to ignore irrelevant color singletons (the interference was still present), but they further argued that when attentional priming effects were maximal, interference from irrelevant distractors was minimal.

A common interpretation of attentional priming effects is that searching for a target automatically creates a representation of that item or feature depending on context [15] which, in turn, influences subsequent attention allocation. This entails the assumption that the processing of the features that are contained in the template is facilitated. Such templates are often thought to be kept in visual working memory—and there is indeed evidence suggesting that visual working memory content can modulate priming [21, 22]. Other accounts of priming involve the dimensional weighting account [23, 24] and the relational encoding account [25] that can both surely account for priming under certain conditions. Note, however, that in their review, Kristjánsson and Campana [26] concluded that priming was so ubiquitous in attentional orienting and occurred on so many levels that no single account would probably ever explain it completely.

1.2 Key Considerations for Studies of Attentional Priming

1.2.1 The Duration of Priming Effects

Attentional priming effects are long-lasting. Maljkovic and Nakayama ([1], Exp. 5) showed that up to at least five trials in the past can cause priming, irrespective of what the target and distractor identities were on the intervening trials. Regrettably, researchers often look only at switches versus repeats in studies of priming rather than cumulative effects over several trials. This is unfortunate for two main reasons: firstly, interesting patterns of cumulative repetition are overlooked, and secondly as the effects are additive over adjacent trials with the same target and distractors, looking only at switches or repeats of target identity may not assess priming effects at their maximum strength.

Priming may have both a transient component and a longer-lasting one [27, 28]. Kruijne et al. [29] concluded however that priming involves a single facilitative memory trace that decays over approximately eight trials (see also [30]).

Maljkovic and Nakayama ([1]; see also [31]) introduced a clever way of assessing how long the priming influence from a single trial lasts independently of what occurs on intervening trials that they called *memory kernel analysis*. The method involves categorizing a given trial as the same or different as the one that appeared on the trial that preceded it, two trials preceding it, three trials preceding it, and so on (or formally, i trials preceding the given trial n). For each trial n, target color on trial $n-i$ can be the same or different as the target color on trial n. But over a large number of trials, the numbers of the same and different color trials between trial n and trial $n-i$ will even out. So, to assess the influence of the color of the target i trials in the past, performance on trials where target color (or any other property that is presumed to be primed) on trial $n-i$ was the same as on trial n and when it was different can be compared. Note that the same procedure can be carried out for *future* trials. Future trials should, of course, cause no priming effect and therefore provide a useful sanity check for the analysis and can also be used as an index of the variability in the data (see [1, 31]).

These memory kernels can then be modeled (see [27, 32]). Martini [27] found, for example, that priming effects from repeated target color and repeated position are well described by the summation of two exponential functions, one with a high gain and fast decay and another with low gain but slower decay (consistent with [28]). Kruijne et al. [29] later concluded that a single temporal function could explain priming of features.

Overall these memory kernels that describe the time course of priming effects highlight an important point, since they show that the priming effects can be subtle and long-lasting and, more importantly, how they can contaminate results in various paradigms. To take one example, unequal numbers of two different targets within blocks can cause a contaminating influence that can bias experimental results.

1.2.2 Disentangling Stimulus Priming and Response Priming

Another key consideration is that what is primed must be disentangled from what is reported to avoid the contaminating influence of response repetition effects. To take one example, Maljkovic and Nakayama [1] ensured that their response variable in their task was independent of any feature-repetition effects to avoid this. That is, while observers looked for an odd-one-out color, they responded to a location of a notch on the oddly colored item. Also, in a present versus absent visual search task tested in Kristjánsson et al. [14], the present-absent judgment was independent of the orientation of targets and distractors that either repeated or not between trials.

1.2.3 Effects of Distractor Repetition

It became clear early on that distractor priming was just as important as target priming [1] and that they have their own influence that can be disentangled in experimental design [10]. It is therefore important not to attribute priming effects to either targets or distractors, unless the effects of each can be convincingly unconfounded by design.

1.2.4 Unwanted Influences of Priming

As mentioned before, the influences of priming can be subtle. To take one example, cues are thought to summon attention, whether by cueing a location or a stimulus feature: a red stimulus may be used to alert observers to an upcoming red stimulus inducing a presumably top-down attention effect. But any benefit from the cue may be confounded with priming since the color cues can *prime* attention shifts [18]. For this reason, we recommend that priming effects should be assessed even if they are not of the main interest unless they are comprehensively ruled out by careful experimental design. Note that this actually touches upon the thorny issue of top-down effects versus priming effects [14, 33–35]. One way to avoid this in cueing studies is to use more symbolic cues (e.g., word cues), to isolate effects of top-down attention.

1.3 Using Role Reversals to Assess Probabilistic Representations of Features in the Environment

Priming effects in visual search reflect information about target and distractors that observers have accumulated over previous trials. Recently, Chetverikov, Campana, and Kristjansson [36] suggested that this allows using priming to assess how observers represent probability distributions of visual features.

The idea that the brain encodes the statistics of the environment and uses them to make inferences is well-established [37–42]. However, a given physical probability distribution of features can be represented by the visual system in different ways. While observers need to know the exact shape of distributions to make optimal inferences, approximate inferences can be made with various simplifications. For example, knowing the feature *range* of a class of objects is enough to say whether or not a new object belongs to this class. But many computational models of vision operate on the premise that a given distribution of physical features will be represented with its shape intact. For example, the core assumption of many ideal observers models is that observers accurately represent the generative model of the environment [43]. It is necessary to distinguish these possibilities to understand what kind of information the brain has access to and what it can use.

Revealing the contents of a representation of a physical probability distribution in the brain is not a trivial problem. It is possible to use traditional psychophysics to assess the representation of a single feature value (e.g., one can ask observers to adjust the orientation of the bar so that it matches a previously shown Gabor patch). It is also possible to ask observers to assess the average value of several stimuli, as has been done in “summary statistics” studies (see review in [44]). However, it is impossible to directly inquire about more complex properties, such as skewness, kurtosis, or even variance. As succinctly put by Kuriki for the case of color distributions, “there is essentially no direct approach to studying the color appearance of a multi-colored patch itself” [45, p. 249]. But by utilizing priming effects, we have found a way of circumventing this problem.

Priming effects in visual search occur because observers accumulate information about targets and distractors in order to solve the task more efficiently. In traditional visual search with known targets and distractors, an ideal observer will make inferences based on the ratio of probabilities that a given measurement originates from a target against that it originates from the distractors [46, 47]. In pop-out search, where observers do not know the targets and distractors beforehand (they simply have to find the odd-one-out), observers have to engage in a costly estimate. First, they have to compute for each stimulus, what would be the potential distractor probability distribution based on the other stimuli, and only then estimate whether or not this stimulus could be a target. Learning features of targets and distractors allow observers to avoid this and to analyze each stimulus separately. That is, an observer can look at a given stimulus and decide whether it is a target or a distractor without the need to analyze other stimuli. This shortcut comes at a cost: if observers have accumulated some knowledge of stimuli distributions, then search efficiency becomes dependent on how well these distributions describe current stimuli. In particular, when the target changes and becomes similar to distractors from previous trials, search efficiency should decrease.

Crucially, decreases in search efficiency when targets become similar to preceding distractors should depend on the degree of similarity between them. For an ideal observer, all other things being equal, search efficiency should be inversely proportional to the probability that a given stimulus belongs to a learned distractor distribution, $\mathrm{RT} \propto \frac{1}{p(x|D)}$, where x is the internal measurement of a stimulus' feature, D are the parameters of the distractor distribution, and RT is response time. Introducing targets with different degrees of similarity to the learned distractors would then enable "probing" the representation at different points. Response times in visual search should then distinguish between different representations of distractor distributions, essentially providing a continuous estimate of an internal probability density function describing the physical distribution of distractors (Fig. 1).

This idea was tested by Chetverikov et al. [36]. They found that the curve describing response times as a function of target orientation (centered on the mean of previous distractors) did indeed follow the shape of a previously presented distractor probability distribution when observers searched for an odd-one-out line target among differently oriented lines. This was observed both for distributions with the same range or standard deviation but different shapes (e.g., uniform vs. Gaussian) and for differently skewed but otherwise identical distributions. Chetverikov and colleagues later described two important limitations of this learning process. First, more complex distributions, such as bimodal ones, require more trials to be learned, while simpler ones, such as Gaussian ones, are already represented relatively accurately after one or two trials

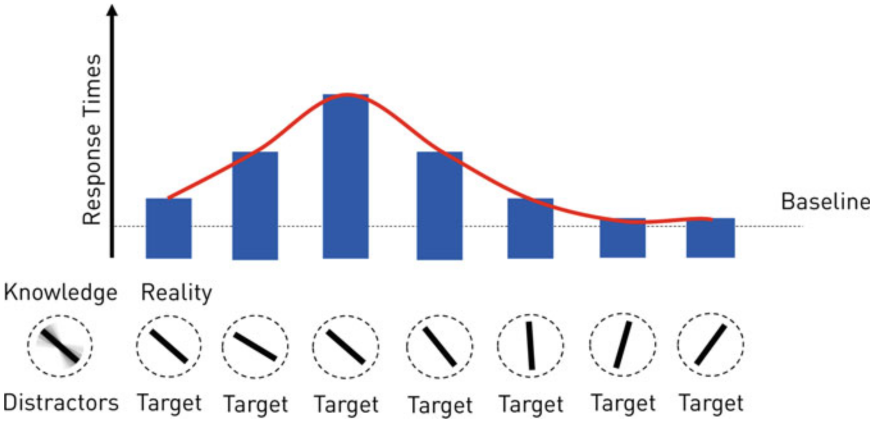

Fig. 1 Hypothetical responses to targets as a function of their similarity to a previously learned distractor distribution. After several visual search trials (here, in the orientation domain), observers obtain knowledge of distractors corresponding to the physical probability distribution of their features (bottom left). It is then possible to "probe" this knowledge by presenting different search targets (bottom central). The response times would be proportional to the degree of similarity between test target and the expected distractors (blue bars). By presenting many different targets, it is then possible to obtain a continuous estimate of a probability density function, corresponding to a representation of distractors (red line)

[48]. Second, in order to learn properties of a distribution, observers need to see a certain minimum number of individual stimuli on each trial [49]. Even mean and variance do not seem to be encoded when only eight lines are presented. Note that this is in sharp contrast to findings on explicit estimates of averages, where observers are able to judge the mean with similar precision regardless of the number of stimuli presented, as long as the overall range stays the same [50]. Finally, this implicit feature distribution learning is not limited to orientation, as similar effects were observed for colored isoluminant diamonds with different hues [51].

2 "Hands-on" Step-by-Step Walk-Through of an Application of the Feature Distribution Learning (FDL) Methodology

The methodological paradigm to assess observers' feature distribution representations has now been used in several different studies [36, 48, 49, 51–53]. Each of these studies probed different aspects of feature distribution learning. While the particular methodological details differ somewhat, they nevertheless all share the same core principle. This involves presenting subjects with a series of odd-one-out visual search trials, where the feature values of the distractors are drawn from a certain type of distribution whose shape and summary statistics stay the same throughout the learning trials. Observers are then presented with a test trial, in which the feature values of the target and the distractor distributions are swapped. This role reversal increases observer's visual search times (as seen previously in a number of studies; [7, 8, 10], and as we have discussed in the section above). When the visual search times obtained from the test trials are plotted as a function of the degree of role reversal, it reveals observer's internal representations of the distractor distribution that was used during the learning trials.

This tutorial will provide detailed information about the three main parts involved in the FDL methodology, which are the design of the visual search task, the structure of the learning trials, and the structure of the test trials. So far, the most studied features with this method are orientation and color. While this tutorial will focus on orientation, the same principles can easily be applied to color (or any other feature space). However, if there are any feature-specific requirements for using color as the main feature, then these are also noted. The next section will focus on the details of the visual search display, while the following ones will focus on how the feature values of the target and distractors are determined.

2.1 Visual Search Task

The main task is an odd-one-out visual search where observers try to find the item whose feature value differs from the rest of the items in the display. When this feature is orientation, the search array includes 36 white lines displayed in a six by six grid on a gray background (Fig. 2b). This method has been tested with smaller set sizes (e.g., 8, 14, 24 lines), but distribution shape learning has only been observed with a larger set size (36 lines, [74]; see discussion above).

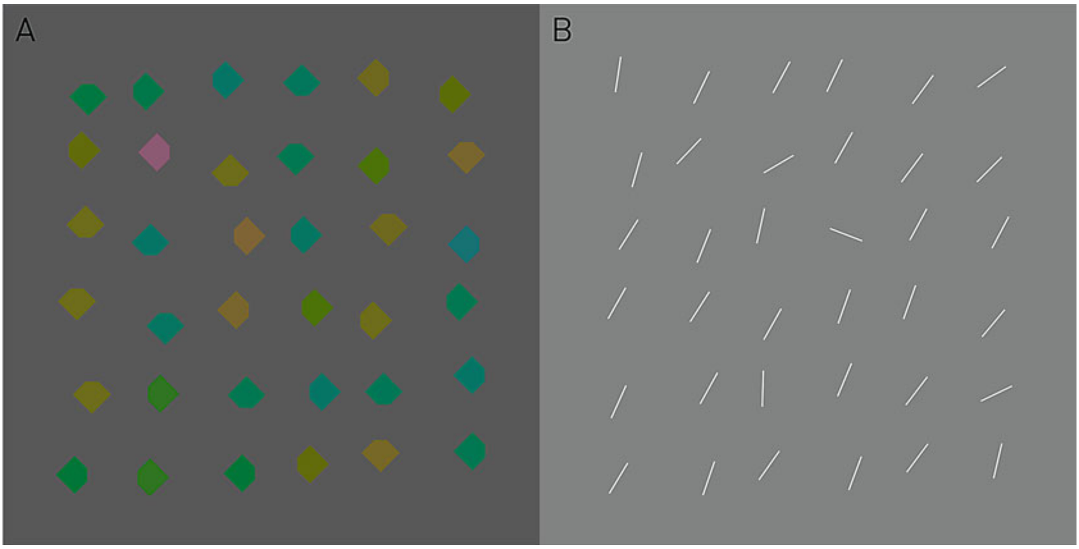

Fig. 2 (**a**) An example visual search display for assessment of the learning of color distributions. Participants search for the oddly colored diamond and report the location of the cutoff on that diamond. (**b**) An example visual search display for testing learning of orientation distributions. Participants search for the oddly oriented line and report whether that line is in the upper or lower half of the display

The size of each line in the search array is set to approximately 1°. The search array subtends approximately 15° × 15° and is positioned at the center of the screen. A random jitter (in the range of ±0.5°) is added to both the vertical and horizontal coordinates of each line. This is done to decrease the precision of the orientation estimate from each individual line, which presumably decreases the viability of serial processing. The position of the target (i.e., the oddly oriented bar) in the search array is randomized.

Participants indicate the location of the target by pressing the "up" button if the target is found in the upper three rows (upper half) and the "down" button if the target is found in the bottom three rows (bottom half) of the search array. The rationale for this choice of response is that the observers respond to the location but not the feature of the target. In other studies, observers responded to a quadrant in which a target was presented [52] or to a shape of a target in a study on colors (see an example below, and further details in [51]). If the participant responds incorrectly, then a feedback display is presented for 1 s, which includes the word "ERROR" in red at the center of the screen. If the response is correct, the search array for the next trial appears immediately. The rationale for not providing feedback after correct responses is to avoid interrupting the between-trial continuity with a feedback display. Feedback screens for incorrect responses slow down the experiment, which, in turn, functions as a motivation for the participant to respond correctly.

However, the main motivating factor for the participant is the score calculated based on accuracy and reaction time. Participants are encouraged to respond as fast as possible to increase their scores. For each trial the score is calculated as follows (where RT is the response time in seconds):

$$\text{For correct answers}: \text{Score} = 10 + (1 - \text{RT}) \times 10$$

$$\text{For incorrect answers}: \text{Score} = -|10 + (1 - \text{RT}) \times 10| - 10$$

On each trial, the score from the previous trial can be shown in one of the corners of the screen. Positive scores are displayed in green and negative in red. When a break is reached during the experiment, the participants' current total score is shown, along with information about what percentage of the experiment has been completed. The only function of keeping score in this experiment is to motivate participants. The particular choice of the score formula is arbitrary, but as is evident from the equation, it is positive for accurate responses faster than 1 s.

When participants perform this odd-one-out visual search task for the first time, their reaction times generally turn out to be too long (>2 s). Therefore, training sessions are needed for naïve participants to get used to the task. The duration of the training might vary, but typically a hundred training trials or more are necessary. The goal of the training is to reduce the average response times and increase accuracy (on a version of the paradigm, such as the one used by [36], for the well-trained observers, the average RT are below 1 s, and accuracy is above 85% correct). The rationale for excluding observers with average response times above a certain threshold is to exclude those who engage in serial processing of the display. As discussed above, that might diminish the learning.

The same principles indicated above can also be applied when color is used as the main feature in visual search. An example display for color search can be seen in Fig. 2a. Instead of lines, participants see 36 diamonds each with a different hue. Each diamond contains a cutoff in any of their four corners. As explained previously, it is important to prevent response repetition from interfering with perceptual priming [1]. Therefore, participants are asked to find the diamond with a hue unlike all the other and report the location of the cutoff (i.e., up, down, left, right).

2.2 Learning Trials

Feature distribution learning experiments consist of blocks. Each block includes streaks of learning trials that are each followed by one or two test trials. Since we are interested in observers' ability to learn the shape of a distribution of feature values, the key aspect of the learning trials is that the shape of the distribution, from where the distractor orientation is drawn, is constant throughout the learning streak.

The length of the learning streaks can be from 1–2 trials up to 10–11 trials depending on the complexity of the distribution used for distractor orientations. For simpler distributions (e.g., Gaussian, uniform), it has been shown that even one to two trials can suffice to uncover learning of distribution shapes. However, when more complex distributions are tested (e.g., bimodal), longer learning streak (7–10 trials) seems to be needed [48]. Even though with simpler distributions one or two trials might be enough to see the learning effects, we recommend keeping the learning streak length to at least three to four trials in order to reduce carry-over effects from preceding streaks. Generally, the length of a learning streak randomly varies during an experiment within a very brief range (e.g., from five to seven trials) in order to break the regularity of the learning and test trials so that participants do not build any expectations about when a block starts and ends. From extensive querying of observers performing this task, they never report having any knowledge of the nature of the sequential trial structure.

The mean of the distractor distribution is randomly determined between −90° and +90° for each learning streak and kept constant within that streak (since this is the distribution that observers are supposed to learn). The target orientation is randomly determined for each trial within a learning streak but is always at least 60° away from the mean of the distractor distribution. This is to ensure that the target is sufficiently dissimilar from distractors, keeping the task relatively easy for observers [54]. In cases where the distractor distribution has high variance and/or has long tails (e.g., Gaussian), some of the distractor lines could turn out to have a similar orientation to the target line. This, in turn, would make the visual search task impossible to carry out. To avoid this, the range of the distractor distribution can be restricted. For example, a Gaussian distractor distribution can be truncated such that any outlier orientation outside of the two-standard deviation range can be removed and then resampled accordingly. The same principles mentioned above for orientation can be applied to color space as well.

2.3 Test Trials

On test trials, the feature values that have been assigned to the target and the distractors during the learning trials are switched in order to reveal observer's internal expectation of distractor orientations (or colors). This, in turn, exposes observer's internal representation of the distractor distribution used in the learning trials. While this is achieved with one test trial, the number of test trials can randomly vary between one and two during an experiment, in order to prevent observers from building expectations about when a block begins and ends. However, only the first test trial is usually included in the data analysis as the effects of the learning streak are expected to dissipate quickly.

The most important consideration for test trials is the selection of the target orientation. The main variable that determines the extent of role-reversal effects is the distance between the target orientation on the test trial and the mean orientation of the previous distractors that are used on the learning trials. This variable, which we refer to as CT-PD ("current target-previous distractor" distance), is essentially the main factor that is manipulated in FDL studies. Plotting reaction times from the test trials as a function of CT-PD reveals observers' internal representation of the previous distractor distribution (Fig. 3). In order to reveal this representation in the range of all possible orientation values, CT-PD values have to uniformly cover this whole range. In order to do that, the orientation space is divided into bins (e.g., 12 bins from −90° to +90° such that each bin covers a range of 15°). Then, a CT-PD value is randomly chosen for each test trial in such way that at the end of the experiment, the number of CT-PD values chosen from each bin would be equal. Once the CT-PD value is determined for a test trial, the target orientation is selected in a way so that the distance between the current target and the previous distractor mean is equal to this CT-PD value (Fig. 3).

Once the target orientation is determined, the mean of the distractor distribution is chosen randomly given that the distance between the target orientation and the mean orientation of the distractor distribution is at least 60°. The distractor distribution on the test trial has to be chosen such that the difficulty of the search on the test trial should be intermediate. If the search is too

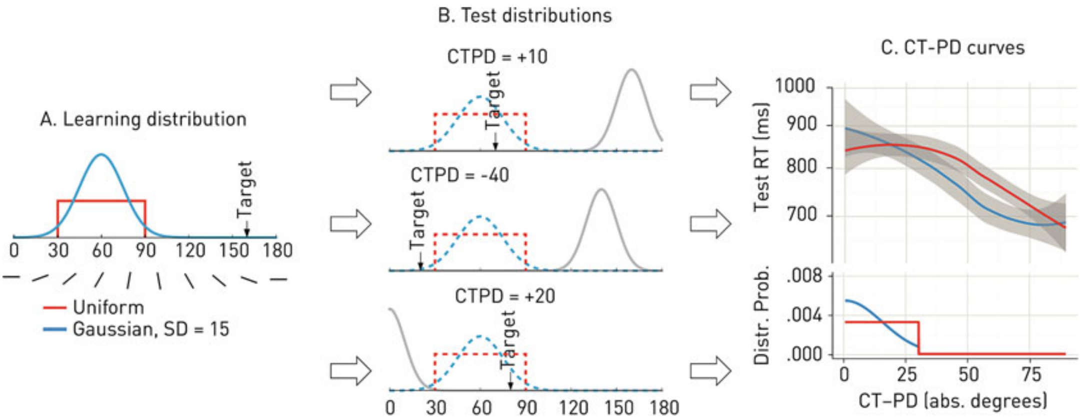

Fig. 3 An example of how CT-PD curves are generated. CT-PD refers to the distance between current target (CT) and previous distractor (PD) distribution mean. After a few trials with distractors drawn randomly from a given learning distribution (**a**), observers are presented with a test trial (**b**). CT-PD distances on test trials are manipulated, and then reaction times obtained from test trials are plotted as a function of these CT-PD distances (**c**, upper). CT-PD curves reveal the observers' internal model of the distractor distribution used on the learning trials and are compared to the physical distribution of stimuli during learning trials (**c**, lower). The CT-PD curves in this figure are based on the results of Experiment 3C from Chetverikov et al. [36] and are replotted from the data available at https:/osf.io/3wcth/

easy or too difficult, floor and ceiling effects might override the role-reversal effects. Previous studies have shown that using a Gaussian distribution with a standard deviation of 10° for the distractor distribution on the test trial provides good testing conditions, but note that this can differ strongly by the aim and characteristics of individual experiments. This Gaussian distribution is also truncated so that no orientation appears as a distractor outside of the two standard deviation ranges.

The number of learning and test streaks used in the experiment is another important factor in this methodology. When the goal is to test whether observers learn the shape of a feature distribution, the experiment should include enough test trials (hence, enough CT-PD values) to fully reveal observers' inner representation of that distribution. We suggest having at least ~350 CT-PD values per observer distributed evenly over the possible feature values. The same principles described here for orientation also apply to color feature space by using just-noticeable differences (JND) as the basic unit (see details on using JND to create a color space for search in [51]).

3 Data Analysis

The goal of the FDL as a method is to infer the characteristics of the ensemble representations based on the search efficiency on test trials. One of the important questions is to analyze whether the visual system represents features as probability distributions rather than only representing the mean and variance (as normally described in summary statistics account of ensemble perception).

In order to reveal observer's internal representation, search times of the first test trial are analyzed. As previously mentioned, the main manipulation factor is the distance between the target on the test trial and the previous distractor distribution mean. Search times of the first test trial are plotted against the distance between the current target and the previous distractor mean (CT-PD) as shown in Fig. 3c. For symmetrical distractor distributions, plotting the absolute distance is sufficient. Figure 3c plots the RT–CT-PD functions using a local regression fit (upper graph) and the probability density function of the distractor distribution (lower graph).

Since search times are analyzed, statistical analyses should be done on log-transformed data due to the skewed search time distributions of the raw data [55–57] and maybe remove outliers due to very slow responses. In addition, only the correct trials are analyzed. One should also make sure that there are enough correct trials for each observer so that the analyses satisfy the usual criterion of more than 85% correct responses.

If role reversals affect search times, then search times of targets within the previous distractor distribution should be slower than search times for targets outside the previous distractor distribution. And, if the actual probabilities of distractors and therefore the distractor distribution shape are represented, the RT function should resemble the shape of the previous distractor distribution.

How can we judge how well the observed RT curve corresponds to the distribution shape? There are several ways for quantitative analysis of RT patterns, including segmented regression and model-fitting. For simple distributions, such as Gaussian or uniform, a useful tool for evaluating these RT functions is to use segmented regression [58, 59]. Following a Gaussian distractor distribution, a monotonically decreasing RT curve is expected, but following, for example, a uniform distractor distribution an RT curve that consists of a flat segment within the distribution range and a steep decrease and faster RT's outside the distribution range is expected, since the probabilities of all feature values within the distribution range are equal (see Fig. 3c). A segmented regression analysis involves searching for significant changes in RT at some particular CT-PD distance. Previous data has shown [36, 48, 49, 51, 53] that the representation of a uniform distribution results in significant breakpoints around the "edge" of the uniform distribution. Search times suddenly change and become faster as the edge of the distribution is reached. The slopes of the individual regression segments are used to support this pattern. Following a uniform distractor distribution, a slope around zero before the breakpoint and a negative slope after the breakpoint resemble the two parts of a uniform distribution. Following a Gaussian distribution, a negative slope of the single segment resembles the monotonic decrease in search time as the distance between the target and the previous distractor mean increases. Statistical tests like the Davies' test [60] compare a two-line model with a single line model that has no breakpoint and provides information about whether the two slopes of the different segments are significantly different. Further analyses could be done on individual subjects' data by comparing the average regression slopes before and after the breakpoint determined on a group basis or based on a priori assumptions (such as the range of the distribution).

A second method of analyzing CT-PD curves is to compare the observed data with pre-defined models that correspond to different distribution shapes. Data can be tested against these pre-defined models of different distribution shapes, and the quality of the different fits can be assessed with the Bayesian Information Criterion. Model fits can be done across subjects or for each subject individually. The model fits seem to be in agreement with segmented regression data in previous studies [51], and they provide a more principled way of testing the hypotheses about encoding of distribution shape. On the other hand, they might lack sensitivity given that the perceptual space might be different from the physical feature space.

We have previously used the following set of models to distinguish between Gaussian and uniform distractor distributions:

1. Half-Gaussian model with a SD $= \sigma$:

$$\mathrm{RT} = c_0 + 2a \times e^{-\frac{\mathrm{CTPD}^2}{2\times\sigma^2}},$$

where a defines the height of the peak, c_0 corresponds to the RT outside the distribution range, and σ corresponds to the standard deviation of the feature distribution that has been used in the experiment (note that the model is half-Gaussian because the Gaussian distribution is symmetric and can be analyzed using the absolute orientation values).

2. Uniform model with a range of $2 \times \sigma$:

$$\mathrm{RT} = \begin{cases} c_0, & \mathrm{CTPD} \le 2 \times \sigma \\ c_1, & \mathrm{CTPD} > 2 \times \sigma \end{cases}$$

where c_0, c_1 determine the RT inside and outside the distribution range, respectively.

3. "Uniform with decrease" model:

$$\mathrm{RT} = \begin{cases} c_0, & \mathrm{CTPD} \le 2 \times \sigma \\ c_0 + b \times \mathrm{CTPD}, & \mathrm{CTPD} > 2 \times \sigma \end{cases}$$

where c_0 determines the RT within the distribution range and b is the decrease outside the distribution range.

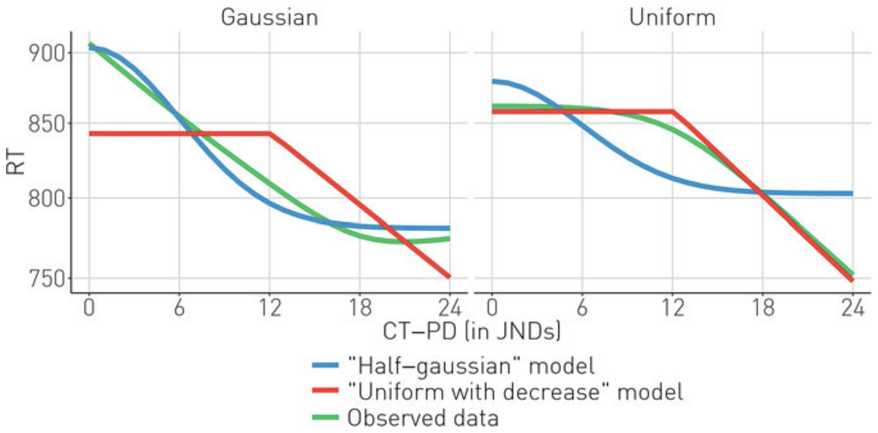

Fig. 4 Observed data and modeling fits for maximum likelihood estimation. Observed data (green) and the best fits of half-Gaussian (blue) and "uniform with decrease" (red) models are based on the results from Chetverikov et al. [51] and are replotted from the data available at https:/osf.io/t2856/

4. Linear model:

$$\mathrm{RT} = c_0 + b \times \mathrm{CTPD}.$$

Each model includes a Gaussian-distributed error term. Different models can be fitted to the data and the best-fitting parameters are obtained using Maximum Likelihood Estimation.

Previous results have shown that search times following a uniform distractor distribution are best fit by a "uniform with decrease" model and following a Gaussian distractor distribution were best fit with either a half-Gaussian or a linear model [51]. Figure 4 shows two of the best-fitting models and the observer data for two different distribution shapes. Together, the segmented regression and the model fitting can provide clear evidence of feature distribution learning.

Note that the goals of data analysis depend on the aims of the study. While here we concentrated on testing distribution shape learning, that is by no means the only possibility. For example, Chetverikov et al. [52] tested whether observers can encode two of the modes of a mixture distribution simultaneously following a single learning block. It is also possible to test, for example, how observers encode the mean or variance of a distribution by analyzing locations of the peaks on a CT-PD curve and its width. Different analytic approaches might be suitable for different research questions.

4 Potential Benefits of FDL

Chetverikov et al. [49] measured whether evidence of FDL leads to higher search efficiency. In that sense it might have a similar effect as contextual cueing. They found some evidence of this, indicating that participants could use the learned distributions, guiding attention across the visual scene. We should also note some preliminary evidence that shows that the distribution learning can be specific to certain portions of the visual scene [61]. But in some sense, the proof is in the pudding, since this paradigm shows how search is affected by the role reversals and that this effect follows the shape of the curves. In other words, the distribution learning improves the efficiency of the search both in terms of speed and accuracy. The paradigm therefore reveals ways in which our visual system makes use of regularities in the visual environment to aid navigation and object recognition.

A notable aspect of feature distribution learning is that it does not require observers to report the feature of interest in any way. This stands in direct contrast to summary statistics studies that require observers to explicitly report or compare some property of the distribution [62–67] which makes FDL more similar to methods relying on other versions of priming or adaptation [68–70]. Importantly, the absence of explicit judgment removes a potential bottleneck in processing of the feature distributions. When making explicit judgments, observers first have to transform the information about the feature likelihood into a single value, i.e., make a “readout” from the information about feature probabilities (akin to how an explicit confidence judgment might involve a readout from probabilistic information on uncertainty [71]). It is also possible that observers do not have explicit access to information about feature probabilities at all and rely on some heuristics when asked to estimate summary statistics [72, 73]. Furthermore, in contrast to other similar methods, FDL allows for the mapping of the feature distribution representations at different points of feature space. In essence, FDL can be considered a behavioral alternative to currently available neurophysiological decoding methods, enabling understanding of the way information is represented in the brain.

5 Summary

Priming effects from repeated target and distractor features are very strong. Their influences last for a long time and they are very hard to willingly overcome [1]. When target repetition effects and distractor repetition effects are combined, role-reversal effects occur that are very large in the context of visual psychophysics and research on visual attention. This makes these effects highly useful since their statistical power enables various applications, including the assessment of more subtle effects, such as how distractor distributions are encoded, just as we described above. But we also emphasize that the priming effects can contaminate experimental results if their influence is not taken into account.

Here we have also shown how priming effects (in particular role reversals) have been used to assess our representations of features in the environment through the implicit assessment of feature distribution learning (FDL). The key insights that our methods have provided involve that our representations are more sophisticated and include far more detail than had previously been assumed in the literature. We believe that this methodology will be of use in assessing many aspects of the function of visual attention, visual working memory, and statistical representations of the environment.

The exact mechanisms of feature distribution learning remain to be studied. However, some characteristics of this process are already known (*see also* [74]). First, as already noted above, there are lower limits on both the number of trials and the number of stimuli necessary for the effect to appear. This suggests that the information about distractors is not accumulated purely locally (otherwise a lower number of stimuli could be compensated for by higher number of trials, which was not the case) and also that this effect involves the accumulation of information over time rather than being a passive aftereffect of stimulus presentation (in the latter case, the bimodal distribution should be represented as unimodal when stimuli from different modes are presented at the same location over several trials). Second, observers can gather information about several different subsets of stimuli on a single trial [52]. This again highlighted the point that information about features is gathered in parallel across the visual field. Third, the generalization from orientation to color suggests that the mechanisms of FDL are dimension-independent. However, whether observers can gather information about the distributions of more complex features (e.g., motion direction) or noncircular ones (e.g., lightness) remains to be studied. There are also a lot of unknowns. The question whether information about feature distribution shape is represented in the cortex or simply readout from a population of differently tuned neurons cannot be answered without neurophysiological studies. A related question is how do observers update the information about the distribution from trial to trial? Do they update the weights of different feature values, or do they update the parameters of the distributions? Also, it is not yet known whether and how the information about different features can be combined. It seems natural to assume that it should be possible, and this combined information can potentially allow for probabilistic object representations. For example, an apple could be a combination of representations of probability distributions for colors and shapes. There is certainly a lot of space for exploring the mechanisms of this unique phenomenon. The method has the additional asset of being an implicit measure of representations of the stimuli in our environment. Statistical representations (such as in the summary statistics literature) are typically assessed with explicit methods, and recently we have provided preliminary evidence that implicit methods may uncover representations of moments of statistical distributions of stimulus ensembles that are masked by explicit methods [75].

Finally, we would like to emphasize that this method demonstrates that observers can learn feature probability distributions in the environment as opposed to a more traditional emphasis on discrete feature values originating in the feature-integration theory of attention [76]. Furthermore, the method shows that such probability distributions are used to guide observer's attention toward (or away from) targets, hindering or facilitating visual search. Recent studies also demonstrate the probabilistic nature of attentional and working memory templates for discrete features [77–81]. Thus, it is likely, in our opinion, that such probabilistic language is more naturally suited to describe representations in the human brain. FDL, in turn, is naturally suited to study representations as it provides an assumption-free continuous description of them. We hope that the addition of this method to the repertoire of commonly used techniques for studying visual attention, memory, or learning processes will bring exciting new discoveries.

Acknowledgments

SHR, ODT, and AK were supported by grant IRF #173947-052 from the Icelandic Research Fund and by a grant from the Research Fund of the University of Iceland. AC was supported by a Radboud Excellence Fellowship.

References

1. Maljkovic V, Nakayama K (1994) Priming of pop-out: I. Role of features. Mem Cognit 22:657–672
2. Bravo MJ, Nakayama K (1992) The role of attention in different visual-search tasks. Percept Psychophys 51:465–472
3. Pascucci D, Mastropasqua T, Turatto M (2012) Permeability of priming of pop out to expectations. J Vis 12:21
4. Shurygina O, Kristjansson Á, Tudge L et al (2019) Expectations and perceptual priming in a visual search task: evidence from eye movements and behavior. J Exp Psychol Hum Percept Perform 45(4):489–499. https://doi.org/10.1037/xhp0000618
5. Sigurdardottir HM, Kristjánsson Á, Driver J (2008) Repetition streaks increase perceptual sensitivity in visual search of brief displays. Vis Cogn 16(5):643–658. https://doi.org/10.1080/13506280701218364
6. Kristjánsson Á, Ásgeirsson ÁG (2019) Attentional priming: recent insights and current controversies. Curr Opin Psychol 29:71–75
7. Wang D, Kristjánsson Á, Nakayama K (2005) Efficient visual search without top-down or bottom-up guidance. Percept Psychophys 67:239–253
8. Lamy DF, Antebi C, Aviani N et al (2008) Priming of Pop-out provides reliable measures of target activation and distractor inhibition in selective attention. Vision Res 48:30–41
9. Saevarsson S, Jóelsdóttir S, Hjaltason H et al (2008) Repetition of distractor sets improves visual search performance in hemispatial neglect. Neuropsychologia 46:1161–1169
10. Kristjánsson Á, Driver J (2008) Priming in visual search: separating the effects of target repetition, distractor repetition and role-reversal. Vision Res 48:1217–1232
11. Chetverikov A, Kristjánsson Á (2015) History effects in visual search for monsters: search times, choice biases, and liking. Atten Percept Psychophys 77:402–412
12. Ásgeirsson ÁG, Kristjánsson Á, Bundesen C (2014) Independent priming of location and

color in identification of briefly presented letters. Atten Percept Psychophys 76:40–48

13. Hillstrom AP (2000) Repetition effects in visual search. Percept Psychophys 62:800–817
14. Kristjánsson Á, Wang D, Nakayama K (2002) The role of priming in conjunctive visual search. Cognition 85:37–52
15. Kristjánsson Á, Ingvarsdóttir Á, Teitsdóttir UD (2008) Object- and feature-based priming in visual search. Psychon Bull Rev 15:378–384
16. Huang L, Holcombe AO, Pashler H (2004) Repetition priming in visual search: episodic retrieval, not feature priming. Mem Cognit 32:12–20
17. Kristjánsson Á (2006) Simultaneous priming along multiple feature dimensions in a visual search task. Vision Res 46:2554–2570
18. Belopolsky AV, Schreij D, Theeuwes J (2010) What is top-down about contingent capture? Atten Percept Psychophys 72:326–341
19. Theeuwes J, van der BE (2011) On the limits of top-down control of visual selection. Atten Percept Psychophys 73:2092–2103
20. Folk CL, Remington RW, Johnston JC (1992) Involuntary covert orienting is contingent on attentional control settings. J Exp Psychol Hum Percept Perform 18:1030–1044
21. Carlisle NB, Kristjánsson Á (2017) How visual working memory contents influence priming of visual attention. Psychol Res 82:833–839
22. Kristjánsson Á, Saevarsson S, Driver J (2013) The boundary conditions of priming of visual search: from passive viewing through task-relevant working memory load. Psychon Bull Rev 20:514–521
23. Muller HJ, Reimann B, Krummenacher J (2003) Visual search for singleton feature targets across dimensions: stimulus- and expectancy-driven effects in dimensional weighting. J Exp Psychol Hum Percept Perform 29:1021–1035
24. Found A, Müller HJ (1996) Searching for unknown feature targets on more than one dimension: investigating a “dimension-weighting” account. Percept Psychophys 58:88–101
25. Becker SI (2010) The role of target-distractor relationships in guiding attention and the eyes in visual search. J Exp Psychol Gen 139:247–265
26. Kristjánsson Á, Campana G (2010) Where perception meets memory: a review of repetition priming in visual search tasks. Atten Percept Psychophys 72:5–18
27. Martini P (2010) System identification in Priming of Pop-Out. Vision Res 50:2110–2115
28. Brascamp JW, Pels E, Kristjánsson Á (2011) Priming of pop-out on multiple time scales during visual search. Vision Res 51:1972–1978
29. Kruijne W, Brascamp JW, Kristjánsson Á et al (2015) Can a single short-term mechanism account for priming of pop-out? Vision Res 115:17–22
30. Kruijne W, Meeter M (2015) The long and the short of priming in visual search. Atten Percept Psychophys 77:1558–1573
31. McPeek RM, Maljkovic V, Nakayama K (1999) Saccades require focal attention and are facilitated by a short-term memory system. Vision Res 39:1555–1566
32. Maljkovic V, Martini P (2005) Implicit short-term memory and event frequency effects in visual search. Vis Res 45(21):2831–2846. https://doi.org/10.1016/j.visres.2005.05.019
33. Theeuwes J, Reimann B, Mortier K (2006) Visual search for featural singletons: no top-down modulation, only bottom-up priming. Vis Cogn 14:466–489
34. Folk CL, Remington RW (2008) Bottom-up priming of top-down attentional control settings. Vis Cogn 16:215–231
35. Wolfe JM, Butcher SJ, Lee C et al (2003) Changing your mind: on the contributions of top-down and bottom-up guidance in visual search for feature singletons. J Exp Psychol Hum Percept Perform 29:483–502
36. Chetverikov A, Campana G, Kristjánsson Á (2016) Building ensemble representations: How the shape of preceding distractor distributions affects visual search. Cognition 153:196–210
37. Girshick AR, Landy MS, Simoncelli EP (2011) Cardinal rules: visual orientation perception reflects knowledge of environmental statistics. Nat Neurosci 14:926–932
38. Rao RP, Olshausen BA, Lewicki MS (2002) Probabilistic models of the brain: perception and neural function. MIT Press, Cambridge, MA
39. Pouget A, Beck JM, Ma WJ et al (2013) Probabilistic brains: knowns and unknowns. Nat Neurosci 16:1170–1178
40. Ma WJ (2012) Organizing probabilistic models of perception. Trends Cogn Sci 16:511–518
41. Fiser J, Berkes P, Orbán G et al (2010) Statistically optimal perception and learning: from behavior to neural representations. Trends Cogn Sci 14(3):119–130
42. Feldman J (2014) Probabilistic models of perceptual features. In: Wagemans J (ed) Oxford handbook of perceptual organization. Oxford University Press, Oxford, pp 933–947

43. Vincent BT (2015) A tutorial on Bayesian models of perception. J Math Psychol 66:103–114
44. Whitney D, Yamanashi-Leib A (2018) Ensemble perception. Annu Rev Psychol 69:105–129
45. Kuriki I (2004) Testing the possibility of average-color perception from multi-colored patterns. Opt Rev 11(4):249–257. https://doi.org/10.1007/s10043-004-0249-2
46. Ma WJ, Navalpakkam V, Beck JM et al (2011) Behavior and neural basis of near-optimal visual search. Nat Neurosci 14:783–790
47. Ma WJ, Shen S, Dziugaite G et al (2015) Requiem for the max rule. Vision Res 116:179–193
48. Chetverikov A, Campana G, Kristjánsson Á (2017) Rapid learning of visual ensembles. J Vis 17:1–15
49. Chetverikov A, Campana G, Kristjánsson Á (2017) Set size manipulations reveal the boundary conditions of distractor distribution learning. Vision Res 140:144–156
50. Utochkin IS, Tiurina NA (2014) Parallel averaging of size is possible but range-limited: a reply to Marchant, Simons, and De Fockert. Acta Psychol (Amst) 146:7–18
51. Chetverikov A, Campana G, Kristjánsson Á (2017) Representing Color Ensembles. Psychol Sci 28:1–8
52. Chetverikov A, Campana G, Kristjánsson Á (2018) Probabilistic rejection templates in visual working memory. Submitted for Review. doi: https://doi.org/10.31234/osf.io/vrbgh. Preprint available at https://psyarxiv.com/vrbgh/
53. Hansmann-Roth S, Chetverikov A, Kristjánsson Á (2019) Representing color and orientation ensembles: can observers learn multiple feature distributions? Submitted for Review
54. Duncan J, Humphreys GW (1989) Visual search and stimulus similarity. Psychol Rev 96:433–458
55. Palmer EM, Horowitz TS, Torralba A et al (2011) What are the shapes of response time distributions in visual search? J Exp Psychol Hum Percept Perform 37:58–71
56. Kristjánsson Á, Jóhannesson ÓI (2014) How priming in visual search affects response time distributions: analyses with ex-Gaussian fits. Atten Percept Psychophys 76:2199–2211
57. Luce RD (1986) Response times: their role in inferring elementary mental organization. Oxford University Press, New York, NY
58. Muggeo VMR (2003) Estimating regression models with unknown break-points. Stat Med 22:3055–3071
59. Muggeo VMR (2008) Segmented: an R package to fit regression models with broken-line relationships. R News 8:20–25
60. Davies RB (1987) Hypothesis testing when a nuisance parameter is present only under the alternative. Biometrika 74:33–43
61. Chetverikov A, Campana G, Kristjánsson Á (2018) Probabilistic perceptual landscapes. J Vis 18:529
62. Atchley P, Andersen GJ (1995) Discrimination of speed distributions: sensitivity to statistical properties. Vision Res 35:3131–3144
63. Morgan MJ, Chubb C, Solomon JA (2008) A "dipper" function for texture discrimination based on orientation variance. J Vis 8:9–9
64. Im HY, Chong SC (2014) Mean size as a unit of visual working memory. Perception 43:663–676
65. Chong SC, Treisman A (2003) Representation of statistical properties. Vision Res 43:393–404
66. Webster J, Kay P, Webster MA (2014) Perceiving the average hue of color arrays. J Opt Soc Am A Opt Image Sci Vis 31:A283–A292
67. Attarha M, Moore CM (2015) The capacity limitations of orientation summary statistics. Atten Percept Psychophys 77:1116–1131
68. Norman LJ, Heywood CA, Kentridge RW (2015) Direct encoding of orientation variance in the visual system. J Vis 15:1–14
69. Michael E, de Gardelle V, Summerfield C (2014) Priming by the variability of visual information. Proc Natl Acad Sci 111:7873–7878
70. Corbett JE, Melcher D (2014) Stable statistical representations facilitate visual search. J Exp Psychol Hum Percept Perform 40:1915–1925
71. Meyniel F, Sigman M, Mainen ZF (2015) Confidence as Bayesian probability: from neural origins to behavior. Neuron 88:78–92
72. Solomon JA (2010) Visual discrimination of orientation statistics in crowded and uncrowded arrays. J Vis 10:19
73. Lau JS, Brady TF (2018) Ensemble statistics accessed through proxies: range heuristic and dependence on low-level properties in variability discrimination. J Vis 18:3
74. Chetverikov A, Campana G, Kristjánsson Á (2017) Learning features in a complex and changing environment: a distribution-based framework for visual attention and vision in general. Prog Brain Res 236:97–120. https://doi.org/10.1016/bs.pbr.2017.07.001
75. Hansmann-Roth S, Kristjansson Á, Whitney D et al (2018) Explicit and implicit judgments of distribution characteristics: Do they lead to different results? Oral presentation at European Conference on Visual Perception 2018, Trieste, Italy. Abstract available at https://guidebook.com/guide/123359/poi/10443998/

76. Treisman AM, Gelade G (1980) A feature-integration theory of attention. Cogn Psychol 12:97–136
77. Won BY, Geng JJ (2018) Learned suppression for multiple distractors in visual search. J Exp Psychol Hum Percept Perform 44:1128–1141
78. Geng JJ, Witkowski P (2019) Template-to-distractor distinctiveness regulates visual search efficiency. Curr Opin Psychol 29:119–125
79. Hout MC, Goldinger SD (2014) Target templates: the precision of mental representations affects attentional guidance and decision-making in visual search. Atten Percept Psychophys 77:128–149
80. Ma WJ, Husain M, Bays PM (2014) Changing concepts of working memory. Nat Neurosci 17:347–356
81. Bays PM (2015) Spikes not slots: noise in neural populations limits working memory. Trends Cogn Sci 19:431–438