Probabilistic rejection templates in visual working memory

Andrey Chetverikov[a,b,c], Gianluca Campana[d,e], & Árni Kristjansson[b,f]

[a] - Center for Cognitive Neuroimaging, Donders Institute for Brain, Cognition and Behavior, Radboud University; [b] - Laboratory for Visual Perception and Visuomotor Control, Faculty of Psychology, School of Health Sciences, University of Iceland, Reykjavik, Iceland; [c] - Cognitive Research Lab, Russian Academy of National Economy and Public Administration, Moscow, Russia; [d] – Department of General Psychology, University of Padova, Padova, Italy; [e] - Human Inspired Technology Research Centre, University of Padova, Italy; [f] - School of Psychology, National Research University Higher School of economics.

Author Note

Correspondence should be addressed to: Andrey Chetverikov, Donders Institute for Brain, Cognition and Behavior, Center for Cognitive Neuroimaging, Kapittelweg 29, 6525 EN Nijmegen, The Netherlands. Email: a.chetverikov@donders.ru.nl. The preprint of this paper is available at https://dx.doi.org/10.17605/OSF.IO/VRBGH The data from the experiments reported in this paper and scripts for simulations and analyses are available at https://osf.io/rg2h8

Word count: 4740.

Abstract

Our interactions with the visual world are guided by attention and visual working memory. Things that we look for and those we ignore are stored as templates that reflect our goals and the tasks at hand. The nature of such templates has been widely debated. A recent proposal is that these templates can most straightforwardly be thought of as probabilistic representations of task-relevant features. We assessed observers' representations by measuring the slowing of visual search when distractor templates unexpectedly match the target. To provide a strong test of the templates' probabilistic nature, distractor stimuli were heterogeneous, randomly drawn on each trial from a bimodal probability distribution. Using two targets on each trial, we tested whether observers encode the full distribution, only one peak of it, or the average of the two peaks. Search was slower when the two targets corresponded to the two modes of previous distractor distributions than when one target was at one of the modes and another between the modes or outside the distribution range. Furthermore, targets on the modes were reported later than targets between the modes that, in turn, were reported later than targets outside this range. These results show that observers represent both distribution modes using templates based on the full probability distribution rather than single features or simple summary statistics. Our findings indicate that visual working memory templates guiding attention are probabilistic and dynamically adapt to task requirements, reflecting the probabilistic nature of the input.

Keywords*: attentional templates, visual working memory, probabilistic representations, visual ensembles, summary statistics, visual search*.

Probabilistic rejection templates in visual working memory

Our senses are constantly bombarded with an overwhelming amount of information that needs to be filtered by the brain to guide action. This information, however, is not completely chaotic. For example, leaves on a tree usually have similar colors, and colors within a single leaf would be more similar to each other than to another leaf. Probabilistic models of vision (Bejjanki, Beck, Lu, & Pouget, 2011; Feldman, 2014; Girshick, Landy, & Simoncelli, 2011; Kersten, Mamassian, & Yuille, 2004; Ma, 2012; Rao, Olshausen, & Lewicki, 2002) suggest that the brain utilizes existing correlations in the environment and uses them in perception. However, some of the incoming information is not relevant for current behavior, and it is important to reject it while processing other stimuli in more detail. Traditionally, the rejection of irrelevant information within a specific feature dimension (e.g., orientation) is thought to be based on specific feature values (Woodman, Carlisle, & Reinhart, 2013). Here we ask whether such rejection can instead be based on probabilistic templates. If this is the case, then probabilistic inference in the brain does not start with perception, but sooner, when to-be-rejected templates are formed (based on previously encountered stimuli) to optimize the prioritization of what is perceived.

Imagine a radiologist looking for signs of tumor in x-ray scans. Malignant signs can take many forms so the targets to look for are diverse. By many accounts, search in this and other contexts is thought to be guided by templates held in visual working memory (Woodman et al., 2013). These templates reflect what one should look for, but may also reflect what should be ignored (Arita, Carlisle, & Woodman, 2012; Won & Geng, 2018). For example, distractors such as the rib cage on a lung scan are salient but not informative and radiologists can therefore ignore it. It is well known that the information about to-be-ignored stimuli or features is kept in memory, but the way they are represented is still unknown.

There are capacity limits in the amount of information that can be stored in visual working memory templates (Bundesen, 1990; Grubert & Eimer, 2013; Vickery, King, & Jiang, 2005), with some authors even suggesting that only one template containing a single feature value can guide attention at any given time (Oberauer, 2002; Olivers, Peters, Houtkamp, & Roelfsema, 2011; van Moorselaar, Theeuwes, & Olivers, 2014). Alternatively, templates could be conceptualized as probabilistic entities of varying precision (Bays, 2015) rather than matches to exact feature values. While previous studies found some support for this, observers typically reported features of single items (Ma, Husain, & Bays, 2014). However, in the real world such isolated features practically never occur. Furthermore, with a few exceptions (Arita et al., 2012; Won & Geng, 2018), templates for ignored information are rarely studied. Here, we provide strong evidence for the probabilistic template view by testing visual working memory templates for rejection.

Our observers searched for two oddly oriented targets among distractors randomly drawn from a bimodal orientation distribution. To expose observers' templates, after a sequence of *learning* trials with distractors randomly drawn from a bimodal distribution, targets on *test* trials could either correspond to regions of feature space previously used for distractors, fall in between the modes of the bimodal distribution, or have feature values outside the previous distribution range. We assume that observers' templates reflect what has been relevant on recent trials. If templates contain features of distractors to be ignored, which then become targets on test trials, search should be slower than otherwise (Chetverikov, Campana, & Kristjánsson, 2016; Kristjánsson & Driver, 2008; Lamy, Antebi, Aviani, & Carmel, 2008; Maljkovic & Nakayama, 1994; Wang, Kristjánsson, & Nakayama, 2005). Crucially, experiments with varied set size and trial numbers show that learning in this paradigm cannot be explained by the sampling of a few items (Chetverikov, Campana, & Kristjánsson, 2017d, 2017b). It also cannot be explained by simple decision rule learning (e.g., all stimuli that have features in a certain range are distractors), because observers response times, on average, reflect the shape of the distractor distribution rather than just a boundary between a target and distractors (Chetverikov et al., 2016, 2017b; Chetverikov, Campana, & Kristjánsson, 2017c; Chetverikov, Hansmann-Roth, Tanrikulu, & Kristjansson, 2019). However, it is not yet clear whether each single set of learning trials can feed observers' templates with the feature probability distribution of distractors, nor is it clear how accurately the information is stored in the templates.

Under the *strong* probabilistic template hypothesis, templates would include information about both peaks of a bimodal distribution. That is, the template would accurately reflect the information about the full probability distribution. Alternatively, templates might include only a single peak (e.g., the attended one), or might reflect only the summary statistics, such as the averages of the whole distribution (Alvarez, 2011). With a two-target search we were able to test whether observers encode both peaks of a distribution following a single learning sequence. The predictions of these models (see Simulations) are qualitatively different regarding both the order in which targets are reported in a two-target search, and search times. If observers accurately encode a bimodal distribution, on trials with a target on a peak and target between peaks, targets between the peaks (associated with a lower distractor probability) should be reported before targets on peaks (associated with the highest distractor probability, Figure 1A). In contrast, if only one peak is encoded or if the whole distribution is averaged, targets on peaks would be associated with a lower distractor probability and should be reported no later than targets between the peaks (associated with lower distractor probability in this case). Notably, while all three hypotheses postulate that observers can use probabilistic *inference*, only the first one assumes that the distractor probability distribution is encoded accurately, that is, that the observers use relatively accurate probabilistic *templates*.

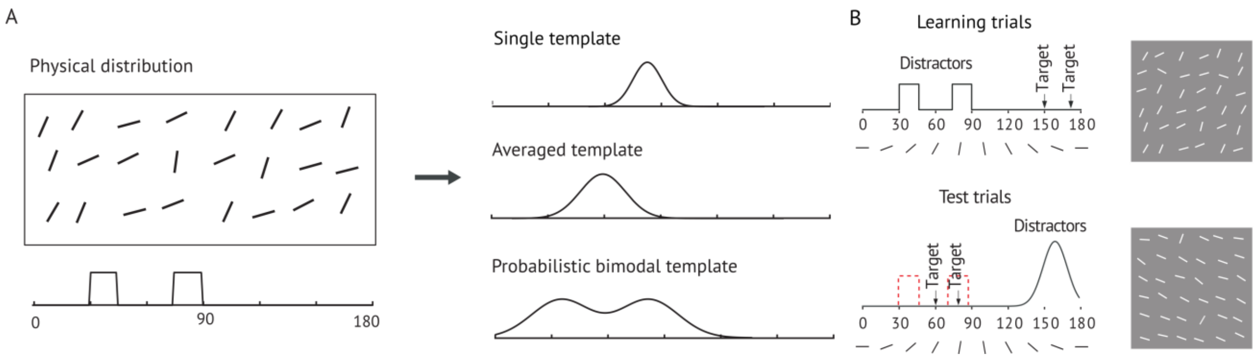

Figure 1. Panel A: The same physical bimodal distribution can be represented in different ways. Panel B: Example learning and test trials with distractor distributions and targets shown on the left.

Experiment

Ethics Statement. The study was approved by the ethics committee of St. Petersburg State University (#75, 21.06.2017). All participants signed a consent form before taking part in the study.

Participants. Fifteen observers (ten female, age M = 25.67) at St. Petersburg State University, Russia, participated voluntarily in a single experimental session lasting approximately 30 min. The data from two observers were excluded because their response times on test trials were too slow (M = 1464 and M = 1871 ms), compared with other observers (M = 1064 ms). Following our previous studies (Chetverikov et al., 2016, 2017b, 2017c, 2017d), the design of this study utilized within-subject comparisons with a relatively small number of trained observers (each observer was trained for at least 100 trials before the main session) performing a large number of trials. The sample size and the trial numbers were similar to those in previous studies using the same paradigm.

Method. We used a task similar to our previous studies (Chetverikov et al., 2016, 2017b). Stimuli were presented on an Acer V193 display (19” with 1280 × 1024 pixel resolution) using PsychoPy 1.84.2 (Peirce, 2007, 2009). Viewing distance was ~ 60 cm. Observers searched for two oddly oriented lines in a 6×6 grid of 36 lines subtending 16°×16° at the centre of a display. The length of each line was 1.41°. Line positions were jittered by randomly adding a value between ±0.5° to both vertical and horizontal coordinates.

Observers were instructed to search for two targets on each trial, with targets being the stimuli that were most different from all the others (“odd-one-out” search (Maljkovic & Nakayama, 1994)). Targets were randomly distributed between the four quadrants of the search display with the constraint that the two targets on a given trial could not appear in the same quadrant. Observers reported the locations of the targets by pressing one of four keys (‘f’, ‘g’, ‘r’, ‘t’ on a standard keyboard) corresponding to the quadrants of the search display. They were informed that two targets would be presented on each trial and were encouraged to respond to each target as soon as they found it and not wait until both targets were found.

Trials were organized in intertwined prime and test ‘streaks’. During prime streaks, distractors were randomly drawn from a bimodal distribution that included two uniform parts with orientations ranging from -30 to -20 and +20 to +30 relative to the overall mean. The distribution mean was the same within streak but chosen randomly between streaks. Target orientations were selected randomly on each trial with the restriction that the distance between target orientation and distractor mean in feature space was 60 degrees at minimum. Prime streak length was set to 6-7 trials (with equal probability) because this streak length is sufficient to learn bimodal distributions with relative accuracy (Chetverikov et al., 2017b).

Within test streaks, distractor orientations were randomly drawn from a truncated Gaussian with SD = 10 deg. and range 20 deg. Test streaks had one or two trials (with equal probability). Different *target* types were used on test trials: targets were either located on a peak of the previous bimodal distribution (“Peak”, at +/-25 deg. relative to the previous distractor mean), between the peaks (“Between”, at 0 deg.) or outside the previous distribution range (“Outside”, at +/- 50 deg.). Four types of test streaks were used: 1) with two targets either on two different peaks (“Peak + Peak”); 2) on a peak and in-between the peaks (“Peak + Between”); 3) on a peak and outside the previous distribution range (“Peak + Outside” – where the “outside” target was always 25 deg. away from the target peak, that is, either the two targets were oriented at +25 and +50 deg. or -25 and -50 relative to the previous distractors’ mean); 4) between the peaks and outside the range (“Outside + Between”). These four *test types* were presented equally often (40 repetitions by participant) in random order. The distractor mean was chosen to be equidistant from both test targets. The second test trial is not analyzed here as the priming effects from the learning streak are not likely to be significant after the first two-target test search. Two-trial test streaks were added for consistency with previous studies and in order to reduce the potential effects of observers’ expectations regarding streak lengths.

Observers participated in one session of approximately 1300 trials. Decision time was not limited but participants were encouraged to respond as quickly and accurately as possible. Feedback based on search time and accuracy on previous trials was shown in the upper-left corner of the screen to motivate participants (see Chetverikov et al., 2016, for details on feedback score calculation). The current trial number and the total number of trials were shown beneath the score. If observers made an error, the word "ERROR" appeared in red letters at display centre for 1 second.

In addition to this two-target search experiment, we also ran a single-target search study (see Supplementary Experiment). The latter was used as a comparison for the single-target search time analyses to ensure that the introduction of a second target and specific conditions of the main experiment did not affect the pattern of results.

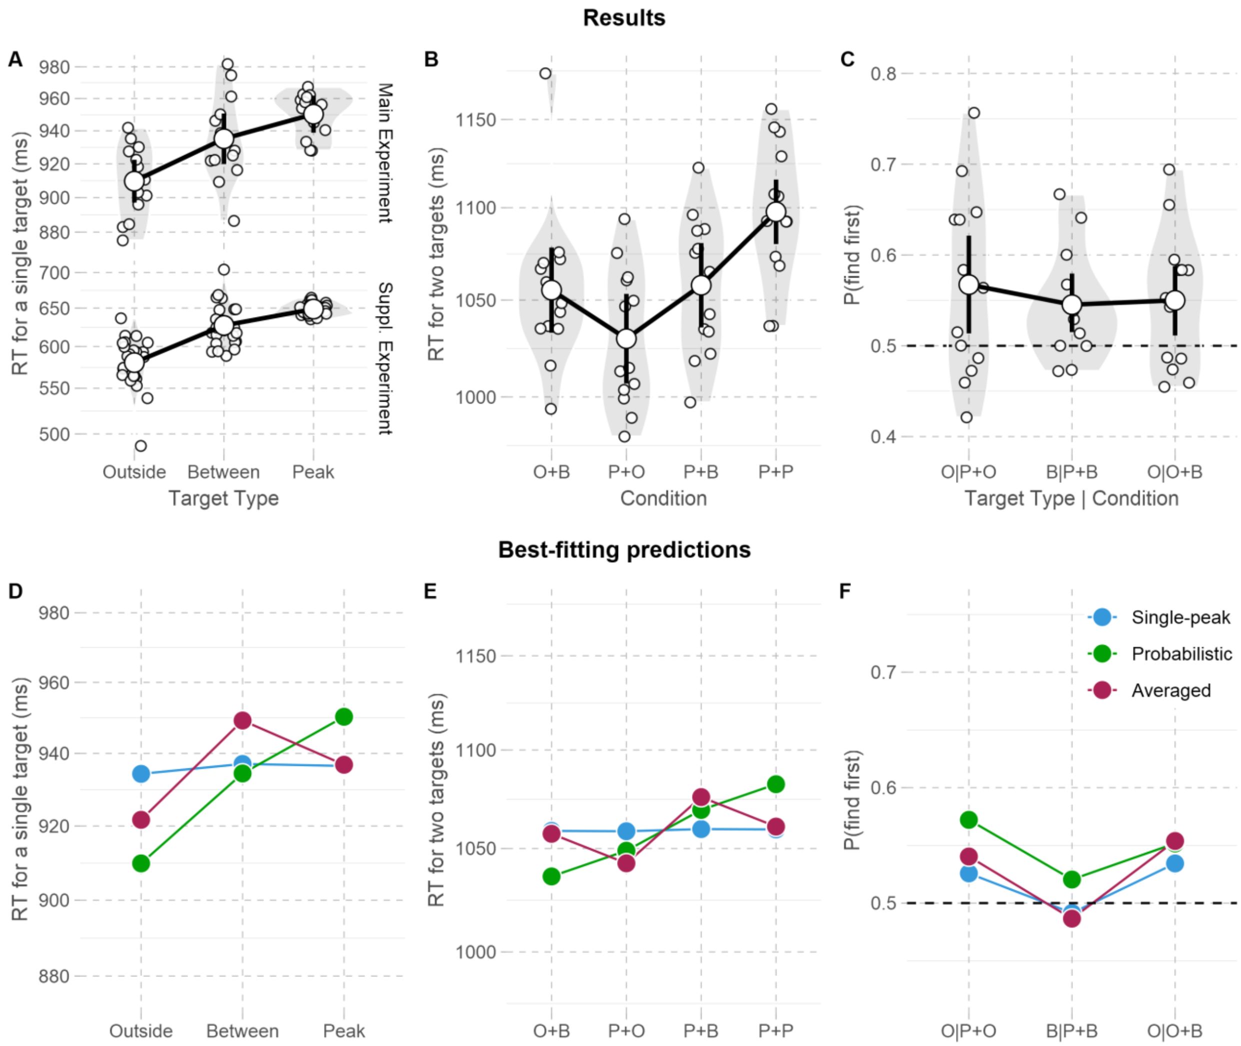

Figure 2. Experimental results and best-fitting predictions of the models (see Simulations). A: Results for different target types from the main Experiment (average search times ignoring the order in which the targets were reported) and the supplementary Experiment where observers searched for only one target on each trial. B: Results for two-target search from the main Experiment. C: Results for the order of target reporting from the main Experiment. D-F: Predictions for single-target search times, search times for two targets, and for the order in which targets would be reported in a two-target search. For A-C large dots show group means, bars show their 95% confidence intervals, smaller dots show individual observers' means, and shaded areas show distributions of individual observers' means. Abbreviations: RT – response times, P - target on a peak, B - target between the peaks, O - target outside the range of previous distractor distribution. The plus sign indicates that two targets of corresponding types are used. For the order of reporting, X|A+B means that target type X was reported first when target types A and B are combined. Note that the slopes of the lines connecting the dots in each plot indicate relative but not absolute difference between conditions because the x-axis is categorical.

Results

Overall performance. On learning trials, observers found both targets in most cases (*M* = 0.72 [0.67, 0.77]), though the share of trials where only one target was reported was high (*M* = 0.27 [0.22, 0.31]; both targets were reported incorrectly on 1% of trials). On test trials, observers reported both targets correctly on *M* = 0.91 [0.89, 0.93] trials (accuracy was comparable to the results of single-target search in the Supplementary Experiment). The delay between the report on the first and the second target was relatively short, but longer on learning than on test trials (*M* = 263 [198, 326] vs. *M* = 176 [130, 233], respectively, $t(12.0) = 4.13$, $p = .001$). Similarly, the first target was reported later on learning than test trials (*M* = 973 [854, 1103] vs. *M* = 826 [753, 904], respectively, $t(12.0) = 5.23$, $p < .001$).

The learning effects were also comparable to those from the single-target search experiment (see supplement). A linear mixed-effects regression with Helmert contrasts (comparing each trial with the average of the following trials) showed that the first trial was slower, ($B = 0.11$, $SE = 0.01$, $t(52.57) = 9.61$, $p < .001$) and less accurate ($B = -0.04$, $SE = 0.02$, $t(13.12) = -2.62$, $p = .021$) than the later trials. The follow-up trials did not differ from one another.

Test trials. Replicating previous results, search times differed depending on target type ($F\ (2, 24) = 8.28$, $p = .003$, $\eta^2_G = .02$, see Figure 2A). Observers search longer for “Peak” targets compared to “Between” targets, which were in turn, found later than “Outside” targets. Crucially, a repeated-measures ANOVA indicated that the time needed to find *both* targets on test trials was affected by the condition ($F(3, 36) = 6.66$, $p = .002$, $\eta^2_G = .02$, Figure 2B). Comparisons between conditions with the same feature difference between the targets showed that performance on "Peak + Peak" trials was slower than on "Outside + Between" trials ($t(12.0) = 3.10$, $p = .009$), while "Peak + Between" trials were not different from "Peak + Outside" ($t(12.0) = -1.68$, $p = .118$) trials. Finally, the "Peak + Peak" condition was also slower than the "Peak + Between" condition ($t(12.0) = 2.58$, $p = .024$).

We then analyzed which type of target was reported first in each condition using a binomial mixed-effects regression. The results showed that targets on peaks were reported after targets between the peaks ($Z = -2.01$, $p = .044$, Figure 2C) or targets outside the preceding distribution range ($Z = -2.43$, $p = .015$), while the latter were reported earlier than targets between the peaks ($Z = 2.08$, $p = .037$).

In sum, search with two targets on the peaks was the most difficult. A comparison of the "Peak + Between" and "Peak + Outside" conditions showed only a numerical difference in total RT. However, in the "Peak + Between" condition, targets on peaks were reported later than targets between the peaks, whereas in the "Outside + Between" condition targets between the peaks were reported later than the "Outside" targets. This shows again that targets on peaks were the most unexpected for observers, followed by targets between the peaks, followed in turn by "Outside" targets that led to the fastest search times.

Simulations. We simulated the predictions from three models (Figure 2E-F; the simulation code is available at https://osf.io/rg2h8). For our main model of interest, the "probabilistic" model, we assumed that the probabilities of different distractors can be represented by two Gaussian templates (for simplicity, we ignore the fact that the stimuli distributions might be more accurately represented by non-Gaussian templates (Chetverikov, Campana, & Kristjánsson, 2017a)) centered on the means of distractor distribution segments. We assumed that observers utilize the knowledge they obtained about distractors and targets optimally. According to a Bayesian ideal observer model in a localization search task, to find a target, observers compare the probability that a given noisy measurement of orientation x at each location L is a target versus the probability that it is a distractor (Ma, Navalpakkam, Beck, van den Berg, & Pouget, 2011; Ma, Shen, Dziugaite, & van den Berg, 2015):

$$p(L|x) \propto p(L)\frac{p(x|T)}{p(x|D)} = p(L)\frac{p(x|s_L)p(s_L|T)ds_L}{p(x|s_L)p(s_L|D)ds_L}$$

where s_L is a true stimulus value at this location, $p(L)$ is the probability that a target is presented at this location, T and D are the parameters of target and distractor distributions, respectively. In our simulations, we assumed that internal representations of target and distractor distributions are independent and response times are inversely proportional to the amount of evidence $p(L|x)$. Given that all locations in our experiment were equiprobable, that is, $p(L)$ is the same for all locations, response times will, on average, be proportional to the probabilities of the test target θ_T under given distractor template parameters:

$$RT \propto p(\theta_T|D)$$

The width of the Gaussian templates was estimated by fitting the model to single-target response time data. To increase the robustness of the estimates, we used an approach similar to bootstrap aggregating (“bagging”), often employed in machine learning (Breiman, 1996). For each model we obtained 500 bootstrapped samples grouped by participant (that is, on each iteration, sampling with replacement was done for each subject and then the samples were combined). We then estimated the template widths for each sample by fitting response times as a linear function of the stimuli probability. For a probabilistic model:

$$RT_{1T} = a + b\left(\frac{1}{2}p(\theta_T|\mu_1,\sigma) + \frac{1}{2}p(\theta_T|\mu_2,\sigma)\right)$$

where $\mu_1 = 25$ and $\mu_2 = -25$, the means of bimodal distractor distribution peaks, and a and b are the scaling parameters necessary to translate the probabilities into response times. The template widths obtained for each sample were then averaged to get the resulting estimates. Estimated template widths were similar for the experiment reported here (18 deg.) and the supplemental experiment (21 deg.).

For the "single-peak" model, we assumed that only one of the two peaks was encoded (with the same approach as with the "probabilistic" model). Given that the peak means are equidistant to the overall distractor mean:

$$RT_{1T} = a + b\big(p(\theta_T|\mu_1, \sigma)\big)$$

The estimated template widths were 27 and 22 deg. for the main and the supplemental experiment.

Finally, the "averaged" model was based on the idea that observers might use a single set of summary statistics to represent the stimuli. Accordingly, we assumed that observers use a single Gaussian template centered at the mean of the overall bimodal distribution:

$$RT_{1T} = a + b\big(p(\theta_T|0, \sigma)\big)$$

The template width was also obtained using ML optimization and bootstrapping. For the main experiment the estimated width was 114 deg., while for the supplemental experiment it was 140 deg. (i.e., almost flat template), already suggesting that this model provides a poor fit to the experimental data.

We then used the estimated template widths to obtain the predictions of the three models for the search times for different target types (Figure 2D), total search time for two targets in different conditions (Figure 2E), and for the order in which the targets should be reported (Figure 2F). For single-target search the equations were the same as when we estimated the template widths, however, we used the data averaged by target type for each subject to reduce the effect of trial-by-trial variability. Two-target search times were assumed to be proportional to a sum of two search times predicted in the same way as for a single target:

$$RT_{2T} = a + b\big(p(\theta_{T1}|D) + p(\theta_{T2}|D)\big)$$

where *D* reflects the distractor distribution parameters for a given model, that is, the template mean(s) and its estimated width(s).

Finally, we assumed that all other things being equal, the order in which the targets are reported would depend on the ratio of the probabilities of observing the test targets under the given distractors template:

$$P(find\ T2\ first) = 0.5 + k\ log\left(\frac{p(\theta_{T1}|D)}{p(\theta_{T2}|D)}\right)$$

with *k* as a scaling constant. The ratio was transformed to logarithm to allow for both positive and negative values.

Figures 2D and 2E show that the probabilistic model provided more accurate predictions for response times than the other models. For single-target response times, it accurately predicted that targets on peaks would be the hardest to find and targets between the peaks would be harder to find than targets outside the range of previously learned distractors. In contrast, the averaged model (ΔBIC = 5.94; here and later ΔBIC refers to the difference in Bayesian Information Criterion compared to the probabilistic model, positive values meaning that the probabilistic model has better fit) suggested that the targets in-between the peaks would be hardest to find, while the single-peak model (ΔBIC = 12.47) predicted relatively similar response times for between targets and targets on peaks. For two-target RTs, the probabilistic model failed to predict longer search for the “outside + between” condition compared to the “peak + outside” condition. Note, however, that this difference was also not significant in the results of our experiment. Speculatively, it might be a result of a higher similarity between the targets in the latter than in the former. Nevertheless, the predictions of the probabilistic model were still better than of the averaged (ΔBIC = 8.04) or the single-peak model (ΔBIC = 6.59).

Crucially, the probabilistic, single-peak, and averaged models gave qualitatively different predictions for the order in which the targets would be found. For both the single-peak and averaged model, the probability of first reporting targets between the peaks when combined with targets on peaks was below 0.5 (Figure 2C). As outlined in the introduction, when observers encode only one peak, on 50% of the trials, the “peak” target on test trials should be on this peak while in the other half of the trials in will be on the non-encoded peak. Depending on the width of the template, the average ratio of the probabilities for a target would vary: with very large or very small template widths, it will be close to 0.5 because targets between the peaks and at the non-encoded peaks will be equally probable, and with intermediate template widths it will be below 0.5 (note that this conclusion is not limited to the specific equation we used for determining the probability of finding one target before another; in fact, it could be shown that this is the case for any monotonic function describing the transformation of a ratio of probabilities of observing the target under a given Gaussian distractor template into average probability of a given reporting order). For the averaged model the target between the peaks should always be reported later than targets on the peaks. In contrast, for the probabilistic model that accurately encodes the probabilities of distractors, the target between the peaks should be reported before the target at the peak. Accordingly, the probabilistic model describes the results better than the single-target (ΔBIC = 13.11) or the averaged model (ΔBIC = 6.86).

Discussion

We assessed the content of templates guiding visual search in the orientation domain, by measuring slowing for targets drawn from a preceding distractor orientation distribution. Crucially, the distribution was bimodal and the searches used to probe the representations involved two simultaneous targets within a trial. Response times were slower when the targets corresponded to the two modes (“peaks”) of previous distractor distributions than when one target was from one of the modes and another from between them, while the latter combination of targets resulted in slower search than when one of the targets was outside the previous distractor range. Furthermore, the order in which the targets were reported on a test trial followed the distractor probabilities observed during prime trials. Targets outside the previous distractor range were reported earlier than the ones between the modes, while the latter were reported before the targets at the modes of previous distractor distribution. When we simulated the predictions of probabilistic, single-template, and averaged template models, we found that the probabilistic model predicts the response times pattern for different target types and different conditions far better than the other models. Moreover, only the probabilistic model was able to accurately predict the order in which the targets were reported. Both the single-template model and the averaged-template model predicted that the target between the peaks should on average be reported after the targets at the peaks, while the reverse was accurately predicted by the probabilistic model. The target between the peaks in the “Peak + Between” condition was on average reported before than the target at one of the peaks. This shows that observers simultaneously represent both modes of distractor distributions. Their representations approximate the physical stimuli, and they fill in the gaps in probability space as demonstrated by slower responses when one of the targets was between the peaks compared to when it was outside the previous distractor range.

Notably, all three models can be considered probabilistic in a sense that they do provide observers with a measure of probability that a certain feature belongs to a distractor class. The difference is in the degree of simplification. The bimodal model reflects the probability distribution accurately (with the assumption of Gaussian approximation). The two other models taken into consideration, however, diverge from an accurate representation in different ways: the "averaged" assumes the use of overall summary statistics, while the "single-peak" assumes the encoding of only one part of the distribution (which could be caused, for example, by biased sampling). While every heuristic or a decision rule can be cast in term of probabilities (e.g., a delta function that assigns probability of 1 for one part of feature space and 0 for the rest), a probabilistic representation as denoted here provides an accurate representation of the probability distribution of the stimuli.

Unlike previous studies assessing how distracting information is stored in visual working memory (Arita et al., 2012; Won & Geng, 2018), the distractors in our studies were heterogeneous and were generated randomly based on a bimodal probability distribution. Nevertheless, observers were able to integrate the information about distractors into an approximate bimodal representation. Speculatively, this demonstrates that using homogeneous distractors may be an artificial limitation, perhaps brought on by earlier technical restrictions on experimental stimuli in pre-modern computer era. In the real world, distracting information is rarely homogeneous, so it may not be particularly surprising that humans are able to form accurate templates representing probability distributions.

Following seminal accounts of priming of pop-out effects (Maljkovic & Nakayama, 1994) we argue that the representations of distractor distributions are kept in visual working memory, rather than long-term memory. Woodman et al (Woodman et al., 2013) have demonstrated that the representation of a single attended target is transferred from VWM to long-term memory in 5 to 7 trials. In contrast, we have previously shown that for simple distractor distributions (such as Gaussian or uniform) one or two trials are enough for observers to develop a probabilistic representation of distractors (Chetverikov et al., 2017b). Representations of more complex distractor distributions take more time (or trials) to develop, but they also progressively change with more repetitions: after one or two trials, bimodal distributions are represented as unimodal, and are only later transformed into bimodal ones. This indicates that more time (trials) is required for sharpening the representation, not for the transfer to long-term memory.

A question of how the probabilistic templates for rejection are stored also taps into a more general question, regarding how working memory templates are stored. Recently, Christophel, Iamschinina and colleagues (2018) demonstrated that while attended stimuli in visual working memory are represented both in parietal and frontal cortex in addition to visual cortex, the latter is not involved in representation of unattended stimuli. It is possible that rejection templates similarly do not involve early visual areas. However, unlike simple unattended items, templates for rejection are actively used by observers to guide attention. As such, their representation might require a level of precision only achievable with the recruitment of sensory areas.

Won and Geng (Won & Geng, 2018) suggested that distractor templates might be more broadly tuned than target templates. This would allow easy generalization of suppression to similar distractors, while for targets such generalization might be harmful as it would lead to an increased number of false alarms. However, the exact costs of generalization for both target and distractor templates depend on the environment. Specific templates are necessary when a target is similar to distractors, but generalization is helpful otherwise. This has indeed been observed by Geng, DiQuattro, and Helm (Geng, DiQuattro, & Helm, 2017): when a target is similar to distractors, its template is sharpened and shifted away from distractors. Moreover, in the real environment we rarely know how exactly the target or distractors would look under a given illumination and point of view, making some degree of generalization essential for efficient search. In contrast, a typical visual search study would require a very narrow distribution of target features, making a narrow template useful. Our results suggest that distractor templates are specific enough to account for bimodality in the distractor distribution. It remains to be studied whether targets or distractors templates are more specific when their physical distributions are equally shaped.

In contrast to our previous studies (Chetverikov et al., 2016, 2017b, 2017d, 2017c), here, we "probed" the distractor representation only at three different points in the feature space. By using targets with a range of features that covered the full feature space, our previous research showed that observers encode the probability distribution of distractors. Here we extend these findings by showing that observers learn the distribution of distractors following a single learning streak. This demonstrates that the previously obtained results are not an artefact of aggregation over multiple trials but rather a true reflection of the templates' content.

We should note that one might interpret our results as simply demonstrating that humans are capable of learning a nonlinear classification rule/decision boundary over a disjoint set in feature space, and can use this to guide visual search. But we think that this alternative proposal is unlikely to hold water because for a simple classifier in this task, learning is not necessary.

There is enough information on each trial to easily tell the target from distractors. Moreover, to include learning in the algorithm, learning of the target would suffice, as the target distribution is constant within the learning streak. The fact that our observers struggle with this shows that they do more than strictly necessary. Second, and perhaps more importantly, we showed in our previous work that observers learn the generative model of the data distribution *on average* rather than learning a simple decision rule (Chetverikov et al., 2016, 2017b, 2017d, 2017c). A decision rule model cannot explain why the response time curves reflect distractor probability both within and outside the distractor distribution range. What determines whether our results support a strong version of the probabilistic templates model, is whether observers learn the distractor distribution on each single learning streak. The aim with the double-target search was to test this specifically.

Conclusions

We found that rejection templates are probabilistic, similarly to items in visual working memory that receive attention (Ma et al., 2014). However, our study also shows that a template for rejection does not need to be a simple bell-shaped curve, as it is typically modelled in working memory studies. In contrast, their representations are dynamically adapted to task requirements, reflecting the probabilistic nature of the input. Whether such flexibility also characterizes templates for attended items remains to be seen. However, our results clearly demonstrate that probabilistic computations start in the brain even before something is perceived, to determine what should be prioritized in perception.

Acknowledgments

Supported by a grant from the Icelandic Research Fund (IRF #173947-052). AC was supported by the Radboud Excellence Initiative. We are grateful to Alena Begler for help with data collection.

Author Contributions

All authors participated equally in conceiving and planning the experiments. AC wrote the experimental scripts, oversaw the data collection, analyzed the results, and wrote the initial version of the manuscript. GC and AK took part in data analyses and interpretation and revised the manuscript.

References

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

Arita, J. T., Carlisle, N. B., & Woodman, G. F. (2012). Templates for rejection: Configuring attention to ignore task-irrelevant features. *Journal of Experimental Psychology: Human Perception and Performance*, *38*(3), 580–584. https://doi.org/10.1037/a0027885

Bays, P. M. (2015). Spikes not slots: Noise in neural populations limits working memory. *Trends in Cognitive Sciences*, *19*(8), 431–438. https://doi.org/10.1016/j.tics.2015.06.004

Bejjanki, V. R., Beck, J. M., Lu, Z.-L., & Pouget, A. (2011). Perceptual learning as improved probabilistic inference in early sensory areas. *Nature Publishing Group*, *14*(5), 642–648. https://doi.org/10.1038/nn.2796

Breiman, L. (1996). Bagging predictors. *Machine Learning*, *24*(2), 123–140. https://doi.org/10.1007/BF00058655

Bundesen, C. (1990). A theory of visual attention. *Psychol Rev*, *97*(4), 523–547. https://doi.org/10.1037/0033-295X.97.4.523

Chetverikov, A., Campana, G., & Kristjánsson, Á. (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

Chetverikov, A., Campana, G., & Kristjánsson, Á. (2017a). Learning features in a complex and changing environment: A distribution-based framework for visual attention and vision in general. In *Progress in Brain Research* (Vol. 236, pp. 97–120). Elsevier. https://doi.org/10.1016/bs.pbr.2017.07.001

Chetverikov, A., Campana, G., & Kristjánsson, Á. (2017b). Rapid learning of visual ensembles.

Journal of Vision, *17*(21), 1–15. https://doi.org/10.1167/17.2.21

Chetverikov, A., Campana, G., & Kristjánsson, Á. (2017c). Representing Color Ensembles. *Psychological Science*, *28*(10), 1–8. https://doi.org/10.1177/0956797617713787

Chetverikov, A., Campana, G., & Kristjánsson, Á. (2017d). Set size manipulations reveal the boundary conditions of distractor distribution learning. *Vision Research*, *140*(April), 144–156. https://doi.org/10.1016/j.visres.2017.08.003

Chetverikov, A., Hansmann-Roth, S., Tanrikulu, Ö. D., & Kristjansson, Á. (2019). Feature Distribution Learning (FDL): A New Method for Studying Visual Ensembles Perception with Priming of Attention Shifts. In *Neuromethods* (pp. 1–21). Springer. https://doi.org/10.1007/7657_2019_20

Christophel, T. B., Iamshchinina, P., Yan, C., Allefeld, C., & Haynes, J. D. (2018). Cortical specialization for attended versus unattended working memory. *Nature Neuroscience, 21*(4), 494–496. https://doi.org/10.1038/s41593-018-0094-4

Feldman, J. (2014). Probabilistic models of perceptual features. In J. Wagemans (Ed.), *Oxford Handbook of Perceptual Organization* (pp. 933–947). https://doi.org/10.1093/oxfordhb/9780199686858.013.049

Geng, J. J., DiQuattro, N. E., & Helm, J. (2017). Distractor probability changes the shape of the attentional template. *Journal of Experimental Psychology: Human Perception and Performance*, *43*(12), 1993–2007. https://doi.org/10.1037/xhp0000430

Girshick, A. R., Landy, M. S., & Simoncelli, E. P. (2011). Cardinal rules: visual orientation perception reflects knowledge of environmental statistics. *Nature Neuroscience*, *14*(7), 926–932. https://doi.org/10.1038/nn.2831

Grubert, A., & Eimer, M. (2013). Qualitative differences in the guidance of attention during single-color and multiple-color visual search: Behavioral and electrophysiological evidence. *Journal of Experimental Psychology: Human Perception and Performance*, *39*(5), 1433–1442. https://doi.org/10.1037/a0031046

Kersten, D., Mamassian, P., & Yuille, A. (2004). Object Perception as Bayesian Inference. *Annual Review of Psychology*, *55*(1), 271–304. https://doi.org/10.1146/annurev.psych.55.090902.142005

Kristjánsson, Á., & Driver, J. (2008). Priming in visual search: separating the effects of target repetition, distractor repetition and role-reversal. *Vision Research*, *48*(10), 1217–1232. https://doi.org/10.1016/j.visres.2008.02.007

Lamy, D. F., Antebi, C., Aviani, N., & Carmel, T. (2008). Priming of Pop-out provides reliable measures of target activation and distractor inhibition in selective attention. *Vision Research*, *48*(1), 30–41. https://doi.org/10.1016/j.visres.2007.10.009

Ma, W. J. (2012). Organizing probabilistic models of perception. *Trends in Cognitive Sciences*, *16*(10), 511–518. https://doi.org/10.1016/j.tics.2012.08.010

Ma, W. J., Husain, M., & Bays, P. M. (2014). Changing concepts of working memory. *Nature Neuroscience*, *17*(3), 347–356. https://doi.org/10.1038/nn.3655

Ma, W. J., Navalpakkam, V., Beck, J. M., van den Berg, R., & Pouget, A. (2011). Behavior and neural basis of near-optimal visual search. *Nature Neuroscience*, *14*(6), 783–790. https://doi.org/10.1038/nn.2814

Ma, W. J., Shen, S., Dziugaite, G., & van den Berg, R. (2015). Requiem for the max rule? *Vision Research*, *116*, 179–193. https://doi.org/10.1016/j.visres.2014.12.019

Maljkovic, V., & Nakayama, K. (1994). Priming of pop-out: I. Role of features. *Memory & Cognition*, *22*(6), 657–672. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/7808275

Oberauer, K. (2002). Access to Information in Working Memory: Exploring the Focus of Attention. *Journal of Experimental Psychology: Learning Memory and Cognition*, *28*(3), 411–421. https://doi.org/10.1037//0278-7393.28.3.411

Olivers, C. N. L., Peters, J., Houtkamp, R., & Roelfsema, P. R. (2011). Different states in visual working memory: When it guides attention and when it does not. *Trends in Cognitive Sciences*, *15*(7), 327–334. https://doi.org/10.1016/j.tics.2011.05.004

Peirce, J. W. (2007). PsychoPy--Psychophysics software in Python. *Journal of Neuroscience Methods*, *162*(1–2), 8–13. https://doi.org/10.1016/j.jneumeth.2006.11.017

Peirce, J. W. (2009). Generating stimuli for neuroscience using PsychoPy. *Frontiers in Neuroinformatics*, *2*. https://doi.org/10.3389/neuro.11.010.2008

Rao, R. P., Olshausen, B. A., & Lewicki, M. S. (2002). *Probabilistic models of the brain: Perception and neural function*. MIT Press.

van Moorselaar, D., Theeuwes, J., & Olivers, C. N. L. (2014). In competition for the attentional template: Can multiple items within visual working memory guide attention? *Journal of Experimental Psychology: Human Perception and Performance*, *40*(4), 1450–1464. https://doi.org/10.1037/a0036229

Vickery, T. J., King, L.-W., & Jiang, Y. (2005). Setting up the target template in visual search. *Journal of Vision*, *5*(1), 8. https://doi.org/10.1167/5.1.8

Wang, D., Kristjánsson, Á., & Nakayama, K. (2005). Efficient visual search without top-down or bottom-up guidance. *Perception & Psychophysics*, *67*(2), 239–253. https://doi.org/10.3758/BF03206488

Won, B.-Y. Y., & Geng, J. J. (2018). Learned suppression for multiple distractors in visual search. *Journal of Experimental Psychology: Human Perception and Performance*, *44*(7), 1128–1141. https://doi.org/10.1037/xhp0000521

Woodman, G. F., Carlisle, N. B., & Reinhart, R. M. G. (2013). Where do we store the memory representations that guide attention? *Journal of Vision*, *13*(3), 1–1. https://doi.org/10.1167/13.3.1

Supplemental experiment

Method

Participants. Fifteen observers (ten female, age *M* = 25.47) at St. Petersburg State University, Russia, took part in a single experimental session lasting approximately 30 min. Two of them were excluded because their response times were very high compared with the other observers (1356 and 1393 ms for excluded observers vs. 761 ms for the remaining sample). One more was excluded because of low accuracy (*M* = 0.71 vs. 0.83 for the remaining sample). The study was approved by the ethics committee of St. Petersburg State University.

Procedure. We used a task similar to our previous studies (Chetverikov et al., 2016, 2017b). Stimuli were presented on an Acer V193 display (19” with 1280 × 1024 pixel resolution) using PsychoPy 1.84.2 (Peirce, 2007, 2009). Viewing distance was ~ 60 cm. Observers searched for an oddly oriented line in a 6×6 grid of 36 lines subtending 16°×16° at the centre of a display. The length of each line was 1.41°. Line positions were jittered by randomly adding a value between ±0.5° to both vertical and horizontal coordinates. Observers indicated whether the target line was in the upper or the lower half of the screen by pressing the ‘i' or ‘j’ keys on a standard keyboard. Trials were organized in intertwined prime and test ‘streaks’. During prime streaks, distractors were randomly drawn from a bimodal distribution that included two uniform parts with orientations ranging from -40 to -20 and +20 to +40 relative to the overall mean. The distribution mean was the same within streak but chosen randomly between streaks. Target orientation was selected randomly on each trial with the restriction that the distance between target orientation and distractor mean in feature space was 60 degrees at minimum. Based on the results of previous studies, prime streak length was set to 6-7 trials because this streak length is sufficient to encode bimodal distributions with relative accuracy (Chetverikov et al., 2017b).

Within test streaks, distractor orientations were randomly drawn from a truncated Gaussian with SD = 10 deg. and range 20 deg. Each test streak had two trials and the targets on these trials (*target type*) were either located on the "peaks" of the previous bimodal distribution (within the +/- 25 to 35 deg. range relative to the previous distractors' mean), in-between the peaks (within 0 to +/-5 deg. range) or outside the previous distribution range (within +/- 55 to 90 deg. range). Three types of test streaks were used with targets on the first and the second test trial either on two different peaks, on a peak and in-between the peaks, or on a peak and outside the previous distribution range. These three conditions were presented equally often in random order. The order of targets within the test trials for each condition was counterbalanced. The distractor mean was chosen randomly with a distance to the target of no less than 60 deg (as on prime trials).

Observers participated in one session of approximately 1244 trials divided into 288 prime and test streaks. Decision time was not limited but participants were encouraged to respond as quickly and accurately as possible. Feedback based on search time and accuracy on previous trials was presented after each trial was shown in the upper-left corner of the screen to motivate participants. The current trial number and the total number of trials were shown beneath the score. If observers made an error, the word "ERROR" appeared in red letters at display centre for 1 second.

Results

Overall performance. Participants were slower ($M = 738$ [683, 794] vs. $M = 615$ [589, 643], $t(11.0) = 5.67$, $p < .001$, $d = 1.64$) and less accurate ($M = 0.78$ [0.75, 0.82] vs. $M = 0.96$ [0.95, 0.97], $t(11.0) = -10.29$, $p < .001$, $d = 2.97$) on learning trials than test trials, due to the fact that learning trials had a broader distribution. Response times decreased while accuracy increased during learning trials: A linear mixed-effects regression indicated that the first trials in each learning sequence were slower, ($B = 0.08$, $SE = 0.02$, $t(11.47) = 4.52$, $p < .001$) and less accurate ($B = -0.04$, $SE = 0.01$, $t(12.52) = -4.50$, $p < .001$) than the later trials.

Test trials. On test trials, observers' performance depended on both target type and condition. Replicating the results of Chetverikov et al. (2017b), on the first trial in a test sequence observers responded more slowly when the target was at one of the peaks of the preceding distractor distribution than when it was in-between the peaks ($t(11.0) = 3.94$, $p = .002$, $d = 1.14$), while responses for the in-between the peaks targets were slower than when they were outside the range of the previous distribution ($t(11.0) = 3.96$, $p = .002$, $d = 1.14$).