

This manuscript has been accepted for publication in Cognition.

Preprint first posted online June 25, 2018.

It is made available under a CC-BY Attribution 4.0 International license. The most current version of the preprint can be found at <https://dx.doi.org/10.17605/OSF.IO/VRBGH>

Probabilistic rejection templates in visual working memory

Andrey Chetverikov^{a,b,c}, Gianluca Campana^{d,e}, & Árni Kristjánsson^{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.

1 **Abstract**

2 Our interactions with the visual world are guided by attention and visual working memory.
3 Things that we look for and those we ignore are stored as templates that reflect our goals and the
4 tasks at hand. The nature of such templates has been widely debated. A recent proposal is that
5 these templates can be thought of as probabilistic representations of task-relevant features.
6 Crucially, such probabilistic templates should accurately reflect feature probabilities in the
7 environment. Here we ask whether observers can quickly form a correct internal model of a
8 complex (bimodal) distribution of distractor features. We assessed observers' representations by
9 measuring the slowing of visual search when target features unexpectedly match a distractor
10 template. Distractor stimuli were heterogeneous, randomly drawn on each trial from a bimodal
11 probability distribution. Using two targets on each trial, we tested whether observers encode the
12 full distribution, only one peak of it, or the average of the two peaks. Search was slower when
13 the two targets corresponded to the two modes of a previous distractor distribution than when
14 one target was at one of the modes and another between them or outside the distribution range.
15 Furthermore, targets on the modes were reported later than targets between the modes that, in
16 turn, were reported later than targets outside this range. This shows that observers use a correct
17 internal model, representing both distribution modes using templates based on the full
18 probability distribution rather than just one peak or simple summary statistics. The findings
19 further confirm that performance in odd-one out search with repeated distractors cannot be
20 described by a simple decision rule. Our findings indicate that probabilistic visual working
21 memory templates guiding attention, dynamically adapt to task requirements, accurately
22 reflecting the probabilistic nature of the input.

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

25

26 **Probabilistic rejection templates in visual working memory**

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

42 Imagine a radiologist looking for signs of tumor in x-ray scans. Malignant signs can take
43 many forms so the targets to look for are diverse. By many accounts, search in this and other
44 contexts is thought to be guided by templates held in visual working memory (Woodman et al.,
45 2013). These templates reflect what one should look for, but may also reflect what should be
46 ignored (Arita, Carlisle, & Woodman, 2012; Won & Geng, 2018). For example, distractors such
47 as the rib cage on a lung scan are salient but not informative and radiologists can therefore ignore
48 them. It is well known that information about to-be-ignored stimuli or features is kept in
49 memory, but the way they are represented is still unknown.


50 There are capacity limits in the amount of information that can be stored in visual
51 working memory templates (Bundesen, 1990; Grubert & Eimer, 2013; Vickery, King, & Jiang,
52 2005), with some authors even suggesting that only one template containing a single feature
53 value can guide attention at any given time (Oberauer, 2002; Olivers, Peters, Houtkamp, &
54 Roelfsema, 2011; van Moorselaar, Theeuwes, & Olivers, 2014). Alternatively, templates could
55 be conceptualized as probabilistic entities of varying precision (Bays, 2015) rather than matches
56 to exact feature values. While previous studies found some support for this, observers typically

57 reported features of single items (Ma, Husain, & Bays, 2014). However, in the real world such
58 isolated features practically never occur. Furthermore, with a few exceptions (Arita et al., 2012;
59 Won & Geng, 2018), templates for ignored information are rarely studied. For any inference
60 based on the probabilistic representations, it is crucial that the internal model used by observers
61 accurately reflects the environment. Here, we provide strong evidence for the probabilistic
62 template view by showing that visual working memory templates for rejection mirror the
63 probability distribution of distractor features.

64 Our observers searched for two oddly oriented targets among distractors randomly drawn
65 from a bimodal orientation distribution. To expose observers' templates, after a sequence of
66 *learning* trials with distractors randomly drawn from a bimodal distribution, targets on *test* trials
67 could either correspond to regions of feature space previously used for distractors, fall in
68 between the modes of the bimodal distribution, or have feature values outside the previous
69 distribution range. We assume that observers' templates reflect what has been relevant on recent
70 trials. If templates contain features of distractors to be ignored, which then become targets on test
71 trials, search should be slower than otherwise (Chetverikov, Campana, & Kristjánsson, 2016;
72 Kristjánsson & Driver, 2008; Lamy, Antebi, Aviani, & Carmel, 2008; Maljkovic & Nakayama,
73 1994; Wang, Kristjánsson, & Nakayama, 2005). Crucially, experiments with varied set size and
74 trial numbers show that learning in this paradigm cannot be explained by the sampling of a few
75 items (Chetverikov, Campana, & Kristjánsson, 2017d, 2017b). It also cannot be explained by
76 simple decision rule learning (e.g., all stimuli that have features in a certain range are
77 distractors), because observers response times, on average, reflect the shape of the distractor
78 distribution rather than just a boundary between a target and distractors (Chetverikov et al., 2016,
79 2017b; Chetverikov, Campana, & Kristjánsson, 2017c; Chetverikov, Hansmann-Roth, Tanrikulu,
80 & Kristjánsson, 2019). However, it is not yet clear whether each single set of learning trials can
81 feed observers' templates with the feature probability distribution of distractors, nor is it clear
82 how accurately the information is stored in the templates.

83 Under the *strong* probabilistic template hypothesis, templates would include information
84 about both peaks of a bimodal distribution. That is, observers would develop an accurate internal
85 model for the task and the template would accurately reflect the information about the full
86 probability distribution. Alternatively, templates might include only a single peak (e.g., the
87 attended one), or might reflect only the summary statistics, such as the averages of the whole

88 distribution (Alvarez, 2011). Using a two-target search task we were able to test whether
 89 observers encode both peaks of a distribution following a single learning sequence. The
 90 predictions of these models (see Simulations) are qualitatively different regarding both the order
 91 in which targets are reported in a two-target search, and search times. If observers accurately
 92 encode a bimodal distribution, on trials with a target on a peak and target between peaks, targets
 93 between the peaks (associated with a lower distractor probability) should be reported before
 94 targets on peaks (associated with the highest distractor probability, Figure 1A). In contrast, if
 95 only one peak is encoded or if the whole distribution is averaged, targets on peaks would be
 96 associated with a lower distractor probability and should be reported no later than targets
 97 between the peaks (associated with lower distractor probability in this case). Notably, while all
 98 three hypotheses postulate that observers can use probabilistic *inference*, only the first one
 99 assumes that the distractor probability distribution is encoded accurately, that is, that the
 100 observers use relatively accurate probabilistic *templates*.



101
 102 **Figure 1.** Panel A: The same physical bimodal distribution can be represented in different ways. Panel B:
 103 Example learning and test trials with distractor distributions and targets shown on the left.

104

105

Experiment

106

107

108

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.

109

110

111

112

113

114

115

116

117

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.

118

119

120

121

122

123

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 $16^\circ \times 16^\circ$ at the centre of a display. The length of each line was 1.41° . Line positions were jittered by randomly adding a value between $\pm 0.5^\circ$ to both vertical and horizontal coordinates.

124

125

126

127

128

129

130

131

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.

132

133

134

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

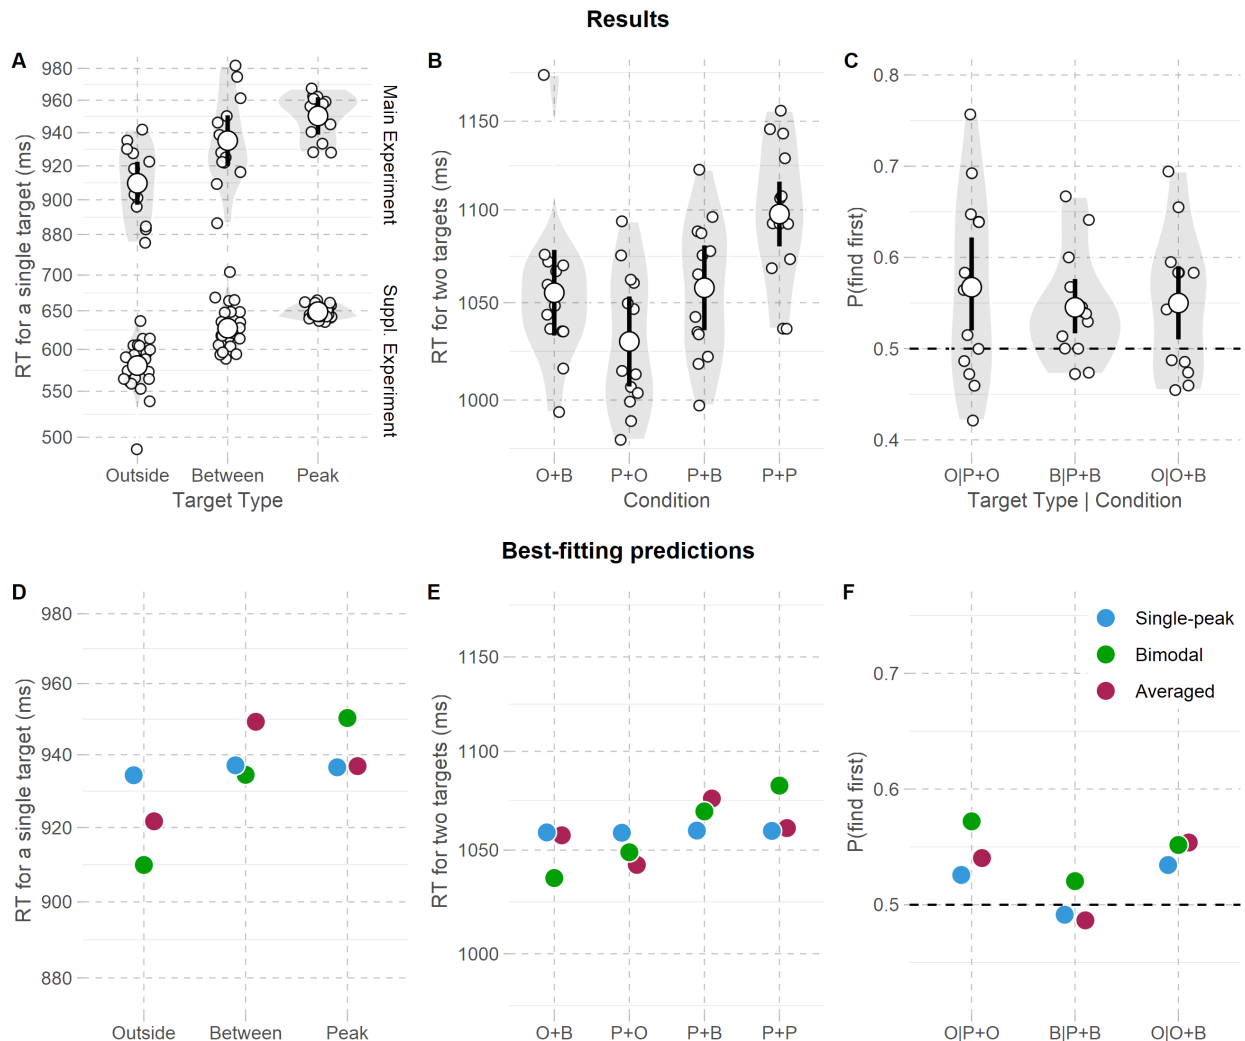
135 distribution mean was the same within streak but chosen randomly between streaks. Target
136 orientations were selected randomly on each trial with the restriction that the distance between
137 target orientation and distractor mean in feature space was 60 degrees at minimum. Prime streak
138 length was set to 6-7 trials (with equal probability) because this streak length is sufficient to learn
139 bimodal distributions with relative accuracy (Chetverikov et al., 2017b).

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

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

163 In addition to this two-target search experiment, we also ran a single-target search study
164 (see Supplementary Experiment). The latter was used as a comparison for the single-target

165 search time analyses to ensure that the introduction of a second target and specific conditions of
 166 the main experiment did not affect the pattern of results.
 167



168
 169 **Figure 2.** Experimental results and best-fitting predictions of the models (see Simulations). A: Results for
 170 different target types from the main Experiment (average search times ignoring the order in which the
 171 targets were reported) and the supplementary Experiment where observers searched for only one target on
 172 each trial. B: Results for two-target search from the main Experiment. C: Results for the order of target
 173 reporting from the main Experiment. D-F: Predictions for single-target search times, search times for two
 174 targets, and for the order in which targets would be reported in a two-target search. For A-C large dots
 175 show group means, bars show their 95% confidence intervals, smaller dots show individual observers'
 176 means, and shaded areas show distributions of individual observers' means. Abbreviations: RT – response
 177 times, P - target on a peak, B - target between the peaks, O - target outside the range of previous distractor
 178 distribution. The plus sign indicates that two targets of corresponding types are used. For the order of
 179 reporting, X|A+B means that target type X was reported first when target types A and B are combined.

180

181 **Results**

182 **Overall performance.** On learning trials, observers found both targets in most cases (M
183 $= 0.72$ [0.67, 0.77]), though the share of trials where only one target was reported was high ($M =$
184 0.27 [0.22, 0.31]; both targets were reported incorrectly on 1% of trials). On test trials, observers
185 reported both targets correctly on $M = 0.91$ [0.89, 0.93] trials (accuracy was comparable to the
186 results of single-target search in the Supplementary Experiment). The delay between the report
187 on the first and the second target was relatively short, but longer on learning than on test trials
188 ($M = 263$ [198, 326] vs. $M = 176$ [130, 233], respectively, $t(12.0) = 4.13$, $p = .001$). Similarly,
189 the first target was reported later on learning than test trials ($M = 973$ [854, 1103] vs. $M = 826$
190 [753, 904], respectively, $t(12.0) = 5.23$, $p < .001$).

191 The learning effects were also comparable to those from the single-target search
192 experiment (see supplement). A linear mixed-effects regression with Helmert contrasts
193 (comparing each trial with the average of the following trials) showed that the first trial was
194 slower, ($B = 0.11$, $SE = 0.01$, $t(52.57) = 9.61$, $p < .001$) and less accurate ($B = -0.04$, $SE = 0.02$,
195 $t(13.12) = -2.62$, $p = .021$) than the later trials. The follow-up trials did not differ from one
196 another.

197 **Test trials.** Replicating previous results, search times differed depending on target type
198 ($F(2, 24) = 8.28$, $p = .003$, $\eta^2_G = .02$, see Figure 2A). Observers search longer for “Peak” targets
199 compared to “Between” targets, which were in turn, found later than “Outside” targets.
200 Crucially, a repeated-measures ANOVA indicated that the time needed to find *both* targets on
201 test trials was affected by the condition ($F(3, 36) = 6.66$, $p = .002$, $\eta^2_G = .02$, Figure 2B).
202 Comparisons between conditions with the same feature difference between the targets showed
203 that performance on “Peak + Peak” trials was slower than on “Outside + Between” trials ($t(12.0)$
204 $= 3.10$, $p = .009$), while “Peak + Between” trials were not different from “Peak + Outside”
205 ($t(12.0) = -1.68$, $p = .118$) trials. Finally, the “Peak + Peak” condition was also slower than the
206 “Peak + Between” condition ($t(12.0) = 2.58$, $p = .024$).

207 We then analyzed which type of target was reported first in each condition using a
208 binomial mixed-effects regression. The results showed that targets on peaks were reported after
209 targets between the peaks ($Z = -2.01$, $p = .044$, Figure 2C) or targets outside the preceding

210 distribution range ($Z = -2.43, p = .015$), while the latter were reported earlier than targets
 211 between the peaks ($Z = 2.08, p = .037$).

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

219 **Simulations.** We simulated the predictions from three models (Figure 2E-F; the
 220 simulation code is available at <https://osf.io/rg2h8>). For our main model of interest, the
 221 “bimodal” model, we assumed that the probabilities of different distractors can be represented by
 222 two Gaussian templates (for simplicity, we ignore the fact that the stimuli distributions might be
 223 more accurately represented by non-Gaussian templates (Chetverikov, Campana, & Kristjánsson,
 224 2017a)) centered on the means of distractor distribution segments. We assumed that observers
 225 utilize the knowledge they obtained about distractors and targets optimally. To find a target in a
 226 localization search task, an ideal observer, would compare the probability that a given noisy
 227 measurement of orientation x at each location L is a target versus the probability that it is a
 228 distractor (Ma, Navalpakkam, Beck, van den Berg, & Pouget, 2011; Ma, Shen, Dziugaite, & van
 229 den Berg, 2015):

$$230 \quad p(L|x) \propto p(L) \frac{p(x|T)}{p(x|D)} = p(L) \frac{p(x|s_L)p(s_L|T)ds_L}{p(x|s_L)p(s_L|D)ds_L}$$

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

$$238 \quad RT \propto p(\theta_T|D)$$

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

$$246 \quad RT_{1T} = a + b \left(\frac{1}{2} p(\theta_T | \mu_1, \sigma) + \frac{1}{2} p(\theta_T | \mu_2, \sigma) \right)$$

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

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

$$255 \quad RT_{1T} = a + b(p(\theta_T | \mu_1, \sigma))$$

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

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

$$261 \quad RT_{1T} = a + b(p(\theta_T | 0, \sigma))$$

262 The template width was also obtained using ML optimization and bootstrapping. For the
 263 main experiment the estimated width was 114 deg., while for the supplemental experiment it was
 264 140 deg. (i.e., almost flat template), already suggesting that this model provides a poor fit to the
 265 experimental data.

266 We then used the estimated template widths to obtain the predictions of the three models
 267 for the search times for different target types (Figure 2D), total search time for two targets in
 268 different conditions (Figure 2E), and for the order in which the targets should be reported (Figure

269 2F). For single-target search the equations were the same as when we estimated the template
 270 widths, however, we used the data averaged by target type for each subject to reduce the effect of
 271 trial-by-trial variability. Two-target search times were assumed to be proportional to a sum of
 272 two search times predicted in the same way as for a single target:

$$273 \quad RT_{2T} = a + b(p(\theta_{T1}|D) + p(\theta_{T2}|D))$$

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

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

$$279 \quad P(\text{find } T2 \text{ first}) = 0.5 + k \log \left(\frac{p(\theta_{T1}|D)}{p(\theta_{T2}|D)} \right)$$

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

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

296 Crucially, the bimodal, single-peak, and averaged models gave qualitatively different
 297 predictions for the order in which the targets would be found. For both the single-peak and
 298 averaged model, the probability of first reporting targets between the peaks when combined with

299 targets on peaks was below 0.5 (Figure 2C). As outlined in the introduction, when observers
300 encode only one peak, on 50% of the trials, the “peak” target on test trials should be on this peak
301 while in the other half of the trials it will be on the non-encoded peak. Depending on the width of
302 the template, the average ratio of the probabilities for a target would vary: with very large or very
303 small template widths, it will be close to 0.5 because targets between the peaks and at the non-
304 encoded peaks will be equally probable, and with intermediate template widths it will be below
305 0.5 (note that this conclusion is not limited to the specific equation we used for determining the
306 probability of finding one target before another; in fact, it could be shown that this is the case for
307 any monotonic function describing the transformation of a ratio of probabilities of observing the
308 target under a given Gaussian distractor template into average probability of a given reporting
309 order). For the averaged model the target between the peaks should always be reported later than
310 targets on the peaks. In contrast, for the bimodal model that accurately encodes the probabilities
311 of distractors, the target between the peaks should be reported before the target at the peak.
312 Accordingly, the bimodal model describes the results better than the single-target ($\Delta\text{BIC} =$
313 13.11) or the averaged model ($\Delta\text{BIC} = 6.86$).

314

315

Discussion

316 Can observers develop an accurate internal model for the probabilities of to-be-ignored items in a
317 visual search task? We assessed the content of templates guiding visual search in the orientation
318 domain, by measuring slowing for targets drawn from a preceding distractor orientation
319 distribution. The distribution was bimodal and the searches used to probe the representations
320 involved two simultaneous targets within a trial. Response times were slower when the targets
321 corresponded to the two modes (“peaks”) of previous distractor distributions than when one
322 target was from one of the modes and another from between them, while the latter combination
323 of targets resulted in slower search than when one of the targets was outside the previous
324 distractor range. Furthermore, the order in which the targets were reported on a test trial followed
325 the distractor probabilities observed during prime trials. Targets outside the previous distractor
326 range were reported earlier than the ones between the modes, while the latter were reported
327 before the targets at the modes of previous distractor distribution. The search times and the order
328 in which targets are reported allowed us to assess the internal model used by observers.

329 We simulated the predictions of a bimodal, single-template, and averaged template
330 models. The first model accurately reflects the actual distribution of distractor features, while the
331 other two oversimplify it in different ways. We found that the bimodal model predicts the
332 response times pattern for different target types and different conditions far better than the other
333 models. Moreover, only the bimodal model could accurately predict the order in which the
334 targets were reported. Both the single-template and the averaged-template model predicted that
335 the target between the peaks should on average be reported after the targets at the peaks, while
336 the reverse was accurately predicted by the bimodal model. The target between the peaks in the
337 “Peak + Between” condition was on average reported before the target at one of the peaks. This
338 shows that observers simultaneously represent both modes of distractor distributions. Their
339 representations approximate the physical stimuli, and they fill in the gaps in probability space as
340 demonstrated by slower responses when one of the targets was between the peaks compared to
341 when it was outside the previous distractor range, or on the peaks.

342 Notably, all three models can be considered probabilistic in a sense that they do provide
343 observers with a measure of probability that a certain feature belongs to a distractor class. The
344 difference is in the degree of simplification. The bimodal model reflects the probability

345 distribution accurately (with the assumption of Gaussian approximation). The two other models
346 taken into consideration, however, diverge from an accurate representation in different ways: the
347 “averaged” assumes the use of overall summary statistics, while the “single-peak” assumes the
348 encoding of only one part of the distribution (which could be caused, for example, by biased
349 sampling). Furthermore, every heuristic or decision rule can be cast in terms of probabilities
350 (e.g., a delta function that assigns probability of 1 for one part of feature space and 0 for the rest).
351 Here we show that the representations used by observers mirror the probability distribution of the
352 stimuli.

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

362 Following seminal accounts of priming of pop-out effects (Maljkovic & Nakayama,
363 1994), we argue that the representations of distractor distributions are kept in visual working
364 memory, rather than long-term memory. Woodman et al (Woodman et al., 2013) have
365 demonstrated that the representation of a single attended target is transferred from VWM to long-
366 term memory in 5 to 7 trials. In contrast, we have previously shown that for simple distractor
367 distributions (such as Gaussian or uniform) one or two trials are enough for observers to develop
368 a probabilistic representation of distractors (Chetverikov et al., 2017b). Representations of more
369 complex distractor distributions take more time (or trials) to develop, but they also progressively
370 change with more repetitions: after one or two trials, bimodal distributions are represented as
371 unimodal, and are only later transformed into bimodal ones. This indicates that more time (trials)
372 is required for sharpening the representation, not for the transfer to long-term memory.

373 A question of how the probabilistic templates for rejection are stored also taps into a
374 more general question, regarding how working memory templates are stored. Recently,
375 Christophel, Iamschinina and colleagues (2018) demonstrated that while attended stimuli in

376 visual working memory are represented both in parietal and frontal cortex in addition to visual
377 cortex, the latter is not involved in representations of unattended stimuli. It is possible that
378 rejection templates similarly do not involve early visual areas. However, unlike simple
379 unattended items, templates for rejection are actively used by observers to guide attention. As
380 such, their representation might require a level of precision only achievable with the recruitment
381 of sensory areas.

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

397 In contrast to our previous studies (Chetverikov et al., 2016, 2017b, 2017d, 2017c;
398 Hansmann-Roth, Chetverikov, & Kristjánsson, 2019), here, we “probed” the distractor
399 representation only at three different points in the feature space. By using targets with a range of
400 features that covered the full feature space, our previous research showed that observers encode
401 the probability distribution of distractors. Here we extend these findings by showing that
402 observers learn the distribution of distractors following a single learning streak. This
403 demonstrates that the previously obtained results are not an artefact of aggregation over multiple
404 trials but rather a true reflection of the templates’ content.

405 Our results agree with previous findings on probabilistic concept learning. Briscoe and
406 Feldman (2011) found that when observers have to form a decision rule based on a multimodal

407 probability distribution, they could do this, although performance became worse with increased
408 mode number. We did not explicitly ask our observers to categorize the stimuli (as distractors
409 and targets), but it is conceivable that they might do so if asked.

410 We should note that one might interpret our results as simply demonstrating that humans
411 are capable of learning a nonlinear classification rule/decision boundary over a disjoint set in
412 feature space, and can use this to guide visual search. But we think that this alternative proposal
413 is unlikely to hold water because for a simple classifier in this task, learning is not necessary.
414 There is enough information on each trial to easily tell the target from distractors. Moreover, to
415 include learning in the algorithm, learning of the target would suffice, as the target distribution is
416 constant within the learning streak. The fact that our observers struggle with this shows that they
417 do more than strictly necessary. Second, and perhaps more importantly, we showed in our
418 previous work that observers learn the correct probabilities of the distractor features *on average*
419 rather than learning a simple decision rule (Chetverikov et al., 2016, 2017b, 2017d, 2017c). A
420 decision rule model cannot explain why the response time curves reflect distractor probability
421 both within and outside the distractor distribution range. By using double-target search we
422 further demonstrate that these results cannot be explained by a combination of different decision
423 rules applied on different test trials.

424 **Conclusions**

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

433

Acknowledgments

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

437

Supplementary materials

438 The data from the experiments reported in this paper and scripts for simulations and analyses are
439 available at <https://osf.io/rg2h8>. The data include trial-by-trial on presented stimuli and
440 observers' responses. The analyses scripts in R provide the full pipeline, including the data
441 preprocessing.

442

Author Contributions

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

447

448

References

- 449 Alvarez, G. A. (2011). Representing multiple objects as an ensemble enhances visual cognition.
450 *Trends in Cognitive Sciences*, 15(3), 122–131. <https://doi.org/10.1016/j.tics.2011.01.003>
- 451 Arita, J. T., Carlisle, N. B., & Woodman, G. F. (2012). Templates for rejection: Configuring
452 attention to ignore task-irrelevant features. *Journal of Experimental Psychology: Human*
453 *Perception and Performance*, 38(3), 580–584. <https://doi.org/10.1037/a0027885>
- 454 Bays, P. M. (2015). Spikes not slots: Noise in neural populations limits working memory. *Trends*
455 *in Cognitive Sciences*, 19(8), 431–438. <https://doi.org/10.1016/j.tics.2015.06.004>
- 456 Bejjanki, V. R., Beck, J. M., Lu, Z.-L., & Pouget, A. (2011). Perceptual learning as improved
457 probabilistic inference in early sensory areas. *Nature Neuroscience*, 14(5), 642–648.
458 <https://doi.org/10.1038/nn.2796>
- 459 Breiman, L. (1996). Bagging predictors. *Machine Learning*, 24(2), 123–140.
460 <https://doi.org/10.1007/BF00058655>
- 461 Briscoe, E., & Feldman, J. (2011). Conceptual complexity and the bias/variance tradeoff.
462 *Cognition*, 118(1), 2–16. <https://doi.org/10.1016/j.cognition.2010.10.004>
- 463 Bundesen, C. (1990). A theory of visual attention. *Psychol Rev*, 97(4), 523–547.
464 <https://doi.org/10.1037/0033-295X.97.4.523>
- 465 Chetverikov, A., Campana, G., & Kristjánsson, Á. (2016). Building ensemble representations:
466 How the shape of preceding distractor distributions affects visual search. *Cognition*, 153,
467 196–210. <https://doi.org/10.1016/j.cognition.2016.04.018>
- 468 Chetverikov, A., Campana, G., & Kristjánsson, Á. (2017a). Learning features in a complex and
469 changing environment: A distribution-based framework for visual attention and vision in
470 general. In *Progress in Brain Research* (Vol. 236, pp. 97–120). Elsevier.
471 <https://doi.org/https://doi.org/10.1016/bs.pbr.2017.07.001>
- 472 Chetverikov, A., Campana, G., & Kristjánsson, Á. (2017b). Rapid learning of visual ensembles.
473 *Journal of Vision*, 17(21), 1–15. <https://doi.org/10.1167/17.2.21>
- 474 Chetverikov, A., Campana, G., & Kristjánsson, Á. (2017c). Representing Color Ensembles.
475 *Psychological Science*, 28(10), 1–8. <https://doi.org/10.1177/0956797617713787>
- 476 Chetverikov, A., Campana, G., & Kristjánsson, Á. (2017d). Set size manipulations reveal the
477 boundary conditions of distractor distribution learning. *Vision Research*, 140(November),

- 478 144–156. <https://doi.org/10.1016/j.visres.2017.08.003>
- 479 Chetverikov, A., Hansmann-Roth, S., Tanrikulu, Ö. D., & Kristjánsson, Á. (2019). Feature
480 Distribution Learning (FDL): A New Method for Studying Visual Ensembles Perception
481 with Priming of Attention Shifts. In *Neuromethods* (pp. 1–21). Springer.
482 https://doi.org/10.1007/7657_2019_20
- 483 Christophel, T. B., Iamshchinina, P., Yan, C., Allefeld, C., & Haynes, J. D. (2018). Cortical
484 specialization for attended versus unattended working memory. *Nature Neuroscience*,
485 *21*(4), 494–496. <https://doi.org/10.1038/s41593-018-0094-4>
- 486 Feldman, J. (2014). Probabilistic models of perceptual features. In J. Wagemans (Ed.), *Oxford*
487 *Handbook of Perceptual Organization* (pp. 933–947).
488 <https://doi.org/10.1093/oxfordhb/9780199686858.013.049>
- 489 Geng, J. J., DiQuattro, N. E., & Helm, J. (2017). Distractor probability changes the shape of the
490 attentional template. *Journal of Experimental Psychology: Human Perception and*
491 *Performance*, *43*(12), 1993–2007. <https://doi.org/10.1037/xhp0000430>
- 492 Girshick, A. R., Landy, M. S., & Simoncelli, E. P. (2011). Cardinal rules: visual orientation
493 perception reflects knowledge of environmental statistics. *Nature Neuroscience*, *14*(7), 926–
494 932. <https://doi.org/10.1038/nn.2831>
- 495 Grubert, A., & Eimer, M. (2013). Qualitative differences in the guidance of attention during
496 single-color and multiple-color visual search: Behavioral and electrophysiological evidence.
497 *Journal of Experimental Psychology: Human Perception and Performance*, *39*(5), 1433–
498 1442. <https://doi.org/10.1037/a0031046>
- 499 Hansmann-Roth, S., Chetverikov, A., & Kristjánsson, Á. (2019). Representing color and
500 orientation ensembles: Can observers learn multiple feature distributions? *Journal of Vision*,
501 *19*(9), 1–17. <https://doi.org/10.1167/19.9.2>
- 502 Kersten, D., Mamassian, P., & Yuille, A. (2004). Object Perception as Bayesian Inference.
503 *Annual Review of Psychology*, *55*(1), 271–304.
504 <https://doi.org/10.1146/annurev.psych.55.090902.142005>
- 505 Kristjánsson, Á., & Driver, J. (2008). Priming in visual search: separating the effects of target
506 repetition, distractor repetition and role-reversal. *Vision Research*, *48*(10), 1217–1232.
507 <https://doi.org/10.1016/j.visres.2008.02.007>
- 508 Lamy, D. F., Antebi, C., Aviani, N., & Carmel, T. (2008). Priming of Pop-out provides reliable

- 509 measures of target activation and distractor inhibition in selective attention. *Vision*
510 *Research*, 48(1), 30–41. <https://doi.org/10.1016/j.visres.2007.10.009>
- 511 Ma, W. J. (2012). Organizing probabilistic models of perception. *Trends in Cognitive Sciences*,
512 16(10), 511–518. <https://doi.org/10.1016/j.tics.2012.08.010>
- 513 Ma, W. J., Husain, M., & Bays, P. M. (2014). Changing concepts of working memory. *Nature*
514 *Neuroscience*, 17(3), 347–356. <https://doi.org/10.1038/nn.3655>
- 515 Ma, W. J., Navalpakkam, V., Beck, J. M., van den Berg, R., & Pouget, A. (2011). Behavior and
516 neural basis of near-optimal visual search. *Nature Neuroscience*, 14(6), 783–790.
517 <https://doi.org/10.1038/nn.2814>
- 518 Ma, W. J., Shen, S., Dziugaite, G., & van den Berg, R. (2015). Requiem for the max rule? *Vision*
519 *Research*, 116, 179–193. <https://doi.org/10.1016/j.visres.2014.12.019>
- 520 Maljkovic, V., & Nakayama, K. (1994). Priming of pop-out: I. Role of features. *Memory &*
521 *Cognition*, 22(6), 657–672. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/7808275>
- 522 Oberauer, K. (2002). Access to Information in Working Memory: Exploring the Focus of
523 Attention. *Journal of Experimental Psychology: Learning Memory and Cognition*, 28(3),
524 411–421. <https://doi.org/10.1037//0278-7393.28.3.411>
- 525 Olivers, C. N. L., Peters, J., Houtkamp, R., & Roelfsema, P. R. (2011). Different states in visual
526 working memory: When it guides attention and when it does not. *Trends in Cognitive*
527 *Sciences*, 15(7), 327–334. <https://doi.org/10.1016/j.tics.2011.05.004>
- 528 Peirce, J. W. (2007). PsychoPy--Psychophysics software in Python. *Journal of Neuroscience*
529 *Methods*, 162(1–2), 8–13. <https://doi.org/10.1016/j.jneumeth.2006.11.017>
- 530 Peirce, J. W. (2009). Generating stimuli for neuroscience using PsychoPy. *Frontiers in*
531 *Neuroinformatics*, 2. <https://doi.org/10.3389/neuro.11.010.2008>
- 532 Rao, R. P., Olshausen, B. A., & Lewicki, M. S. (2002). *Probabilistic models of the brain:*
533 *Perception and neural function*. MIT Press.
- 534 van Moorselaar, D., Theeuwes, J., & Olivers, C. N. L. (2014). In competition for the attentional
535 template: Can multiple items within visual working memory guide attention? *Journal of*
536 *Experimental Psychology: Human Perception and Performance*, 40(4), 1450–1464.
537 <https://doi.org/10.1037/a0036229>
- 538 Vickery, T. J., King, L.-W., & Jiang, Y. (2005). Setting up the target template in visual search.
539 *Journal of Vision*, 5(1), 8. <https://doi.org/10.1167/5.1.8>

- 540 Wang, D., Kristjánsson, Á., & Nakayama, K. (2005). Efficient visual search without top-down or
541 bottom-up guidance. *Perception & Psychophysics*, *67*(2), 239–253.
542 <https://doi.org/10.3758/BF03206488>
- 543 Won, B.-Y. Y., & Geng, J. J. (2018). Learned suppression for multiple distractors in visual
544 search. *Journal of Experimental Psychology: Human Perception and Performance*, *44*(7),
545 1128–1141. <https://doi.org/10.1037/xhp0000521>
- 546 Woodman, G. F., Carlisle, N. B., & Reinhart, R. M. G. (2013). Where do we store the memory
547 representations that guide attention? *Journal of Vision*, *13*(3), 1–1.
548 <https://doi.org/10.1167/13.3.1>
- 549
- 550

551 **Supplemental experiment**552 **Method**

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

559 **Procedure.** We used a task similar to our previous studies (Chetverikov et al., 2016,
560 2017b). Stimuli were presented on an Acer V193 display (19" with 1280×1024 pixel resolution)
561 using PsychoPy 1.84.2 (Peirce, 2007, 2009). Viewing distance was ~ 60 cm. Observers searched
562 for an oddly oriented line in a 6×6 grid of 36 lines subtending $16^\circ \times 16^\circ$ at the centre of a display.
563 The length of each line was 1.41° . Line positions were jittered by randomly adding a value
564 between $\pm 0.5^\circ$ to both vertical and horizontal coordinates. Observers indicated whether the target
565 line was in the upper or the lower half of the screen by pressing the 'i' or 'j' keys on a standard
566 keyboard. Trials were organized in intertwined prime and test 'streaks'. During prime streaks,
567 distractors were randomly drawn from a bimodal distribution that included two uniform parts
568 with orientations ranging from -40 to -20 and $+20$ to $+40$ relative to the overall mean. The
569 distribution mean was the same within streak but chosen randomly between streaks. Target
570 orientation was selected randomly on each trial with the restriction that the distance between
571 target orientation and distractor mean in feature space was 60 degrees at minimum. Based on the
572 results of previous studies, prime streak length was set to 6-7 trials because this streak length is
573 sufficient to encode bimodal distributions with relative accuracy (Chetverikov et al., 2017b).

574 Within test streaks, distractor orientations were randomly drawn from a truncated
575 Gaussian with $SD = 10$ deg. and range 20 deg. Each test streak had two trials and the targets on
576 these trials (*target type*) were either located on the "peaks" of the previous bimodal distribution
577 (within the ± 25 to 35 deg. range relative to the previous distractors' mean), in-between the
578 peaks (within 0 to ± 5 deg. range) or outside the previous distribution range (within ± 55 to 90
579 deg. range). Three types of test streaks were used with targets on the first and the second test trial
580 either on two different peaks, on a peak and in-between the peaks, or on a peak and outside the

581 previous distribution range. These three conditions were presented equally often in random
582 order. The order of targets within the test trials for each condition was counterbalanced. The
583 distractor mean was chosen randomly with a distance to the target of no less than 60 deg. (as on
584 prime trials).

585 Observers participated in one session of approximately 1244 trials divided into 288 prime
586 and test streaks. Decision time was not limited but participants were encouraged to respond as
587 quickly and accurately as possible. Feedback based on search time and accuracy on previous
588 trials was presented after each trial was shown in the upper-left corner of the screen to motivate
589 participants. The current trial number and the total number of trials were shown beneath the
590 score. If observers made an error, the word "ERROR" appeared in red letters at display centre for
591 1 second.

592 Results

593 **Overall performance.** Participants were slower ($M = 738$ [683, 794] vs. $M = 615$ [589,
594 643], $t(11.0) = 5.67$, $p < .001$, $d = 1.64$) and less accurate ($M = 0.78$ [0.75, 0.82] vs. $M = 0.96$
595 [0.95, 0.97], $t(11.0) = -10.29$, $p < .001$, $d = 2.97$) on learning trials than test trials, due to the fact
596 that learning trials had a broader distribution. Response times decreased while accuracy
597 increased during learning trials: A linear mixed-effects regression indicated that the first trials in
598 each learning sequence were slower, ($B = 0.08$, $SE = 0.02$, $t(11.47) = 4.52$, $p < .001$) and less
599 accurate ($B = -0.04$, $SE = 0.01$, $t(12.52) = -4.50$, $p < .001$) than the later trials.

600 **Test trials.** On test trials, observers' performance depended on both target type and
601 condition. Replicating the results of Chetverikov et al. (2017b), on the first trial in a test
602 sequence observers responded more slowly when the target was at one of the peaks of the
603 preceding distractor distribution than when it was in-between the peaks ($t(11.0) = 3.94$, $p = .002$,
604 $d = 1.14$), while responses for the in-between the peaks targets were slower than when they were
605 outside the range of the previous distribution ($t(11.0) = 3.96$, $p = .002$, $d = 1.14$).