



Short Communication

Feature distribution learning by passive exposure

David Pascucci^{a,*}, Gizay Ceylan^a, Árni Kristjánsson^b^a Laboratory of Psychophysics, Ecole Polytechnique Fédérale de Lausanne (EPFL), Switzerland^b Vision Sciences Laboratory, School of Health Sciences, University of Iceland, Reykjavik, Iceland.

ARTICLE INFO

Keywords:

Feature distribution learning
 Passive viewing
 Visual search
 Pop-out search

ABSTRACT

Humans can rapidly estimate the statistical properties of groups of stimuli, including their average and variability. But recent studies of so-called *Feature Distribution Learning* (FDL) have shown that observers can quickly learn even more complex aspects of feature distributions. In FDL, observers learn the full shape of a distribution of features in a set of distractor stimuli and use this information to improve visual search: response times (RT) are slowed if the target feature lies inside the previous distractor distribution, and the RT patterns closely reflect the distribution shape. FDL requires only a few trials and is markedly sensitive to different distribution types. It is unknown, however, whether our perceptual system encodes feature distributions automatically and by passive exposure, or whether this learning requires active engagement with the stimuli. In two experiments, we sought to answer this question. During an initial exposure stage, participants passively viewed a display of 36 lines that included one orientation singleton or no singletons. In the following search display, they had to find an oddly oriented target. The orientations of the lines were determined either by a Gaussian or a uniform distribution. We found evidence for FDL only when the passive trials contained an orientation singleton. Under these conditions, RT's decreased as a function of the orientation distance between the target and the mean of the exposed distractor distribution. These results suggest that passive exposure to a distribution of visual features can affect subsequent search performance, but only if a singleton appears during exposure to the distribution.

1. Introduction

Humans can extract meaningful information from complex visual scenes in a fraction of a second. Most of the information that is available in the immediate 'gist', takes advantage of statistical regularities and redundancies typically found in the visual world. Similar objects are often arranged into groups defined by distributions of low-level features, like color, orientation, size, and location. This provides the visual system with the opportunity to form coarse, global representations, without recognizing individual details: "there is a big basket of small red apples on the left side of the grocery shop". But how coarse are these global representations that are automatically and effortlessly available?

A large body of evidence indicates that humans can rapidly extract basic statistical summaries, like the average and variability of an ensemble of similar stimuli (Ariely, 2001; Whitney & Leib, 2018). Recent studies, however, demonstrate that the perceptual system can quickly learn even more complex and detailed aspects of stimulus distributions, such as the whole shape of a distribution of visual features. In studies of *feature-distribution learning* (FDL; Chetverikov, Campana, &

Kristjánsson, 2016, 2017b, 2017a), for instance, observers are presented with a sequence of visual search trials containing a singleton target embedded in a set of distractors (see Fig. 1). After only a few exemplars of distractors drawn from a particular distribution of features, the visual system learns the shape of the distribution and uses this information to guide future search: searching for a new target becomes easier if the target lies outside the distribution of previous distractors and harder if the target lies inside the previous distractor distribution (a role-reversal effect; Kristjánsson & Driver, 2008). In this way, the response time functions reflect the shape and characteristics of the distractor distribution. This suggests that even when ensembles are not directly relevant for behavior (e.g., a set of distractors in a search task), the brain encodes detailed representations of their properties, beyond summary statistics, which can ultimately aid the detection of task-relevant outliers. This form of learning has been observed for both colors (Chetverikov et al., 2017b) and orientation (Chetverikov et al., 2016), for various distributions types (even bimodal ones, Chetverikov et al., 2017a) and occurs implicitly, since observers cannot explicitly judge distribution shapes, even though the learning strongly affects their search performance

* Corresponding author.

E-mail address: david.pascucci@epfl.ch (D. Pascucci).<https://doi.org/10.1016/j.cognition.2022.105211>

Received 4 November 2021; Received in revised form 20 June 2022; Accepted 21 June 2022

Available online 1 July 2022

0010-0277/© 2022 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

(Hansmann-Roth, Kristjánsson, Whitney, & Chetverikov, 2021).

An important question concerning this learning is whether such detailed statistical representations are encoded automatically and independently of a task —i.e., by simple exposure. In FDL, observers are engaged in a search task that requires them to segment out a potential target from a set of task-irrelevant features. In most of our everyday routines, however, we are not constantly engaged in visual search. Yet, learning the distribution of visual features can help to maintain an accurate representation of our visual environment while maximizing our ability to identify outliers and relevant changes in the environment. In line with this, research has shown that our attentional and perceptual systems constantly learn the characteristics of the visual input, even while idling, giving rise to phenomena of latent learning and plasticity (Turatto, Bonetti, Pascucci, & Chelazzi, 2018; Won & Geng, 2020). In studies of the habituation of the attentional capture response, for example, simple exposure to a repetitive onset stimulus can prepare the attentional system for resisting capture by the same onset in the context of a task (Turatto et al., 2018; Turatto, Bonetti, & Pascucci, 2017; Turatto & Pascucci, 2016).

While feature distribution learning from active search has been found to aid visual performance, this mechanism would be more useful if it can also operate more generally, outside this task-specific context. In the present work, we, therefore, asked whether passive exposure to sets of oriented lines coming from distributions of orientations can induce feature distribution learning effects that modulate visual search during a

subsequent active search task. We used a modified FDL paradigm, presenting a sequence of five displays containing 36 oriented lines. All the lines were oriented according to either a Gaussian or a uniform distribution (see Methods and Fig. 1), except for one oddly oriented line resembling the target of a typical FDL display. In a departure from typical FDL methods, within short blocks with sequences of trials with oriented lines, observers were asked to passively view each display without performing any explicit task. After the sequence, a single test display was presented, requiring observers to perform a search task, reporting the location of the singleton target (upper vs. lower half of the screen). For the test display, that probed the signatures of FDL, search times were analyzed as a function of the angular distance between the mean of the prior distractor distribution and the orientation of the test target (the Current-Target/Previous-Distractor distance, CT-PD) (Chetverikov, Hansmann-Roth, Tanrikulu, & Kristjánsson, 2019).

2. Methods

2.1. Participants

A total of 40 healthy participants (24 in Experiment 1 and 16 in Experiment 2, age range of 19–35 years, 18 females), from the EPFL and the University of Lausanne, participated in the study for a monetary reward (20 CHF/h). All participants had normal or corrected-to-normal vision and were naïve as to the purpose of the experiments. Before the

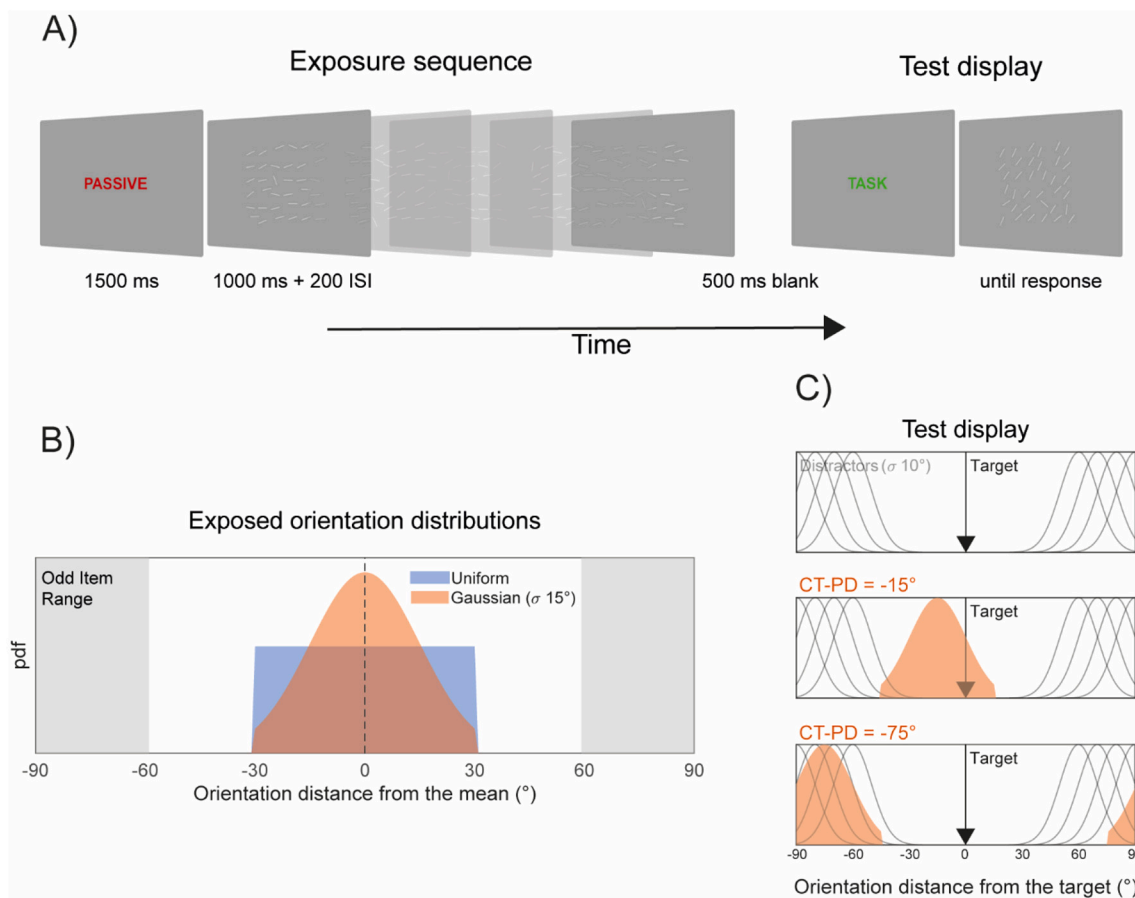


Fig. 1. Trial structure and main variables. A) Example of a sequence of displays in one trial. Participants were instructed by a written cue to passively view the display (exposure sequence) or to perform a visual search (test display). In the visual search task, they had to report the location of the oddly oriented target (upper vs. lower quadrant). B) In the exposure sequence, the orientations in the ensemble of lines were drawn according to two different distributions, a truncated Gaussian, and a uniform distribution. One oddly oriented line with orientation distance of $\pm 60\text{--}90^\circ$ from the ensemble mean, was included in the exposed displays of Experiment 1, but not in Experiment 2. C) In the test display, distractors followed a Gaussian distribution with 10° of standard deviation and a mean located between $\pm 60^\circ$ and $\pm 90^\circ$ away from the target orientation (gray curves). The target orientation was presented at different distances from the mean of the exposed distractors distribution (current-target previous-distractor distances, CT-PD; orange Gaussian distribution in the example).

experiments, visual acuity was tested with the Freiburg Acuity test (Bach, 1996). A value of 1 or above reached with both eyes open was used as the selection criterion. The sample sizes were selected based on previous FDL studies. All participants gave written informed consent, and the study was approved by the local ethics committee under the Declaration of Helsinki (apart from preregistration).

2.2. Apparatus

Stimuli were presented on a gamma-corrected VG248QE monitor (resolution: 1920 × 1080 pixels, refresh rate: 120 Hz) and were generated with custom-made scripts written in Matlab (R2013a) and the Psychophysics Toolbox (Brainard, 1997), running on a Windows-based machine. Experiments were performed in a darkened room and participants sat at 57 cm from the computer screen, with their head positioned on a chin rest.

2.3. Stimuli and procedure

In both experiments, trials were arranged into short mini-blocks comprising an exposure sequence of five displays (1000 ms each, separated by a 200 ms blank screen), followed by a single test display (see Fig. 1). All stimuli were presented on a gray background. A colored written cue (1500 ms) indicated the upcoming sequence of exposed displays (the word 'PASSIVE', written in red) or the test display ('TASK' written in green). Each display contained 36 white lines of 1° length, arranged in a 6-by-6 array within a 14° square centred at the fovea. A small jitter was added to the location of each line within each cell of the 6-by-6 array (jitter range: ±0.5°) to prevent uniform appearance.

During the exposure stage, the orientation of each line was drawn from a truncated Gaussian distribution (standard deviation (σ): 15°, cut-off at ±30° [2σ] from the mean) or a uniform distribution (range: ±30° from the mean), in separate and intermixed mini-blocks consisting of 5 learning trials and a test trial. A 'seed' mean orientation of both distributions was randomly selected for each mini-block and remained the same for the 5 learning trials of each mini-block. In Experiment 1, the exposure displays always contained a singleton line, oriented between ±60° and ±90° away from the mean of the distribution. In Experiment 2, no singleton line was present in the passively viewed displays during the exposure sequence.

On test trials, the orientation of the target line was drawn from among a set of 12 angular distances (from -75° to +75° in steps of 15°) from the seed mean of the previous sequence of exposed distributions. The distance between the target line orientation and the previous distribution mean defined the Current-Target/Previous-Distractor distance (CT-PD). The distractor distribution of orientations in the test display was always Gaussian with σ of 10° and the mean located between ±60° and ±90° away from the target orientation. Note that the true circular mean of the orientations shown in each display of the learning trials could slightly differ from their seed mean. This is because orientations were randomly redrawn in each display, sampled from truncated Gaussian or uniform distributions with the chosen seed mean. However, these variations were negligible (standard deviation of the difference between the true circular mean and the seed mean across trials, subjects, and experiments: <1.60°) and well below the 15° steps of the CT-PD variable used to model changes in search times (see Analysis).

The target line in the test display could be presented in one of the 36 possible locations. Participants were asked to report the location (upper vs. lower half of the screen) of the target line by pressing the 'i' (upper) or 'j' (lower) keys of a computer keyboard as quickly as possible. The test display remained on the screen until a response was made. Two types of feedback were used to ensure that participants engaged in the task and maintained a high accuracy rate. First, error feedback (the word 'error', written in black) was presented following incorrect responses. Second, a score was calculated on each mini-block (as in previous work, Chetverikov et al., 2019), and the average score was shown after every 33

mini-blocks.

At the beginning of each experiment, participants were provided with written and verbal instructions and performed a sequence of practice trials under the supervision of the experimenter. Practice trials were not analyzed further but served to ensure that participants understood the task. The experiments consisted of 264 mini-blocks with 8 breaks, lasting approximately 1 h.

2.4. Analysis

Trials with errors and response times slower than 1 s or faster than 200 ms were excluded from the analysis, following guidelines from previous work (Chetverikov et al., 2019) (34.6% in total for Experiment 1, 23.8% for Experiment 2). Participants performed with an average accuracy of 94 ± 4% in Experiment 1 and 93 ± 4% in Experiment 2. In Experiment 2, one participant was excluded due to average response times larger than 2 s, well beyond the threshold recommended in previous FDL studies (Chetverikov et al., 2019).

To evaluate the effect of the exposed distributions on search times, the relationship between single-trial log-transformed response times and the CT-PD variable was modeled with a set of five different models. These included a model assuming no effect of CT-PD on the log-transformed search times, defined as a model with only a 'constant' parameter (e.g., an intercept) and no parameter associated with the CT-PD variable (see Table 1). The other models assumed effects of various forms. In particular, the Linear model predicted a linear relationship between CT-PD and search times, the Uniform model predicted a step-like shift of search times in the CT-PD range, whereas the remaining models predicted effects that followed Gaussian and truncated-Gaussian distributions. Note that the truncated-Gaussian and uniform distributions corresponded to the one used to generate the exposure distributions. The explicit mathematical form of each model is reported in Table 1. All models were fitted by minimizing the negative log-likelihood of the data given the parameter values, using a quasi-Newton optimization algorithm (MATLAB's *fminunc* function). For each model, the Bayesian information criterion (BIC; Schwarz, 1978) was computed as $BIC = -2 * \log L + k * \log(n)$, where $\log L$ is the maximum log likelihood value under each model, k is the number of parameters in each model (see Table 1) and n is the total number of data points. Model comparison was performed for the two types of exposed distributions separately, by subtracting all BIC values from the largest one. Differences in BIC (ΔBIC) larger than 2 are considered positive evidence against the model with the higher BIC (Raftery, 1995).

Model comparison was paralleled with a two-way repeated-measures ANOVA with factors: Distribution Type (2 levels) and CT-PD (6 levels, from 0° to 75° in steps of 15°, considering the absolute value of the original CT-PD levels). For the ANOVA analysis, the response times of each participant were averaged across CT-PD levels. In the location priming analysis of Experiment 1, we compared response times as a function of whether the location of the singleton in the last exposure display was the same or different from the location of the test singleton, using paired *t*-test analysis.

3. Results

3.1. Experiment 1

In Experiment 1, where observers passively viewed arrays of oriented lines where one item was an orientation singleton, there was strong learning of the orientation distributions beyond the average or range. But there was little evidence that observers learned differences between the Gaussian and uniform distributions.

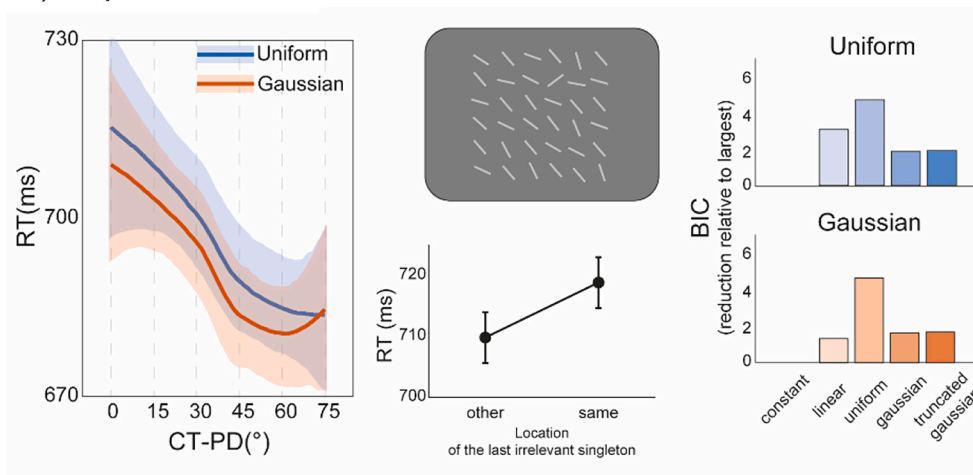
To evaluate whether the orientation distribution in the exposure displays affected performance on the test trial, we compared how well a set of models fit the shape of response times (RT) as a function of CT-PD (see Table 1). The set of models included a model with only a constant (i.

Table 1

The set of models used in both experiments and the BIC of each model for trials with uniform and Gaussian distributions presented during exposure. The Δ BIC against the model with the higher BIC in each model comparison is reported below the BIC value. The model with the higher BIC is indicated by $\Delta = 0$. In all equations, y stands for the log-transformed, single-trial search times, x is the CT-PD variable in degrees (see Methods), a models the intercept, b is either the slope of the Linear model, the offset of the Uniform model or the amplitude of the Gaussian models (Gaussian and Truncated-Gaussian) and σ is the standard deviation of the Gaussian function.

Experiment	Function	1		2	
		Distribution		Distribution	
Model	Function	Gaussian BIC	Uniform BIC	Gaussian BIC	Uniform BIC
Constant	$y = a$	-619.24 $\Delta = 0$	-528.72 $\Delta = 0$	-67.73 $\Delta = 6.65$	-2.31 $\Delta = 6.26$
Linear	$y = a + b x $	-620.57 $\Delta = 1.32$	-531.81 $\Delta = 3.09$	-62.25 $\Delta = 1.17$	3.32 $\Delta = 0.61$
Uniform	$y = \begin{cases} a, & x \leq 2\sigma \\ b, & x > 2\sigma \end{cases}$	-623.9 $\Delta = 4.65$	-533.44 $\Delta = 4.71$	-62.03 $\Delta = 0.95$	3.94 $\Delta = 0$
Gaussian	$y = a + 2be \left(-\frac{ x ^2}{2\sigma^2} \right)$	-620.87 $\Delta = 1.62$	-530.59 $\Delta = 1.87$	-61.10 $\Delta = 0.02$	3.82 $\Delta = 0.12$
Truncated-Gaussian	$y = \begin{cases} a, & x \leq 2\sigma \\ a + 2be \left(-\frac{ x ^2}{2\sigma^2} \right), & x > 2\sigma \end{cases}$	-620.93 $\Delta = 1.68$	-530.65 $\Delta = 1.92$	-61.08 $\Delta = 0$	3.82 $\Delta = 0.12$

A) Experiment 1



B) Experiment 2

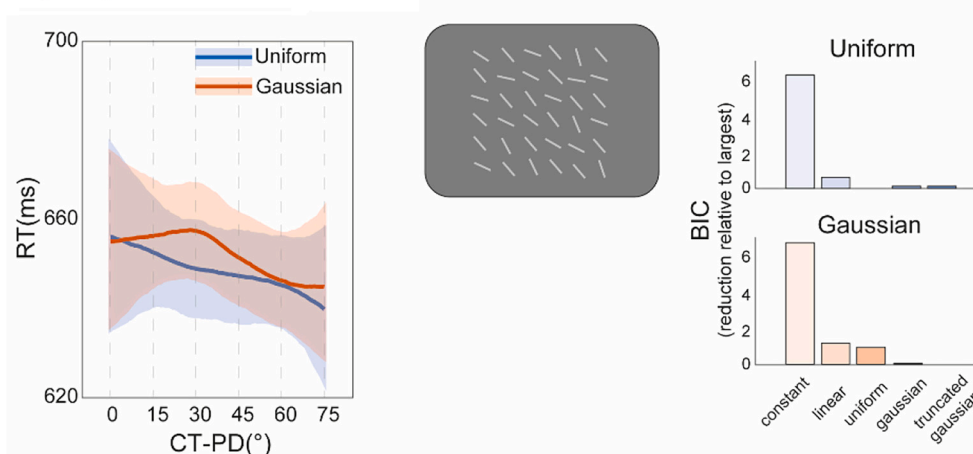


Fig. 2. Results of Experiment 1 and 2. A) In Experiment 1, the exposed set of distractors contained a singleton orientation (as in the example display). Despite the absence of an explicit task, participants learned the distractors distribution during exposure, as evident from the decrease of response times (RT) as the angular distance between the previous distractors mean and the current target orientation increased (current target – previous distractor, CT-PD), a typical hallmark of FDL. According to Bayesian Information Criterion analysis (BIC), all models assuming an effect of the exposed distractors distribution performed better than a model assuming no effect (right panel, see also Table 1). This pattern was comparable across exposure sequences in which a uniform (blue bars) and a Gaussian distribution (red bars) were presented. The singleton line in Experiment 1 captured attention, causing significant negative location priming effects (response times were slower when the test target appeared at the same location of the last exposed singleton, line plot in the bottom-central panel). This involuntarily attentional capture likely caused the implicit learning of the distractors’ distribution. B) In Experiment 2, in which the singleton line was removed during exposure (as in the example display), FDL did not occur and a model assuming no effect of the previous distractor distribution performed better than all the other models (BIC plot, right panel). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

e., no effect of the prior distractor distribution), a linear and a Gaussian model, and a uniform and a truncated-Gaussian model that resembled those used in generating the original distributions of the exposed orientations (see Fig. 2A and Table 1).

All models assuming an effect of the prior distractor distribution performed better than the model assuming no effect (see Table 1; minimum ΔBIC across model comparisons against the constant = 1.32). In particular, model comparison through the Bayesian Information Criterion (BIC; Schwarz, 1978), revealed positive evidence favoring the uniform over the constant and over the other models (all $\Delta\text{BIC} > 2$, except for the comparison ‘uniform vs. linear’ in the condition with uniform distractor distributions: $\Delta\text{BIC} = 1.62$).

This pattern was similar for the two types of distributions observers were exposed to during the learning trials. This means that observers’ search times were clearly affected by the distribution of orientations during the exposure stage, and while this result was further supported by a significant main effect of CT-PD, in a two-way repeated-measures ANOVA with the factors Distribution Type X CT-PD (main effect of CT-PD: $F(5,115) = 6.86, p < 0.001, \eta_p^2 = 0.23$), the main effect of Distribution Type and interaction were not significant. While we would normally expect a main effect of Distribution Type, the results from the model comparison show that observers can passively learn the characteristics of a distribution of features, beyond the simple average. However, the learned characteristics did not allow sufficient resolution to distinguish between the shape of the uniform and the Gaussian distributions.

One potential explanation for this form of learning is that the presence of a singleton, although not task-relevant, per se, might have triggered automatic attentional capture mechanisms that lead to the implicit learning of the distractor features. That is, the singleton may have captured attention because it was an outlier in the orientation distribution and was implicitly represented and learned. To evaluate this possibility, we inspected potential priming effects due to the location of the irrelevant singleton in the last exposure display before the test trial. A small but significant negative priming effect revealed that search times were slower when the singleton during the last learning trials was in the same quadrant as the target singleton ($t(23) = -2.24, p = 0.035$). This confirmed that observers were not completely immune to the presence of a singleton at the exposure stage, and therefore, FDL might have been due to the automatic segmentation of a singleton from the distribution of features. In other words, the feature contrast causes singleton detection which is, in turn also connected with the learning of distributions. To address this point more directly, we performed Experiment 2 in which no singleton appeared during the exposure stage.

3.2. Experiment 2

As in the first experiment, observers passively viewed the display during the exposure trials and were tested on a single test trial. No evidence of an effect of the exposed distributions from the learning trials was found in this experiment (see Fig. 2B and Table 1), with all the models performing worse than the constant only model (constant model against each of the other models with higher BIC, all $\Delta\text{BIC} > 4$), and no significant main effect or interaction in the two-way repeated-measures ANOVA with factors Distribution Type X CT-PD (all p 's > 0.05). This means that the presence of a singleton during the learning trials in Experiment 1, was crucial for triggering FDL by passive viewing.

4. Discussion

Recent work has shown that human observers can rapidly learn rich and detailed representations of distributions of visual features (Chetverikov et al., 2016, 2017a), above what is proposed in summary statistics accounts (Cohen, Dennett, & Kanwisher, 2016). Here we asked whether this form of learning, in which the attentional system represents the full shape of a distribution of visual features, occurs during passive

viewing or whether it requires active visual search.

We found that passive exposure to distributions of visual features in the absence of an explicit search task can affect future search performance, provided that the passively viewed displays contain a singleton element. In previous work, observers also performed a singleton search on the learning trials, and the RT patterns following learning tracked the precise shape of the learned distributions. It was possible to distinguish, for instance, learned representations that resembled the shape of Gaussian, uniform, and even more complex distributions (Chetverikov et al., 2016, 2017b, 2019). In Experiment 1, observers’ performance was affected by the distribution of features during the exposure stage, beyond the simple average, but the resulting RT patterns did not allow us to distinguish between the Gaussian and the uniform distributions.

One explanation for the inability to distinguish different distribution types is that this form of learning is eventually shaped and refined by the repeated attentional selection of a target singleton during active search. The detection of a deviant ‘outlier’ during singleton search may indeed be mediated by an implicit representation of the statistical properties of the whole display (Haberman & Whitney, 2012). Once the active process of segmenting out a deviant from the same distribution of features is reiterated a few times, the shape of the distribution is learned and temporarily stored in detail. Several findings reinforce this idea, suggesting that statistical representations are the building blocks of segmentation and categorization processes (Im, Tiurina, & Utochkin, 2021; Khayat & Hochstein, 2018; Utochkin, 2015). In line with this, the most likely explanation for our results is that the irrelevant singleton on passive trials triggered an unsolicited and involuntary attentional capture, which in turn, led to the automatic segmentation of a deviant from the distribution, and a coarse representation of the distribution’s shape. This is supported by the absence of learning after removing the singleton from the passive trials of Experiment 2, and even more, by the negative location priming effect found in Experiment 1. The negative location priming suggests that participants did encode the irrelevant singleton (and its location) during passive viewing. This, in turn, impaired performance when the target appeared in the same location as the previous singleton, a form of negative priming typically observed at the location of previously irrelevant and distracting stimuli (Fox, 1995). We therefore speculate that the involuntary detection of a deviant singleton left a coarse trace of the statistical distribution of visual features from which the singleton was an outlier. In more active conditions, where the singleton is the target of a visual search, the characteristics of the distribution may be learned and retained more in detail (Chetverikov et al., 2016).

An alternative explanation is that search times were affected by the difference between the singleton orientation and the search target: search times could have been faster when the search target was similar to the singleton and slower as their difference increased. Since the odd item range was always far from the exposed distribution mean, this could have led to the gradual decrease in search times as the CT-PD distance increased (and the similarity between the singleton and the target increased). Because the CT-PD and the singleton-target difference in orientation were highly negatively correlated, due to the experimental design, it is not possible to fully disentangle between the role of one or the other, which could be an interesting question for future research. Nevertheless, both explanations are consistent with the idea that the presence of a singleton triggers an automatic capture and segmentation of the display, which could be mediated by activity in saliency maps that operate as early as in primary visual cortex (Li, 1999, 2002).

Our results add to the emerging field of research showing latent forms of short-term plasticity in the attentional system. These studies have provided evidence of basic and non-associative forms of learning, like habituation of the attentional capture response, after passive and task-free exposure to single stimuli (Turatto et al., 2018; Turatto & Pascucci, 2016; Won & Geng, 2020). Here we show that involuntary attentional capture and singleton segmentation can foster more complex

forms of learning in which the properties of a distribution of visual features are latently and automatically learned. The dependence of passive FDL on the presence of an outlier might be functionally meaningful: constantly learning the entire statistics of irrelevant features in the visual world might be redundant and resource-consuming; learning the distribution of features that signal a ‘surprise’, can aid the attentional system in directing resources more efficiently when such surprise becomes relevant for behavior.

Data availability

The datasets generated during and/or analyzed during the current study are available at <https://doi.org/10.5281/zenodo.6669062>.

Authors contribution

Conceptualization: D.P., A.K. Data curation: D.P., G.C. Investigation: D.P., G.C. Resources: D.P., A.K. Formal analysis: D.P. Writing - original draft: D.P. Writing - review & editing: G.C., A.K., D.P

Acknowledgments

This research was supported by funding from the Swiss National Science Foundation (grant no. PZ00P1_179988 to DP), from the Icelandic Research Fund (#207045-052 to AK & DP) and the Research Fund of the University of Iceland (to AK). The funders had no role in study design, data collection and analysis.

References

- Ariely, D. (2001). Seeing sets: Representation by statistical properties. *Psychological Science*, 12(2), 157–162.
- Bach, M. (1996). The Freiburg Visual Acuity Test-automatic measurement of visual acuity. *Optometry and Vision Science*, 73(1), 49–53.
- Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, 10, 433–436.
- Chetverikov, A., Campana, G., & Kristjánsson, Á. (2016). Building ensemble representations: How the shape of preceding distractor distributions affects visual search. *Cognition*, 153, 196–210.
- Chetverikov, A., Campana, G., & Kristjánsson, Á. (2017a). Rapid learning of visual ensembles. *Journal of Vision*, 17(2), 21.
- Chetverikov, A., Campana, G., & Kristjánsson, Á. (2017b). Representing color ensembles. *Psychological Science*, 28(10), 1510–1517.
- Chetverikov, A., Hansmann-Roth, S., Tanrikulu, Ö. D., & Kristjánsson, Á. (2019). Feature distribution learning (FDL): A new method for studying visual ensembles perception with priming of attention shifts. In *Spatial learning and attention guidance* (pp. 37–57). Springer.
- Cohen, M. A., Dennett, D. C., & Kanwisher, N. (2016). What is the bandwidth of perceptual experience? *Trends in Cognitive Sciences*, 20(5), 324–335.
- Fox, E. (1995). Negative priming from ignored distractors in visual selection: A review. *Psychonomic Bulletin & Review*, 2(2), 145–173.
- Haberma, J., & Whitney, D. (2012). Ensemble perception: Summarizing the scene and broadening the limits of visual processing. In *From perception to consciousness: Searching with Anne Treisman* (pp. 339–349).
- Hansmann-Roth, S., Kristjánsson, Á., Whitney, D., & Chetverikov, A. (2021). Dissociating implicit and explicit ensemble representations reveals the limits of visual perception and the richness of behavior. *Scientific Reports*, 11(1), 1–12.
- Im, H. Y., Tiurina, N. A., & Utochkin, I. S. (2021). An explicit investigation of the roles that feature distributions play in rapid visual categorization. *Attention, Perception, & Psychophysics*, 83(3), 1050–1069.
- Khayat, N., & Hochstein, S. (2018). Perceiving set mean and range: Automaticity and precision. *Journal of Vision*, 18(9), 23.
- 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.
- Li, Z. (1999). Contextual influences in V1 as a basis for pop out and asymmetry in visual search. *Proceedings of the National Academy of Sciences*, 96(18), 10530–10535.
- Li, Z. (2002). A saliency map in primary visual cortex. *Trends in Cognitive Sciences*, 6(1), 9–16.
- Raftery, A. E. (1995). Bayesian model selection in social research. *Sociological Methodology*, 111–163.
- Schwarz, G. (1978). Estimating the dimension of a model. *The Annals of Statistics*, 6(2), 461–464.
- Turatto, M., Bonetti, F., & Pascucci, D. (2017). Filtering visual onsets via habituation: A context-specific long-term memory of irrelevant stimuli. *Psychonomic Bulletin & Review*, 1–7.
- Turatto, M., Bonetti, F., Pascucci, D., & Chelazzi, L. (2018). Desensitizing the attention system to distraction while idling: A new latent learning phenomenon in the visual attention domain. *Journal of Experimental Psychology: General*, 147(12), 1827.
- Turatto, M., & Pascucci, D. (2016). Short-term and long-term plasticity in the visual-attention system: Evidence from habituation of attentional capture. *Neurobiology of Learning and Memory*, 130, 159–169.
- Utochkin, I. S. (2015). Ensemble summary statistics as a basis for rapid visual categorization. *Journal of Vision*, 15(4), 8.
- Whitney, D., & Leib, A. Y. (2018). Ensemble perception. *Annual Review of Psychology*, 69, 105–129.
- Won, B.-Y., & Geng, J. J. (2020). Passive exposure attenuates distraction during visual search. *Journal of Experimental Psychology: General*, 149(10), 1987.