

Representing Color Ensembles



Andrey Chetverikov^{1,2,3}, Gianluca Campana^{4,5}, and
Árni Kristjánsson¹

¹Laboratory for Visual Perception and Visuomotor Control, Faculty of Psychology, School of Health Sciences, University of Iceland; ²Cognitive Research Lab, Russian Presidential Academy of National Economy and Public Administration; ³Department of Psychology, Saint Petersburg State University; ⁴Dipartimento di Psicologia Generale, Università degli Studi di Padova; and ⁵Human Inspired Technology Research Centre, Università degli Studi di Padova

Psychological Science
1–8
© The Author(s) 2017
Reprints and permissions:
sagepub.com/journalsPermissions.nav
DOI: 10.1177/0956797617713787
www.psychologicalscience.org/PS
 SAGE

Abstract

Colors are rarely uniform, yet little is known about how people represent color distributions. We introduce a new method for studying color ensembles based on intertrial learning in visual search. Participants looked for an oddly colored diamond among diamonds with colors taken from either uniform or Gaussian color distributions. On test trials, the targets had various distances in feature space from the mean of the preceding distractor color distribution. Targets on test trials therefore served as probes into probabilistic representations of distractor colors. Test-trial response times revealed a striking similarity between the physical distribution of colors and their internal representations. The results demonstrate that the visual system represents color ensembles in a more detailed way than previously thought, coding not only mean and variance but, most surprisingly, the actual shape (uniform or Gaussian) of the distribution of colors in the environment.

Keywords

color perception, visual search, ensemble representation, summary statistics, priming of pop-out, open data, open materials

Received 9/5/16; Revision accepted 5/16/17

The natural environment is rich with colors. Even simple objects contain a multitude of hues that depend on texture, shape, and reflections from other objects and light sources. How are such color ensembles represented in the brain? Most previous studies have focused on the perception of uniformly colored patches or average impressions of a few colors (Kuriki, 2004; Maule & Franklin, 2015, 2016; McDermott, Malkoc, Mulligan, & Webster, 2010; Sunaga & Yamashita, 2007; Webster, Kay, & Webster, 2014). This is perhaps not surprising since common psychophysical methods are well suited to provide estimates of single values in perceptual space. For example, color matches for average hues within ensembles can be measured (e.g., Kuriki, 2004), but asking participants to match the colors of whole ensembles is more difficult. How would one reply when asked about the color of flowers in a botanical garden? Simple labels and single matched colors are unlikely to capture the richness of perceptual experience.

We introduce a new method based on intertrial learning in visual search that reveals continuous

representations of colors in multicolored ensembles and does not depend on participants' explicit judgments. Following Bayesian models of perception, we consider representations as probabilistic models (Ma, 2012; Rao, Olshausen, & Lewicki, 2002). When searching for an unknown target among distractors, participants form representations of targets and distractors in short-term memory after a few repetitions, and search becomes more efficient (Wang, Kristjánsson, & Nakayama, 2005; see reviews in Kristjánsson & Campana, 2010, and Lamy & Kristjánsson, 2013). Crucially, target stimuli can be used as probes into distractor representations: The more likely it is that a distractor has a given feature value, the slower search will be when a target is defined by this feature (Kristjánsson & Driver, 2008). It is then possible to assess distractor representations by varying

Corresponding Author:

Andrey Chetverikov, Faculty of Psychology, University of Iceland, Saemundargata 2, 101 Reykjavík, Iceland
E-mail: andrey@hi.is

target features relative to preceding distractors and measuring response time (RT) variations.

Using this approach in previous research, we found that participants have surprisingly detailed representations of ensembles of differently oriented lines, and these representations depend on the probability density function (PDF) determining the distribution shape (Chetverikov, Campana, & Kristjánsson, 2016). Such fine-grained representations of feature distributions suggest that perceptual mechanisms utilize probabilistic inferences. Although probabilistic inferences are sometimes considered universal (e.g., Clark, 2013), orientation might be a special case, and humans might rely more on approximations in other feature domains. We used intertrial learning in visual search to investigate how color ensembles are represented by human participants: whether representations of color ensembles involve only mean color and variance or whether finer-grain statistics, such as distribution probability density, are encoded.

Method

Participants

Ten participants (6 female, 4 male; mean age = 32.50 years) at the University of Iceland took part in a single experimental session lasting approximately 45 min, after first completing a training session lasting for at least 100 trials. All participants signed informed consent, participating voluntarily without monetary reward. Sample size and trial numbers were based on the results of previous studies (Chetverikov et al., 2016).

Materials and procedure

Participants sat in a darkened room in front of a 24-in. Asus (Taipei, Taiwan) VX248H display (1,920- × 1,080-pixel resolution; viewing distance was 60 cm). The display was precalibrated using a Cambridge Research Systems (Rochester, England) ColorCal MK2 photometer. The experiment was run using the Psychophysics Toolbox (Brainard, 1997) in MATLAB (The MathWorks, Natick, MA). Forty-eight isoluminant hues were used (luminance: 30 cd/m², background: $x = 0.31$, $y = 0.34$, $Y = 30.0$; all hues were defined in Commission Internationale de l'Éclairage, CIE, 1931, color space on the basis of the values obtained from an isoluminant circle in Derrington, Krauskopf, and Lennie, DKL, color space; see Fig. 1). Feature space was corrected for inequalities in average sensitivity to different colors so that adjacent hues were separated by 1 group-averaged just-noticeable-difference (JND) unit. Data on average JNDs were provided by Christoph Witzel (personal communication,

November 25, 2015). Witzel and Gegenfurtner (2013, 2015) found that JNDs for different hues in isoluminant DKL color space were relatively consistent between participants. For example, for all participants, JNDs were higher for greenish hues than for orange ones. Accordingly, information about average JNDs can be used to partly correct for anisotropies of color space by dividing a color circle into equally discriminable steps. The validity of this correction has been demonstrated in previous studies of color ensembles (Maule & Franklin, 2015, 2016).

The task was an odd-one-out visual search. Participants viewed a set of 36 diamonds, each of which had one corner cut off, arranged on a jittered grid (Fig. 2); on each trial, their task was to find the diamond with a hue most unlike all the others (the target) and report, using the arrow keys, whether the left, right, top, or bottom part of the target diamond was cut off. The structure of this task therefore dissociated response repetition from perceptual priming (Maljkovic & Nakayama, 1994). The experiment consisted of 272 blocks of four to six trials each, with alternating learning sequences (3 or 4 consecutive trials per block) and test sequences (1 or 2 consecutive trials per block; 1,360 trials per participant in total). On each trial, stimuli appeared on the screen and lasted until the observer's response. The sequence of trials was uninterrupted, and participants did not know whether the present trial was a learning or a test trial, nor were they explicitly informed about the features of the target or distractors. During learning sequences, distractor colors were randomly drawn from either a uniform distribution with a range of 24 JNDs or a Gaussian distribution with a standard deviation equal to 6 JNDs, in which outliers with values more than 2 standard deviations above or below the mean were removed, so that on each trial both distributions had the same maximal possible range of 24 JNDs (Fig. 2). The mean of the distributions was constant within learning sequences (but not test sequences) and changed randomly between all sequences.

Target color was chosen randomly, with the restriction that the distance from the target to the distractors' mean in color space was between 18 and 24 JNDs. The means of the distractor distribution were chosen randomly for each learning sequence, and the shape of the distribution and its mean were kept constant within each learning sequence.

On test trials, target orientation was chosen randomly, as was the distractor distribution mean, with the restriction that the target-to-distractor distance ranged from 18 to 24 JNDs (as during the learning sequences). Test trials were used to measure RTs as a function of the distance between the target's color and the mean color of the previous distractor distribution (for convenience,




Fig. 1. Hues used in the study, arranged in Derrington, Krauskopf, and Lennie (DKL) color space. The colored disks were isoluminant, and adjacent hues were separated by an average of 1 just-noticeable-difference (JND) step. The x -axis represents the contrast between L cones and M cones ($L - M$), roughly corresponding to the “red-green” dimension. The y -axis represents the variation in S-cone excitation as $L + M$ activation is constant for isoluminant stimuli, roughly corresponding to the “blue-yellow” dimension. Note that because of differences in sensory thresholds, colors in some parts of the circle are more distant from each other than in other parts, though differences in JNDs remain the same. Color numbers correspond to the scale shown in Figure 2.

we will refer to this as current-target-to-previous-distractors distance, or CT-PD). To keep difficulty constant, we always used a Gaussian distractor distribution with a standard deviation of 3 JNDs for the test sequence.

Decision time was unlimited, but participants were encouraged to respond as quickly and accurately as possible. Feedback based on search time and accuracy was presented after each trial to motivate participants. The current trial number and the total number of trials were shown beneath their score. If participants made an error, the word “ERROR” appeared in red letters at the center of the display for 1 s.

Data analysis

RTs were log-transformed. Linear mixed-effects regression was used to analyze repetition effects from the mean and shape of the distractor distribution during learning sequences. To analyze effects of distractor distributions during learning sequences on subsequent performance on test trials (in particular, to assess differences in the shape of the RT~CT-PD function), we used segmented regression (Muggeo, 2008). Following the approach we used in previous research (Chetverikov et al., 2016), we excluded both errors (4%) and

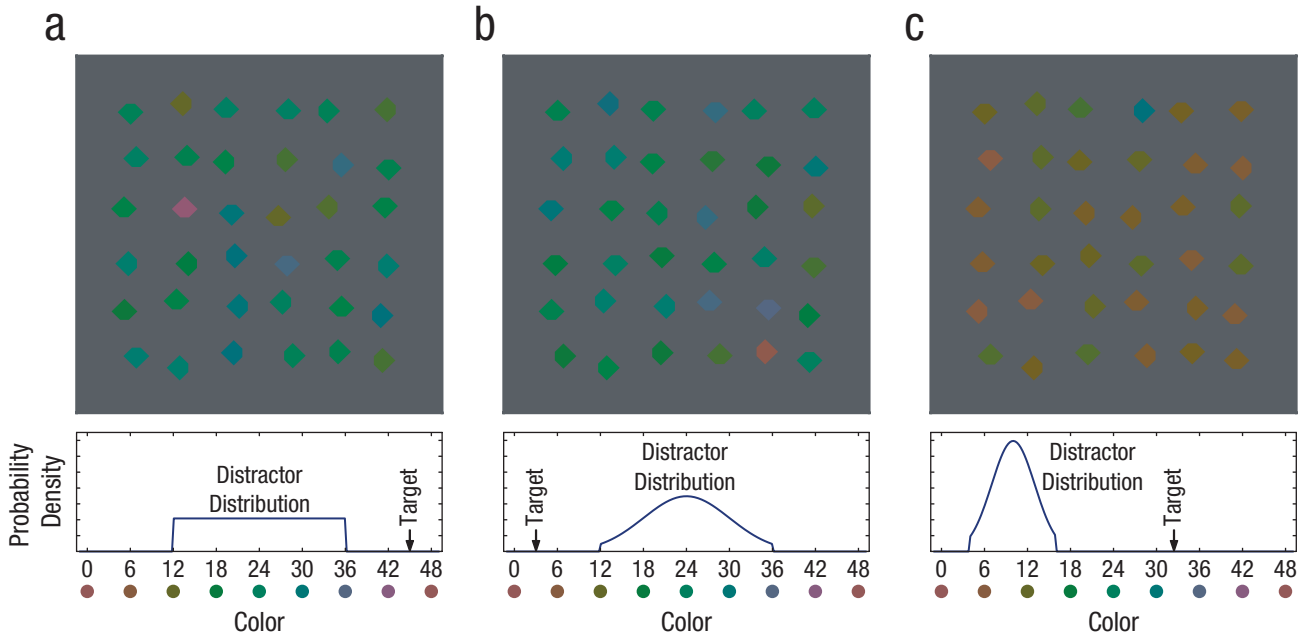


Fig. 2. Example of the three types of display used in the experiment (top) with the distributions from where the distractors were drawn (bottom). In learning trials, distractor colors were drawn from either (a) a uniform distribution with a mean distractor color of 24 just-noticeable-difference (JND) steps and a range of 24 JNDs or (b) a Gaussian distribution with a mean distractor color of 24 JNDs ($SD = 6$), with values more than 2 standard deviations above or below the mean removed, which gave this distribution the same range as the uniform distribution. In test trials (c), distractor colors were drawn from a Gaussian distribution with a mean distractor color of 24 JNDs ($SD = 3$), again with values more than 2 standard deviations above or below the mean truncated. The mean—the position of the distribution on the color wheel—was constant within learning sequences and changed randomly between all sequences and within test sequences.

posterror responses (5%) from analyses on test trials. Posterror responses were excluded because errors may lead to abrupt increases in conscious control over task performance (*posterror slowing*), which might mask any carryover effects from learning trials. Given the results of previous studies, we expected that both uniform and Gaussian PDFs could be approximated as consisting of two segments. For a uniform distribution, the probability of a stimulus with a given value is equal within the distribution range and zero elsewhere. So if participants can learn the shape of the distractor distribution and use it to aid visual search by ignoring hues that have higher probabilities of belonging to that distribution, the first segment of the RT~CT-PD function on test trials following a uniform distribution should be flat, and the second should have a negative slope corresponding to a gradually decreasing probability of observing a distractor in that position in the color space (an abrupt change in the PDF results in a gradual decrease in RT~CT-PD; see Chetverikov et al., 2016). For a Gaussian distribution, the probability monotonically decreases within the distribution range and is also zero elsewhere. Hence, the first segment of the RT~CT-PD function on test trials should have a negative slope, while the second segment should be

flat. An alternative analysis based on model fitting instead of segmented regression is provided in the Supplemental Material available online.

Results

Performance during learning sequences

During learning sequences, search for a target among distractors from the uniform distribution was slightly more difficult than among distractors from the Gaussian distribution, as indexed both by RT ($M = 995$ ms, $SD = 206$ vs. $M = 917$ ms, $SD = 178$, respectively), $t(9) = 7.77$, $p < .001$, and accuracy ($M = 0.94$ ms, $SD = 0.03$ vs. $M = 0.95$ ms, $SD = 0.03$, respectively), $t(9) = 2.31$, $p = .046$. Search times decreased during the first few repetitions, reaching a plateau approximately after the second trial. Accuracy also increased after the first trial (Fig. 3). Linear mixed-effects regression with Helmert contrasts (comparing RT or accuracy on each trial with the average on subsequent trials) showed that participants responded more slowly, $b = 0.15$ (0.01), $t = 20.71$, and less accurately, $b = -0.35$ (0.10), $Z = 3.69$,¹ on the first trial than on later trials within a block.

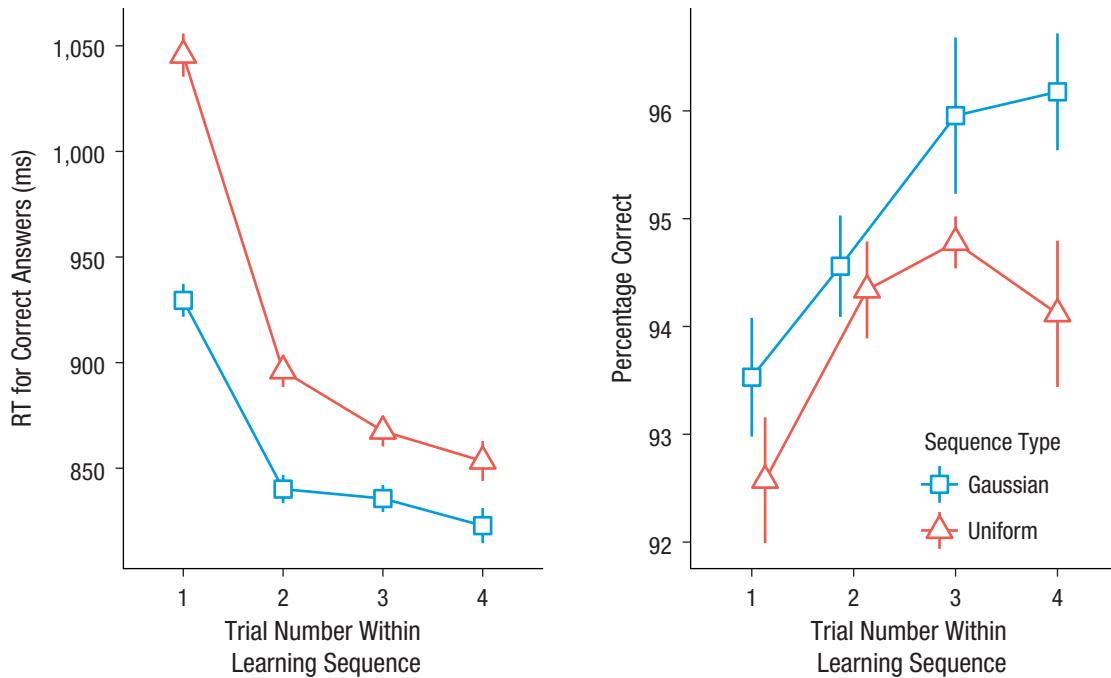


Fig. 3. Mean response time (RT) for correct answers (left) and mean percentage of correct responses (right) as a function of trial number within learning sequences and distribution type. Error bars show ± 1 SEM.

Performance during test sequences

Most important, we found that the shape of the distribution during the learning sequence was reflected in the shape of the RT~CT-PD function. Figure 4 shows that when the distractor distribution during the learning sequence was uniform, RTs during the test sequence were equally slow as long as the target was within the range of the preceding distribution, but RTs decreased when the target was outside that range. A segmented regression confirmed this: Following the uniform learning distribution, RTs can be described as a two-part linear function with a break point at 9 JNDs away from the mean of the learning distribution (95% confidence interval, CI = [5.4, 12.6]). The slope of the first part did not differ significantly from zero, $b = 3.86$, 95% CI = [-5.31, 13.03] (values represent the slope and CI for untransformed search times; the actual analysis was done using log-transformed RTs and yielded similar results), while for the second part, the slope was significantly negative, $b = -8.84$, 95% CI = [-12.31, -5.36]. A Davies test comparing a two-line model that had a break point with a linear model that had no break point (Muggeo, 2008) confirmed that the difference in slopes for the two parts was indeed significant, $p < .001$.

In contrast, following the learning sequence with a Gaussian distractor distribution, RTs during the test sequence (a) monotonically decreased with increasing

CT-PD distance when the target was within the range of the preceding distribution but (b) became flat when the target fell outside the range of the preceding distractor distribution. Again, these observations were confirmed by segmented regression. Following the Gaussian distribution, there was a significant break point as well, but this time it was located 17 JNDs away from the mean of the learning distribution (95% CI = [12.9, 21.4], Davies's $p = .017$). In contrast to RTs following a uniform distribution, here the segment before the break point had a negative slope ($b = -10.08$, 95% CI = [-13.69, -6.46]), while the second had a flat slope ($b = 2.73$, 95% CI = [-7.40, 12.86]).

In sum, following the uniform distribution, the break point showed that participants responded equally slowly when target color was within the range of the previous distractor distribution, whereas following the Gaussian distribution, the break point indicates that participants responded equally fast when target color was well outside the range of the previous distractor distribution. It is, in fact, remarkable how closely the RT~CT-PD function follows the actual PDFs for the two different learning-sequence distributions (Fig. 4, lower panel).

To ensure that the segmented regression result was not an artifact of data aggregation, we also fitted individual models with break points obtained from a segmented model built on the raw data. We then compared slopes before (β_0) and after (β_1) the break point as a

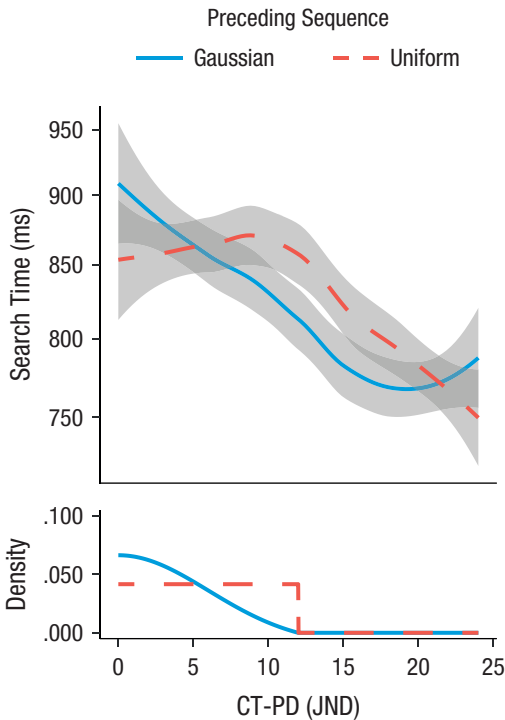


Fig. 4. Mean search time on test trials (top panel) and probability density functions during learning trials (lower panel) as a function of the shape of the distribution during the preceding learning trials (Gaussian vs. uniform) and the distance between the target on test trials and the mean of the distractor distribution on learning trials (current-target-to-previous-distractors distance, or CT-PD, shown in just-noticeable-difference, JND, steps). Curves for search time show mean values smoothed with local polynomial regression. The shaded areas indicate 95% confidence intervals.

function of the learning-sequence distribution. A repeated measures analysis of variance showed a significant interaction between slope type (flat vs. negative) and type of distribution during the previous sequence, $F(1, 9) = 14.13$, $p = .004$, generalized η^2 (η^2_G) = .38 for a break point at 9 JNDs, and $F(1, 9) = 19.29$, $p = .002$, $\eta^2_G = .51$, for a break point at 17 JNDs. In both cases, the slope before the break point was shallower following the uniform than the Gaussian distribution, while following the break point, the opposite was observed (Fig. 5). The analyses using different break points for each distribution (a break point at 9 JNDs for the uniform distribution and at 17 JNDs for the Gaussian distribution) yielded the same results. Individual t tests further showed that for the break point at 9 JNDs, which best fit the raw data, the slope of the first segment following the uniform distribution did indeed not differ from zero, $t(9.0) = 1.25$, $p = .243$, while the slope of the second segment was negative, $t(9.0) = -8.37$, $p < .001$. In contrast, following a Gaussian distribution, the first of the segments with a break point at 17 JNDs was negative, $t(9.0) = -6.67$, $p < .001$, while the second did not differ from zero, $t(9.0) = 0.15$, $p = .880$.

Discussion

Using intertrial learning in visual search, we showed that participants can obtain remarkably detailed representations of color distributions in multihued ensembles of distractors after only three to four repetitions. There was a striking correspondence between RT curves on test trials and the shape of preceding color distributions. RTs following exposure to uniform distributions were constant when test targets were within the range of the learned distribution. In contrast, following exposure to Gaussian distributions, there was a monotonic decrease in RTs that paralleled the monotonic decrease in the probability density of that distribution (Fig. 4). When the target fell within a preceding distractor distribution, RTs were slowed because these distributions contained features previously associated with distractors (Kristjánsson & Driver, 2008). Subjectively, two color distributions were very much alike, as Figure 2 demonstrates. The subtle difference was nevertheless learned.

Our findings show the validity and usefulness of using intertrial learning in visual search to study internal representations of feature distributions in the environment. Along with our previous results (Chetverikov et al., 2016), these data show that learning of feature distributions occurs for different feature domains. Processing of color and orientation in the brain become separated at early stages (Zeki & Shipp, 1988), although they are not independent (e.g., Clifford, Spehar, Solomon, Martin, & Qasim, 2003). Similar sensitivity to learned distributions may suggest that detailed representations are ubiquitous in perceptual processing, which lends support to empirical theories of vision, including Bayesian models and the empirical ranking theory (Howe, Lotto, & Purves, 2006; Rao et al., 2002; Seriès & Seitz, 2013; Yuille & Kersten, 2006).

Previous studies on color perception have mostly used homogeneous patches of colors. Notable exceptions are recent studies of averaged colors (Kuriki, 2004; Sunaga & Yamashita, 2007; Webster et al., 2014) and of color variance (Maule & Franklin, 2015, 2016; Michael, de Gardelle, & Summerfield, 2014), which demonstrate participants' ability to estimate average hues. Instead, we aimed to understand representations of color distributions. Interestingly, we found no evidence of preferential encoding of the mean of the uniform distribution (no differences in RT were seen for any value—mean included—within the range of the learned distribution). Participants may not encode mean color specifically but rather may compute it later when asked, basing judgments on the represented distribution. Moreover, unlike some previous studies (Michael et al., 2014), this research found no priming effect of distribution variance or of the shape of the previous distribution per se (see

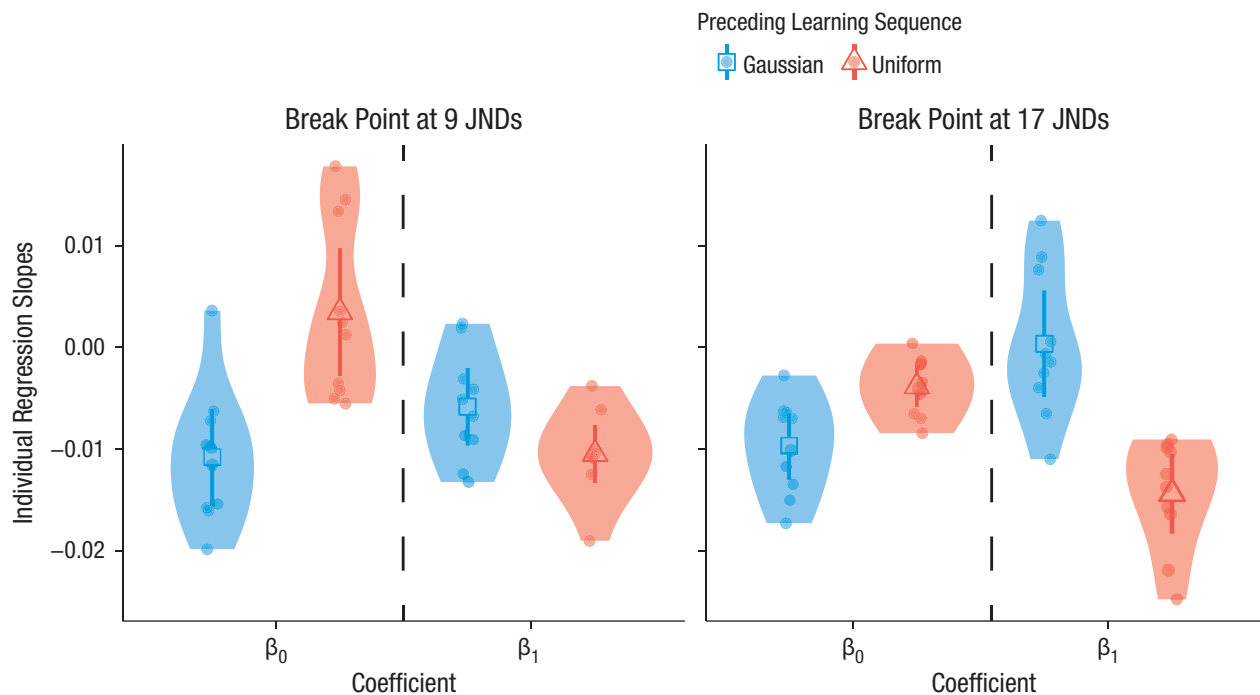


Fig. 5. Individual regression slopes before (β_0) and after (β_1) the break point as a function of the distribution during the learning sequence, separately for break points at 9 and 17 just-noticeable-difference (JND) steps. Squares and triangles show between-participants means, error bars show 95% confidence intervals, dots show regression slopes from individual participants, and shaded areas show the distribution of individual participants' values (the size of the shaded areas along the x -axes indicate the frequency of the corresponding slope value on the y -axes).

the Supplemental Material). That is, RTs in test sequences were not affected by the variance or shape of distractor distributions in learning sequences aside from the changes in RT curves described in the Results. This may indicate different distribution encoding from when participants explicitly judge distribution properties.

There is increasing evidence that perceptual systems estimate natural statistics when assessing colors. For example, many phenomena related to color perception (e.g., variation of hue discrimination or contextual influences on perceived colors) can be explained by statistical distributions in natural images (Long, Yang, & Purves, 2006; Yang & Purves, 2004). If human participants utilize natural image statistics, they must first learn these statistics (Lotto, 2004). Our results show that such learning occurs rapidly and in surprising detail. Similarly, just as uniformly colored patches do not reflect the diversity of hues in the real world, illumination is rarely uniform, and participants take this into account (e.g., Brainard et al., 2006). Knowing the shape of color distributions can aid with estimating the illuminant if participants know that an object has uniform reflectance. In sum, the natural environment is rich with colors, and color representations are detailed enough to reflect this.

Action Editor

Edward S. Awh served as action editor for this article.

Author Contributions

All the authors participated equally in conceiving and planning the experiment. A. Chetverikov wrote the experimental scripts, oversaw the data collection, analyzed the results, and wrote the initial version of the manuscript. G. Campana and Á. Kristjánsson took part in data analyses and interpretation and revised the manuscript.

Acknowledgments

We are grateful to Christoph Witzel for providing perceptually corrected feature space and valuable suggestions on experimental procedure and analyses.

Declaration of Conflicting Interests

The authors declared that they had no conflicts of interest with respect to their authorship or the publication of this article.

Funding

This research was supported by the Russian Foundation for Humanities (Grant No. 15-36-01358A2), the Icelandic Research Fund (Grant No. 152427), and the research fund of the University of Iceland.

Supplemental Material

Additional supporting information can be found at <http://journals.sagepub.com/doi/suppl/10.1177/0956797617713787>

Open Practices



All data and materials have been made publicly available via the Open Science Framework and can be accessed at <https://osf.io/t2856/>. The complete Open Practices Disclosure for this article can be found at <http://journals.sagepub.com/doi/suppl/10.1177/0956797617713787>. This article has received the badges for Open Data and Open Materials. More information about the Open Practices badges can be found at <http://www.psychologicalscience.org/publications/badges>.

Note

1. Wald tests are typically used for binomial regression because variance is known; hence we report Z .

References

- Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision, 10*, 433–436.
- Brainard, D. H., Longèrè, P., Delahunt, P. B., Freeman, W. T., Kraft, J. M., & Xiao, B. (2006). Bayesian model of human color constancy. *Journal of Vision, 6*(11), Article 10. doi:10.1167/6.11.10
- Chetverikov, A., Campana, G., & Kristjánsson, Á. (2016). Building ensemble representations: How the shape of preceding distractor distributions affects visual search. *Cognition, 153*, 196–210. doi:10.1016/j.cognition.2016.04.018
- Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral & Brain Sciences, 36*, 181–204. doi:10.1017/S0140525X12000477
- Clifford, C. W. G., Spehar, B., Solomon, S. G., Martin, P. R., & Qasim, Z. (2003). Interactions between color and luminance in the perception of orientation. *Journal of Vision, 3*(2), Article 1. doi:10.1167/3.2.1
- Howe, C. Q., Lotto, R. B., & Purves, D. (2006). Comparison of Bayesian and empirical ranking approaches to visual perception. *Journal of Theoretical Biology, 241*, 866–875. doi:10.1016/j.jtbi.2006.01.017
- Kristjánsson, Á., & Campana, G. (2010). Where perception meets memory: A review of repetition priming in visual search tasks. *Attention, Perception, & Psychophysics, 72*, 5–18. doi:10.3758/APP.72.1.5
- Kristjánsson, Á., & Driver, J. (2008). Priming in visual search: Separating the effects of target repetition, distractor repetition and role-reversal. *Vision Research, 48*, 1217–1232. doi:10.1016/j.visres.2008.02.007
- Kuriki, I. (2004). Testing the possibility of average-color perception from multi-colored patterns. *Optical Review, 11*, 249–257. doi:10.1007/s10043-004-0249-2
- Lamy, D. F., & Kristjánsson, Á. (2013). Is goal-directed attentional guidance just intertrial priming? A review. *Journal of Vision, 13*(3), Article 14. doi:10.1167/13.3.14
- Long, F., Yang, Z., & Purves, D. (2006). Spectral statistics in natural scenes predict hue, saturation, and brightness. *Proceedings of the National Academy of Sciences, USA, 103*, 6013–6018. doi:10.1073/pnas.0600890103
- Lotto, R. B. (2004). Visual development: Experience puts the colour in life. *Current Biology, 14*, 619–621. doi:10.1016/j.cub.2004.07.045
- Ma, W. J. (2012). Organizing probabilistic models of perception. *Trends in Cognitive Sciences, 16*, 511–518. doi:10.1016/j.tics.2012.08.010
- Maljkovic, V., & Nakayama, K. (1994). Priming of pop-out: I. Role of features. *Memory & Cognition, 22*, 657–672.
- Maule, J., & Franklin, A. (2015). Effects of ensemble complexity and perceptual similarity on rapid averaging of hue. *Journal of Vision, 15*(4), Article 6. doi:10.1167/15.4.6
- Maule, J., & Franklin, A. (2016). Accurate rapid averaging of multihue ensembles is due to a limited capacity subsampling mechanism. *Journal of the Optical Society of America A: Optics, Image Science, and Vision, 33*, A22–A29. doi:10.1364/JOSAA.33.000A22
- McDermott, K. C., Malkoc, G., Mulligan, J. B., & Webster, M. A. (2010). Adaptation and visual salience. *Journal of Vision, 10*(13), Article 17. doi:10.1167/10.13.17
- Michael, E., de Gardelle, V., & Summerfield, C. (2014). Priming by the variability of visual information. *Proceedings of the National Academy of Sciences, USA, 111*, 7873–7878. doi:10.1073/pnas.1308674111
- Muggeo, V. M. R. (2008). segmented: An R package to fit regression models with broken-line relationships. *R News, 8*(1), 20–25.
- Rao, R. P., Olshausen, B. A., & Lewicki, M. S. (2002). *Probabilistic models of the brain: Perception and neural function*. Cambridge, MA: MIT Press.
- Seriès, P., & Seitz, A. R. (2013). Learning what to expect (in visual perception). *Frontiers in Human Neuroscience, 7*, Article 668. doi:10.3389/fnhum.2013.00668
- Sunaga, S., & Yamashita, Y. (2007). Global color impressions of multicolored textured patterns with equal unique hue elements. *Color Research and Application, 32*, 267–277. doi:10.1002/col.20330
- Wang, D., Kristjánsson, Á., & Nakayama, K. (2005). Efficient visual search without top-down or bottom-up guidance. *Perception & Psychophysics, 67*, 239–253. doi:10.3758/BF03206488
- Webster, J., Kay, P., & Webster, M. A. (2014). Perceiving the average hue of color arrays. *Journal of the Optical Society of America A: Optics, Image Science, and Vision, 31*, A283–A292. doi:10.1364/JOSAA.31.00A283
- Witzel, C., & Gegenfurtner, K. R. (2013). Categorical sensitivity to color differences. *Journal of Vision, 13*(7), Article 1. doi:10.1167/13.7.1
- Witzel, C., & Gegenfurtner, K. R. (2015). Categorical facilitation with equally discriminable colors. *Journal of Vision, 15*(8), Article 22. doi:10.1167/15.8.22
- Yang, Z., & Purves, D. (2004). The statistical structure of natural light patterns determines perceived light intensity. *Proceedings of the National Academy of Sciences, USA, 101*, 8745–8750. doi:10.1073/pnas.0402192101
- Yuille, A., & Kersten, D. (2006). Vision as Bayesian inference: Analysis by synthesis? *Trends in Cognitive Sciences, 10*, 301–308. doi:10.1016/j.tics.2006.05.002
- Zeki, S., & Shipp, S. (1988). The functional logic of cortical connections. *Nature, 335*, 311–317. doi:10.1038/335311a0