

Individual differences in response precision correlate with adaptation bias

Marcelo G. Mattar

Department of Psychology, University of Pennsylvania,
Philadelphia, PA, USA

Present address: Princeton Neuroscience Institute,
Princeton University, Princeton, NJ, USA



Marie V. Carter

Department of Neurology, University of Pennsylvania,
Philadelphia, PA, USA

Marc S. Zebrowitz

Department of Neurology, University of Pennsylvania,
Philadelphia, PA, USA

Sharon L. Thompson-Schill

Department of Psychology, University of Pennsylvania,
Philadelphia, PA, USA

Geoffrey K. Aguirre

Department of Neurology, University of Pennsylvania,
Philadelphia, PA, USA

The internal representation of stimuli is imperfect and subject to bias. Noise introduced at initial encoding and during maintenance degrades the precision of representation. Stimulus estimation is also biased away from recently encountered stimuli, a phenomenon known as adaptation. Within a Bayesian framework, greater biases are predicted to result from poor precision. We tested for this effect on individual difference measures. Through an online experiment, 202 subjects contributed data. During separate face and color blocks, they performed three different tasks: an immediate stimulus match, a delayed match-to-sample, and a delayed match following 5 s of adaptation. The stimulus spaces were circular, and subjects entered their responses on a color/face wheel. Bias and precision of responses were extracted while accounting for the probability of random guesses. We found that the adaptation manipulation induced the expected bias in responses, and the magnitude of this bias varied reliably and substantially between subjects. Across subjects, there was a negative correlation between mean precision and bias. This relationship was replicated in a new experiment with 192 subjects. This result is consistent with a Bayesian observer model, in which the precision of perceptual representation influences the magnitude of perceptual bias.

Introduction

Perception is our “best guess” as to what is in the world (von Helmholtz, 1867). When sensory information is imperfect or incomplete, this best guess relies more heavily on expectations and prior beliefs (Knill & Richards, 1996). Prior knowledge influences perception continuously and automatically, and may occasionally lead to perceptual illusions (Bar, 2004; Summerfield & Egner, 2009; Lafer-Sousa, Hermann, & Conway, 2015). To achieve statistically optimal inference, sensory evidence and prior beliefs are quantitatively combined using Bayes’s rule (Bayes & Price, 1763).

While many aspects of human perception are found to be consistent with Bayesian observer models (Knill & Richards, 1996; Stocker & Simoncelli, 2006a), it has been more difficult to apply this framework to the ubiquitous phenomenon of sensory adaptation. Recent sensory history tends to bias perception “away” from observed stimuli (Levinson & Sekuler, 1976; Clifford, 2002; although see Gibson & Radner, 1937; Fischer & Whitney, 2014), which seems at odds with a Bayesian account of perception that is biased toward prior beliefs. Recent theoretical work has attempted to explain perceptual adaptation as arising from either an asymmetric likelihood (Wei & Stocker, 2015) or a modification of the prior (Chopin & Mamassian, 2012).

Citation: Mattar, M. G., Carter, M. V., Zebrowitz, M. S., Thompson-Schill, S. L., & Aguirre, G. K. (2018). Individual differences in response precision correlate with adaptation bias. *Journal of Vision*, 18(13):18, 1–12, <https://doi.org/10.1167/18.13.18>.



	Completed registration	Completed screening	Completed experiment
Number of subjects	530	422	202
Age ($M \pm SD$)	36 ± 12	36 ± 11	35 ± 11
Number male/female	252/277	191/231	92/110
Number left-/right-handed	18/508	14/405	3/197

Table 1. Subject enrollment and exclusion in the main experiment. *Notes:* Totals may not match the total number of subjects due to missing responses from some subjects.

In either case, larger perceptual biases should be observed in the setting of greater imprecision.

Here we tested this key prediction of Bayesian inference. Using a web-based visual-adaptation experiment with color and face stimuli, we investigated the relationship across individuals between variability in response precision and in adaptation biases. We found that precision and bias are relatively stable measures of an individual, and that biases for color stimuli are correlated with biases for face stimuli. We then investigated the relationship between bias and precision across individuals and found that biases for both materials are larger when precision is lower. These results conform with predictions of Bayesian observer models, whereby perception is biased away from recently observed stimuli when sensory information is uncertain.

Methods

Participants

We recruited 1,002 people through Amazon Mechanical Turk: 530 for the main experiment (Table 1) and 472 for the replication experiment (Tables 2 and 3). This research was reviewed and deemed exempt from oversight by the University of Pennsylvania Institutional Review Board, and therefore informed consent was not collected. Information on the home page of the

	Completed registration	Completed screening	Completed experiment
Number of subjects	191	149	98
Age ($M \pm SD$)	37 ± 12	37 ± 12	39 ± 12
Number male/female	110/81	87/62	60/38
Number left-/right-handed	8/176	4/139	0/95

Table 2. Subject enrollment and exclusion in the replication experiment with colors. *Notes:* Totals may not match the total number of subjects due to missing responses from some subjects.

	Completed registration	Completed screening	Completed experiment
Number of subjects	281	234	94
Age ($M \pm SD$)	36 ± 12	37 ± 12	37 ± 12
Number male/female	164/117	134/100	48/46
Number left-/right-handed	19/262	13/221	3/91

Table 3. Subject enrollment and exclusion in the replication experiment with faces. *Notes:* Totals may not match the total number of subjects due to missing responses from some subjects.

web-based experiment indicated the research nature of the project. No information that could identify participants was collected. All subjects received a fixed minimum compensation of \$0.25 for their participation in addition to a performance-based bonus of up to \$12.00. The full experiment took approximately 75 min to complete, and subjects that reached the end received an average bonus of \$7.68 ($SD = \0.51).

Stimuli and materials

The experiment was programmed in JavaScript and hosted on a custom-built website that subjects accessed using their own computers (<https://cfn.upenn.edu/iadapt>). Stimuli consisted of synthetic faces generated with FaceGen Main SDK (Inversions, 2012) and colored squares.

The face set used in the main experiment varied in age and gender. Two base stimuli were generated by varying the gender of an identity-neutral face from male to female, and another two base stimuli were generated by varying the age of an identity-neutral face from 15 to 65 years. Based on these four stimuli, a set of 360 faces were generated in a circular space by varying both dimensions, with main axes corresponding to age and gender.

The color set varied in hue, with no nominal variation in lightness and saturation. A set of 360 color values were generated in HSL space with L^* held fixed at 25 and saturation equal to 7. This saturation value was determined in pilot experiments to produce stimuli whose response variance approximately matched that of face stimuli. The set of HSL values was then converted to sRGB space.

For the replication experiment, new sets of face and color stimuli were generated. The face stimuli were again generated based on four stimuli that varied by apparent age (15–65 years) and gender (male–female). To produce more distinctive faces than the previous set, the four base stimuli also varied in internal facial features. This was realized by displacing each base face stimulus along a random vector direction within the

face-stimulus space, orthogonal to both age and gender. The color stimuli were again generated in HSL space with L^* held fixed at 25, but now with saturation equal to 20 (the maximum value that produced sRGB values within the 0–255 range displayable in regular monitors), producing more distinctive colors than in the main experiment.

Experimental procedure

Subjects recruited through Mechanical Turk were redirected to the experiment website, where they entered their responses using mouse and keyboard input. Upon completion, subjects received an anonymous code that they provided to Mechanical Turk for payment. Each experimental session started with a basic description of the procedures, followed by a demographic questionnaire. Subjects then completed a screen-calibration procedure and a color-perception test, followed by the main experimental blocks.

In the screen-calibration procedure, subjects were presented with a set of discrete color gradients organized in rows, each ranging from black to a distinct, saturated color value. They were then asked to adjust the screen settings or the angle of their laptop screen to allow them to simultaneously distinguish between neighboring colors on both ends of the each gradient. In the color-perception test, subjects completed eight trials of the Ishihara test, a test for congenital color deficiencies. Subjects proceeded to the main experimental blocks only if at least seven responses were correct. This online test was not capable of detecting subtle color-vision deficits but did ensure that the subjects and their computers could support a baseline level of color-discrimination ability.

Subjects then completed a battery of tests designed to estimate representation precision, precision after a delay, and precision and bias after adaptation. Data were collected over two to six blocks of each of the following experiments: stimulus match, delayed match, and adaptation (Figure 1a). Results from the delayed-match task are not further considered here. Each block consisted of 15–30 trials (each lasting between 2 and 5 min) in which subjects were instructed to report the value of a target stimulus—a color or a face—by clicking and moving the mouse cursor around a stimulus wheel (Figure 1b), allowing a fine adjustment of their responses. Prior to performing each type of experiment for the first time, subjects were presented with detailed instructions, a mini-quiz containing three questions with three alternatives each about the instructions, and five practice trials. If any answer to the quiz was incorrect, subjects were repeatedly presented with the instructions and the quiz until all three answers were simultaneously correct. Similarly, subjects repeated the practice exper-

iment as many times as necessary until all five responses were within 5° of the target. Together, these approaches ensured that subjects comprehended the experiment instructions and were able to adequately perform the experiment on their computers.

Experimental blocks in the main experiment were completed in the following order: color stimulus match; face stimulus match; color delayed match; face delayed match; color adaptation; face adaptation; color delayed match; face delayed match; color adaptation; face adaptation; color stimulus match; face stimulus match. In the replication experiment, blocks (only one stimulus class) were completed in the following order: stimulus match; delayed match; adaptation (six blocks); delayed match; stimulus match.

We calculated subject accuracy on every trial by linearly mapping from errors (0° – 90°) to accuracy (100%–0%); at the end of each block, we calculated the average accuracy for that block (100% being perfect performance). The compensation accumulated by subjects increased at the end of each experimental block by an amount proportional to the average accuracy. Subjects were presented with their average accuracy in the finished block, the corresponding dollar amount accumulated, and the total compensation accumulated in the experimental session up to that point. If the accuracy on any block ended up below 20%, the experimental session was terminated and subjects were directed back to Amazon Mechanical Turk to receive their payment. Only subjects who maintained accuracy above 20% in all blocks were able to reach the end of the experiment; those subjects received twice the regular compensation. Out of the 1,002 subjects recruited across both the main experiment and the replication experiment, 197 were excluded for either abandoning or not passing the color-perception test and 411 for not maintaining accuracy above 20% throughout the entire session. Only data from the remaining 394 subjects were included in the analyses described in this article (Tables 1, 2, 3).

Stimulus-match experiment

Each block consisted of 15 trials, with target values sampled uniformly (24° spacing) from the circular space; each target was presented once per block. On each trial, subjects were presented with a target stimulus on both the left and right sides of the black background screen, along with a stimulus wheel containing eight equally spaced thumbnails with representative stimuli from the circular space (Figure 1c, top). Because we could not control the properties of the display used to present the stimulus to each subject, we could not ensure that the perceptual similarity between equally spaced stimuli was uniform around the

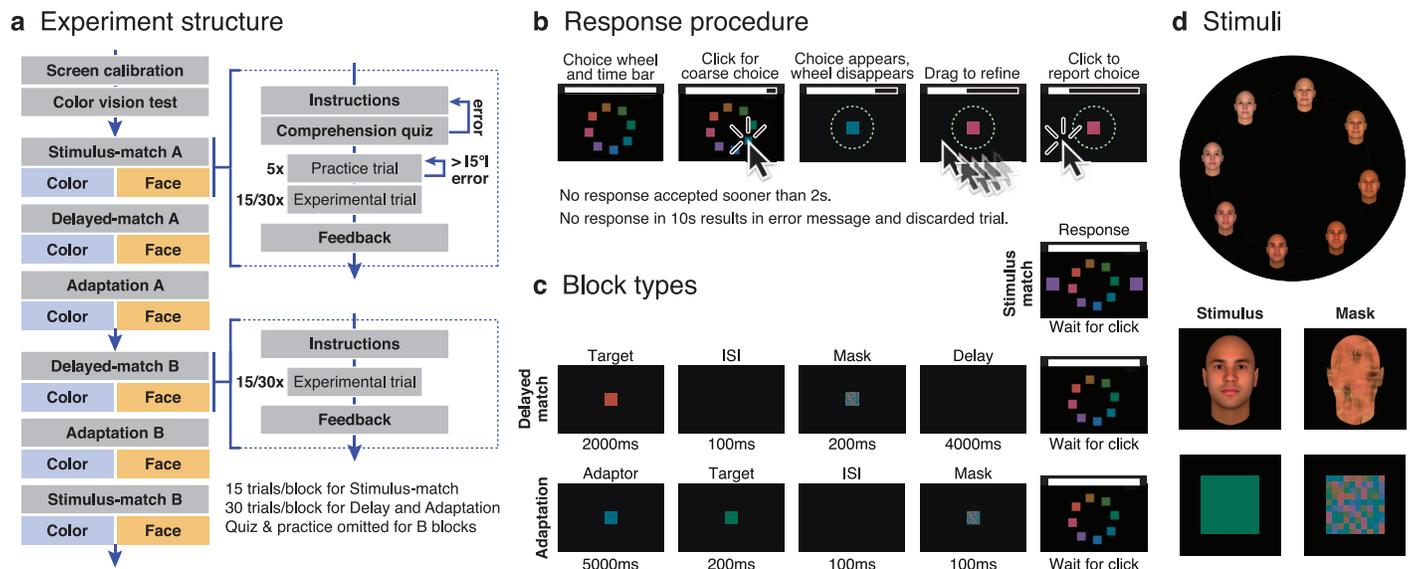


Figure 1. Experiment setup and methods. (a) Experiment structure. Each subject completed two blocks of each experiment, for both color and face stimuli. Prior to each experiment, subjects were presented with instructions and a comprehension quiz; the first time they performed each block of a given experiment, they also completed five practice trials. Each block of the stimulus-match experiment had 15 trials. Each block of the delayed-match and adaptation experiments had 30 trials. (b) Response procedure. Subjects were presented with eight equally spaced representative stimuli around the wheel. To enter their responses, they first performed a coarse selection by clicking on the region of the wheel that approximately matched the target. They then adjusted their selection more precisely by dragging the pointer around the wheel and clicking a second time to confirm their responses. (c) Block types. Top: On each trial of the stimulus-match experiment, subjects were instructed to match to a stimulus presented on both sides of the screen; the target stimulus remained on the response screen while they made the match. Middle: On each trial of the delayed-match experiment, subjects were instructed to match to a target stimulus following a 4-s interval. The data from this task are not further considered in this article. Bottom: On each trial of the adaptation experiment, subjects were instructed to match to a target stimulus that was presented after a 5-s adaptation period and followed by a mask. An interval of 50 ms separated adaptor and target (not shown), and an interval of 100 ms separated target and mask. (d) Stimuli. A circular space with 360 stimuli was used for both color and face stimuli. Color stimuli were generated to vary in hue but not saturation or luminance (HSL space). Face stimuli varied in age and gender, each forming one axis of the space. Color masks were a checkerboard composed of various stimuli randomly sampled from the color space. The eight stimuli that are shown in the figure and in the response screen are equally spaced examples from the entire set of 360, and were selected at random on each trial. Face masks were created using the steerable pyramids method to match various low-level visual properties (Portilla & Simoncelli, 2000).

wheel. Further, the perception of the stimuli may be subject to long-term priors over the stimulus space. However, by sampling uniformly from the wheel and providing test stimuli in both the clockwise and counterclockwise directions, we expected to ameliorate either of these forms of stimulus bias on average. Moreover, the specific thumbnails, their position, and the mapping of stimulus value to screen position varied randomly on each trial (i.e., there was no systematic relationship between stimulus values and screen location). Subjects were instructed to click once with the cursor positioned on the region of the screen corresponding to the target location. The selected stimulus was then presented in the center of the screen, and subjects were allowed to fine-tune their response by moving the cursor around the stimulus wheel before confirming their selection with a second mouse click (Figure 1b).

During the fine-tuning phase, the stimulus presented in the center of the screen varied (in steps of 1°), to allow subjects to precisely match their responses to the target stimulus. The second (confirmation) click was registered only if it occurred within 2–10 s after the trial onset. If no response was entered for 10 s, a dialog box was displayed warning the subject to pay attention and click the “OK” button to continue. On these trials, an accuracy of 0% was registered (for the purpose of calculating the average block accuracy), though they were not included in the main analyses. Similarly, if responses were more than 90° away from the target, a modal dialog box was displayed warning the subject to pay attention. These measures ensured that subjects maintained continuous attention throughout the entire block and slowed down subjects who attempted to rush through the experiment without care.

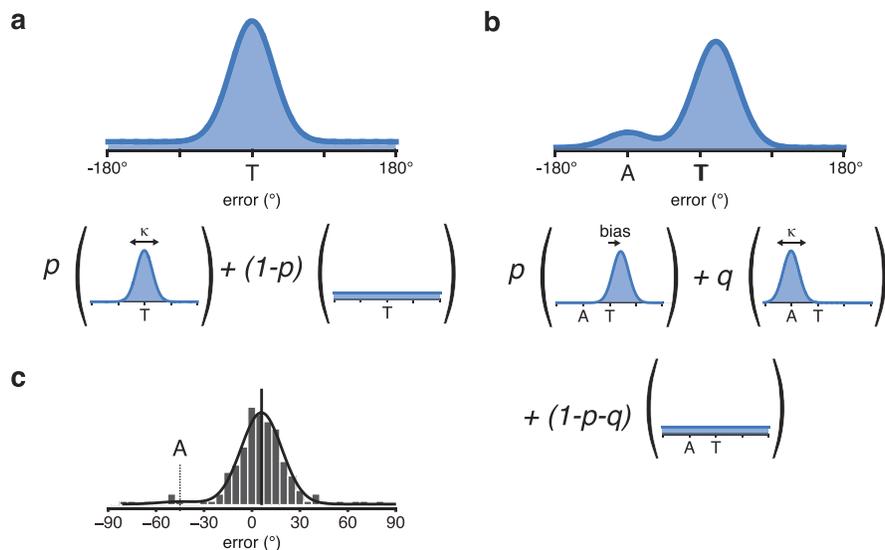


Figure 2. Mixture model-fitting approach. The distribution of errors was modeled as a superposition of probability distributions to account for different types of responses. (a) In the stimulus-match and delayed-match experiments, responses could be concentrated near the target value (von Mises distribution centered at the target value) or fall randomly in any position of the space with equal probability (von Mises distribution with concentration parameter equal to 0, equivalent to a uniform distribution). Two parameters were estimated: the concentration parameter of the responses near the target and the probability of random responses. (b) In the adaptation experiment, responses could be concentrated near the target value (von Mises distribution displaced from the target by a fixed amount), be concentrated near the adaptor value (a stimulus that the subjects were instructed to ignore, modeled as a von Mises distribution centered at the adaptor value), or fall randomly in any position of the space with equal probability (von Mises distribution with concentration parameter equal to 0). Four parameters were estimated: the magnitude of the displacement of responses near the target (bias), the concentration parameter of the responses near the target, the probability of responses near the adaptor, and the probability of random responses. (c) Example of a good fit for a subject in the adaptation experiment. Data are collapsed across both blocks of the adaptation experiment and indicate the existence of a positive (repulsive) bias in relation to the adaptor.

Delayed-match experiment

Each block consisted of 30 trials, with target values sampled uniformly (12° spacing) from the circular space—that is, each target was presented once per

block. On each trial, subjects were presented with a target stimulus in the center of the screen for 2,000 ms, followed by a mask stimulus at the same location for 200 ms (interstimulus interval: 100 ms), followed by a 4,000-ms interval of a blank screen during which no response was allowed (Figure 1c, middle). After this

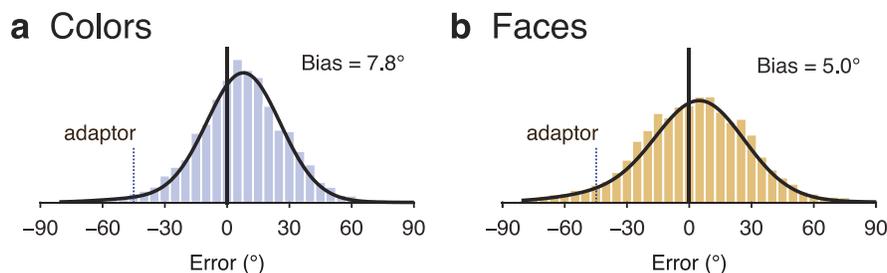


Figure 3. Estimation of adaptation biases at the population level. (a) Group adaptation results for all color trials. The bars provide the histogram of error (5° bins) in responses across all trials and all subjects in the adaptation experiment that used color stimuli. If subjects reported the color of the target stimulus perfectly, all trials would have zero error. Each trial in the adaptation experiment featured a 5-s adaptor stimulus, which in this plot has a relative location of -45° and is indicated with the blue dotted line. The black curve shows the fit to the data provided by the model shown in Figure 2b. As can be seen, the peak of the distribution of error responses is shifted to the right of 0, indicating that subjects had a bias ($M = 7.8^\circ$) in reporting the value of the target stimulus. (b) Similar results for all face trials. As can be seen, the peak of the distribution of error responses is again shifted to the right of 0, indicating that subjects had a bias ($M = 5.0^\circ$) in reporting the value of the target stimulus.

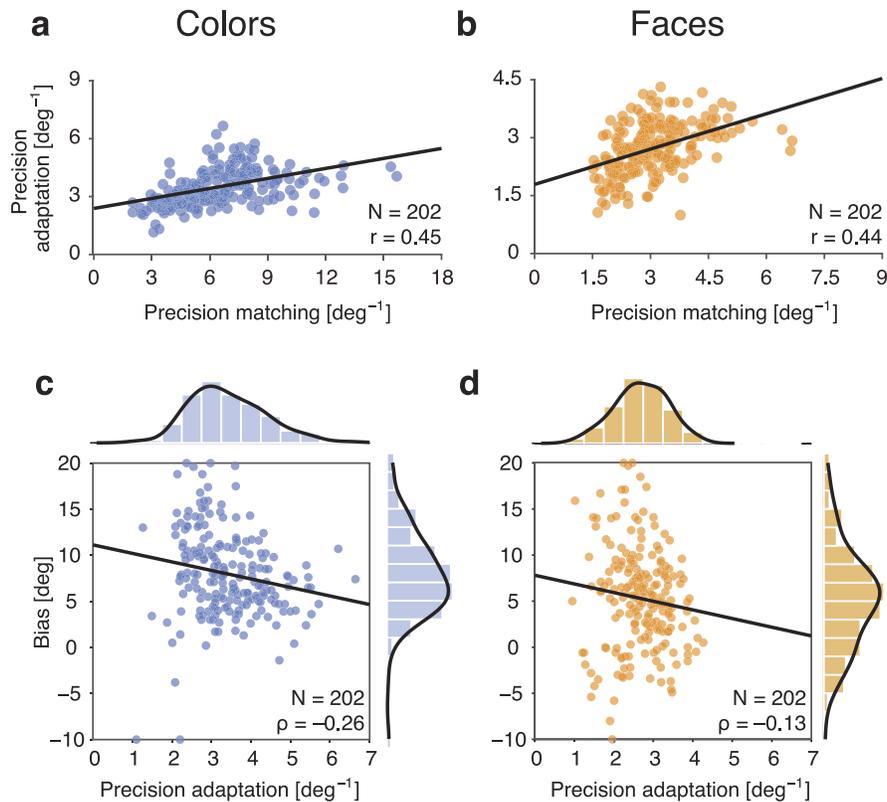


Figure 4. Individual differences in representation precision are negatively correlated with adaptation bias. (a) We fitted data from each subject collapsed across both blocks of the color stimulus-match experiment and across both blocks of the color adaptation experiment, using the mixture-model approach described for Figure 2. The correlation between representation precision in both experiments was $r = 0.45$. (b) Similar to (a), for face stimuli. The correlation between representation precision in the stimulus-match and adaptation experiments was $r = 0.44$. (c) We fitted data from each subject collapsed across both blocks of the color adaptation experiment, using the mixture-model approach described for Figure 2. The Spearman's rank-correlation coefficient between representation precision and adaptation bias was $\rho = -0.26$. (d) Similar to (c), for face stimuli. The Spearman's rank-correlation coefficient between representation precision and adaptation bias was $\rho = -0.13$.

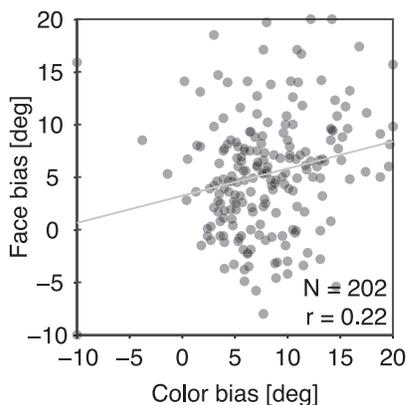


Figure 5. Relationship between color bias and face bias. We fitted data from each subject collapsed across both blocks of the color and face adaptation experiments, using the mixture-model approach described for Figure 2. The Pearson correlation coefficient between adaptation bias for colors and faces was $r = 0.22$.

interval, subjects entered their responses using the same procedure described previously (Figure 1b). The same measures described previously were used to ensure that subjects maintained continuous attention throughout the entire block.

Color masks were checkerboards composed of various colors randomly sampled from within the stimulus set. Face masks were created by synthesizing textures based upon the original face stimuli using the steerable pyramid approach (Portilla & Simoncelli, 2000; Figure 1d). The resulting textures were placed within the outlines of the face stimuli, resulting in a mask that had similar low-level feature properties to the faces (e.g., spatial frequency, line curvature) but was not recognizable as a face.

Adaptation experiment

Each block consisted of 30 trials, with target values sampled uniformly (12° spacing) from the circular

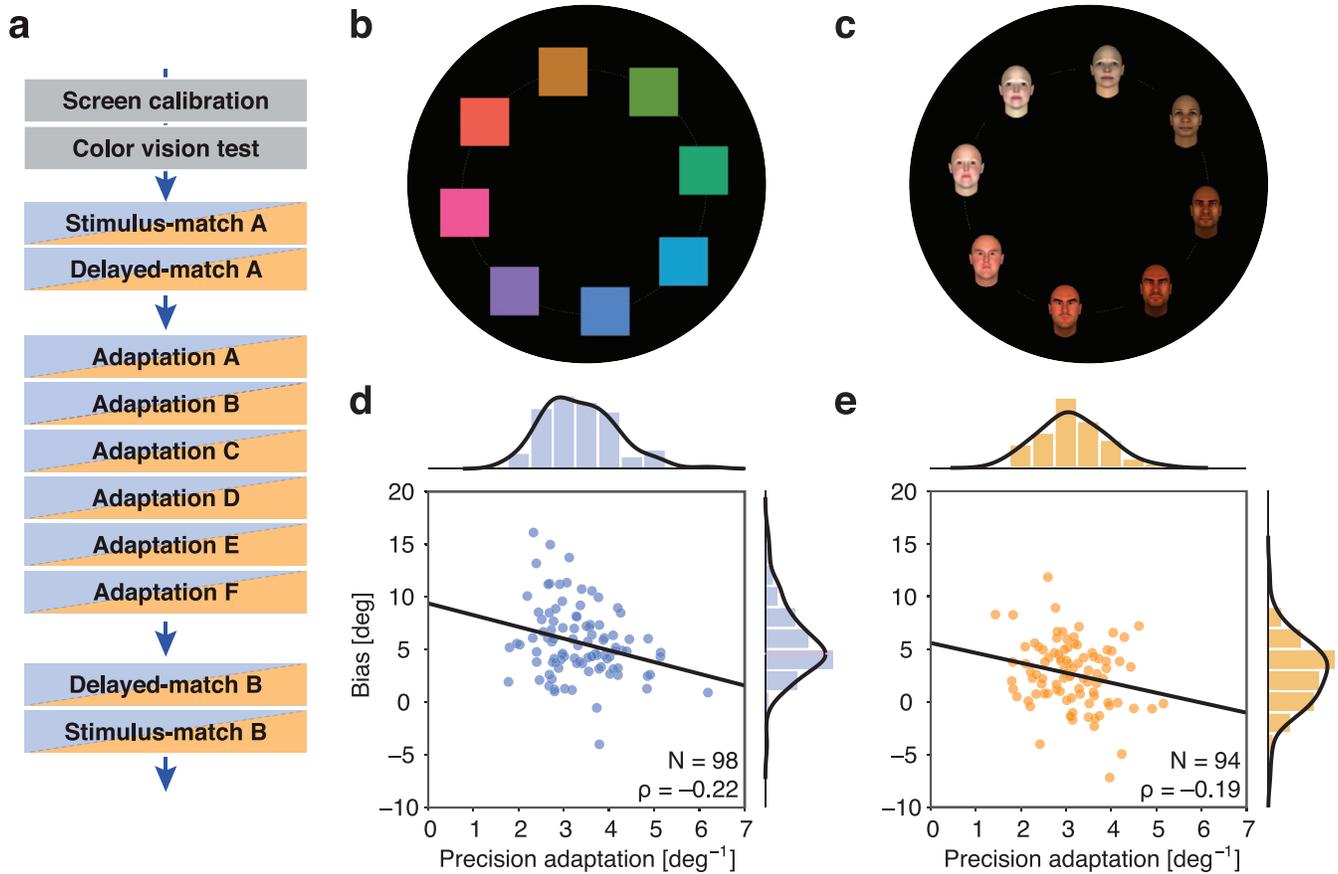


Figure 6. Relationship between representation precision and adaptation bias in the replication experiment. (a) Replication-experiment structure. Subjects completed a battery of tests designed to estimate representation precision, precision after a delay, and precision and bias after adaptation. Each subject completed two blocks of the stimulus-match and delayed-match experiments and six blocks of the adaptation experiment, for either colors or faces. Prior to each experiment, they were presented with instructions and a comprehension quiz, and the first time they performed each block of a given experiment, they also completed five practice trials. Each block of the stimulus-match experiment comprised 15 trials. Each block of the delayed-match and adaptation experiments comprised 30 trials. (b) Color stimuli. A circular space with 360 stimuli was used. Colors were generated to vary in hue but not saturation or luminance (HSL space), though saturation was higher than in the first experiment. (c) Face stimuli. A circular space with 360 stimuli was used. Face stimuli varied in age and gender, each forming one axis of the space. To maximize stimulus differences, the extreme points on each axis also varied in identity. (d) We fitted data from each subject collapsed across all blocks of the color adaptation experiment, using the mixture-model approach described for Figure 2. The Spearman's rank-correlation coefficient between representation precision and adaptation bias was $\rho = -0.22$. (e) Similar to (d), for face stimuli. The Spearman's rank-correlation coefficient between representation precision and adaptation bias was $\rho = -0.19$.

space. Each target was presented once per block. On each trial, subjects were presented with an adaptor stimulus in the center of the screen for 5000 ms, drawn from a position in the stimulus space $\pm 45^\circ$ from the target (adaptors at $+45^\circ$ and -45° were intermixed within a block). The adaptor was followed by a target stimulus for 200 ms (interstimulus interval: 50 ms), which was then followed by a mask stimulus at the same location for 100 ms (interstimulus interval: 100 ms; Figure 1c, bottom). After the mask offset, subjects entered their responses using the same procedure described previously (Figure 1b). Subjects received a warning if their responses were within 10° of the adaptor stimulus position. In these cases, a modal

dialog box was presented indicating that subjects should ignore the adaptor and report the value of the target.

Data analysis

We used statistical methods for circular data (Fisher, 1995). We calculated the error on each trial as the angular deviation on the stimulus wheel between the target value and the response entered. We then used maximum-likelihood estimation to fit the distribution of error values in the circular space. In both stimulus-match and delayed-match experiments, the distribution

of error values was decomposed into two parameters representing a mixture of a uniform distribution (corresponding to random responses) and a von Mises distribution—the circular analog of the Gaussian distribution on a line—centered on the target value. The parameters fitted by this procedure correspond to the probability of guesses, which is proportional to the height of the uniform distribution, and the precision of responses, which is the inverse of the standard deviation of the best-fitting Gaussian distribution (which is in turn related to the concentration parameter of the fitted von Mises distribution; Figure 2a).

In the adaptation experiment, the distribution of error values was decomposed into four parameters representing a mixture of a uniform distribution (corresponding to random responses), a von Mises distribution centered on the adaptor value (corresponding to responses where the subject mistakenly attempted to report the value of the adaptor, due to not seeing or not paying attention to the target), and a von Mises distribution with equal concentration parameter centered *near* the target value. The parameters fitted by this procedure correspond to the probability of guesses, the precision of responses (the inverse of the standard deviation of the best-fitting Gaussian distribution), and the bias, which is the mean of the von Mises distribution centered near the target value (Figure 2b).

Results

We investigated the relationship between variation in precision and adaptation bias across 530 individuals. Each subject performed a series of psychophysics experiments on their personal computers and received compensation for their time (Figure 1a). On each experimental trial, subjects were instructed to report the value of a target stimulus—a color or a face—by clicking and dragging the mouse pointer around a stimulus wheel, allowing a fine adjustment of their responses (Figure 1b). Wheels comprised 360 distinct stimuli varying in hue (colors) or in age and gender (faces; Figure 1d).

Subjects completed three types of experiments for each stimulus class. In the stimulus-match experiment (15 trials), designed to obtain a baseline response precision for each subject, subjects were instructed to select a value on the wheel matching the target stimulus (Figure 1c, top). In the delayed-match experiment (30 trials), designed to estimate subjects' working-memory precision, a target stimulus was presented on the center of the screen for 2.0 s, followed by a 4.0 s delay, after which subjects were to select a value on the wheel matching the target stimulus (Figure 1c, middle). In the adaptation experiment (30 trials), designed to estimate

the magnitude of adaptation biases, an adaptor stimulus was presented in the center of the screen for 5 s followed by a target stimulus $\pm 45^\circ$ away from the adaptor for 200 ms and by a brief mask, after which subjects were to select a value on the wheel matching the target stimulus (Figure 1c, bottom). Subjects performed two separate blocks of each experiment in a session, allowing for tests of measure reliability, and the trials within each experiment were sampled uniformly and in random order from the circular space (Figure 1d). Throughout the entire session, subjects were only allowed to proceed to the next block of the experiment if their accuracy remained above a minimum threshold (see Methods). A total of 328 subjects either were excluded or abandoned the experiment (108 at the color-vision test and 220 during the main experimental trials), leaving 202 subjects for the main analyses (Table 1).

We calculated the error on each trial as the difference between the target value and the response entered ($\varepsilon = \theta_{\text{target}} - \theta_{\text{response}}$), and we fitted the distribution of error values for each subject using a superposition of distributions defined over a circular support ($-180^\circ < \theta < +180^\circ$; Zhang & Luck, 2008; Bays, Catalao, & Husain, 2009). This procedure simultaneously estimates the precision of the error distribution (i.e., the inverse of the standard deviation of the best-fitting Gaussian distribution; Jammalamadaka & Sengupta, 2001) and the probability of random guesses (Figure 2a).

In the adaptation experiment, biases expected to result from adaptors located at $+45^\circ$ and -45° away from the target should have equal magnitude and opposite sign. Thus, we combined the distribution of error values from trials where the adaptor was located at -45° from the target (ε) with the distribution of reflected error values ($-\varepsilon$) from trials where the adaptor was located at $+45^\circ$ from the target. In these experiments, the fitting procedure also estimates two additional parameters: the mean of the error distribution (i.e., the bias induced by the paradigm) and the probability that the selected response matches the adaptor stimulus and not the target stimulus (Figure 2b and 2c).

Adaptation produces a repulsive bias

To confirm the effectiveness of our adaptation paradigm in inducing repulsive biases, we first fitted the adaptation data from all subjects combined (12,120 trials per stimulus class) with a mixture of distributions as described. We observed a positive (repulsive) bias of 7.8° and 5.0° for color and face stimuli, respectively. We also estimated that subjects responded randomly in 2.6% and 4.1% of the trials, and that their responses

	Bias	Concentration	$P(\text{target})$	$P(\text{adaptor})$	$P(\text{random})$
Colors	7.8°	10.3	95.1%	2.3%	2.6%
Faces	5.0°	7.0	90.1%	5.2%	4.1%

Table 4. Parameters from group fit of the adaptation data in the main experiment.

matched the adaptor stimulus in 2.3% and 5.2% of the trials, with the remaining 95.1% and 90.8% of responses concentrated around the target value (Figure 3). These results confirm that the paradigm induced the typical *repulsive* aftereffects, in which subject responses to target stimuli tend to be biased *away* from the preceding, adapting stimulus (Table 4).

Individual differences in adaptation bias and representation precision are negatively correlated

We then fitted the response-bias data for each individual subject. We found that adaptation bias did not significantly differ between blocks (repeated-measures analysis of variance [ANOVA] with block and stimulus class as within-subject factors), $F(1, 201) = 0.0160$, $p = 0.9$, suggesting that the magnitude of adaptation biases is a stable individual characteristic for the duration of the experiment. For that reason, we again fitted the response bias using data from both blocks of each experiment. We then performed a repeated-measures ANOVA to identify the sources of variability in adaptation bias. We observed a small but significant effect of stimulus class, $F(1, 201) = 34.9$, $p < 0.001$, explaining only 6% of the total variance. The effect of mean subject bias was more substantial, $F(201, 201) = 1.56$, $p < 0.001$, accounting for 57% of the total variance in bias values. The distribution of individual subject bias was well fitted with a Gaussian with a mean of 8.0° for colors (95% confidence interval [CI] [0.2°, 18.8°]) and 5.3° for faces (95% CI [−4.9°, 17.1°]). These results are consistent with an individual difference in adaptation bias that is present across face and color stimuli.

We next extracted the width (i.e., precision) of the error distribution for each subject, separately for each block of each experiment. Again, because precision did not significantly differ between blocks (repeated-measures ANOVA with block and stimulus class as within-subject factors), $F(1, 201) = 0.86$, $p = 0.36$, we combined data from both blocks of each experiment. We then performed a repeated-measures ANOVA to identify the sources of variability in response precision. We observed a significant effect of stimulus class, $F(1, 201) = 170.1$, $p < 0.001$, accounting for 16.5% of the total variance, with response precision being higher for colors. We also observed a significant effect of mean

subject precision, $F(201, 201) = 3.28$, $p < 0.001$, accounting for 64% of the total variance. Furthermore, response precision estimated from stimulus-match trials was well correlated with precision estimated from the adaptation experiment (color: $r = 0.45$, $p < 0.001$; faces: $r = 0.44$, $p < 0.001$; Figure 4a and 4b). This correlation is possibly due to shared sources of variability between experiments. Note, however, that the variance associated with reporting the target on stimulus-match trials is smaller than that associated with reporting the target in the adaptation experiment, possibly due to the fact that the latter is presented briefly and involves a short-term-memory component.

We then asked if subjects with a lower representation precision are more or less prone to adaptation biases. We investigated the statistical relationship between average representation precision and adaptation bias, both estimated from the adaptation experiment. We observed that, across subjects, there was a negative correlation between mean precision and bias (Spearman's rank correlation)—color: $\rho = -0.26$, jackknife resampling 95% CI [−0.28, −0.25], $p < 0.001$; faces: $\rho = -0.13$, jackknife resampling 95% CI [−0.15, −0.12], $p < 0.001$ (Figure 4c and 4d). This suggests that subjects with lower representation precision exhibit larger biases away from the adapting stimulus, consistent with predictions from a Bayesian observer model.

We also examined the relationship between color bias and face bias across subjects (Figure 5), finding a weak positive correlation in these scores (Pearson correlation), $r = 0.22$, jackknife resampling 95% CI [−0.21, −0.23], $p < 0.001$. A possible source for this relationship is the shared effect of individual differences in precision. We cannot, however, be certain that this is the entire explanation, as the correlation between face and color bias values remains in partial correlations that attempt to account for individual variation in precision (partial Pearson correlation), $r = 0.18$, jackknife resampling 95% CI [−0.17, −0.19], $p < 0.001$. We considered the possibility that individual differences in reaction time could account for individual differences in adaptation magnitude. The magnitude of perceptual bias decays following cessation of the adaptor (Greenlee, Georgeson, Magnussen, & Harris, 1991). Perhaps subjects who respond quickly tend to experience more bias due to less decay. However, this prediction was not confirmed. For faces, we found no significant relationship between response bias and average response time (Spearman's rank correlation), $\rho = 0.07$, $p = 0.30$. For colors, the opposite pattern was found: Subjects who responded quickly experienced less bias (Spearman's rank correlation), $\rho = 0.20$, $p = 0.0047$. We do not have a specific mechanism (neural or otherwise) to offer as the basis for this small correlation

in induced adaptation bias between materials across subjects.

Replication experiment

We wished to replicate the observed relationship between representation precision and adaptation bias with more trials per subject. An additional group of 472 people were recruited through Amazon Mechanical Turk. From this set, 89 were excluded or abandoned at the color-vision test and 191 during the main experimental trials, leaving 192 subjects for the replication analyses (Tables 2 and 3). Subjects performed a slightly modified version of our experiment: two blocks of the stimulus-match experiment (30 trials), two blocks of the delayed-match experiment (60 trials), and six blocks of the adaptation experiment (180 trials). Each subject performed the experiments for only one of the two stimulus classes (faces: 94 subjects; colors: 98 subjects; Figure 6a). A slightly more saturated version of the stimuli was used in an attempt to increase overall performance (Figure 6b and 6c).

After fitting the data using the same approach as described previously (Figure 2), we analyzed the sources of variability in response precision. Across both color and face experiments, experimental block explained 0.9% of variance, $F(5, 950) = 3.72$, $p = 0.002$, and subject explained 53.5%, $F(190, 950) = 5.873$, $p < 0.001$. We then analyzed the sources of variability in adaptation bias. Similarly, experimental block explained 0.8% of variance, $F(5, 950) = 2.75$, $p = 0.02$, and subject explained 47.1%, $F(190, 950) = 4.51$, $p < 0.001$. To investigate whether the inverse relationship between representation precision and adaptation biases was replicated, we extracted the bias and precision values for each subject using data from all blocks and calculated the correlation between these two quantities. We again observed that, across subjects, there was a negative correlation between mean precision and bias—color: $\rho = -0.22$, jackknife resampling 95% CI $[-0.25, -0.20]$, $p < 0.001$; faces: $\rho = -0.19$, jackknife resampling 95% CI $[-0.22, -0.17]$, $p < 0.001$ (Figure 6d and 6e).

Discussion

We investigated the relationship between individual differences in adaptation bias and response precision for colors and faces. In two cohorts of 202 and 192 subjects recruited through Amazon Mechanical Turk, precision and adaptation bias were found to be stable properties of the observer, with substantial variance in the measurements arising from between-subjects dif-

ferences. Across experiments and materials, we found that lower response precision in an individual was associated with greater perceptual bias.

Two hypotheses have been proposed to reconcile repulsive biases with the Bayesian framework. First, adaptation may induce asymmetries in the likelihood function (Stocker & Simoncelli, 2006b). Combined with a symmetric prior around the adaptor, an asymmetric likelihood can produce estimates that are shifted away from an observer's prior (Wei & Stocker, 2015). Alternatively, adaptation may affect the prior itself, according to the expectation that stimulus statistics in recent history match the statistics accumulated over a longer, more distant past (Chopin & Mamassian, 2012). After a prolonged stimulus presentation, a Bayesian observer model would predict biases “away” from this stimulus to maintain balanced statistics in recent history. Note that in the extreme case of infinite variance, the first proposal predicts attraction toward the adaptor (because the likelihood would no longer be asymmetric, estimates would converge again toward the prior, which is centered around the adaptor), while the second continues to predict a bias away from the adaptor and toward the long-term stimulus prior.

A notable finding of our study was the substantial individual variation in the magnitude of perceptual adaptation. This measurement was reproducible across blocks within a testing session. As we did not measure across testing sessions, our measurement likely also contains a component of state variation (although prior studies of individual variation in blur adaptation suggest this component is small; Vera-Diaz, Woods, & Peli, 2010). Stable variation across individuals explained almost an order of magnitude more variance in bias measures than did variation within subject across stimulus type (face and color). This indicates a shared mechanism of variation in adaptation magnitude. We find that individual differences in sensory precision provide one such mechanism.

Our results are consistent with a model in which each subject is a Bayesian observer, each of whom differs in the fidelity with which they represent sensory input. We estimated each subject's precision by measuring response variability. A Bayesian interpretation assumes that response variability in turn reflects (to some degree) individual differences in the precision of sensory encoding. Although quite reproducible, the magnitude of correlation between response precision and adaptation bias was small (accounting for approximately 5% of between-subjects variability in adaptation bias). Response precision is therefore an imperfect proxy for sensory precision—or other factors contribute to the substantial between-subjects variation in the adaptation effect that is shared across stimulus types.

We recruited subjects through the Amazon Mechanical Turk platform and conducted our experiments using a custom-built website. In addition to allowing a larger sample size, web-based experiments improve subject diversity (Woods, Velasco, Levitan, Wan, & Spence, 2015). A major challenge of web-based data collection is that subjects may be motivated not to provide a high level of performance but instead to complete the tasks as quickly as possible to obtain payment. To meet this challenge, our online test was designed so that it would be completed most rapidly if the subject produced accurate responses. Additional measures to improve data quality included quizzes to ensure comprehension of the instructions and paying proportionally large bonuses for compliant subjects. While the exclusion of a large proportion of subjects with low accuracy limited the range of precision values we could have measured from our population, we regarded this as an acceptable compromise to exclude subjects who made no attempt to achieve the goals of the measurement.

To conclude, our study provides support for the view that adaptation biases conform with a key property of Bayesian inference: Greater perceptual bias should be expected in conditions of lower response precision. While we cannot adjudicate between the conflicting proposals that adaptation affects the prior (Chopin & Mamassian, 2012) or the likelihood (Wei & Stocker, 2015)—or both—our results nonetheless rule out two possibilities: first that adaptation is not modulated by precision (which would suggest that adaptation bias is not a result of optimal inference), and second that the perceptual prior is centered on the adaptor when noise is symmetric (which would predict smaller biases in conditions of low precision). Future studies explicitly manipulating expectations independently from sensory evidence, or testing for asymmetries on response distribution, may be conducted to provide further mechanistic evidence.

Keywords: visual adaptation, individual differences, face perception, color perception, Bayesian inference, Amazon Mechanical Turk

Acknowledgments

This research was supported by National Institutes of Health Grants EY021717 (STS and GKA). We thank Alan Stocker for many helpful discussions. Author contributions: Conceptualization: MGM, GKA, and SLT-S; methodology: MGM and GKA; software and formal analysis: MGM, MVC, and MSZ; visualization: MGM and GKA; supervision: GKA and SLT-S; original draft: MGM and GKA;

review and editing: MGM, GKA, and SLT-S; funding acquisition: GKA and SLT-S.

Commercial relationships: none.

Corresponding author: Marcelo G. Mattar.

Email: mmattar@princeton.edu.

Address: Princeton Neuroscience Institute, Princeton University, Princeton, NJ, USA.

References

- Bar, M. (2004). Visual objects in context. *Nature Reviews Neuroscience*, 5(8), 617–629.
- Bayes, T. (1763). An essay towards solving a problem in the doctrine of chances. *Philosophical Transactions of the Royal Society of London*, 53, 370–418.
- Bays, P. M., Catalao, R. F., & Husain, M. (2009). The precision of visual working memory is set by allocation of a shared resource. *Journal of Vision*, 9(10):7, 1–11, <https://doi.org/10.1167/9.10.7>. [PubMed] [Article]
- Chopin, A., & Mamassian, P. (2012). Predictive properties of visual adaptation. *Current Biology*, 22(7), 622–626.
- Clifford, C. W. (2002). Perceptual adaptation: Motion parallels orientation. *Trends in Cognitive Sciences*, 6(3), 136–143.
- Fischer, J., & Whitney, D. (2014). Serial dependence in visual perception. *Nature Neuroscience*, 17(5), 738–743.
- Fisher, N. I. (1995). *Statistical analysis of circular data*. Cambridge, UK: Cambridge University Press.
- Gibson, J. J., & Radner, M. (1937). Adaptation, after-effect and contrast in the perception of tilted lines. I: Quantitative studies. *Journal of Experimental Psychology*, 20(5), 453–467.
- Greenlee, M. W., Georgeson, M. A., Magnussen, S., & Harris, J. P. (1991). The time course of adaptation to spatial contrast. *Vision Research*, 31(2), 223–236.
- Jammalamadaka, S. R., & Sengupta, A. (2001). *Topics in circular statistics* (Vol. 5). City, ST: World Scientific.
- Knill, D. C., & Richards, W. (1996). *Perception as Bayesian inference*. City, ST: Cambridge University Press.
- Lafer-Sousa, R., Hermann, K. L., & Conway, B. R. (2015). Striking individual differences in color perception uncovered by “the dress” photograph. *Current Biology*, 25(13), R545–R546.
- Levinson, E., & Sekuler, R. (1976). Adaptation alters

- perceived direction of motion. *Vision Research*, 16(7), 779–781.
- Portilla, J., & Simoncelli, E. P. (2000). A parametric texture model based on joint statistics of complex wavelet coefficients. *International Journal of Computer Vision*, 40(1), 49–70.
- Singular Inversions. (2010). FaceGen [software]. Toronto, ON: Author.
- Stocker, A. A., & Simoncelli, E. P. (2006a). Noise characteristics and prior expectations in human visual speed perception. *Nature Neuroscience*, 9(4), 578–585.
- Stocker, A., & Simoncelli, E. P. (2006b). Sensory adaptation within a Bayesian framework for perception. *Advances in Neural Information Processing Systems*, 18, 1289–1296.
- Summerfield, C., & Egner, T. (2009). Expectation (and attention) in visual cognition. *Trends in Cognitive Sciences*, 13(9), 403–409.
- Vera-Diaz, F. A., Woods, R. L., & Peli, E. (2010). Shape and individual variability of the blur adaptation curve. *Vision Research*, 50(15), 1452–1461.
- von Helmholtz, H., & Southall, J. P. (1924). *Treatise on physiological optics. Vol. III*. New York, NY: Dover. (Original work published 1866)
- Wei, X.-X., & Stocker, A. A. (2015). A Bayesian observer model constrained by efficient coding can explain “anti-Bayesian” percepts. *Nature Neuroscience*, 18(10), 1509–1517.
- Woods, A. T., Velasco, C., Levitan, C. A., Wan, X., & Spence, C. (2015). Conducting perception research over the internet: a tutorial review. *PeerJ*, 3, e1058.
- Zhang, W., & Luck, S. J. (2008, May 8). Discrete fixed-resolution representations in visual working memory. *Nature*, 453(7192), 233–235.