Perceptual learning increases orientation sampling efficiency

Denise Moerel

Donders Institute for Brain, Cognition and Behavior, Radboud University, Nijmegen, The Netherlands

Sam Ling

Department of Psychological and Brain Sciences, Boston University, Boston, MA, USA

Janneke F. M. Jehee

Donders Institute for Brain, Cognition and Behavior, Radboud University, Nijmegen, The Netherlands

Visual orientation discrimination is known to improve with extensive training, but the mechanisms underlying this behavioral benefit remain poorly understood. Here, we examine the possibility that more reliable task performance could arise in part because observers learn to sample information from a larger portion of the stimulus. We used a variant of the classification image method in combination with a global orientation discrimination task to test whether a change in information sampling underlies training-based benefits in behavioral performance. The results revealed that decreases in orientation thresholds with perceptual learning were accompanied by increases in stimulus sampling. In particular, while stimulus sampling was restricted to the parafoveal, inner portion of the stimulus before training, we observed an outward spread of sampling after training. These results demonstrate that the benefits of perceptual learning may arise, in part, from a strategic increase in the efficiency with which the observer samples information from a visual stimulus.

Introduction

Perceptual learning refers to a training-related improvement in performance on a perceptual task (Gibson, 1963), which is sometimes specific to certain stimulus factors, such as retinal location (Fiorentini & Berardi, 1981; Kapadia, Gilbert, & Westheimer, 1994; Karni & Sagi, 1991), stimulus orientation (Fahle, 1997; Fahle & Edelman, 1993; Fiorentini & Berardi, 1980; Jehee, Ling, Swisher, van Bergen, & Tong, 2012; Schoups, Vogels, & Orban, 1995; Shiu & Pashler, 1992), spatial frequency (De Valois, 1977; Fiorentini & Berardi, 1980), or motion direction (Ball & Sekuler, 1982, 1987; Zanker, 1999). The mechanisms underlying such training-based improvements in behavioral performance, however, remain poorly understood.

One mechanism that might underlie perceptual learning is a sharpening of tuning at the neural population level, which would result in more reliable perceptual estimates. While some evidence supporting training-based changes in receptive field structure has been found (Crist, Li, & Gilbert, 2001; Jehee et al., 2012; Schoups, Vogels, Qian, & Orban, 2001), an alternative (though not mutually exclusive) possibility is that perceptual learning increases the efficiency with which sensory evidence is sampled. That is, a more reliable estimate of a stimulus could be obtained by increasing the portion of the stimulus from which the observer samples information in order to make a decision. In a more general context, this selective sampling of sensory evidence was proposed by Dakin (2001), who tested the effect of external orientation noise on perceptual thresholds to determine the size of the sampling region in a global orientation discrimination task. The external orientation noise was obtained by varying the local orientation of individual Gabor elements within a stimulus array, which had an overall global orientation. The results suggested that subjects' decision-making relied on only portions of the orientation stimulus, the size of which was determined by the total number of patches within the array.

Although the study by Dakin (2001) showed that subjects rely on a subsample of a stimulus when making perceptual decisions, it remains unclear whether the sampling efficiency is improved with training. In the present study, we developed a novel paradigm to investigate whether a change in subsampling efficiency could be a possible mechanism underlying perceptual learning on an orientation discrimination task. We predicted that subsampling would become more efficient after training, as shown by an increase in the
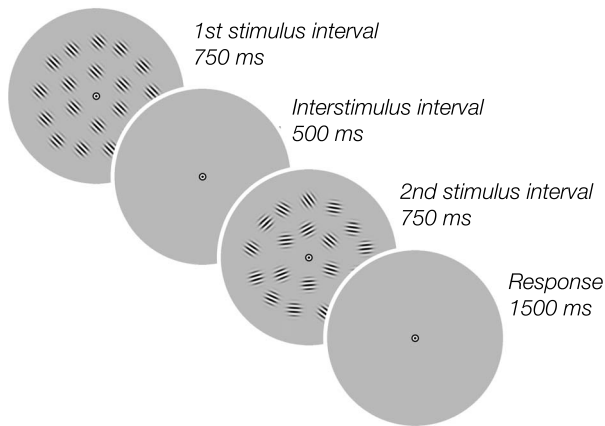
Figure 1. Schematic representation of a single trial. Participants maintained fixation on a central bull's-eye throughout the entire trial, and in a given trial were shown a pair of stimuli, in sequence. The stimulus was composed of Gabor elements, which varied in their local orientation variability, as well as their global mean orientation. In the first interval, all Gabor elements had zero local variability, such that all elements were aligned perfectly along the global orientation (45° or 135°). In the second interval, each Gabor element had its local orientation sampled randomly from a Gaussian distribution with a standard deviation of 22°, and with a mean global orientation offset that was staircased from trial to trial. Participants reported whether the global orientation of the second stimulus was rotated clockwise or counterclockwise relative to the first. The size and contrast of the Gabor patches in the Figure are exaggerated purely for illustrative purposes, and not drawn to scale.

size of the sub-sample, and a decrease of orientation thresholds after training. To test this prediction, we made use of the external orientation noise of the stimulus and behavioral responses of the participant, in order to reconstruct a spatial map, referred to as a decision template, revealing the perceptual decision weights of individual elements within the stimulus. This technique, a variant of the classification image method (Eckstein & Ahumada, 2002; Gold, Murray, Bennett, & Sekuler, 2000; Nagai, Bennett, & Sekuler 2007), allowed us to infer which elements of the stimulus array drove the perceptual decision, thereby enabling us to compare the pattern of sampling before and after training on an orientation discrimination task.

Methods

Participants

Six healthy adult volunteers (aged 22–28 years, all female) with normal or corrected-to-normal vision

participated in this experiment. All participants gave informed written consent. The experiment was approved by the Radboud University Institutional Review Board, and was conducted in accordance with the Declaration of Helsinki. Four of the participants were naive to the aims of the study; one of the remaining two subjects is an author (DM).

Apparatus

The participants viewed the stimuli binocularly on a 21-in. gamma linearized CRT monitor at a resolution of 800 × 600 pixels, and a distance of 57 cm, in a dark room. The mean luminance of the display was 60.15 cd/m^2. In order to prevent the use of the rectangular frame of the monitor as a reference for orientation discrimination, a black annulus was used to cover the edges of the frame, resulting in a circular field of view with a diameter of 29 cm. A chin rest was used to stabilize the participants' heads and prevent motion. The stimuli were created using the Psychophysics Toolbox (PTB3) extensions of Matlab software (Brainard, 1997; Kleiner et al., 2007; Pelli, 1997).

Stimuli

Each stimulus consisted of 63 full-contrast Gabor patches, placed on a gray background (Figure 1). The Gabor patches were positioned in an annulus around fixation. Each individual Gabor element consisted of a 4 cycles per degree sinusoidal grating, windowed by a circular Gaussian with a full width at half the maximum height of 0.5° of visual angle. The initial stimulus array had an outer diameter of 17°, and an inner diameter of 3°. Individual patch positions within the array were defined in terms of polar coordinates, forming a polar grid consisting of four rings (with a thickness of 3.5°) around the fixation dot (Figure 1). We first placed the patches on this grid, and then introduced random spatial jitter, so that spatial predictability was decreased while ensuring an even distribution of the patches and minimizing overlap between patches. The random jitter consisted of uniformly distributed noise that was applied to the two polar coordinates independently.

Procedure

The same task was performed in all sessions. Observers were instructed to maintain fixation on a central bull's-eye target throughout the trial. In each trial, two successive stimulus arrays were briefly presented (750 ms), separated by a 500-ms interval

(Figure 1). After the presentation of the second stimulus array, participants were given 1500 ms to report whether the global orientation of the second (comparison) stimulus was rotated clockwise or counter clockwise with respect to the first (base) stimulus. All Gabor patches comprising the first stimulus array had identical orientations and served as an external discrimination reference. In contrast, the second stimulus array contained orientation noise. That is, the orientations of the individual patches within the second array were drawn from a Gaussian distribution with a different mean (i.e., global orientation) from the orientation of the first stimulus, and with a standard deviation of 22° (see Girshick, Landy, & Simoncelli, 2011, for a similar procedure). The task of the participants was to judge whether the global orientation of the second stimulus was rotated clockwise or counterclockwise relative to the first stimulus. While the level of orientation noise was kept constant, the (global) orientation of the comparison stimulus was adjusted by means of a two-down one-up staircase procedure, converging on global orientation thresholds of 70.7% accuracy (Leek, 2001; Treutwein, 1995). The use of a staircase ensured a stable task difficulty over all sessions, therefore ruling out changes in sampling strategy due to task difficulty. For each session, the initial staircase value was based on the threshold acquired in the previous session (i.e., the previous threshold +7°; this latter value was added to ensure that the staircases readily converged to the orientation discrimination threshold). Each 1-hr session consisted of 12 staircases (composed of 50 trials per staircase). Given the large standard deviation of the Gaussian noise distribution, some patches had a rotation larger than 90° on individual trials (this occurred on 0.25% of all trials). Excluding these potentially ambiguous trials from our analyses did not greatly affect any of our results.

Each participant performed three different experimental stages in succession: (1) a pretraining thresholding stage, (2) a training stage, and (3) a posttraining thresholding stage. These stages, which were spread over at least 16 days, were preceded by a practice session on the first day in order to get acquainted with the task. Both the pre- and posttraining thresholding stages consisted of three sessions, and at least 10 training sessions were performed. All sessions were planned on consecutive days, and at approximately the same time each day.

Thresholding sessions

In the pre- and posttraining thresholding sessions, we measured orientation discrimination thresholds and decision templates for two different base orientations (at 45° and 135°). The order of these two base orientations was counterbalanced over blocks, so that one block (i.e., staircase) consisted of a single base orientation. Subjects completed three pretraining and three posttraining thresholding sessions, with a total of 1,800 trials completed over three sessions, resulting in 900 trials (18 staircases) per base orientation. On 2.6% of all trials, participants failed to respond before the end of the response window. These trials were excluded from further analysis.

Training sessions

The task used for the training sessions was identical to the task used in the thresholding sessions, with the exception that participants performed the task on only one of the two base orientations. This trained orientation was counterbalanced over participants. The participants received auditory feedback for every correct response in the training sessions. For motivational purposes, participants were presented with information about their threshold at the end of each block, and their progress across different sessions was shown at the end of every session. The participants trained for a minimum of 10 days, resulting in a total of at least 6,000 training trials per participant.

Thresholds

Just noticeable differences (JNDs), the smallest detectable differences in degrees between the base stimulus and the comparison stimulus, were used as an orientation discrimination threshold measure. For the thresholding sessions, the data of three sessions was combined in order to get a single JND measure per base orientation; we used the data from three sessions combined because of the fairly large number of trials needed to obtain reliable decisional weight estimates. A three-way repeated measures ANOVA on the JNDs with session (1, 2, or 3), orientation (trained or untrained), and training (pre or post) as factors, indicated that no reliable learning occurred over the course of these three thresholding sessions (main effect of session $F(2, 4) = 0.505$, $p = 0.638$). For the training sessions, one JND per session was acquired. In order to obtain these JNDs, a cumulative Gaussian psychometric function was fit to the behavioral data, using a maximum likelihood criterion (Figure 2). The standard deviation of this fitted psychometric function served as the JND. Computing JNDs from the last n reversals in each staircase gave similar results (where n was obtained by first removing the first three to four reversals of each staircase and then averaging across the remaining data).

Figure 2. Example psychometric function from one naive participant, across conditions. The relative size of the data points indicates the number of responses combined in each point. In order to obtain these JNDs, a cumulative Gaussian psychometric function was fit to the behavioral data, using a maximum likelihood criterion. The standard deviation of this fitted psychometric function served as the JND.

Decision templates

To measure decision templates, we used a variant of the classification image method (Eckstein & Ahumada, 2002; Gold et al., 2000). This technique uses reverse correlation between the behavioral choices of the participant and the properties of the noisy stimulus in order to infer which specific elements of the stimulus determine the decision of the observer. The rationale behind this procedure is as follows. Because the orientations of the individual patches within the comparison stimulus were drawn from a Gaussian distribution with a large standard deviation (22°), on a given trial some patches in the comparison stimulus were shifted in a direction opposite to the global orientation rotation. If the observer relied most on these patches for their decision, then they might lead the participant to erroneously believe that the global stimulus orientation was shifted in the opposite direction, causing an incorrect response. Such a bias will reveal itself when comparing the presented orientations between correct and incorrect trials, providing insight into which elements of the stimulus display consistently drove decisions over trials. To arrive at these decision templates, we first subtracted,

for each trial, the global orientation from individual patch orientations (this ensured an unbiased comparison of the decision template between conditions). We then calculated the mean orientation across trials for each individual (recentered) Gabor patch, and compared between correct and incorrect trials. This was done for the different base orientations (45° and 135°), and rotations (clockwise and counterclockwise) independently, resulting in four decision templates. After correcting for the direction of the effect, the decision templates were averaged over different rotations. This procedure resulted in two decision templates (i.e., one for each base orientation) for each participant. Because the trained orientation was counterbalanced over participants, the decision templates were subsequently flipped when subjects trained on a base orientation of 135°, resulting in final decision templates centered on a (trained) base orientation of 45° for all subjects. Similar procedures were performed for the untrained base orientation.

Some of our analyses focused on a training-based change in the size of the decision window, quantified as the slope of an exponential decay function. To determine significance for the altered slope due to training, we used a permutation test implemented as follows. For each subject and eccentricity, we randomly shuffled the data between pre- and post-training thresholding sessions, fitted the exponential decay function, calculated the difference in slope between the two sessions, and repeated this 10,000 times. This procedure effectively cancels out any potential effect of training, revealing the range of changes in slope that is expected without any effects of training. This randomization procedure was performed for the trained and untrained orientation separately.

Results

Perceptual performance improved substantially over the course of training, with orientation discrimination thresholds decreasing to roughly two-thirds of their original value (Figure 3). To quantify the effects of training, we obtained discrimination thresholds for the trained as well as an orthogonal, untrained orientation, both before and after the training sessions. A repeated-measures ANOVA performed on these thresholds, with training (pre- vs. posttraining) and orientation (trained vs. untrained orientation) as factors, revealed a significant main effect for training, $F(1, 5) = 12.332$, $p = 0.017$. No significant effect was found for orientation, $F(1, 5) = 2.558$, $p = 0.171$, nor for the interaction between training and orientation, $F(1, 5) = 0.154$, $p = 0.711$. Further analyses indicated significant training-

Figure 3. Orientation discrimination thresholds over time. Mean orientation thresholds changed substantially with training for both the trained and untrained orientation, $t(5) = 3.967$, $p = 0.011$, and $t(5) = 2.486$, $p = 0.055$, respectively. Shaded area and error bars correspond to ± 1 *SEM*. In this and subsequent figures, $\triangle$ = Nonnaive author, $\triangleright$ = Nonnaive participant.

based benefits for the trained orientation, $t(5) = 3.967$, $p = 0.011$, and a trend for the untrained orientation, $t(5) = 2.486$, $p = 0.055$.

What mechanisms might underlie this improvement? Here, we examine the hypothesis that training-based benefits might arise, in part, due to increased spatial efficiency with which the observer samples from the stimulus in order to make a decision. To quantify the observer's decisional efficiency across space, we used a variant of the classification image method (Eckstein & Ahumada, 2002; Gold et al., 2000). Specifically, we reverse correlated the behavioral choices of the participant with the properties of the noisy stimulus (see Methods). This allowed us to infer which specific elements of the stimulus display consistently drove decisions over trials. These elements were captured in a spatial decision template, which quantified the decisional weight placed on individual stimulus elements in the observer's decision. Decision templates were obtained separately for the trained and untrained orientations, and for pre- and posttraining sessions.

Previous work has shown that observers tend to sample information from only a small portion of an image when making decisions. What changes, if any, occur in this decisional template with training? Before training, patches with the highest weight values were mostly restricted to the inner portions of the stimulus,

indicating that subjects relied mostly on a small portion of patches (Dakin, 2001) located close to the fovea. Over the course of training, however, the spatial distribution of decision weights appeared to spread to more peripheral locations, increasing the spatial extent across which a perceptual decision seemed to be driven by (Figure 4a). To quantify the effect of training on the decision templates, we averaged the templates across iso-eccentric rings. Specifically, we first divided the stimulus into four bins, or rings, of comparable eccentricity (Ring 1: 3°–6.5°, Ring 2: 6.5°–10°, Ring 3: 10°–13.5°, Ring 4: 13.5°–17°), and then averaged the decision template values across all patches within each eccentricity ring. This resulted in a single orientation difference score, referred to as the decision weight, per ring. To quantify the drop-off in decisional sampling as a function of eccentricity, an exponential decay function was fitted to the decision weights per ring averaged across all observers, with the slope of this function used as a measure of eccentricity (Figure 4b). A permutation test revealed a significant decrease in the slope of the fitted functions after training, both for the trained orientation ($p < 0.001$) and the untrained orientation ($p < 0.001$). Thus, there was a reliable outward spread of the decision template with training, consistent with the training-based reductions in orientation discrimination threshold.

How does the training-based change in decision template emerge across training sessions? To evaluate the effect over the course of training, we calculated decision templates for each of two consecutive training sessions and assessed the templates over time (Figure 5a). This revealed a gradual spread in the decision template with training. The effect of eccentricity is summarized in Figure 5b, which shows the decision weights per ring over consecutive training sessions. The slope of the fitted exponential decay function became consistently shallower with each training session, consistent with the gradual decrease in behavioral discrimination thresholds. Indeed, correlating the training-based benefits in behavioral performance (calculated relative to the pretraining thresholds, for each of two consecutive sessions) with the slope of the exponential decay function revealed a significant link between the two ($r = 0.975$, $p < 0.001$).

Discussion

Previous studies have shown that the visual system is limited in terms of processing efficiency, making it plausible that training would target this efficiency. For example, change blindness studies have shown that observers are often not able to detect changes to a scene between views, even though the observers typically feel

Figure 4. Training-based improvements in decision templates. (A) Decision templates for both the trained and untrained orientation, pre- and posttraining. In this and subsequent figures, each pixel reflects the decisional weight placed on the image element at that location (see Methods), while red depicts stimulus regions that are congruent, and blue those that are incongruent with the participant's response. (B) Mean patch orientation averaged across rings of equal eccentricity. Bars indicate the mean decision weight per ring. The exponential decay function is shown in red. Training significantly reduced the slope of the fitted function for both the trained orientation ($M = -2.8$, $SE = 0.9$; $p < 0.001$), and the untrained orientation ($M = -1.7$, $SE = 0.9$; $p < 0.001$).

as though they are able to see the scene in great detail (Rensink, O'Regan, & Clark, 1997; Simons & Levin, 1997; Simons & Rensink, 2005). This suggests that instead of forming a complete representation of their surroundings, observers tend to select only certain parts of the visual scene for further processing. In the present study, we sought to directly examine the spatial sampling efficiency with which individual visual stimuli are sampled from. In particular, we investigated whether an increase in the spread of spatial sampling of a stimulus could be one of the mechanisms underlying

perceptual learning. The classification image method (Eckstein & Ahumada, 2002; Gold et al., 2000) was used to construct the decision templates of the participants before and after training on an orientation discrimination task, by correlating the noisy stimulus and the response of the participant. While the orientation thresholds showed a significant decrease with training, the decision templates showed an increase in sample size. The decision weight was restricted to the inner ring before training, but this distribution of weights spread outwards as training


Figure 5. (A) Decision templates over training sessions. The decision weights gradually spread with training. Because the number of trials in one training session was insufficient to obtain a reliable decision template, the data was combined across two training sessions. As some participants performed more than 10 training sessions, all sessions after eight were combined into a single decision template. (B) Data summarized for eccentricity. The slope of the fitted exponential decay function gradually becomes shallower over time (slopes over time, averaged across observers: -1.71, -1.53, -1.33, -1.23, and -1.06).

progressed. These results show that perceptual learning on this orientation discrimination task is accompanied by an increase in spatial sampling efficiency. After training, more patches were used to make a decision, resulting in improved perceptual performance.

The observation that the decision weights were largely restricted to the inner ring before training is in line with the well-established eccentricity effect (Carrasco, Evert, Chang, & Katz, 1995; Carrasco & Frieder, 1997). It has been shown that targets presented near the fixation point are processed more accurately than peripheral targets. Furthermore, orientation discriminations thresholds have been shown to depend on stimulus eccentricity, showing an advantage for orientation processing in the fovea (Paradiso & Carney, 1988). Therefore, it makes sense that this part of the stimulus has a larger influence on the decision, although our results demonstrate that subjects can learn to depend less on the foveal presentation with training, leveraging peripheral information to aid in performing a task. This finding is consistent with work by Dobres and Seitz (2010), which used a similar classification image approach to demonstrate that perceptual training can increase the spatial extent with which a stimulus is sampled from for perceptual decision making. Further work will be needed to examine the degree with which applying a cortical magnification factor to the stimulus array may change the spread of decisional sampling, as well as how perceptual learning may interact with such a stimulus.

Perceptual learning is generally assumed to be orientation specific (Ahissar & Hochstein, 1993; Ahissar & Hochstein, 1996; Poggio, Fahle, & Edelman, 1992; Sigman & Gilbert, 2000), but not all studies have found orientation specificity of training (McGovern, Webb, & Peirce, 2012; Schoups et al., 1995; Xiao et al., 2008; Zhang et al., 2010). The reverse hierarchy theory of perceptual learning (Ahissar & Hochstein, 2000, 2004; Hochstein & Ahissar, 2002) could potentially provide an explanation for these inconsistent findings. According to this theory, learning is a top-down process starting at high-level visual areas that perform global processing. When the task is too difficult or specialized for these high-level areas, more local or precise processing is obtained by using low-level areas. Because global processing seems such an integral component of the task used here, the use of high-level areas might have been sufficient. It also appears that in the current experiment, little would be gained from a training-based change in lower level visual areas: narrowing neural tuning to better represent the orientation of individual patch elements would presumably have only a small effect on behavioral thresholds because of the high degree of external orientation noise in our stimuli. If learning indeed occurred predominantly in higher level areas, then this

could explain the transfer of the effect of learning to the untrained orientation because of the broader tuning curves observed for neurons in these high-level areas as compared to early visual areas (David, Hayden, & Gallant, 2006; Desimone & Schein, 1987; Hubel & Wiesel, 1968). This explanation is also in line with the findings of McGovern et al. (2012), who found that the benefits of training transferred across those tasks requiring a more global processing of the stimulus. In this context, it is also interesting to note that a recent neuroimaging study demonstrated optimized decision templates in higher level, posterior occipitotemporal regions after prolonged training on a shape discrimination task (Kuai, Levi, & Kourtzi, 2013), with no training-based modulation of activity in early visual cortex.

Previous studies of perceptual learning have suggested that sampling efficiency may play a role in the benefits of perceptual learning. For instance, a study by Lu and Dosher (2004) showed a significant reduction in orientation thresholds with learning for a high external noise condition, but not for a noiseless condition. These findings suggested that participants were able to exclude external noise better after training, which was explained by a retuning of the perceptual template after training. Interestingly, a change in tuning of the perceptual template is consistent with a change in sampling efficiency. Although the spatial sampling hypothesis was not directly investigated in the Lu and Dosher study, it is in line with another study by Li, Levi, and Klein (2004) that investigated this hypothesis. This study, as well as a study by Kurki and Eckstein (2014), correlated the observer's decision with positional noise of the stimulus, allowing the authors to infer which parts of the stimulus consistently drove the perceptual decision. The results of both studies showed a change in sampling after training on a position discrimination task. After training, the observers used a larger sample of stimulus elements in order to make a position judgment. Although the studies by Li et al. (2004) and Kurki and Eckstein (2014) showed a consistent increase in sampling efficiency with learning, this was only shown for training on a position discrimination task. It remained unclear whether increased sampling efficiency is a general feature of training that would transfer to other perceptual features, such as orientation.

In conclusion, the present study has shown an increase in the efficiency with which a stimulus is sampled from after perceptual learning on an orientation discrimination task. Because an increase in sample size is directly related to a decrease in orientation threshold (Dakin, 2001), altered sampling may be one of the core mechanisms underlying training-based improvements in perceptual performance.

References

Ahissar, M., & Hochstein, S. (1993). Attentional control of early perceptual learning. *Proceedings of the National Academy of Sciences*, 90(12), 5718–5722.

Ahissar, M., & Hochstein, S. (1996). Learning pop-out detection: Specificities to stimulus characteristics. *Vision Research*, 36(21), 3487–3500.

Ahissar, M., & Hochstein, S. (2000). The spread of attention and learning in feature search: Effects of target distribution and task difficulty. *Vision Research*, 40(10), 1349–1364.

Ahissar, M., & Hochstein, S. (2004). The reverse hierarchy theory of visual perceptual learning. *Trends in Cognitive Sciences*, 8(10), 457–464.

Ball, K., & Sekuler, R. (1982). A specific and enduring improvement in visual motion discrimination. *Science*, 218(4573), 697–698.

Ball, K., & Sekuler, R. (1987). Direction-specific improvement in motion discrimination. *Vision Research*, 27(6), 953–965.

Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, 10, 433–436.

Carrasco, M., Evert, D. L., Chang, I., & Katz, S. M. (1995). The eccentricity effect: Target eccentricity affects performance on conjunction searches. *Perception & Psychophysics*, 57(8), 1241–1261.

Carrasco, M., & Frieder, K. S. (1997). Cortical magnification neutralizes the eccentricity effect in visual search. *Vision Research*, 37(1), 63–82.

Crist, R. E., Li, W., & Gilbert, C. D. (2001). Learning to see: Experience and attention in primary visual cortex. *Nature Neuroscience*, 4(5), 519–525.

Dakin, S. C. (2001). Information limit on the spatial integration of local orientation signals. *Journal of the Optical Society of America A*, 18(5), 1016–1026.

David, S. V., Hayden, B. Y., & Gallant, J. L. (2006). Spectral receptive field properties explain shape selectivity in area V4. *Journal of Neurophysiology*, 96(6), 3492–3505.

De Valois, K. K. (1977). Spatial frequency adaptation can enhance contrast sensitivity. *Vision Research*, 17(9), 1057–1065.

Desimone, R., & Schein, S. J. (1987). Visual properties of neurons in area V4 of the macaque: Sensitivity to stimulus form. *Journal of Neurophysiology*, 57(3), 835–868.

Dobres, J., & Seitz, A. R. (2010). Perceptual learning of oriented gratings as revealed by classification images. *Journal of Vision*, 10(13):8, 1–11, doi:10.1167/10.13.8. [PubMed] [Article]

Eckstein, M. P., & Ahumada, A. J. (2002). Classification images: A tool to analyze visual strategies. *Journal of Vision*, 2(1): i, doi:10.1167/2.1.i. [PubMed] [Article]

Fahle, M. (1997). Specificity of learning curvature, orientation, and vernier discriminations. *Vision Research*, 37(14), 1885–1895.

Fahle, M., & Edelman, S. (1993). Long-term learning in vernier acuity: Effects of stimulus orientation, range and of feedback. *Vision Research*, 33(3), 397–412.

Fiorentini, A., & Berardi, N. (1980). Perceptual learning specific for orientation and spatial frequency. *Nature*, 287, 43–44.

Fiorentini, A., & Berardi, N. (1981). Learning in grating waveform discrimination: Specificity for orientation and spatial frequency. *Vision Research*, 21, 1149–1158.

Gibson, E. J. (1963). Perceptual learning. *Annual Review of Psychology*, 14(1), 29–56.

Girshick, A. R., Landy, M. S., & Simoncelli, E. P. (2011). Cardinal rules: Visual orientation perception reflects knowledge of environmental statistics. *Nature Neuroscience*, 14(7), 926–932.

Gold, J. M., Murray, R. F., Bennett, P. J., & Sekuler, A. B. (2000). Deriving behavioural receptive fields for visually completed contours. *Current Biology*, 10(11), 663–666.

Hochstein, S., & Ahissar, M. (2002). View from the top: Hierarchies and reverse hierarchies in the visual system. *Neuron*, 36(5), 791–804.

Hubel, D. H., & Wiesel, T. N. (1968). Receptive fields and functional architecture of monkey striate cortex. *The Journal of Physiology*, 195(1), 215–243.

Jehee, J. F., Ling, S., Swisher, J. D., van Bergen, R. S., & Tong, F. (2012). Perceptual learning selectively refines orientation representations in early visual cortex. *The Journal of Neuroscience*, *32*(47), 16747–16753.

Kapadia, M. K., Gilbert, C. D., & Westheimer, G. (1994). A quantitative measure for short-term cortical plasticity in human vision. *Journal of Neuroscience*, *14*, 451–457.

Karni, A., & Sagi, D. (1991). Where practice makes perfect in texture discrimination: Evidence perceptual learning refines orientation representation for primary visual cortex plasticity. *Proceedings of the National Academy of Sciences*, *88*(11), 4966–4970.

Kleiner, M., Brainard, D., & Pelli, D. (2007). What's new in Psychtoolbox-3. *Perception*, *36*, ECVP Abstract Supplement.

Kuai, S. G., Levi, D., & Kourtzi, Z. (2013). Learning optimizes decision templates in the human visual cortex. *Current Biology*, *23*(18), 1799–1804.

Kurki, I., & Eckstein, M. P. (2014). Template changes with perceptual learning are driven by feature informativeness. *Journal of Vision*, *14*(11):6, 1–18, doi:10.1167/14.11.6. [PubMed] [Article]

Leek, M. R. (2001). Adaptive procedures in psychophysical research. *Perception & Psychophysics*, *63*(8), 1279–1292.

Li, R. W., Levi, D. M., & Klein, S. A. (2004). Perceptual learning improves efficiency by retuning the decision "template" for position discrimination. *Nature Neuroscience*, *7*(2), 178–183.

Lu, Z. L., & Dosher, B. A. (2004). Perceptual learning retunes the perceptual template in foveal orientation identification. *Journal of Vision*, *4*(1):5, 44–56, doi:10.1167/4.1.5. [PubMed] [Article]

McGovern, D. P., Webb, B. S., & Peirce, J. W. (2012). Transfer of perceptual learning between different visual tasks. *Journal of Vision*, *12*(11):4, 1–11, doi:10.1167/12.11.4. [PubMed] [Article]

Nagai, M., Bennett, P. J., & Sekuler, A. B. (2007). Spatiotemporal templates for detecting orientation-defined targets. *Journal of Vision*, *7*(8):11, 1–16, doi:10.1167/7.8.11. [PubMed] [Article]

Paradiso, M. A., & Carney, T. (1988). Orientation discrimination as a function of stimulus eccentricity and size: Nasal/temporal retinal asymmetry. *Vision Research*, *28*(8), 867–874.

Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, *10*, 437–442.

Poggio, T., Fahle, M., & Edelman, S. (1992). Fast perceptual learning in visual hyperacuity. *Science*, *256*(5059), 1018–1021.

Rensink, R. A., O'Regan, J. K., & Clark, J. J. (1997). To see or not to see: The need for attention to perceive changes in scenes. *Psychological Science*, *8*(5), 368–373.

Schoups, A. A., Vogels, R., & Orban, G. A. (1995). Human perceptual learning in identifying the oblique orientation: Retinotopy, orientation specificity and monocularity. *The Journal of Physiology*, *483*(Pt 3), 797–810.

Schoups, A., Vogels, R., Qian, N., & Orban, G. (2001). Practising orientation identification improves orientation coding in V1 neurons. *Nature*, *412*(6846), 549–553.

Shiu, L. P., & Pashler, H. (1992). Improvement in line orientation discrimination is retinally local but dependent on cognitive set. *Perception & Psychophysics*, *52*(5), 582–588.

Sigman, M., & Gilbert, C. D. (2000). Learning to find a shape. *Nature Neuroscience*, *3*(3), 264–269.

Simons, D. J., & Levin, D. T. (1997). Change blindness. *Trends in Cognitive Sciences*, *1*(7), 261–267.

Simons, D. J., & Rensink, R. A. (2005). Change blindness: Past, present, and future. *Trends in Cognitive Sciences*, *9*(1), 16–20.

Treutwein, B. (1995). Adaptive psychophysical procedures. *Vision Research*, *35*(17), 2503–2522.

Xiao, L. Q., Zhang, J. Y., Wang, R., Klein, S. A., Levi, D. M., & Yu, C. (2008). Complete transfer of perceptual learning across retinal locations enabled by double training. *Current Biology*, *18*(24), 1922–1926.

Zanker, J. M. (1999). Perceptual learning in primary and secondary motion vision. *Vision Research*, *39*(7), 1293–1304.

Zhang, J. Y., Zhang, G. L., Xiao, L. Q., Klein, S. A., Levi, D. M., & Yu, C. (2010). Rule-based learning explains visual perceptual learning and its specificity and transfer. *The Journal of Neuroscience*, *30*(37), 12323–12328.