

Contents lists available at [SciVerse ScienceDirect](http://SciVerse.ScienceDirect.com)

Vision Research

journal homepage: www.elsevier.com/locate/visres

Motion aftereffect duration is not changed by perceptual learning: Evidence against the representation modification hypothesis

Alexander A. Petrov*, Nicholas M. Van Horn

Ohio State University, Columbus, OH 43210, United States

ARTICLE INFO

Article history:

Received 5 April 2011

Received in revised form 4 August 2011

Available online xxx

Keywords:

Perceptual learning

Motion aftereffect

Representation modification

Selective reweighting

Task specificity

ABSTRACT

The representation modification hypothesis of perceptual learning attributes the practice-induced improvements in sensitivity and/or discriminability to changes in the early visual areas. We used motion aftereffects (MAE) to probe the representations of motion direction. In two experiments, four practice sessions on a fine direction-discrimination task caused large stimulus-specific improvements in d' but no significant stimulus-specific changes in either static or dynamic MAE duration at posttest relative to a pretest. Power analysis indicated that the data were approximately 100 times more likely given the hypothesis of no MAE change than the hypothesis of a 10% relative change. In light of converging evidence in the MAE literature, this suggests that little or no change occurred in the cortical representations of visual motion up to and including area MT. The task specificity of the learning effect challenges the representation modification hypothesis and supports an alternative—selective reweighting.

© 2011 Elsevier Ltd. All rights reserved.

1. Introduction

Visual perceptual learning is defined as practice-induced improvement in visual tasks (see Fahle & Poggio, 2002; Fine & Jacobs, 2002, for reviews). It has been documented in orientation discrimination (e.g., Doshier & Lu, 1998), Vernier acuity (e.g., Fahle & Edelman, 1993), visual search (e.g., Ahissar & Hochstein, 1997), texture discrimination (e.g., Karni & Sagi, 1991) face identification (e.g., Gold, Bennett, & Sekuler, 1999), and motion detection and discrimination (e.g., Ball & Sekuler, 1987; Huang et al., 2008; Law & Gold, 2008; Liu, 1999; Petrov & Hayes, 2010; Watanabe, Náñez, & Sasaki, 2001). The learning effects are typically long-lasting and (partially) specific to the particular stimuli used in training (e.g., Ahissar & Hochstein, 1996, 1997; Ball & Sekuler, 1987; Crist et al., 1997; Fahle & Edelman, 1993; Liu, 1999).

The mechanisms of perceptual learning are still poorly understood and are a topic of active research (e.g., Gilbert, Sigman, & Crist, 2001; Lu et al., 2009). Two prominent hypotheses in the field are *representation modification (RM)* and *selective reweighting (SRW)*. The representation modification hypothesis attributes the behavioral improvement to changes in the early visual representations (e.g., Gilbert, Sigman, & Crist, 2001; Karni & Sagi, 1991; Schoups et al., 2001). It is advanced on the basis of the stimulus specificity of the learning effect, which is consistent with the stimulus-specific tuning of neurons in the early sensory areas. The RM hypothesis has been

implemented in two related models (Schwabe & Obermayer, 2005; Teich & Qian, 2003) that have accounted for some neurophysiological correlates of perceptual learning (Schoups et al., 2001; Yang & Maunsell, 2004). In other sensory modalities, abundant evidence for training-dependent cortical map plasticity has been found in the primary somatosensory and auditory cortices (see, e.g., Buonomano & Merzenich, 1998; Das, 1997, for reviews).

However, an alternative explanation is equally consistent with the stimulus specificity of learning (Petrov, Doshier, & Lu, 2005). It is possible that the system learns which features of the redundant, multifaceted early representations are most diagnostic for the task at hand and strengthens the read-out connections from the units encoding these features (Doshier & Lu, 1998; Mollon & Danilova, 1996). This selective reweighting hypothesis has been implemented in numerous models that account for a range of behavioral (e.g., Doshier & Lu, 1999; Lu, Liu, & Doshier, 2010; Petrov, Doshier, & Lu, 2005, 2006; Seung & Sompolinsky, 1993; Sotiropoulos, Seitz, & Seriès, 2011; Vaina, Sundaeswaran, & Harris, 1995) and neurophysiological (Law & Gold, 2008; Meinhardt, 2002; Shiu & Pashler, 1992). However, most of these studies involve pairs of tasks that depend on completely unrelated stimulus dimensions. For example, Shiu and Pashler (1992) used brightness discrimination and orientation discrimination with line stimuli that varied in both brightness and orientation. Such experimental designs demonstrate the role of attention in learning (Ahissar & Hochstein, 2002), but provide little information about the locus of plasticity. This is because the two tasks depend on stimulus dimensions that are encoded by non-overlapping populations of neurons. The use of the same stimulus set for both tasks does not guarantee that the same sensory representations are engaged in both (Petrov, Doshier,

* Corresponding author. Address: Department of Psychology, 200B Lazenby Hall, Ohio State University, Columbus, OH 43210, United States.

E-mail address: apetrov@alexpetrov.com (A.A. Petrov).

URL: <http://alexpetrov.com> (A.A. Petrov).

& Lu, 2005). When the representations overlap little, both RM and SRW hypotheses predict little transfer across tasks.

Our goal in this article is to test the RM hypothesis. Two conditions are required for a stringent test (Petrov, Doshier, & Lu, 2005): First, we need tasks that engage the same (or at least strongly overlapping) sensory representations. Second, the tasks must depend on different (or at most weakly overlapping) read-out connections. These two requirements are hard to satisfy simultaneously. We are aware of only three studies that satisfy them to some degree (Crist et al., 1997; Fahle, 1997; Saffell & Matthews, 2003). Two experiments demonstrated lack of transfer of learning across orientation discrimination and Vernier discrimination (Crist et al., 1997; Fahle, 1997). While a good case can be made that Vernier discrimination depends on the orientation-selective neurons in V1 (e.g., Weiss, Edelman, & Fahle, 1993; Wilson, 1986), this remains an assumption and Vernier hyperacuity can also be modeled on different principles (e.g., Cao & Grossberg, 2005). The third experiment (Saffell & Matthews, 2003) demonstrated lack of transfer of learning across motion direction discrimination and speed discrimination. Convergent evidence indicates that most neurons sensitive to visual motion are tuned conjunctively for both direction and speed (see, e.g., Britten, 2004, for review). Thus, both tasks engage the same population of motion-sensitive neurons and the lack of transfer poses a problem for the RM hypothesis. However, these data cannot rule out one prominent form of representation modification—selective increase of the slope of the tuning curves (Schoups et al., 2001). It is possible that direction-discrimination training sharpens the direction tuning of certain critical neurons without affecting their speed tuning, whereas speed-discrimination training sharpens the latter but not the former. More generally, selective sharpening seems compatible with task-specific learning of any two stimulus dimensions. A more stringent test of the RM hypothesis requires a pair of dissociable tasks based on a single dimension.

The main idea of this article is to use adaptation to probe the early representations of visual motion direction. Adaptation is a valuable tool for studying the visual system (see, e.g., Clifford, 2002; Clifford et al., 2007, for reviews). Here, we test whether practicing a direction-discrimination task affects the strength of the motion aftereffect in the trained direction relative to a control direction.

The motion aftereffect (MAE) is a well-known visual illusion (see Mather, Verstraten, & Anstis, 1998; Mather et al., 2008, for reviews). Prolonged exposure to motion in a given direction causes a subsequent illusory percept of motion in the opposite direction. The aftereffect is stimulus-specific (see Thompson, 1998, for review). For example, it shows both spatial- and temporal-frequency tuning (e.g., Bex, Verstraten, & Mareschal, 1996; Cameron, Baker, & Boulton, 1992; Schofield, Ledgeway, & Hutchinson, 2007). Converging evidence from psychophysics (e.g., Nishida & Sato, 1995; Verstraten et al., 1999), single-unit recording (e.g., Kohn & Movshon, 2004), brain imaging (e.g., Taylor et al., 2000), transcranial magnetic stimulation (e.g., Theoret et al., 2002), and other methods (see Mather et al., 2008, for review) indicates that MAE is not a monolithic phenomenon but “an amalgam of neural adaptation at several visual cortical sites” (Mather et al., 2008, p. 481). Our present experiments use two types of aftereffects: *static* (sMAE) and *dynamic* (dMAE). The adapting stimulus—moving filtered-noise texture—is the same for both types. They differ in the test stimulus during the post-adaptation phase: a static texture frame for sMAE versus dynamic visual noise for dMAE. These two types of aftereffects have different properties and seem to arise at different levels of the motion-processing pathway (e.g., Nishida & Ashida, 2000; Nishida & Sato, 1995; Verstraten et al., 1999); see (Culham et al., 1998; Mather et al., 2008, for reviews). In particular, areas V1, V2, and V3 are implicated in static MAE (e.g., Maruya, Watanabe,

& Watanabe, 2008; Taylor et al., 2000) and MT is implicated in both static and dynamic MAE (e.g., Kohn & Movshon, 2004; Theoret et al., 2002; Tootell et al., 1995). We tested both types to probe for representation modification across these areas.

Various theoretical explanations of the MAE have been proposed (see Mather & Harris, 1998; Mather et al., 2008; Vidnyánszky, Blaser, & Papathomas, 2002, for reviews). More generally, there is an extensive literature on the mechanisms of motion adaptation and related phenomena such as the direction aftereffect (see, e.g., Clifford, 2002; Kohn, 2007, for reviews). Some of these topics are discussed briefly in Section 5 below. While the details differ, all models agree that motion adaptation impacts the neuronal populations involved in processing and representing visual motion. In the closely related domain of orientation processing, the influential model of Teich and Qian (2003) proposes a common mechanism for both adaptation and perceptual learning—sharpening of the orientation tuning curves by means of changing the lateral connections in a recurrent network. According to this proposal, “adaptation in behaving subjects may be viewed as a short-term form of learning” (Teich & Qian, 2003, Abstract).

In sum, motion direction discrimination and MAE seem to rely on strongly overlapping sensory representations and are based on a single stimulus dimension. The representation modification hypothesis predicts that extensive discrimination practice near a particular direction will change the neuronal representations for this direction. We further hypothesize that this representational change will cause a detectable change in the strength of the MAE in the trained direction relative to an untrained control.

Turning to the selective reweighting hypothesis, it is important to choose the two experimental tasks that rely on different read-out connections. Only then do the SRW and RM predictions diverge. The MAE task is very similar to detection.¹ We use reported MAE duration as an index of strength (Pantle, 1998). The observers pressed a key when the illusory motion was no longer detectable. There is mounting psychophysical (Hol & Treue, 2001; Jazayeri & Movshon, 2007; Phinney, Bowd, & Patterson, 1997; Regan & Beverley, 1985), neurophysiological (e.g., Purushothaman & Bradley, 2005; Raiguel et al., 2006), and computational (e.g., Jazayeri & Movshon, 2006; Petrov, Doshier, & Lu, 2005; Seung & Sompolinsky, 1993) evidence that detection and *fine*² discrimination rely on different read-outs. Thus, we use fine discrimination training in our experiments. The observers practiced to discriminate small differences in motion direction from trial to trial. We hypothesize that the fine-discrimination read-out differs from the MAE read-out. Therefore, if perceptual learning occurs via selective reweighting of the read-out connections from unchanging representations, fine discrimination practice should have no effect on the MAE.

The RM and SRW hypotheses thus make opposite predictions about the interaction of fine-discrimination practice and MAE. These predictions were tested in two experiments.

2. Experiment 1

The first experiment uses static and dynamic MAE to probe for practice-induced changes at different levels of the visual hierarchy. We use a moving filtered-noise texture for the adapting stimulus in both cases. The test stimulus is a static texture frame on sMAE trials and multiple frames of dynamic visual noise on dMAE trials. The experimental schedule begins with a MAE pretest session, followed by four discrimination sessions practicing one particular direction, followed by a MAE posttest session. A final discrimina-

¹ There are some subtle but important differences between detecting real and illusory motion. They are discussed in Section 5.

² As opposed to *coarse* discrimination of motion in opposite directions, which is similar to detection (Petrov & Hayes, 2010).

tion session with the orthogonal direction verifies the specificity of the learning effect on the fine discrimination task. This experimental design produces the following three binary factors: Session (pretest vs. posttest) \times Direction (trained vs. control) \times Type (static vs. dynamic). All three factors are crossed within each participant. The Direction factor is counterbalanced between participants.

The question of main interest is whether there is a statistically significant interaction between the Session and Direction factors. The RM hypothesis predicts such interaction whereas the SRW hypothesis predicts no interaction. We use analysis of variance (ANOVA) to test for statistical significance. A technical difficulty arises at this point because the SRW prediction amounts to asserting the null hypothesis in the ANOVA (Cohen, 1992; Keppel & Wickens, 2004). To make such null result meaningful, the test must have sufficient statistical power to detect the effect if, in fact, the RM hypothesis is correct and practice does induce direction-specific changes in the MAE. This concern is addressed in four ways: First, the adapting stimuli on the MAE sessions are as similar as possible to the training stimuli on the discrimination sessions. Second (and most important), the statistical test is performed individually for each participant. This leverages the power across multiple participants. The probability to fail to detect a true effect (“Type II error”) decreases exponentially with the number of independent tests. Thus, even though the individual tests may have modest power, the power of the combined test increases dramatically. Third, we use *a priori* planned contrasts to maximize the power of each individual test to detect the specific kind of interaction predicted by the RM hypothesis—see Section 2.1.6 for details. Fourth, we calculate explicitly the smallest effect size that our leveraged test is expected to detect at the conventional significance level ($p < .05$).

2.1. Method

2.1.1. Observers

The participants were 11 students at the Ohio State University with normal or corrected-to-normal vision. They were naïve to the purposes of the experiment and were paid \$6 per hour plus a bonus contingent on their accuracy.

2.1.2. Stimuli and apparatus

The stimuli were filtered-noise textures moving coherently at a constant speed of 12 deg/s behind a circular aperture (Fig. 1). The filter had a Gaussian cross-section along the frequency axis in the Fourier domain and was radially symmetric at all directions. The spectral power peaked at 3 cpd and the bandwidth was 4 octaves (full width at half height). A new texture was generated on every trial by applying the same filter to a fresh sample of independent, identically distributed Gaussian noise. Consecutive, overlapping frames were cut out from a larger texture patch and presented at 96 Hz on a 21" NEC AccuSync 120 color CRT using the Psychophysics Toolbox (Brainard, 1997; Pelli, 1997) in The MathWorks (2009). The direction of motion was manipulated by rotating the individual frames via OpenGL calls to the ATI Radeon HD2600 Pro graphics card. The aperture was implemented by a circular mask that was fully transparent at radius $R = 9.5$ cpd, fully opaque at $R = 10.5$ cpd, and ramped down linearly in between. A small red fixation dot (14.2 min in diameter) marked the stimulus center at all times. The monitor gamma function was estimated via a psychophysical matching procedure (cf., Colombo & Derrington, 2001) and was verified with a Minolta 1° luminance meter. A software lookup table defined 255 evenly spaced luminance levels between $L_{min} = 1.5$ cd/m² and $L_{max} = 32.5$ cd/m². The display was the only light source in the room and was viewed binocularly with the natural pupil from a chin rest located ≈ 93 cm away. At that dis-




Fig. 1. The stimuli were filtered-noise textures moving behind a circular aperture. On discrimination trials, the texture moved in one of four possible directions (depicted by arrows) relative to an implicit reference direction (dotted line). The adapting stimuli on motion-aftereffect trials were the same, except that they moved for 10 s rather than 397 ms. [Note: The fixation dot at the center and the angles between the directions are exaggerated for visibility.]

tance, 1° of visual angle subtended ≈ 42 pixels (1024 \times 768 resolution).

2.1.3. Discrimination task and procedure

The fine direction discrimination task was defined with respect to a *reference direction* θ that was set implicitly for each block. The actual motion direction took four possible values: $(\theta - 3.5)$, $(\theta - 2)$, $(\theta + 2)$, and $(\theta + 3.5)$ degrees from vertical. Each block presented 120 stimuli of each kind. The instructions designated the first two directions as “counterclockwise” and the other two as “clockwise” and the observers made a binary discrimination choice by pressing one of two keys on the computer keyboard.

Each trial began with a brief beep. The texture appeared 500 ms later, moved for 397 ms, and disappeared. The beep onset always preceded the texture onset by exactly 500 ms, and thus could serve as a reliable attentional cue. A bonus system helped to motivate the observers and provided feedback. The reward for each correct response was a bonus point. The penalty for each error was the loss of a bonus point, an unpleasant beep, and the addition of 250 ms to the 800-ms intertrial interval. The cumulative bonus was displayed prominently at all times and was converted to dollars and cents after the last session.

2.1.4. MAE task and procedure

Each MAE trial consisted of three phases: adaptation, MAE test, and reset. The adaptation duration was 10 s, the MAE duration was the dependent variable, and the reset duration was adjusted to make the trial total exactly 30 s. The adapting stimuli were the same as in the discrimination task, except that their duration was ≈ 25 times longer. The adaptation duration (10 s) was determined in a pilot experiment and was designed to avoid ceiling and floor effects. The observers were instructed to fixate the stationary red dot in the middle of the display throughout the adaptation period. As an incentive for doing so, a simple attentional task provided an opportunity to score bonus points. The motion direction alternated between $(\theta - 3)$ and $(\theta + 3)$ at random intervals and the participants pressed a key whenever they detected a change. A variable number of such changes occurred on a trial and the bonus points scored for detecting them were visible throughout the adaptation period.

The test phase presented stimuli of two kinds: static and dynamic. On static (sMAE) trials, the test stimulus was a stationary frame of the same texture (Fig. 1). It appeared to move in direction $-\theta$ due to the aftereffect. On dynamic (dMAE) trials, the test stimulus was dynamic visual noise. That is, each frame was a mosaic of small square tiles of size 4×4 pixels. (Individual-pixel noise tended to blend into near-uniform gray.) The intensity of each tile was drawn independently from a Gaussian distribution and resampled 96 times per second. The semitransparent circular mask was applied throughout. This stimulus contained no globally coherent motion energy and yet it appeared to move due to the aftereffect. The observers were asked to press a key when the apparent motion stopped. The re-appearance of the bonus (which was not displayed during the test phase) signaled that the response had been registered.

The third and final phase was identical for both sMAE and dMAE trials. Dynamic noise was presented for the remainder of the duration in an effort to reset the motion processing system and minimize adaptation carry-overs across trials (von Grünau, 2002). It also removed the incentive to report short MAE durations because all trials lasted for 30 s regardless of the participant's response.³ A brief alert beep signaled the beginning of the next trial.

2.1.5. Experimental design and presentation schedule

The experiment involved two reference directions: $\theta = -50^\circ$ and $+40^\circ$ from vertical. We avoided the diagonal directions to discourage verbalizable decision strategies. For six participants, the trained direction was -50 and the control direction was $+40$; these values were reversed for the other five participants. Each participant completed 2 MAE sessions (on days 1 and 6) and five discrimination sessions (on days 2–5 and 7). The MAE sessions consisted of 84 trials, whereas the discrimination sessions consisted of 960 trials. Each MAE session was divided into seven blocks and the observers were encouraged to rest between the blocks. The MAE duration was measured repeatedly for each combination of reference direction (trained vs. control) and type (static vs. dynamic). Each block presented three replications of each combination, in random order. The discrimination sessions were divided into two blocks, which were further subdivided into “miniblocks.” The reference direction θ was the same throughout the four practice sessions—this is what made it the trained direction for this individual. Each miniblock presented a counterbalanced, randomized sequence of trials with small clockwise and counterclockwise deviations around θ (see Section 2.1.3). After the MAE posttest on day 6, there was one final discrimination session. The first miniblock (120 trials) of this session revisited the trained reference direction and transitioned back to the discrimination task. Then θ was switched to the orthogonal direction to assess the specificity of the discrimination learning effect.

2.1.6. Power analysis

The question of main interest is whether the MAE duration differs significantly between the trained and control directions at posttest. The conventional way to answer this question is to test the Session \times Direction interaction in a within-subject ANOVA. However, the statistical power of this test would be relatively low. We needed to maximize the power to give the representation modification hypothesis a fair hearing. We conducted a detailed power analysis to make a null effect interpretable. Power calculations required an explicitly specified alternative hypothesis (Cohen, 1992; Keppel & Wickens, 2004). This gave a distinctly

Bayesian flavor to our frequentist analysis. We examined a family of such hypotheses defined by an effect-size parameter c .

The crossing of the Direction and Session factors produce four conditions: “trained at pretest,” “control at pretest,” “trained at posttest,” and “control at posttest.” The first two are not expected to differ because the reference directions are (essentially) symmetrical and counterbalanced between participants. Let M denotes the MAE duration for both pretest conditions as shown in the top section of Table 1. Let us temporarily⁴ assume that the MAE in the control direction does not change between pretest and posttest. Thus, the control MAE duration at posttest is also M .

The trained direction at posttest is the critical condition. The RM hypothesis predicts a change there, whereas SRW predicts the same MAE across the board (Table 1). It is convenient to express the change in relative units by introducing the multiplicative parameter c . It can be positive or negative. For example, $c = -.10$ represents the hypothesis that perceptual learning causes a 10% drop in MAE duration.

A conventional ANOVA would partition the variance in our data into two main effects and one interaction effect. Each of them has 1 degree of freedom in our design and thus can be expressed as a planned contrast (Keppel & Wickens, 2004). The corresponding coefficients are listed in Table 1. The rightmost column shows the predictions of the RM hypothesis. The predicted MAE change is dispersed across all three tests, which dilutes the statistical power of any one of them. To focus the power, we replaced tests 2 and 3 by the *simple-effect* tests labeled 4 and 5 in Table 1. This concentrated the predicted change into a single test—the one that compares the MAE durations for the trained and control directions at posttest. The coefficients for this critical test are listed in the last row in the table. Note that contrasts 1, 4, and 5 form a mutually orthogonal set. Thus, we can use the same error term and the same significance level as the conventional partitioning (Keppel & Wickens, 2004). The Type factor (static vs. dynamic) is orthogonal to all these contrasts. Although not shown in Table 1, it was included in the analysis and the associated variance was subtracted from the error term. We adopted significance level $\alpha = .05$ for all power calculations.

To increase the power still further, a separate ANOVA was performed for each individual participant. These analyses capitalized on the fact that each cell of our experimental design contained 21 replications per observer. A group-level ANOVA would use the cell means only and thus miss the valuable information contained in these multiple replications. The outcomes of the individual tests were combined using the Bernoulli formula for repeated independent trials (Feller, 1957).

2.2. Results and discussion

2.2.1. Discrimination data

Two discriminability values (d' Macmillan & Creelman, 2005) were calculated for each observer in each block: for the easy ($\theta \pm 3.5$) and difficult ($\theta \pm 2$) pair of stimuli. Fig. 2 plots the group-averaged learning curves. The classic pattern was clearly replicated—the performance improved with practice but the improvement was partially specific to the trained direction, hence the drop after the switch to the orthogonal (control) direction on the last session.

The response times (RTs) also improved with practice. The RT distributions were fitted by a model describing the process of accumulating sensory evidence over time and making a binary decision. We used the *diffusion model* (Ratcliff, 1978; Ratcliff & McKoon, 2008). This analysis was based entirely on the discrimination data

³ On a few (<2%) of the trials, no response was made before the deadline. These cases were attributed to lapses of attention and omitted from the analysis. Many of these lapses came from one particular observer.

⁴ The final analysis does not depend on this assumption; it is introduced here for expository convenience only.

Table 1

Top section: Motion-aftereffect (MAE) duration patterns predicted by the selective reweighting (SRW) and representation modification (RM) hypotheses of perceptual learning. M denotes the MAE duration at pretest and c is an effect-size parameter. Middle and bottom sections: Coefficients of the planned linear contrasts for the two types of analysis of variance (ANOVA) discussed in Section 2.1.6.

	Pretest		Posttest		Predictions	
	Trained	Control	Trained	Control	SRW	RM
MAE duration predicted by						
Selective reweighting (SRW)	M	M	M	M		
Representation modification (RM)	M	M	$(1 + c)M$	M		
Conventional ANOVA						
(1) Main effect of Session	-1	-1	+1	+1	0 ^a	$cM/4^a$
(2) Main effect of Direction	+1	-1	+1	-1	0	$cM/4$
(3) Session \times Direction	-1	+1	+1	-1	0	$cM/4$
ANOVA using simple effects						
(1) Main effect of Session	-1	-1	+1	+1	0 ^a	$cM/4^a$
(4) Direction at pretest	+1	-1	0	0	0	0
(5) Direction at posttest	0	0	+1	-1	0	$cM/2$

^a Note: The predictions for the main effect of Session assume (incorrectly) that the response criterion does not change from pretest to posttest.




Fig. 2. Learning curves for the discrimination task. Group average of the 11 participants in Experiment 1. The error bars are 90% within-subject confidence intervals. The shaded areas mark the two motion aftereffect sessions.

and is published separately (Petrov, Van Horn, & Ratcliff, 2011, summarized in Section 4). The decision-making aspects of the discrimination task are not relevant for the interpretation of the MAE data. Fig. 2 contains the information needed for our present purposes. The training manipulation was successful—by the time of the MAE posttest (depicted by the shaded area), the d' for the trained direction seemed to increase by approximately 50% relative to its initial level.

2.2.2. MAE data

Fig. 3 plots the mean MAE durations for the group average (large panel) and for the 11 individual observers (small panels). As expected, there were no significant differences between the two reference directions at pretest. The simple effect of Direction at pretest (contrast 4 in Table 1) was not significant in the group-level data ($F(1, 10) < 1$, n.s.). Nor was it significant in 10 of the 11 individual ANOVAs. It did reach significance ($F(1, 160) = 4.21$, $p = .042$) for 1 observer, but this appears to be a Type I error. Recall that the two directions ($\theta = -40$ and $+50$) were symmetrical with respect to the vertical and there is no reason to expect MAE differences prior to the discrimination training.

The main question is whether this symmetry was broken after training with one direction but not the other. The data in Fig. 3 show no evidence for any significant asymmetry at posttest. The two directions (plotted with o's and x's, respectively) continued to elicit very similar MAEs. The simple effect of Direction at post-

test (contrast 5 in Table 1) was not significant in the group-level analysis ($F(1, 10) < 1$, n.s.). More importantly, none of the 11 individual tests revealed any significant differences either. The individual F values for this critical test are printed above their corresponding panels in Fig. 3.

The average statistical power of the individual tests was estimated as follows. The mean square error (MSE) of each individual ANOVA was divided by the (squared) grand mean MAE for the corresponding observer. The median of these normalized error terms was .10 in our sample (mean = .13, SD = .10). The effect-size parameter c in Table 1 is expressed in the same normalized units (Section 2.1.6). We calculated the power for a range of effect sizes on the basis of the median normalized MSE using standard formulas (Keppel & Wickens, 2004). For the critical test (contrast 5 in Table 1), we obtained power estimates $P = \{.050, .109, .304, .586, .827\}$ for effect sizes $c = \{0, \pm.05, \pm.10, \pm.15, \pm.20\}$, respectively.

This is the main result of the present article, so let us examine it more closely. Consider the hypothesis that the MAE did, in fact, change by 10% ($c = .10$). According to the above estimate, there is approximately 30% chance to detect a change of this size in the data from one (median) observer. In other words, the majority of the tests (70%) will fail to detect it (Type II error). Clearly, a negative outcome on one individual test does not warrant any strong conclusions. (The null result of the group-level analysis is similarly inconclusive.) However, the probability of observing 11 such failures on 11 independent tests is only $p = .0185$. Thus it is very unlikely that the collective pattern in Fig. 3 could occur if the hypothesis $c = .10$ were true. The leveraged test has enough statistical power to reject this hypothesis at the 5% significance level. Bigger changes ($c \geq .15$) can be rejected strongly ($p < .0001$). Very small MAE changes ($c = .05$) are compatible with these data ($p = .28$).

Consider now the selective-reweighting prediction that MAE did not change at all ($c = 0$). The probability to incorrectly reject a true null hypothesis (Type I error) is $\alpha = .05$ for each individual test. The probability of 0 rejections on 11 tests is $(1 - \alpha)^{11} = .569$. Thus, our data pattern is 30 times more likely to occur under selective reweighting than under representation modification involving 10% MAE change.

The main effect of Session (contrast 1 in Table 1) was significant for eight individuals. The MAE duration at posttest compared to pretest increased for four of them and decreased for the other four. Because of the cancelation of these inconsistent shifts, the effect did not reach significance in the group-level ANOVA ($F(1, 10) = 0.33$). We attribute these individual differences to the subjective choice of the response criterion. Recall that the task was to press a key when the apparent motion stopped. It seems that the observers had

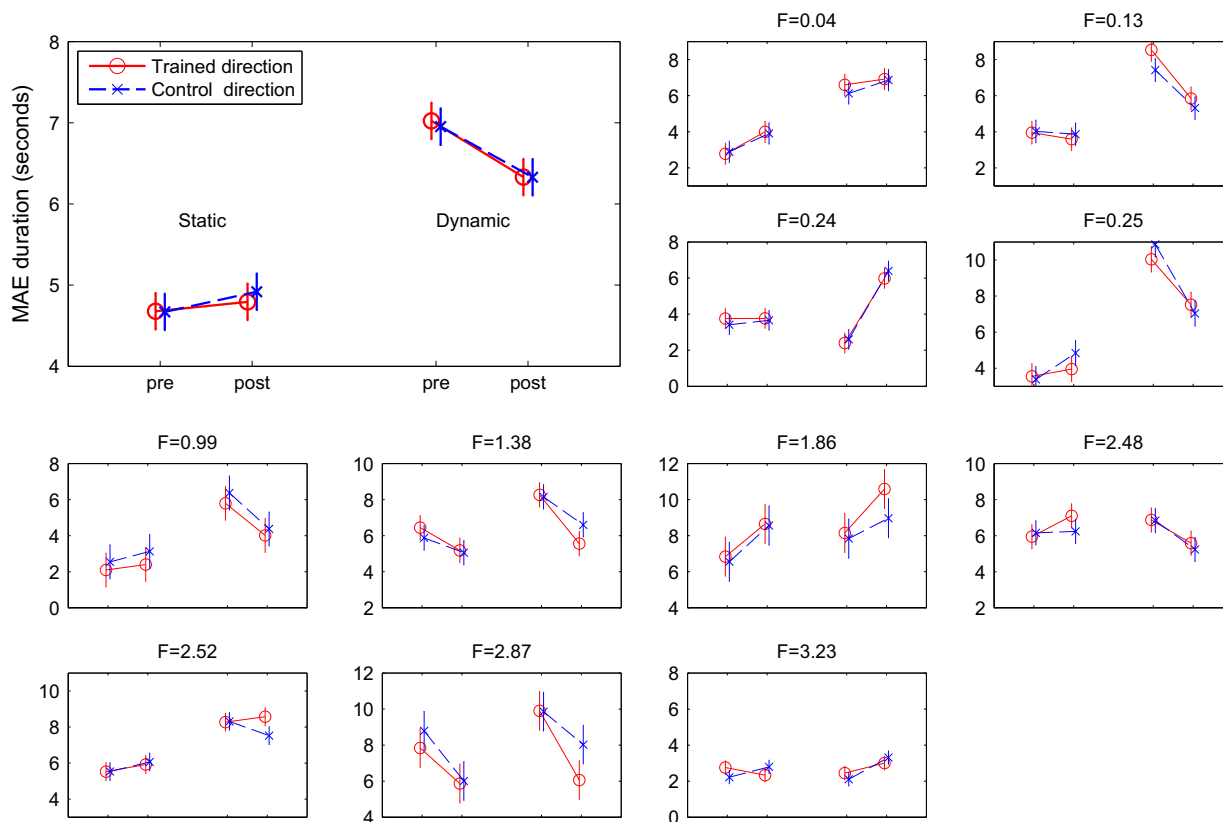


Fig. 3. Mean motion aftereffect (MAE) durations in Experiment 1. The large panel plots the group average. The error bars are 90% within-subject confidence intervals. The *pre* and *post* labels on the horizontal axis denote the pre- and posttest MAE sessions, respectively. The small panels plot the mean MAEs for each individual observer (using the same graphical layout). The error bars are 90% confidence intervals estimated from 21 replications per condition. The *F* statistic above each panel tests whether the trained direction (solid line) differs from the control direction (dashed line) at posttest. None of the tests is significant (critical $F(1, 160) = 3.90$).

adopted different criteria and changed them across sessions in idiosyncratic ways. The lack of significant Session \times Direction interaction (as tested by contrasts 4 and 5 above) indicates that these criteria were applied consistently for both directions.

The Type factor had a significant main effect for 9 individuals. Moreover, it was the one effect that reached significance at the group level ($F(1, 10) = 18.2, p < .002$). The MAE lasted significantly longer on dynamic trials than on static ones (Fig. 3). We attribute this to the random fluctuations in the dynamic-noise stimuli. The Session \times Type interaction was significant for six observers but not for the group ($F(1, 10) = 2.01, n.s.$). The general pattern was that static and dynamic MAE differed more at pretest than at posttest. One possible interpretation of this interaction is that the response criteria became better calibrated as observers gained experience with the MAE task. The Direction \times Type interaction was not significant for any individual participant, again suggesting consistent response criteria for both directions. The three-way interaction ($S \times D \times T$) was significant for only 1 of the 11 observers ($F(1, 160) = 4.21, p = .042$, probably a Type I error).

In conclusion, the data from Experiment 1 suggest that the MAE duration did not change (or changed very little) with discrimination practice. This challenges the representation modification hypothesis. There is, however, one important theoretical possibility that was not addressed in this experiment. The test stimulus in MAE appears to move in the direction that is opposite to the direction of the adapting stimulus. Given that the trained direction was used only for the adapting stimulus in Experiment 1, the illusory motion was never aligned with the trained direction. It was opposite to it. Previous research (e.g., Ball & Sekuler, 1987) indicates that the improvement in motion direction discrimination does not transfer to the opposite direction. This opens the possibility that representa-

tion modification did occur but was not detected because Experiment 1 did not test the MAE in the relevant direction. We conducted a second experiment to evaluate this possibility.

3. Experiment 2

This experiment was identical to Experiment 1 except that the MAE duration was measured in four directions: $\theta = -50, +40, +130$, and -140 relative to vertical. Also, all dynamic tests were dropped to make room for the static tests in the new directions.

3.1. Method

3.1.1. Observers

Sixteen new participants were recruited from the same population and were paid the same hourly rate and bonus as in Experiment 1.

3.1.2. Stimuli, task, and procedure

The stimuli and apparatus were the same as in Experiment 1. The fine discrimination task and procedure were also the same. The MAE sessions contained static trials only—21 replications in each of the four directions. There were no dynamic trials and thus the MAE session length (and hence statistical power) was the same as in Experiment 1.

3.1.3. Experimental design and presentation schedule

Half of the participants trained with $\theta = -50$ and the other half with $+40$. There was a switch to the orthogonal *upward* direction on the last session, as in Experiment 1. The two *downward* directions ($+130$ and -140) were never used for discrimination. They were only used for the adapting stimuli on MAE trials.

3.1.4. Power analysis

Because the Direction factor has four levels (and hence 3 degrees of freedom), it cannot be formulated as a single linear contrast. A set of three such contrasts would be necessary. This negates the advantages of the alternative analysis that was used in Experiment 1. Here we used the conventional partitioning of the variance into main effects and interactions. An ANOVA was performed for each individual, the power was calculated for a range of effect sizes using standard formulas (Keppel & Wickens, 2004), and the outcomes of the individual tests were combined as in Experiment 1.

3.2. Results and discussion

3.2.1. Discrimination data

The d' learning curves for the discrimination task are plotted in Fig. 4. It closely replicates the results of Experiment 1 (Fig. 2). The d' for the trained direction (solid symbols) improved with practice but the improvement transferred little to the control direction (open symbols).

3.2.2. MAE data

Fig. 5 shows the mean aftereffect durations for the group average (large panel) and for the 16 individual observers (small panels). The results of Experiment 1 were replicated. Most importantly, the Direction \times Session interaction was not statistically significant for 15 of the 16 observers. The individual F values for this critical test are printed above the corresponding panels in Fig. 5. One observer did show a significant interaction ($F(3,45) = 3.22, p < .025$) but their data (lower right-hand corner in Fig. 5) suggest that it was due to an anomalously low MAE for the control direction at pretest—an apparent Type I error. The group-level test did not reach significance ($F(3,45) = 0.84$).

The normalized mean square error of each individual ANOVA was calculated as in Experiment 1. The median of the new sample was $MSE = .11$ (mean = .12, $SD = .06$), replicating our earlier results. On the basis of this error term, we calculated the power to detect the Direction \times Session interaction in the data from one (median) observer. The power estimates were $P = \{.050, .083, .202, .387, .605\}$ for effect sizes $c = \{0, \pm.05, \pm.10, \pm.15, \pm.20\}$, respectively. The probability of obtaining exactly 1 significant outcome on n independent tests was calculated by the formula $p = nP(1 - P)^{n-1}$ (Feller, 1957). The results were $p \leq \{.37, .36, .11, .004, .00001\}$ for $n = 16$ and the effect sizes listed above.

Experiment 2 thus had slightly less statistical power than Experiment 1. The hypothesis that MAE changes by 15% or more can be rejected strongly ($p < .004$) but the hypothesis $c = .10$ cannot ($p = .11$). Still, the data pattern is 3.4 times more likely to occur under the null hypothesis ($c = 0$) than under a 10% change. This reinforces the main conclusion of Experiment 1: The MAE duration did not change (or changed very little) with discrimination practice. The new result is that this was demonstrated for adapting stimuli moving not only in the trained direction but in the opposite direction as well.

The lack of main effect of Session at the group level was replicated too ($F(1, 15) = 1.55, p = .23$). Averaged across all four directions, the MAE duration at posttest compared to pretest increased significantly for two individuals and decreased significantly for nine others. This replicated the finding that different observers adopt different response criteria and change them in idiosyncratic ways. Finally, the Direction factor did not have a significant main effect for the group ($F(3, 45) = 0.16$) but did for eight individuals, mostly due to differences between the upward and downward directions.

4. Decision-making aspects

The present article focused on the MAE, whereas a companion article (Petrov, Van Horn, & Ratcliff, 2011) focused on the discrimi-

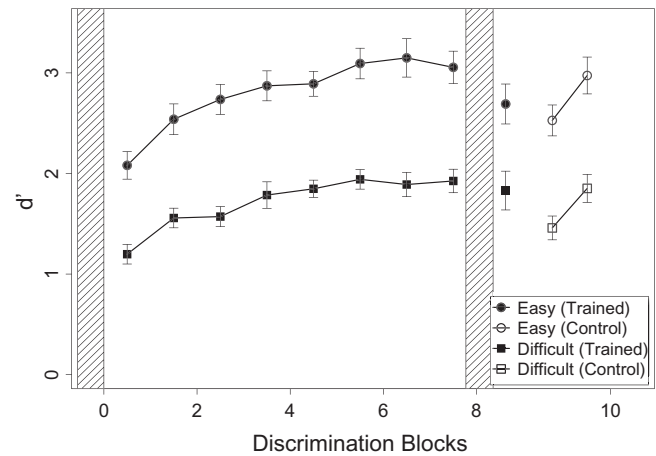


Fig. 4. Learning curves for the discrimination task. Group average of the 16 participants in Experiment 2. The error bars are 90% within-subject confidence intervals. The shaded areas mark the two motion aftereffect sessions.

nation sessions. Because of their homogeneity, the data from Experiments 1 and 2 were analyzed together. A learning index (Fine & Jacobs, 2002) was calculated to quantify the increase of d' relative to its initial level: $LI = (d'_8 - d'_1)/d'_1$, where the subscripts denote block numbers (cf. Figs. 2 and 4). It was⁵ $LI = .55 \pm .08$ for the average learning curve in the combined sample. A specificity index (Ahissar & Hochstein, 1997) quantified the disruption caused by the switch to the control direction: $SI = (d'_8 - d'_9)/(d'_8 - d'_1)$. It was $SI = .60 \pm .10$ (Petrov, Van Horn, & Ratcliff, 2011). In words, the average d' improved by approximately 55% after 4 days of discrimination practice but 60% of this improvement were specific to the trained direction.

As the response times (RTs) also improved with practice, the d' indices underestimate the true learning effect. The diffusion model (DM) is a valuable tool for analyzing joint accuracy and RT data (Ratcliff, 1978; see Ratcliff & McKoon, 2008, for review). Just as signal detection theory (Macmillan & Creelman, 2005) converts hits and false alarms into theoretically motivated estimates of discriminability and bias, the DM converts hits, false alarms, and RT distribution statistics into estimated parameters of various processing components. The DM accounted very well for the detailed discrimination data and re-expressed the regularities in these data in terms of seven parameters per observer per block (Petrov, Van Horn, & Ratcliff, 2011). The drift rate parameter is especially important for our present purposes because it quantifies the sensory evidence feeding into the decision process. By the end of training, the average drift rates were two times greater than their initial level ($LI = .99 \pm .23$). This accounted for both the observed d' increase (Figs. 2 and 4) and part⁶ of the RT decrease. The learning effect was partially specific to the trained direction ($SI = .68 \pm .09$).

5. General discussion

The present study tested a prediction of the representation modification hypothesis of perceptual learning. According to this hypothesis, practicing a direction-discrimination task induces stimulus-specific changes in the cortical representation of visual motion. In turn, these changes predict stimulus-specific changes in the motion aftereffect. In two experiments, fine-discrimination practice caused large stimulus-specific improvements in d' but no significant stimulus-specific changes in MAE duration. Power analysis indi-

⁵ $\pm 80\%$ bootstrap confidence interval, see Petrov, Van Horn, and Ratcliff (2011), for details.

⁶ The DM analysis suggested a second learning mechanism that improved the timing of the decision-process onset relative to the stimulus onset. It is not discussed here because it does not affect the interpretation of the MAE data.

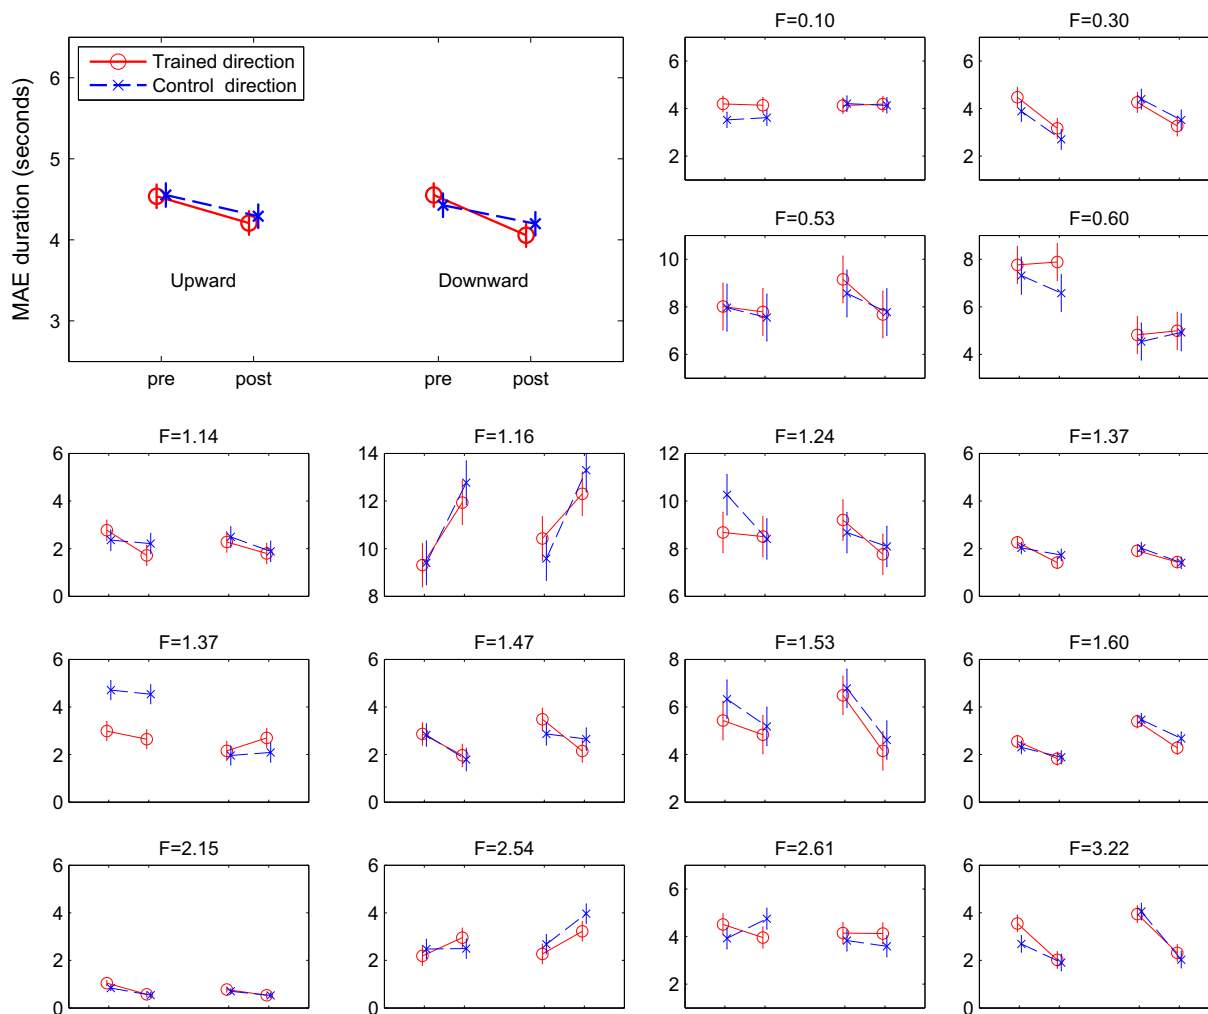


Fig. 5. Mean motion aftereffect (MAE) durations in Experiment 2. The large panel plots the group average. The error bars are 90% within-subject confidence intervals. The *pre* and *post* labels on the horizontal axis denote the pre- and posttest MAE sessions, respectively. The small panels plot the mean MAEs for each individual observer (using the same graphical layout). The error bars are 90% confidence intervals estimated from 21 replications per condition. The *F* statistic above each panel tests the Direction \times Session interaction. Only 1 of the 16 tests is significant (critical $F(3, 160) = 2.66$).

cated that the data were approximately 100 times⁷ more likely given the hypothesis of no MAE change than the hypothesis of a 10% relative change.

Section 1 cited convergent evidence that the MAE arises in the cortical areas that process and represent visual motion. Recall that areas V1, V2, and V3 are implicated in static MAE (e.g., Maruya, Watanabe, & Watanabe, 2008; Taylor et al., 2000) and MT is implicated in both static and dynamic MAE (e.g., Kohn & Movshon, 2004; Theoret et al., 2002; Tootell et al., 1995). The absence of significant stimulus-specific changes in either static or dynamic MAE suggests that little or no change occurred in the cortical representations of visual motion up to and including area MT. This challenges the RM hypothesis.

This inference depends on two premises: (i) that our tasks engaged the same population of motion-sensitive neurons and (ii) that modifying these internal representations would cause detectable changes in the MAE. We consider (and reject) two alternative interpretations of our results that negate these premises in turn.

The first alternative interpretation points out, correctly, that overlap at the level of cortical areas may not necessarily imply overlap in the specific neurons that determine the behavioral out-

come in our two experimental tasks. There is mounting evidence (e.g., Hol & Treue, 2001; Jazayeri & Movshon, 2007; Purushothaman & Bradley, 2005) that the neurons that are most diagnostic for fine discrimination are those whose direction preferences are shifted away from the discrimination boundary. The tuning curve of these *flanking neurons* has the steepest slope at the boundary (Jazayeri & Movshon, 2006; Petrov, Doshier, & Lu, 2005; Seung & Sompolinsky, 1993). The psychophysical demonstrations of this (Clifford et al., 2001; Hol & Treue, 2001; Phinney, Bowd, & Patterson, 1997; Regan & Beverley, 1985) are particularly relevant here because they involve adaptation (see Clifford, 2002, for review). Fine-discrimination thresholds increase by up to 60% for motion directions 20–30° on either side of the adaptor (Hol & Treue, 2001; Phinney, Bowd, & Patterson, 1997) and for orientation angles 10–15° on either side (Clifford et al., 2001; Regan & Beverley, 1985). The discrimination threshold at the adapted direction or orientation decreases by up to 20% (Clifford et al., 2001; Phinney, Bowd, & Patterson, 1997; Regan & Beverley, 1985). The latter improvement has been attributed to disinhibition of the flanker neurons when the neurons tuned for the boundary are suppressed by adaptation (Phinney, Bowd, & Patterson, 1997). Other explanations are also possible (e.g., Clifford et al., 2001; Seriès, Stocker, & Simoncelli, 2009). In one way or another, the effect seems to arise from the lateral interactions among the tuned mechanisms.

⁷ These odds (or Bayes factors, Kass & Raftery, 1995) combine multiplicatively across data sets. They were 30 and 3.4 in Experiments 1 and 2, respectively.

The detection thresholds, on the other hand, are most elevated at the adapted direction (e.g., Hol & Treue, 2001) or orientation (e.g., Regan & Beverley, 1985). This unimodal pattern is qualitatively different from the Mexican-hat pattern of change in the fine-discrimination thresholds. This is why Section 1 cited these studies in support of the claim that the two tasks in our experiments rely on different read-out connections.

In light of this evidence, why should we expect an interaction between fine-discrimination training and the MAE under the representation modification hypothesis? For concreteness, consider the preeminent RM proposal: Perceptual learning sharpens the tuning curves of the flanker neurons on either side of the trained reference direction (Schwabe & Obermayer, 2005; Schoups et al., 2001; Teich & Qian, 2003; Yang & Maunsell, 2004). The neurons tuned for the reference direction itself need not be modified.⁸ Given that the adapting stimuli on our MAE trials were very close to the reference direction, why should we expect a change in MAE triggered by modifications in the flanker neurons 20–30° away?

The answer is threefold: First, there are well documented lateral interactions among the neurons in the early sensory areas (e.g., Gilbert, 1992). The horizontal collaterals in V1, for example, contribute to orientation sensitivity (e.g., Nelson et al., 1994; Somers, Nelson, & Sur, 1995), contrast gain control (e.g., Heeger, 1992), and context integration (e.g., Gilbert et al., 2000). As discussed above, release from lateral inhibition is one likely mechanism of the improvement of the discrimination threshold at the adapted direction (Phinney, Bowd, & Patterson, 1997). The RM models (Schwabe & Obermayer, 2005; Teich & Qian, 2003) explicitly attribute perceptual learning to plasticity in the recurrent lateral connections. Second, adaptation affects the flanker neurons too—a phenomenon known as *flank adaptation* (e.g., Kohn & Movshon, 2004; see Kohn, 2007, for review). The effects of adaptation are not confined to the neurons that respond most strongly to the adaptor. Thus, if practice had modified significantly the response properties of the flanker neurons on both sides of the reference direction in our experiments, it is not unreasonable to expect a significant modification in the strength and/or time course of adaptation along this direction as well.

Finally, the detection of *illusory* motion is characterized by additional smoothing and integration compared to real motion. There is strong evidence that “the perceived global motion direction during motion aftereffects results from *local vector averaging* of the co-localized motion-direction signals induced by adaptation” (Vidnyánszky, Blaser, & Pappathomas, 2002, Abstract, emphasis added). A straightforward demonstration of this is that adaptation to bivectorial transparent motion normally leads to univectorial non-transparent MAE in the direction opposite to the vector sum of the adapting directions (e.g., Verstraten, Fredericksen, & van de Grind, 1994). This phenomenon reveals the inadequacy of the classic opponent-process (or ratio) model (Sutherland, 1961), which attributes the MAE to an adaptation-induced imbalance between the responses of the units tuned to the adapted direction and to its opposite. This problem motivated the development of the distribution shift model (Mather, 1980) and its successor, the automatic gain-control model (Grunewald, 1996; Grunewald & Lankheet, 1996; van de Grind, Lankheet, & Tao, 2003; van de Grind, van de Smagt, & Verstraten, 2004). The details of these models are beyond the scope of this article (see Mather & Harris, 1998; van de Grind, Lankheet, & Tao, 2003, for reviews). It suffices for our purposes that they all agree that the MAE involves multiple channels and that the mutual (dis)inhibition among them plays a major role. In the case of

bivectorial adaptation, the inhibitory interactions fuse the aftereffect into one single direction. Now, the RM hypothesis postulates changes in the channels flanking the discrimination boundary. Given the stimulus-specificity of perceptual learning, the response properties of these flanker channels should be different from those of the surrounding channels. This two-prong modification pattern induced by training is qualitatively similar to the two-prong suppression pattern induced by bivectorial adaptation. This gives reason to expect that the strength of the (univectorial) MAE in the direction opposite to the trained boundary would differ from the untrained control. This also suggests an interesting follow-up experiment that uses bivectorial transparent motion (e.g., $\theta \pm 25^\circ$) for the adapting stimulus on MAE trials.

In summary, the first alternative interpretation of our results does not seem convincing upon close examination. It appears that our two experimental tasks did engage highly overlapping populations of neurons, including the lateral connections between them. It must be acknowledged that the main conclusions of this article depend on this conjecture and that further research is needed to clarify the role of the lateral connectivity. Note also that the same connectivity is exactly the plasticity site proposed by the RM hypothesis.

The other alternative interpretation disputes the second premise of our reasoning. Why should we expect that modifying these representations would cause detectable MAE changes? Suppose, for the sake of argument, that there really are common plasticity mechanisms for adaptation and learning (Teich & Qian, 2003). Then the MAE posttest in the trained direction reflects the combined effect of learning plus adaptation, whereas that in the control direction reflects the effect of adaptation alone. If the adaptation effect were much stronger than the learning effect, adaptation plus learning would not differ significantly from adaptation alone. This would explain the observed lack of significant differences in the MAE data.

To answer this question we need to compare the effects of adaptation and learning. We use the change in discrimination threshold as common currency to estimate the approximate relative strengths of these two factors. Recall that the learning index for the average d' data is $\approx 55\%$. To a first approximation (cf. Fine & Jacobs, 2002), this corresponds to a commensurate decrease in discrimination threshold. On the other hand, the two studies that measured the effect of adaptation on motion discrimination (Hol & Treue, 2001; Phinney, Bowd, & Patterson, 1997) report $\approx 60\%$ increases in the fine discrimination thresholds for the flanker directions (and up to 20% decrease for the adapted direction). Note that the adaptation protocols in these studies were much stronger than ours. Phinney, Bowd, and Patterson (1997), for example, used an initial adaptation period of 5 min followed by top-up periods of 10 s interspersed between the discrimination trials. In contrast, our adaptation protocol was calibrated to avoid ceiling and floor effects, and included reset periods at the end of each trial. Our adaptation effects are thus probably smaller than 60%. Still, a conservative estimate is that the behavioral effects of adaptation and learning have approximately equal strength. This casts serious doubt on the second alternative interpretation of our results.

In conclusion, while we cannot rule out the possibility that a small amount of representation modification did occur but went undetected, our experimental technique seems powerful enough to rule out modifications that could account for a non-trivial portion of the large behavioral improvement. Most of the increase in d' (and diffusion drift rate) apparently stems from some other plasticity mechanism(s).

This conclusion agrees with the results of a recent study (Law & Gold, 2008) that recorded from approximately 250 individual MT neurons from two monkeys before and during extensive training on a motion direction-discrimination task. Despite dramatic improvements in the behavioral thresholds, no significant changes in firing rates or tuning properties were observed in the MT sample,

⁸ Indeed, no such modifications have been reported in single-cell recordings to date, except that (Ghose, Yang, & Maunsell, 2002) found slightly fewer V1 neurons whose optimal orientation was near the trained orientation. This resulted in a small but significant decrease in the V1 population response to the trained orientation at the trained location compared to a location in the opposite, untrained hemifield.

contrary to the prediction of the RM hypothesis. To our knowledge, Law and Gold (2008) is the only single-cell recording study of motion perceptual learning, but there are five similar studies of orientation-discrimination learning (Crist, Li, & Gilbert, 2001; Ghose, Yang, & Maunsell, 2002; Raiguel, Vogels, Mysore, & Orban, 2006; Schoups et al., 2001; Yang & Maunsell, 2004). One of them (Schoups et al., 2001) found modest changes in V1 but this result did not replicate in two subsequent attempts (Ghose, Yang, & Maunsell, 2002, 2004). Two studies Raiguel et al., 2006; Yang & Maunsell, 2004) found changes in area V4. While stronger than those in V1, they were still insufficient to account for the massive behavioral improvements (see Petrov, Doshier, & Lu, 2005; Raiguel et al., 2006, for reviews). Again, these data seem inconsistent with the RM hypothesis.

The selective reweighting (SRW) hypothesis, on the other hand, offers a straightforward and natural account. It attributes the behavioral improvement to synaptic plasticity in the read-out connections to higher areas. The pattern of activation (or the tuning properties) in early representational areas such as V1, V2, or MT is not expected to change. Rather, such changes are predicted for higher areas, particularly those involved in determining the response. This is precisely the pattern observed by Law and Gold (2008)—no significant changes in MT coupled with pronounced changes in LIP, a sensory-motor area. Moreover, the changes in LIP accumulated as learning progressed and their time course correlated with the time course of the behavioral improvement (Law & Gold, 2008). A neural network model accounted for the detailed neurophysiological recordings in terms of selective reweighting of the connections between MT and LIP (Law & Gold, 2009). A closely related model had been developed on the basis of human behavioral data (Doshier & Lu, 1998, 1999; Lu, Liu, & Doshier, 2010; Petrov, Doshier, & Lu, 2005, 2006; Sotiropoulos, Seitz, & Seriès, 2011; Vaina, Sundaeswaran, & Harris, 1995).

Our present data lends further support to the SRW hypothesis. On this account, the d' learning curves in Figs. 2 and 4 stem from selective reweighting of the read-out connections. Areas V1, V2, and MT do not change and neither does the motion aftereffect. The signal-to-noise ratio of the sensory input to the decision-making areas does change, however, because task-relevant inputs gain importance, whereas task-irrelevant inputs lose it (Petrov, Doshier, & Lu, 2005). This accounts for the twofold increase in the drift rate parameter of the diffusion-model fit to the RT data (Petrov, Van Horn, & Ratcliff, 2011). The MAE duration is not affected by the discrimination practice because it relies on a different set of read-out connections, as discussed above. These non-overlapping read-outs explain the task specificity of the learning effect.

To our knowledge, the present result is the first demonstration that MAE changes little (or not at all) after practicing a direction-discrimination task. There is one prior study (Vidnyánszky & Sohn, 2005) that did find significant MAE changes after practicing an unrelated task. Vidnyánszky and Sohn (2005) aimed to investigate the effect of practice on the efficiency of attentional selection. To that end, they constructed bivectorial transparent motion displays in which red dots moved horizontally and green dots moved vertically (or vice versa). The observers were instructed to attend one of the motion planes and ignore the other. The task was to detect brief increases of the attended dot population's luminance. The MAE duration was measured before and after practice on the luminance task. Four of the six participants showed significant decreases in the MAE evoked by the motion signal which was neglected during the seven practice sessions. Vidnyánszky and Sohn (2005, Abstract) concluded that "attentional suppression of task-irrelevant stimuli becomes more efficient with practice." The effect size varied between 20% and 70% across the four observers. This indicates that MAE duration is a sensitive measure that can be affected by practice manipulations. Our experiment had more than enough statistical power to detect a MAE change of this magnitude.

In conclusion, the present study establishes a new empirical constraint on theories of perceptual learning. Any model of motion direction-discrimination learning in particular must be able to produce 55% d' improvement (and 99% drift-rate improvement, Petrov, Van Horn, & Ratcliff, 2011) with less than 10% change in either static or dynamic MAE. The selective reweighting hypothesis predicts this pattern, whereas it poses a great challenge to the representation modification hypothesis.

References

- Ahissar, M., & Hochstein, S. (1993). Attentional control in early perceptual learning. *Proceedings of the National Academy of Sciences, USA*, 90, 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. (1997). Task difficulty and the specificity of perceptual learning. *Nature*, 387, 401–406.
- Ahissar, M., & Hochstein, S. (2002). The role of attention in learning simple visual tasks. In M. Fahle & T. Poggio (Eds.), *Perceptual learning* (pp. 253–272). Cambridge, MA: MIT Press.
- Ball, K., & Sekuler, R. (1987). Direction-specific improvement in motion discrimination. *Vision Research*, 27(6), 953–965.
- Bex, P. J., Verstraten, F. A. J., & Mareschal, I. (1996). Temporal and spatial frequency tuning of the flicker motion aftereffect. *Vision Research*, 36, 2721–2727.
- Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, 10, 433–436.
- Britten, K. H. (2004). The middle temporal area: Motion processing and the link to perception. In L. M. Chalupa & J. S. Werner (Eds.), *The visual neurosciences* (pp. 1203–1216). Cambridge, MA: MIT Press.
- Buonomano, D. V., & Merzenich, M. M. (1998). Cortical plasticity: From synapses to maps. *Annual Review of Neuroscience*, 21, 149–186.
- Cameron, E. L., Baker, C. L., & Boulton, J. C. (1992). Spatial frequency selective mechanism underlying the motion aftereffect. *Vision Research*, 32, 561–568.
- Cao, Y., & Grossberg, S. (2005). A laminar cortical model of stereopsis and 3D surface perception: Closure and da vinci stereopsis. *Spatial Vision*, 18, 515–578.
- Clifford, C. W. G. (2002). Perceptual adaptation: Motion parallels orientation. *Trends in Cognitive Sciences*, 6(3), 136–143.
- Clifford, C. W. G., Webster, M. A., Stanley, G. B., Stocker, A. A., Kohn, A., Sharpee, T. O., et al. (2007). Visual adaptation: Neural, psychological and computational aspects. *Vision Research*, 47(25), 3125–3131.
- Clifford, C. W. G., Wyatt, A. M., Arnold, D. H., Smith, S. T., & Wenderoth, P. (2001). Orthogonal adaptation improves orientation discrimination. *Vision Research*, 41(2), 151–159.
- Cohen, J. (1992). Statistical power analysis. *Current Directions in Psychological Science*, 1(3), 98–101.
- Colombo, E., & Derrington, A. (2001). Visual calibration of CRT monitors. *Displays*, 22, 87–95.
- Crist, R. B., Li, W., & Gilbert, C. D. (2001). Learning to see: Experience and attention in primary visual cortex. *Nature Neuroscience*, 4(5), 519–525.
- Crist, R. E., Kapadia, M. K., Westheimer, G., & Gilbert, C. D. (1997). Perceptual learning of spatial location: Specificity for orientation, position, and context. *Journal of Physiology*, 78, 2889–2894.
- Culham, J., Nishida, S., Ledgeway, T., Cavanagh, P., von Grünau, M. W., Kwas, M., et al. (1998). Higher-order effects. In G. Mather, F. A. J. Verstraten, & S. Anstis (Eds.), *The motion aftereffect: A modern perspective* (pp. 85–124). MIT Press.
- Das, A. (1997). Plasticity in adult sensory cortex: A review. *Network: Computation in Neural Systems*, 8, R33–R76.
- Doshier, B. A., & Lu, Z.-L. (1998). Perceptual learning reflects external noise filtering and internal noise reduction through channel reweighting. *Proceedings of the National Academy of Sciences, USA*, 95(23), 13988–13993.
- Doshier, B. A., & Lu, Z.-L. (1999). Mechanisms of perceptual learning. *Vision Research*, 39(19), 3197–3221.
- 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 feedback. *Vision Research*, 33(3), 397–412.
- Fahle, M., & Poggio, T. (Eds.). (2002). *Perceptual learning*. Cambridge, MA: MIT Press.
- Feller, W. (1957). *An introduction to probability theory and its applications* (2nd ed.). New York: John Wiley and Sons.
- Fine, I., & Jacobs, R. A. (2002). Comparing perceptual learning across tasks: A review. *Journal of Vision*, 2(2), 190–203.
- Ghose, G. M., Yang, T., & Maunsell, J. H. R. (2002). Physiological correlates of perceptual learning in monkey V1 and V2. *Journal of Neurophysiology*, 87, 1867–1888.
- Gilbert, C. D. (1992). Horizontal integration and cortical dynamics. *Neuron*, 9, 1–13.
- Gilbert, C. D., Ito, M., Kapadia, M., & Westheimer, G. (2000). Interactions between attention, context and learning in primary visual cortex. *Vision Research*, 40(10), 1217–1226.
- Gilbert, C. D., Sigman, M., & Crist, R. E. (2001). The neural basis of perceptual learning. *Neuron*, 31(5), 681–697.
- Gold, J. M., Bennett, P. J., & Sekuler, A. B. (1999). Signal but not noise changes with perceptual learning. *Nature*, 402(6758), 176–178.
- Grunewald, A. (1996). A model of transparent motion and non-transparent motion aftereffects. In D. Touretzky, M. C. Mozer, & M. Hasselmo (Eds.), *Advances in*

- neural information processing systems (Vol. 8, pp. 837–843). MA, USA: MIT Press Cambridge.
- Grunewald, A., & Lankheet, M. J. M. (1996). Orthogonal motion after-effect illusion predicted by a model of cortical motion processing. *Nature*, *384*(6607), 358–360.
- Heeger, D. J. (1992). Normalization of cell responses in cat striate cortex. *Visual Neuroscience*, *9*, 181–197.
- Hol, K., & Treue, S. (2001). Different populations of neurons contribute to the detection and discrimination of visual motion. *Vision Research*, *41*(6), 685–834.
- Huang, X., Lu, H., Tjan, B. S., Zhou, Y.-F., & Liu, Z. (2008). Motion perceptual learning: When only task-relevant information is learned. *Journal of Vision*, *7*(10), 1–10.
- Jazayeri, M., & Movshon, J. A. (2006). Optimal representation of sensory information by neural populations. *Nature Neuroscience*, *9*, 690–696.
- Jazayeri, M., & Movshon, J. A. (2007). A new perceptual illusion reveals mechanisms of sensory decoding. *Nature*, *446*, 912–915.
- Karni, A., & Sagi, D. (1991). Where practice makes perfect in texture discrimination: Evidence for primary visual cortex plasticity. *Proceedings of the National Academy of Sciences, USA*, *88*, 4966–4970.
- Kass, R. E., & Raftery, A. E. (1995). Bayes factors and model uncertainty. *Journal of the American Statistical Association*, *90*, 773–795.
- Keppel, G., & Wickens, T. D. (2004). *Design and analysis: A researcher's handbook* (4th ed.). Upper Saddle River, NJ: Prentice Hall.
- Kohn, A. (2007). Visual adaptation: Physiology, mechanisms, and functional benefits. *Journal of Vision*, *7*(5), 3155–3164.
- Kohn, A., & Movshon, J. A. (2004). Adaptation changes the direction tuning of macaque MT neurons. *Nature Neuroscience*, *7*, 764–772.
- Law, C.-T., & Gold, J. I. (2008). Neural correlates of perceptual learning in a sensory-motor, but not a sensory, cortical area. *Nature Neuroscience*, *11*(4), 505–513.
- Law, C.-T., & Gold, J. I. (2009). Reinforcement learning can account for associative and perceptual learning on a visual-decision task. *Nature Neuroscience*, *12*(5), 655–663.
- Liu, Z. (1999). Perceptual learning in motion discrimination that generalizes across motion directions. *Proceedings of the National Academy of Sciences, USA*, *96*(24), 14085–14087.
- Lu, Z.-L., Liu, J., & Doshier, B. A. (2010). Modeling mechanisms of perceptual learning with augmented Hebbian re-weighting. *Vision Research*, *50*, 375–390.
- Lu, Z.-L., Yu, C., Watanabe, T., Sagi, D., & Levi, D. (2009). Perceptual learning: Functions, mechanisms, and applications. *Vision Research*, *49*, 2531–2534.
- Macmillan, N. A., & Creelman, C. D. (2005). *Detection theory: A user's guide* (2nd ed.). New York: Cambridge University Press.
- Maruya, K., Watanabe, H., & Watanabe, M. (2008). Adaptation to invisible motion results in low-level but not high-level aftereffects. *Journal of Vision*, *8*(11), 7:1–7:11.
- Mather, G. (1980). The movement aftereffect and a distribution-shift model for coding the direction of visual movement. *Perception*, *9*, 379–392.
- Mather, G., & Harris, J. (1998). Theoretical models of the motion aftereffect. In G. Mather, F. A. J. Verstraten, & S. Anstis (Eds.), *The motion aftereffect: A modern perspective* (pp. 157–185). MIT Press.
- Mather, G., Pavan, A., Campana, G., & Casco, C. (2008). The motion aftereffect reloaded. *Trends in Cognitive Sciences*, *12*(12), 481–487.
- Mather, G., Verstraten, F. A. J., & Anstis, S. (Eds.). (1998). *The motion aftereffect: A modern perspective*. Cambridge, MA: MIT Press.
- The MathWorks. 2009. *MATLAB user's guide*. Natick, MA.
- Meinhardt, G. (2002). Learning to discriminate simple sinusoidal gratings is task specific. *Psychological Research*, *66*(2), 143–156.
- Mollon, J. D., & Danilova, M. V. (1996). Three remarks on perceptual learning. *Spatial Vision*, *10*(1), 51–58.
- Nelson, S., Toth, L., Sheth, B., & Sur, M. (1994). Orientation selectivity of cortical neurons during intracellular blockade of inhibition. *Science*, *265*, 774–777.
- Nishida, S., & Ashida, H. (2000). A hierarchical structure of motion system revealed by interocular transfer of flicker motion aftereffects. *Vision Research*, *40*, 265–278.
- Nishida, S., & Sato, T. (1995). Motion aftereffect with flickering test patterns reveals higher stages of motion processing. *Vision Research*, *35*(4), 477–490.
- Pantle, A. (1998). How do measures of the motion aftereffect measure up. In G. Mather, F. A. J. Verstraten, & S. Anstis (Eds.), *The motion aftereffect: A modern perspective* (pp. 25–39). MIT Press.
- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, *10*, 437–442.
- Petrov, A. A., Doshier, B. A., & Lu, Z.-L. (2005). The dynamics of perceptual learning: An incremental reweighting model. *Psychological Review*, *112*(4), 715–743.
- Petrov, A. A., Doshier, B. A., & Lu, Z.-L. (2006). Perceptual learning without feedback in non-stationary contexts: Data and model. *Vision Research*, *46*(19), 3177–3197.
- Petrov, A. A., & Hayes, T. R. (2010). Asymmetric transfer of perceptual learning of luminance- and contrast-modulated motion. *Journal of Vision*, *10*(14:11), 1–22.
- Petrov, A. A., Van Horn, N. M., & Ratcliff, R. (2011). Dissociable perceptual learning mechanisms revealed by diffusion-model analysis of the patterns of specificity. *Psychonomic Bulletin & Review*, *18*(3), 490–497.
- Phinney, R. E., Bowd, C., & Patterson, R. (1997). Direction-selective coding of stereoscopic (cyclopean) motion. *Vision Research*, *34*(7), 865–869.
- Purushothaman, G., & Bradley, D. C. (2005). Neural population code for fine perceptual decisions in area MT. *Nature Neuroscience*, *8*(1), 99–106.
- Raiguel, S., Vogels, R., Mysore, S. G., & Orban, G. A. (2006). Learning to see the difference specifically alters the most informative V4 neurons. *Journal of Neuroscience*, *26*(24), 6589–6602.
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, *85*, 59–108.
- Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: Theory and data for two-choice decision tasks. *Neural Computation*, *20*, 873–922.
- Regan, D., & Beverley, K. I. (1985). Postadaptation orientation discrimination. *Journal of the Optical Society of America A*, *2*(2), 147–155.
- Saffell, T., & Matthews, N. (2003). Task-specific perceptual learning on speed and direction discrimination. *Vision Research*, *43*, 1365–1374.
- Schofield, A. J., Ledgeway, T., & Hutchinson, C. V. (2007). Asymmetric transfer of the dynamic motion aftereffect between first- and second-order cues and among different second-order cues. *Journal of Vision*, *7*(8:1), 1–12.
- Schoups, A. A., Vogels, R., Qian, N., & Orban, G. A. (2001). Practising orientation identification improves orientation coding in V1 neurons. *Nature*, *412*, 549–553.
- Schwabe, L., & Obermayer, K. (2005). Adaptivity of tuning functions in a generic recurrent network model of a cortical hypercolumn. *Journal of Neuroscience*, *25*(13), 3223–3332.
- Seriès, P., Stocker, A. A., & Simoncelli, E. P. (2009). Is the homunculus “aware” of sensory adaptation? *Neural Computation*, *21*, 3271–3304.
- Seung, H. S., & Sompolinsky, H. (1993). Simple models for reading neuronal population codes. *Proceedings of the National Academy of Sciences, USA*, *90*(22), 10749–10753.
- 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.
- Somers, D. C., Nelson, S. B., & Sur, M. (1995). An emergent model of orientation selectivity in cat visual cortical simple cells. *Journal of Neuroscience*, *15*(8), 5448–5465.
- Sotiropoulos, G., Seitz, A. R., & Seriès, P. (2011). Perceptual learning in visual hyperacuity: A reweighting model. *Vision Research*, *51*(6), 585–599.
- Sutherland, N. S. (1961). Figural after-effects and apparent size. *Quarterly Journal of Experimental Psychology*, *13*, 222–228.
- Taylor, J. G., Schmitz, N., Ziemons, K., Grosse-Ruyken, M.-L., Gruber, O., Mueller-Gaertner, H.-W., et al. (2000). The network of brain areas involved in the motion aftereffect. *NeuroImage*, *11*(4), 257–270.
- Teich, A. F., & Qian, N. (2003). Learning and adaptation in a recurrent model of V1 orientation selectivity. *Journal of Neurophysiology*, *89*(4), 2086–2100.
- Théoret, H., Kobayashi, M., Ganis, G., Di Capua, P., & Pascual-Leone, A. (2002). Repetitive transcranial magnetic stimulation of human area MT/V5 disrupts perception and storage of the motion aftereffect. *Neuropsychologia*, *40*(13), 2280–2287.
- Thompson, P. (1998). Tuning of the motion aftereffect. In G. Mather, F. A. J. Verstraten, & S. Anstis (Eds.), *The motion aftereffect: A modern perspective* (pp. 41–55). MIT Press.
- Tootell, R. B., Reppas, J. B., Dale, A. M., Look, R. B., Sereno, M. I., Malach, R., et al. (1995). Visual motion aftereffect in human cortical area MT revealed by functional magnetic imaging. *Nature*, *375*, 139–141.
- Vaina, L. M., Sundareswaran, V., & Harris, J. G. (1995). Learning to ignore: Psychophysics and computational modeling of fast learning of direction in noisy motion stimuli. *Cognitive Brain Research*, *2*(3), 155–163.
- van de Grind, W. A., Lankheet, M. J., & Tao, R. (2003). A gain-control model relating nulling results to the duration of dynamic motion aftereffects. *Vision Research*, *43*, 117–133.
- van de Grind, W. A., van de Smagt, M. J., & Verstraten, F. A. J. (2004). Storage for free: A surprising property of a simple gain-control model of motion aftereffects. *Vision Research*, *44*, 2269–2284.
- Verstraten, F. A. J., Fredericksen, R. E., & van de Grind, W. A. (1994). Movement aftereffect of bi-vectorial transparent motion. *Vision Research*, *34*(3), 349–358.
- Verstraten, F. A. J., van der Smagt, M. J., Fredericksen, R. E., & van de Grind, W. A. (1999). Integration after adaptation to transparent motion: Static and dynamic test patterns result in different aftereffect directions. *Vision Research*, *39*, 803–810.
- Vidnyánszky, Z., Blaser, E., & Pappathomas, T. V. (2002). Motion integration during motion aftereffects. *Trends in Cognitive Sciences*, *4*(157–167).
- Vidnyánszky, Z., & Sohn, W. (2005). Learning to suppress task-irrelevant visual stimuli with attention. *Vision Research*, *45*(6), 677–685.
- von Grünau, M. W. (2002). Bivectorial transparent stimuli simultaneously adapt mechanisms at different levels of the motion pathway. *Vision Research*, *42*(5), 577–587.
- Watanabe, T., Náñez, J., & Sasaki, Y. (2001). Perceptual learning without perception. *Nature*, *413*, 844–848.
- Weiss, Y., Edelman, S., & Fahle, M. (1993). Models of perceptual learning in vernier hyperacuity. *Neural Computation*, *5*, 695–718.
- Wilson, H. R. (1986). Responses of spatial mechanisms can explain hyperacuity. *Vision Research*, *26*(3), 453–469.
- Yang, T., & Maunsell, J. H. R. (2004). The effect of perceptual learning on neuronal responses in monkey visual area V4. *Journal of Neuroscience*, *24*(7), 1617–1626.