Vision Research 166 (2020) 72–80

Contents lists available at ScienceDirect

Vision Research

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

Contingent adaptation in masking and surround suppression

Hörmet Yiltiz[a], David J. Heeger[a,b], Michael S. Landy[a,b,*]

[a] Department of Psychology, New York University, New York, NY, United States
[b] Center for Neural Science, New York University, New York, NY, United States

ARTICLE INFO

Keywords:
Contingent adaptation
Masking
Surround suppression
Homeostasis
Hebbian learning
Divisive normalization

ABSTRACT

Adaptation is the process that changes a neuron's response based on recent inputs. In the traditional model, a neuron's state of adaptation depends on the recent input to that neuron alone, whereas in a recently introduced model (Hebbian normalization), adaptation depends on the structure of neural correlated firing. In particular, increased response products between pairs of neurons leads to increased mutual suppression. We test a psychophysical prediction of this model: adaptation should depend on 2nd-order statistics of input stimuli. That is, if two stimuli excite two distinct sub-populations of neurons, then presenting those stimuli simultaneously during adaptation should strengthen mutual suppression between those subpopulations. We confirm this prediction in two experiments. In the first, pairing two gratings synchronously during adaptation (i.e., a plaid) rather than asynchronously (interleaving the two gratings in time) leads to increased effectiveness of one pattern for masking the other. In the second, pairing the gratings in a center-surround configuration results in reduced apparent contrast for the central grating when paired with the same surround (as compared with a condition in which the central grating appears with a different surround at test than during adaptation). These results are consistent with the prediction that an increase in response covariance leads to greater mutual suppression between neurons. This effect is detectable both at threshold (masking) and well above threshold (apparent contrast).

1. Introduction

Adaptation describes the change in neural responses based on recent input. It is prevalent throughout the nervous system, in sensory systems such as vision (Clifford et al., 2007), audition (Condon & Weinberger, 1991), and in non-sensory areas as well (Louie, Grattan, & Glimcher, 2011). One consequence of adaptation is a reduction of sensitivity of neurons selective for the adapting stimulus. This was originally characterized as a kind of neural fatigue in which highly stimulated neurons reduce their gain (Maffei, Fiorentini, & Bisti, 1973; Ohzawa, Sclar, & Freeman, 1982; Sekuler & Pantle, 1967; Shapley & Enroth-Cugell, 1984).

However, neural responses to adaptation are not limited to gain control (Solomon & Kohn, 2014). Adaptation also modifies tuning curves, primarily by repelling tuning curves away from the adapter, and can reduce correlation in neural responses (Dragoi, Sharma, & Sur, 2000; Müller, Metha, Krauskopf, & Lennie, 1999). The consequences of adaptation include homeostatic regulation of neural firing rates and pairwise response covariance (Benucci, Saleem, & Carandini, 2013). As a result, some adaptation effects are stimulus-specific, i.e., responses are reduced more in all neurons for stimuli similar to the adapter than for stimuli that differ from it. However, in the gain-control model of adaptation, reduction in gain is specific to each neuron, rather than to a specific stimulus.

We recently introduced a new model of adaptation, the Hebbian normalization model (Westrick, Heeger, & Landy, 2016). In this model, each neuron has a feed-forward drive determined by the stimulus and the neuron's tuning curve; that drive is normalized by (i.e., divisively suppressed by) a weighted sum of the responses of other neurons to generate the response for that neuron. In the model, adaptation leads to a change in the normalization weights. The weight update is, in a sense, Hebbian: the normalization strength is increased between neurons that fire together more often than expected. But, since that weight is effectively suppressive, the net effect is anti-Hebbian: if neurons fire together, they suppress one another more. This Hebbian normalization model is similar to a previous model of Barlow and Földiák (Barlow & Földiák, 1989; Földiák, 1990), except that in our model the suppression is divisive and the learning rule in their model tries to drive neural correlations to zero. The Hebbian normalization model predicts many post-adaptation effects on neural responses including reductions in gain, shifts of tuning curves (Fig. 1), and homeostasis of neural response covariance. In the classic gain-control model, adaptation of a neuron

* Corresponding author at: Department of Psychology, New York University, 6 Washington Place, New York, NY 10003, United States.
E-mail address: landy@nyu.edu (M.S. Landy).

https://doi.org/10.1016/j.visres.2019.11.004
Received 31 March 2019; Received in revised form 6 November 2019; Accepted 8 November 2019

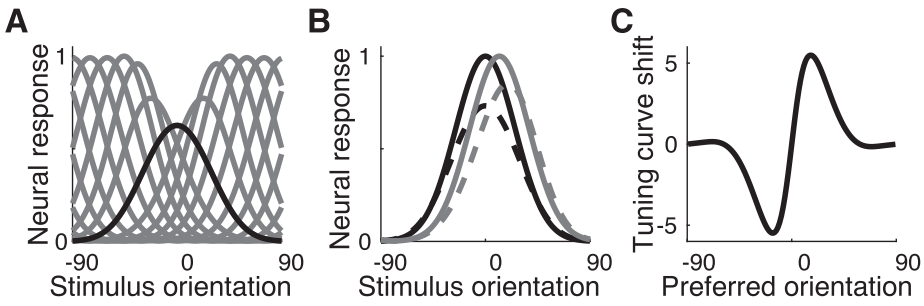

Fig. 1. The Hebbian normalization model. Before adaptation, neurons had identical tuning curves varying in preferred orientation. **A**. Tuning curves after adaptation. Black curve: Post-adaptation tuning curve of the neuron selective for the 0-deg adapter. **B**. Pre- (solid) and post-adaptation (dashed) tuning curves of neurons tuned to the adapter (black) and a neighboring orientation (gray). Note the repulsive shift of the post-adaptation tuning curve of this latter neuron. **C**. Post-adaptation tuning-curve shift as a function of pre-adaptation preferred orientation. Note the repulsive shifts for nearby orientations on either side of the adapter. Figure reprinted with permission from Westrick et al. (2016).

depends only on the recent history of input to that neuron. A distinguishing feature of the new model is that adaptation of a neuron depends on the recent response history of all neurons in the neuron's normalization pool, rather than just on its own input.

The Hebbian normalization model predicts that adaptation should depend on the 2nd-order statistics of input stimuli. That is, if two stimuli excite distinct sub-populations of neurons, then presenting those stimuli *simultaneously* during adaptation should increase suppression between those subpopulations, an effect we call *contingent adaptation*. Using pairs of gratings differing in orientation, we have confirmed the prediction that contingent adaptation strengthens normalization in cortex (Aschner, Solomon, Landy, Heeger, & Kohn, 2018).

Adaptation effects in the nervous system have robust consequences for perception. Reduced neural gain suggests that sensitivity will be reduced for the adapted stimulus, and indeed, threshold elevation for stimuli similar to the adapter is a hallmark of adaptation (Blakemore & Campbell, 1969; Blakemore & Nachmias, 1971; Pantle & Sekuler, 1968). This occurs in other senses as well, e.g., vibro-tactile detection (Hahn, 1968). In addition to effects at threshold, reduced neural gain also predicts effects on appearance due to the altered population response after adaptation. For example, in the visual system, a prolonged exposure to a slightly clockwise-tilted grating results in repulsion: a vertical grating is perceived as tilted counter-clockwise, a result known as the tilt after-affect (Gibson & Radner, 1937). And, there is an analogous spatial frequency after-effect (Blakemore & Sutton, 1969). Similar effects are present in other senses as well (Calzolari, Azañón, Danvers, Vallar, & Longo, 2017; Li, Chan, Iqbal, & Goldreich, 2017).

Here, we test the prediction of the effects of contingent adaptation on perception. In the first experiment, adapting to two overlapping gratings (i.e., a plaid) leads to strengthened masking as compared to adaptation to those same gratings presented asynchronously. In the second, adapting to two gratings in a center-surround configuration leads to reduced apparent contrast for the center grating when paired with the same surround as in the adapter, as compared to a center-surround pairing that differs between adaptation and test. Thus, contingent adaptation leads to perceptual effects both at and well above detection threshold.

2. Materials and methods

2.1. Participants

Four graduate students (three females; age range: 22–26) from New York University participated in Experiment 1. Two observers were naive to the purposes of the study. Three NYU graduate students participated in Experiment 2 (two females; age range: 22–26; two naive). All participants had normal or corrected-to-normal vision. The study was approved by the NYU Committee on Activities Involving Human Subjects. All participants gave informed consent prior to the experiment.

2.2. Apparatus

We presented stimuli using Psychtoolbox-3 for MATLAB (Kleiner et al., 2007) on a gamma-corrected, 20-inch HP p1230 CRT monitor (Hewlett–Packard, Palo Alto, CA) with a resolution of 1024 x 768 pixels and 60 Hz refresh rate. The monitor was controlled by a Mac Mini 2014 (Apple, Cupertino, CA) using an Intel Iris 1536 graphics card (Intel, Santa Clara, CA). To ensure that observers maintained steady fixation throughout a session, we monitored eye position using an SR Research Eyelink-1000 desktop eye tracker (SR Research Ltd., Ottawa, Ontario, Canada) with a sampling rate of 1000 Hz, controlled using the Eyelink Matlab Toolbox (Cornelissen, Peters, & Palmer, 2002). Observers viewed the stimuli from a viewing distance of 80 cm (so that resolution was 40 pixel/deg) with the head stabilized using a chin rest in a dimly lit room.

2.3. Stimuli

2.3.1. Experiment 1: Detection

Throughout the experiment, a white fixation dot (100% Weber contrast, radius: 0.1 deg) appeared at the center of the screen. Background luminance was 40 cd/m^2. Adapter and test stimuli were displayed 3 deg left and right of fixation. Stimuli were circular patches of sinusoidal grating (radius: 0.75 deg) with contrast

$$G(x, y) = c \cdot \sin(2\pi f\,(x\cos(\theta) + y\sin(\theta)) + \varphi), \quad (1)$$

where c, θ, f and φ were the contrast, orientation, spatial frequency and phase of the grating, respectively. All gratings had a spatial frequency of 3 cycle/deg, and phase was uniformly sampled between 0 and 2π. Grating orientations were ± 15 or ± 45 deg. Grating edges were windowed with a raised cosine (period: 0.15 deg).

There were two adaptation conditions: *asynchronous* and *contingent*. An adapter consisted of two 50%-contrast gratings with opposite orientations relative to the vertical. In the asynchronous condition, the two orientations alternated at 5 Hz (Fig. 2, top). Thus, the two different orientations were never presented simultaneously. Identical grating patches were displayed left and right of fixation. For the contingent adapters, the 50%-contrast gratings with the two orientations were summed to form a plaid (Fig. 2, middle). This plaid was alternated with a uniform, mean-luminance display at 5 Hz. As a result, time-averaged contrast powers of the asynchronous and contingent adapters were identical.

Test stimuli consisted of a 60%-contrast masker grating presented in both locations, plus an additional target grating added to one of the maskers (Fig. 2, bottom). The target orientation was the opposite of the masker orientation. The target location and masker/target orientations were randomly varied across trials.

2.3.2. Experiment 2: Appearance

Adapter and test stimuli consisted of two circular patches of sinusoidal grating ('center', radius: 0.75 deg) with a surrounding ring of

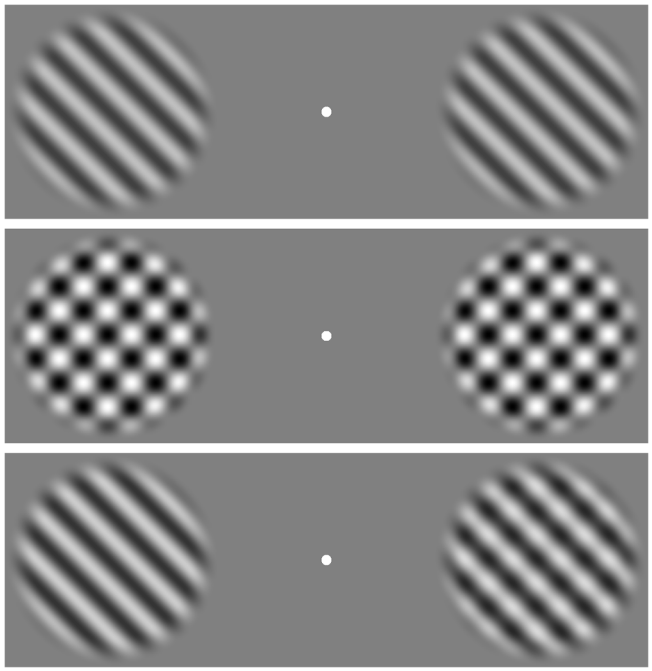

Fig. 2. Stimuli for Experiment 1. Top: Asynchronous adapter. The gratings (contrast: 50%) alternated between ± 15 or ± 45 deg (as shown here) at 5 Hz. Middle: Contingent adapter. The two gratings were summed to form a plaid pattern and alternated with a blank, mean-gray, uniform field at 5 Hz. Bottom: Test stimuli. One grating (the masker) was displayed in both locations (contrast: 60%). In one location, chosen randomly, the target grating (with the opposite orientation) was added to the masker (here, on the right).

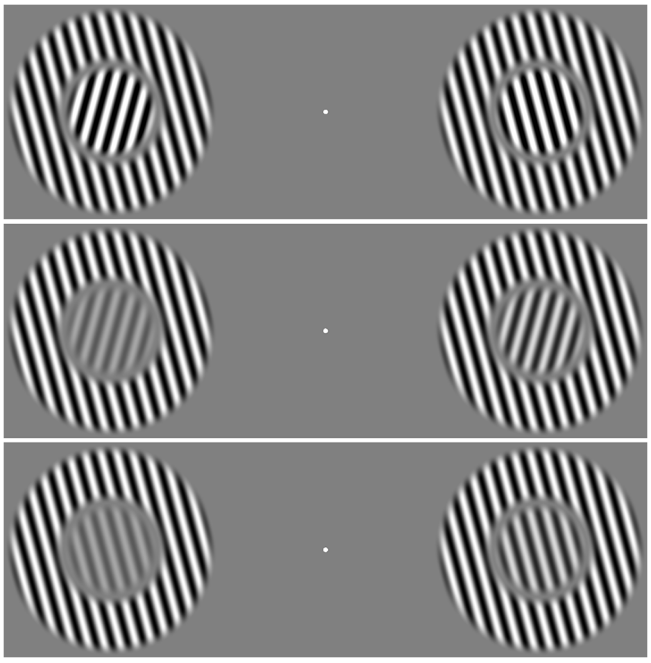

Fig. 3. Stimuli for Experiment 2. Top: Adapters. Pictured is an opposite-orientation configuration on the left and a same-orientation configuration on the right. The opposite configuration was used in half the sessions. Pictured is one of the two displays for this configuration; the other swapped the orientations of all gratings. These two displays alternated at 5 Hz. Middle: Test stimulus with the opposite-orientation configuration on both sides. The 30% contrast standard stimulus is the center patch on the left. Bottom: Test stimulus in the same-orientation configuration with the standard patch on the right. These adapters and test stimuli are from the ± 15 deg condition.

grating ('surround', inner radius: 1 deg, outer radius: 2 deg). Centers and surrounds were windowed by a raised cosine (period: 0.15 deg). There were two types of adapters: *same-orientation*, where orientations of the center and surround gratings were identical, and *opposite-orientation*, where the orientations were opposite. An adapter consisted of one same-orientation and one opposite-orientation configuration (Fig. 3, top); one was displayed 3 deg to the left and one 3 deg to the right of fixation. The choice of which adapter appeared on the left was varied across subjects. Adapters had 100% contrast. The phases of the center and surround adapters were chosen randomly and independently from a uniform distribution from 0 to 2π. Adapters alternated orientation at 5 Hz. The test stimuli had several key differences from the adapters(Fig. 3, middle and bottom): (1) test stimuli were either same-orientation on both sides or opposite-orientation on both sides (varied within a session); (2) test stimulus surrounds had 65% contrast; and (3) one center grating (the standard) had a fixed contrast of 30%. The other (the comparison) had a contrast that was varied across trials by a staircase procedure.

2.4. Experimental design and statistical analysis

2.4.1. Experiment 1: Detection

2.4.1.1. Procedure. We measured contrast threshold for detecting a target grating superimposed on a mask grating. Target and mask had opposite orientations relative to vertical (± 15 or ± 45 deg). We used a two-alternative forced-choice (2AFC) design. There were two adaptation conditions (asynchronous and contingent). In both adaptation conditions, gratings with the same orientations as the mask and test each appeared 50% of the time. They either appeared simultaneously as a plaid (contingent adaptation) or alternated in time (asynchronous). Because the two gratings differed in orientation, they stimulated partially non-overlapping populations of neurons. Those neurons that typically respond only to the counter-clockwise-rotated grating and those that respond only to the clockwise orientation will alternate activity in the asynchronous condition, but will have correlated responses in the contingent condition. The Hebbian normalization model thus predicts an increase in mutual suppression between these neural populations in response to contingent adaptation.

There were four types of sessions that varied in the grating orientations that were used (± 15 or ± 45 deg) and the type of adaptation (asynchronous or contingent). Sessions lasted about 30 min and were separated by at least one day. In the beginning of each session, we showed observers all possible stimuli at every contrast and the adapters that would be used in the experiment to familiarize them with the stimulus set. Observers dark-adapted to the dimly lit room for at least 3 min. Each session consisted of four runs of 88 trials separated by a brief break. Each run began with a 60 s presentation of the adapter followed by the experimental trials. Each trial began with a 2 s top-up adapter, a 150 ms inter-stimulus interval, and a 250 ms stimulus (Fig. 4). An auditory tone signaled the start of the stimulus interval. Observers indicated by key press which patch (left or right) contained the target grating. The response initiated the subsequent trial after an inter-trial-interval of 250 ms. On the rare occasions when fixation was broken before the stimulus appeared, the stimulus was delayed until fixation was reacquired. Feedback after each trial indicated whether the response was correct.

Four randomly interleaved staircases (Levitt, 1971) controlled target contrast. Two 1-up-2-down staircases targeting 71% correct started from a contrast of 5% and two 1-up-3-down staircases targeting 79% correct started from a contrast of 10%. One staircase of each type was associated with a target appearing on the left, and the other with a target appearing on the right. Contrast values were restricted to lie between 1.5 and 40% with 20 logarithmically spaced steps. Each staircase ran for 20 trials. The session also included eight easy trials (40% contrast, 4 each with target on left or right) at

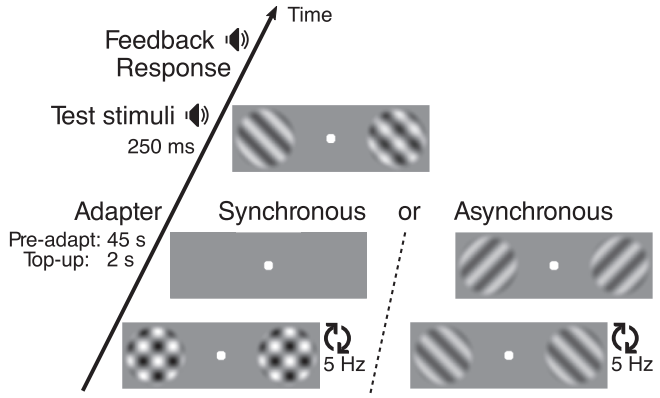

Fig. 4. Experiment 1: Trial sequence. Adapters alternated between two images at 5 Hz for 60 s (pre-adapter) with 2 s top-up adaptation at the beginning of each trial. The top-up adapter was followed by a 150 ms uniform field and then a 250 ms stimulus. Observers indicated which patch (left or right) contained the target grating. Adapters had two gratings with opposite orientations. These gratings were presented simultaneously as a plaid interleaved with a uniform field (contingent) or were interleaved (asynchronous), illustrated on the left and right respectively.

randomly chosen times to maintain observer motivation and better constrain estimates of lapse rate, for a total of 88 trials/run. The orientation of the masker (clockwise or counterclockwise) and test (the opposite orientation) were randomized across trials. In a separate initial session, observers practiced with feedback without adapters for approx. 30 min or until thresholds stabilized. Observers were required to maintain fixation throughout each run. Subjects ran a total of 8–12 sessions on separate days.

2.4.1.2. Analysis. There were four conditions based on adapter type (asynchronous vs. contingent adapter) and orientation ($\pm$ 15 and $\pm$ 45 deg). We discarded runs where accuracies for the catch trials were lower than 80% (only one run for one subject met this criterion). In the resulting data, we have between 1 and 4 runs for each subject in each condition (at least 2 runs in all but one combination of subject and condition). We fit psychometric curves for each condition (proportion correct as a function of target contrast) by maximum-likelihood estimation. With the assumption of log-normal internal noise, the psychometric function has the following form:

$$\Psi(c;\ \mu,\ \sigma,\ \lambda) = \frac{1}{2} + \left(\frac{1}{2} - \lambda\right)\Phi\left(\frac{\ln c - \mu}{\sigma}\right), \tag{2}$$

where c is the contrast of the target grating, μ and σ are the location and scale parameters in log contrast of the cumulative normal distribution Φ, and λ is the lapse rate. ½ is the guess rate for our 2AFC task. We determined the parameter values (μ, σ and λ) that maximized the likelihood of the data for each run within each session, then averaged the parameter estimates across runs that shared a common adaptation condition. Note that due to a moderate between-session variability in the slope of the psychometric curves, we did not pool data across runs and sessions.

We used parametric bootstrapping to estimate confidence intervals for the model parameters for each observer (Maloney, 1990). For each condition, we simulated the experiment 10,000 times assuming the fitted version of the psychometric function as an observer, and randomly sampled the same number of trials following the same procedure as in the experimental data. We re-fit each bootstrapped dataset by maximum likelihood to obtain sample parameter estimates. Based on the empirical distribution of these samples, we computed 95% bootstrap confidence intervals of threshold contrast μ. We tested for the significance of the threshold elevation effects due to contingent adaptation with parametric bootstraps. From each of the empirical distributions for asynchronous and contingent adapters, we re-sampled without replacement pairs of threshold contrast estimates. The difference of these estimates is a measure of effect size. We computed a *p*-value by counting the proportion of paired threshold estimates for which the difference in thresholds (asynchronous minus contingent) was less than zero. We tested separately for each observer and $\pm$ 15 and $\pm$ 45 deg orientations.

2.4.2. Experiment 2: Appearance

2.4.2.1. Procedure. We measured apparent contrast of a grating in a two-spatial-alternative forced-choice task. Observers adapted to one of two contingencies between the orientations of center and surround on the left of fixation, and the opposite contingency on the right. For two subjects, the same-orientation adapter was on the left; for the other subject it was on the right. Each session began with a 60 s presentation of the adapter followed by the experimental trials. Each trial began with a 2 s top-up adapter, a 150 ms inter-stimulus interval, and a 250 ms stimulus, followed by an inter-trial interval of 250 ms. In the test stimulus, there were identical contingencies on both sides (both same-orientation, or both opposite-orientation). Observers indicated which center patch appeared to have higher contrast. One center patch (the standard) had a fixed contrast of 30%. The other patch (the comparison) had contrast that varied across trials controlled by a staircase procedure. A session consisted of 8 staircases, each run for 20 trials. These corresponded to all combinations of two configurations of test stimuli (same- vs. opposite-orientation configurations), position of the standard stimulus (left or right of fixation), and type of staircase (1-up-2-down starting from 20% contrast and 2-up-1-down starting from 40% contrast). In addition, for each configuration and standard location there were two easy trials (comparison contrast: 80%) for a total of 168 trials. Observers participated in two sessions on one day separated by at least 15 min. The sessions differed in which orientations were used ($\pm$ 15 and $\pm$ 45 deg), but had the same adapter configuration for both (either same-orientation on the left or on the right). The staircases were restricted to range from 15 to 80% contrast in 20 discrete log-spaced steps.

2.4.2.2. Analysis. Each session comprised four conditions based on whether the test stimulus was a same- or opposite-orientation pair, and whether the comparison stimulus matched its corresponding adapter on the same side of fixation. If the comparison and the adapter matched, the theory predicts increased surround suppression (contingent adaptation), reducing apparent contrast of the comparison stimulus's center patch. This should, as a result, require higher contrast for the comparison to appear to have the same contrast as the standard stimulus. Recall that the standard stimulus had the same orientation configuration as the test, and the adapter had opposite configurations to the left and right of fixation, respectively. If the comparison and adapter did not match, the comparison should require less contrast to appear to have the same contrast as the standard (Fig. 5).

We reduced the data from four conditions to two by pooling the two conditions in which the comparison matched its corresponding adapter (*matching case*), and pooling the two conditions in which they did not match (*non-matching case*). For each of the resulting two conditions, we fit a psychometric function by maximum likelihood to the data, i.e., the probability of choosing the comparison stimulus as a function of comparison contrast. With the assumption of normally distributed internal noise for appearance judgments, the psychometric function had the following form:

$$\Psi(c;\ \mu,\ \sigma, \lambda) = \lambda + (1 - 2\lambda)\Phi\left(\frac{c - \mu}{\sigma}\right). \tag{3}$$

μ and σ are the location and scale parameters in physical contrast. We again fit the data separately for each observer and condition. The μ parameter corresponds to the 50% point of the psychometric function

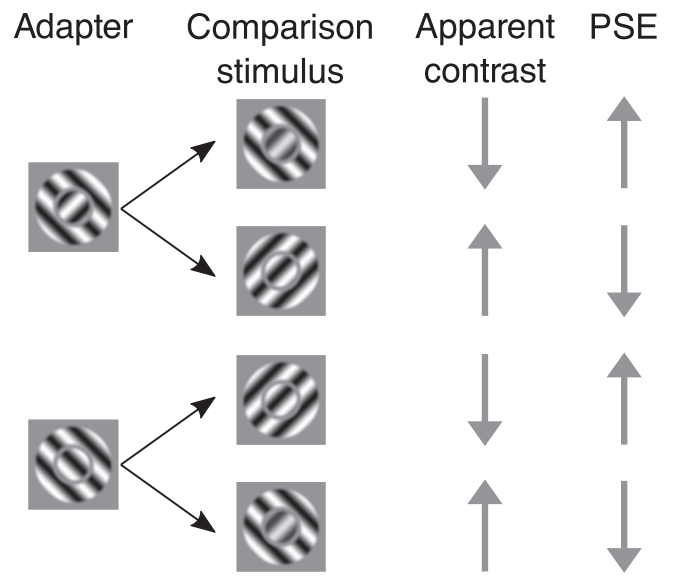

Fig. 5. Experiment 2: Conditions and predictions. Icons represent stimulus configurations for the adapter (first column) and the corresponding comparison stimulus on the same side of fixation (column 2). In the first two rows, the adapter at the position of the comparison stimulus had opposite orientations in the center patch and surrounding ring; in the last two rows, these orientations matched. In the first and third rows, the center-surround orientation contingencies in the comparison stimulus matched those of the corresponding adapter; in the second and fourth rows, they did not match. When those contingencies matched, the theory predicts adaptation will lead to an increase in suppression between surround and center for the comparison stimulus, leading to reduced apparent contrast of the center patch. Therefore, the comparison will require increased physical contrast to match the standard stimulus. Similarly, in the non-matching cases (rows 2 and 4), the comparison will require less physical contrast to match the standard. (Note: in this design, when the comparison matches the corresponding adapter, the standard does not match its corresponding adapter, and vice versa.)

(the point of subjective equality, PSE) at which the comparison and standard appeared to have equal contrast. We estimated confidence limits on the PSEs using parametric bootstraps.

2.5. Model

We explored whether the Hebbian normalization model (Westrick et al., 2016) can predict our experimental findings. We simulated a population of orientation-selective neurons with orientation tuning preferences that were equally spaced and tiled the full range of orientations in a fully-connected feed-forward network. A learning rule adjusted the normalization weights between these neurons to find a balance between maintaining weights near their initial values and homeostatically maintaining the responses products of the corresponding pairs of neurons. The model for the detection task in Experiment 1 is illustrated in Fig. 6. Model source code is available at https://github.com/hyiltiz/contingent-adaptation.

The tuning curves of the model neurons were raised cosines. The tiling was arranged so that the response energy (i.e., the sum of squared neural responses) was exactly equal to 1 for a full-contrast grating of any orientation (Simoncelli, Freeman, Adelson, & Heeger, 1992). The feed-forward responses were then normalized by a weighted sum of neural responses. It is these weights that were modified by the learning rule. To simulate human behavior, the model was exposed to the same stimuli as the human observers, and the learning rule was applied after each stimulus. For test trials, we ran the model until simulated activity due to adaptation was stable. For Experiment 1, we determined the probability of correct response assuming late additive Gaussian noise for the population response and an ideal decoder responding to the

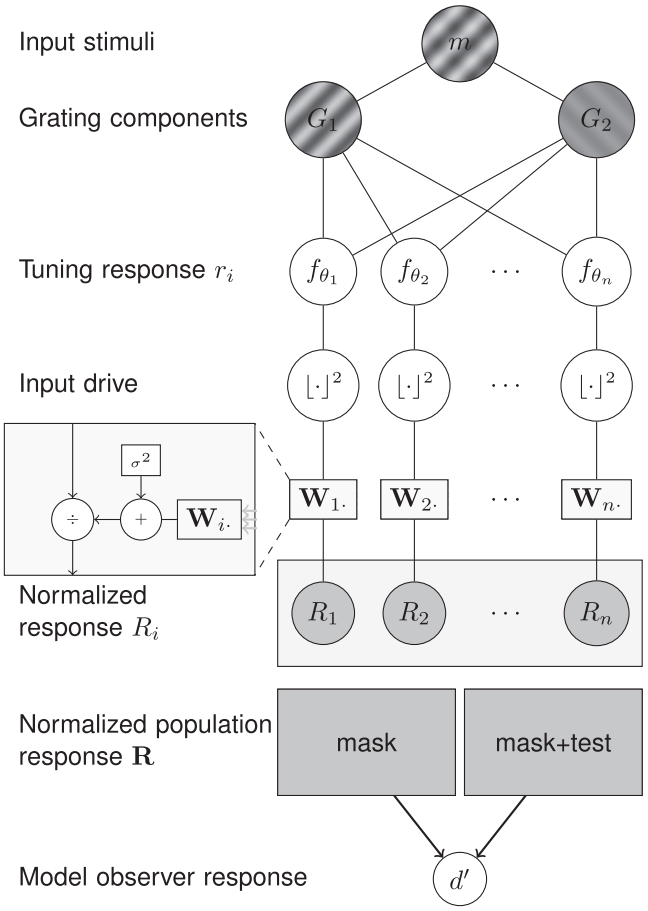

Fig. 6. Model architecture for Experiment 1. The model is a feed-forward neural network. The input image m (adapter or test stimulus) was analyzed by a set of neurons with tuning curves of identical shape tiling the orientation space (from 0 to 180 deg). The squared outputs of these units were then normalized by (i.e., divided by) a weighted sum of other neurons' squared input drives plus a normalization constant (illustrated in the zoomed-in box on the left). These weights were modified by a Hebbian learning rule (not shown). For modeling the detection task, we assumed independent late noise so that d' is proportional to the Euclidean distance between the output vectors of the two test stimuli.

simulated activity of the full population of model neurons. For Experiment 2, we took the mean of the population response as an estimate of the apparent contrast, and numerically solved for the PSEs.

2.5.1. Experiment 1: Detection

The network consisted of $N = 180$ fully connected neurons with raised cosine tuning curves and preferred orientations φ_i spaced every 1 deg. The neurons shared the same receptive field location centered on the stimuli. The feed-forward response of the i^{th} neuron to a set of gratings θ_j was:

$$r_i = \sum_j c_j \cdot f_i(\theta_j) + b, \text{ where} \tag{4}$$

$$f_i(\theta) = \begin{cases} \frac{1}{2}\left(\cos\left(\frac{2\pi}{B}(\theta - \varphi_i)\right) + 1\right) & -B \leqslant \theta - \varphi_i \leqslant B \\ b & \text{otherwise,} \end{cases} \tag{5}$$

where θ_j was the orientation of the j^{th} stimulus component with corresponding contrast c_j. Note that if b were zero, the sum of squared tuning curves would have been precisely uniform. The feed-forward responses were then subject to divisive normalization. The resulting normalized response R_i^t at time t was

$$R_i^t = \frac{\lfloor r_i(t) \rfloor^2}{\sigma^2 + \sum_j W_{ij}^t \lfloor r_j(t) \rfloor^2 / N}, \quad (6)$$

where W_{ij}^t was the normalization weight from neuron j to neuron i, and σ controlled the semi-saturation constant of divisive normalization. $\lfloor \cdot \rfloor$ denotes half-wave rectification, also known as a rectified linear unit (ReLU). We divided the normalization term by the number of neurons to make the normalization strength effectively independent of network size. The normalization weights were updated based on a Hebbian learning rule:

$$W_{ij}^{t+1} = W_{ij}^t + \alpha(R_i R_j - C_{ij}) + \beta(W_{ij}^t - C_{ij}). \quad (7)$$

The second term is the Hebbian learning rule, which increases normalization weights if pre- and post-synaptic firing rates are both high, with learning rate α. To reach the stable state faster, we started with a learning rate $\alpha = 10^3$ that quickly decayed exponentially to 1. The third term pushed weights back toward their initial values at leakage rate β. The same values C_{ij} were used as the initial normalization weights, the value to which they decayed (the third term) and the value of response product homeostatically controlled by the second term. These values were set to the average response products of the network in response to a uniform distribution of grating orientations for a network with fixed weights $W_{ij} = 1$.

To simulate the experiments, we presented the adapter for a given condition, and updated the weights repeatedly until the normalization weights, and therefore simulated responses, were stable. Then, we froze the normalization weights at the current values and determined the network response vector $\mathbf{R}$ for the mask alone as well as for the mask plus the target at each presented target contrast. To convert network responses to predicted behavioral response, we assumed the responses were perturbed by independent Gaussian noise with zero mean and standard deviation σ_n, and that the decision was made by an ideal observer as predicted by signal detection theory (Wickens, 2001). Thus, for each target stimulus we computed

$$d'_{\text{ideal}}(c) = \frac{||\mathbf{R}_{\text{mask}} - \mathbf{R}_{\text{mask+target}}||}{\sigma_n}. \quad (8)$$

Assuming a symmetric, unbiased decision criterion, the predicted percentage correct for our two-alternative, forced-choice task is

$$PC = \Phi\left(\frac{d'}{\sqrt{2}}\right), \quad (9)$$

where Φ is the standard, cumulative normal distribution. We repeated the above procedure for each candidate set of values for the parameters b, B, σ and β. Observers shared all of the model parameters, but we allowed σ_n to vary across observers to account for observed individual differences in threshold elevation. Since this noise parameter effectively scaled d', we did not require a parameter for the absolute contrast sensitivity (i.e., the peak height of the tuning curves). We selected the parameters by maximum likelihood using a fine-resolution grid search. Using these best-fit parameters, we simulated the stimulation paradigm of Experiment 1 for orientations ranging from ± 10 to ± 80 and generated model predictions of threshold contrast.

2.5.2. Experiment 2: Appearance

The model for Experiment 2 was similar to the model above with modifications to account for the center-surround structure of the stimuli and the contrast-matching task. Specifically, we included an additional population of “surround” neurons that were tuned to the surround grating. This “surround” population had the same size and tuning preferences as the “center” population used for Experiment 1, and only differed in receptive field location. Experiment 2 measured PSEs rather than thresholds. We took the mean of the center neurons’ responses as an estimate of apparent contrast. Observers shared all model parameters except for the semi-saturation constant σ, which varied to account for individual differences in PSE. For each set of model parameters, we computed the root mean squared error (RMSE) of predicted vs. observed PSEs over all observers. We selected the set of parameters that led to the smallest RMSE using a fine-resolution grid search.

3. Results

3.1. Experiment 1: Detection

Fig. 7A shows responses (binned) and fitted psychometric functions for all observers and orientations for each of the two adaptation conditions. Estimated contrast threshold (Fig. 7C, black points) was greater for contingent relative to asynchronous adaptation for all subjects in all conditions. The increased threshold for contingent adaptation is consistent with our hypothesis: contingent adaptation led to greater cross-orientation suppression, which in this task implied greater masking efficacy and hence elevated threshold. This effect was significant pooling over all 8 results ($t_7 = 3.85$, $p < .01$), but was not significant for each condition individually (± 45 deg condition: $t_3 = 2.91$, $p = .06$; ±15 deg condition: $t_3 = 2.22$, $p > .05$). At the individual subject level, this effect was significant only for two observers (GL, SL) in the ± 45 deg condition (parametric bootstrap for ± 15 deg condition: $p = 0.096$, 0.395, 0.106 and 0.19; ± 45 deg condition: $p = 0.012$, 0.197, 0.269 and 0.022), although the threshold in the contingent condition was greater than in the asynchronous condition in all eight data sets (4 subjects, 2 orientation pairs).

We simulated a neural network that followed the Hebbian normalization model with the stimulation paradigm of Experiment 1 (see Methods). The gratings of both adapters and test stimuli were input to a set of neurons with tuning curves spanning all orientations. The best-fit parameters for the network were $b = .01$, $\sigma = .19$, $B = 84$ and $\beta = 0$. The best-fitting value of parameter B, 84 deg, resulted in a half-height full-width bandwidth of 28 deg. The output noise scaling parameters σ_n for our four observers (GL, SL, HSI and HY) were 11.4, 12.1, 14.3 and 15.0, respectively. The RMSEs were 0.0044, 0.0041, 0.0040 and 0.0035. These parameter settings resulted in threshold elevations similar to those demonstrated experimentally (Fig. 7C, red points and Fig. 7E). Thus, our results are consistent with the predictions of the Hebbian normalization model.

With these model parameters, we simulated the stimulation paradigm for a wider range of orientations to extend our experimental findings. Fig. 7F shows predictions for observer GL (predictions for other observers were similar).

A single-neuron gain-control model would fail to predict these results. Each of the two orientations in each adapter was present 50% of the time. Thus, a single-neuron gain-control model would predict equivalent amounts of adaptation in the two conditions. Furthermore, if one includes cross-orientation suppression in such a gain-control model, contingent adaptation should lead to some cross-orientation suppression when the plaid adapter is displayed, weakening its effectiveness as an adapter. This predicts the opposite of what we found.

3.2. Experiment 2: Appearance

Fig. 7B shows responses (binned) and fitted psychometric functions for all observers and adaptation conditions. The estimated PSEs are shown in Fig. 7D (black points). As predicted, when the comparison matched the corresponding adapter (solid curves and circles), the comparison underwent greater surround suppression and thus required more contrast to appear to have the same contrast as the standard as compared to the condition when the comparison did not match the corresponding adapter (dashed curves and triangles). This effect is significant when pooled over all the data ($t_5 = 4.23$, $p < .01$), but not when testing for differences between thresholds separately for each condition (± 45 deg condition: $t_2 = 2.29$, $p = .15$; ±15 deg condition: $t_2 = 3.36$, $p = .08$). The result is significant for 5 of the 6 combinations

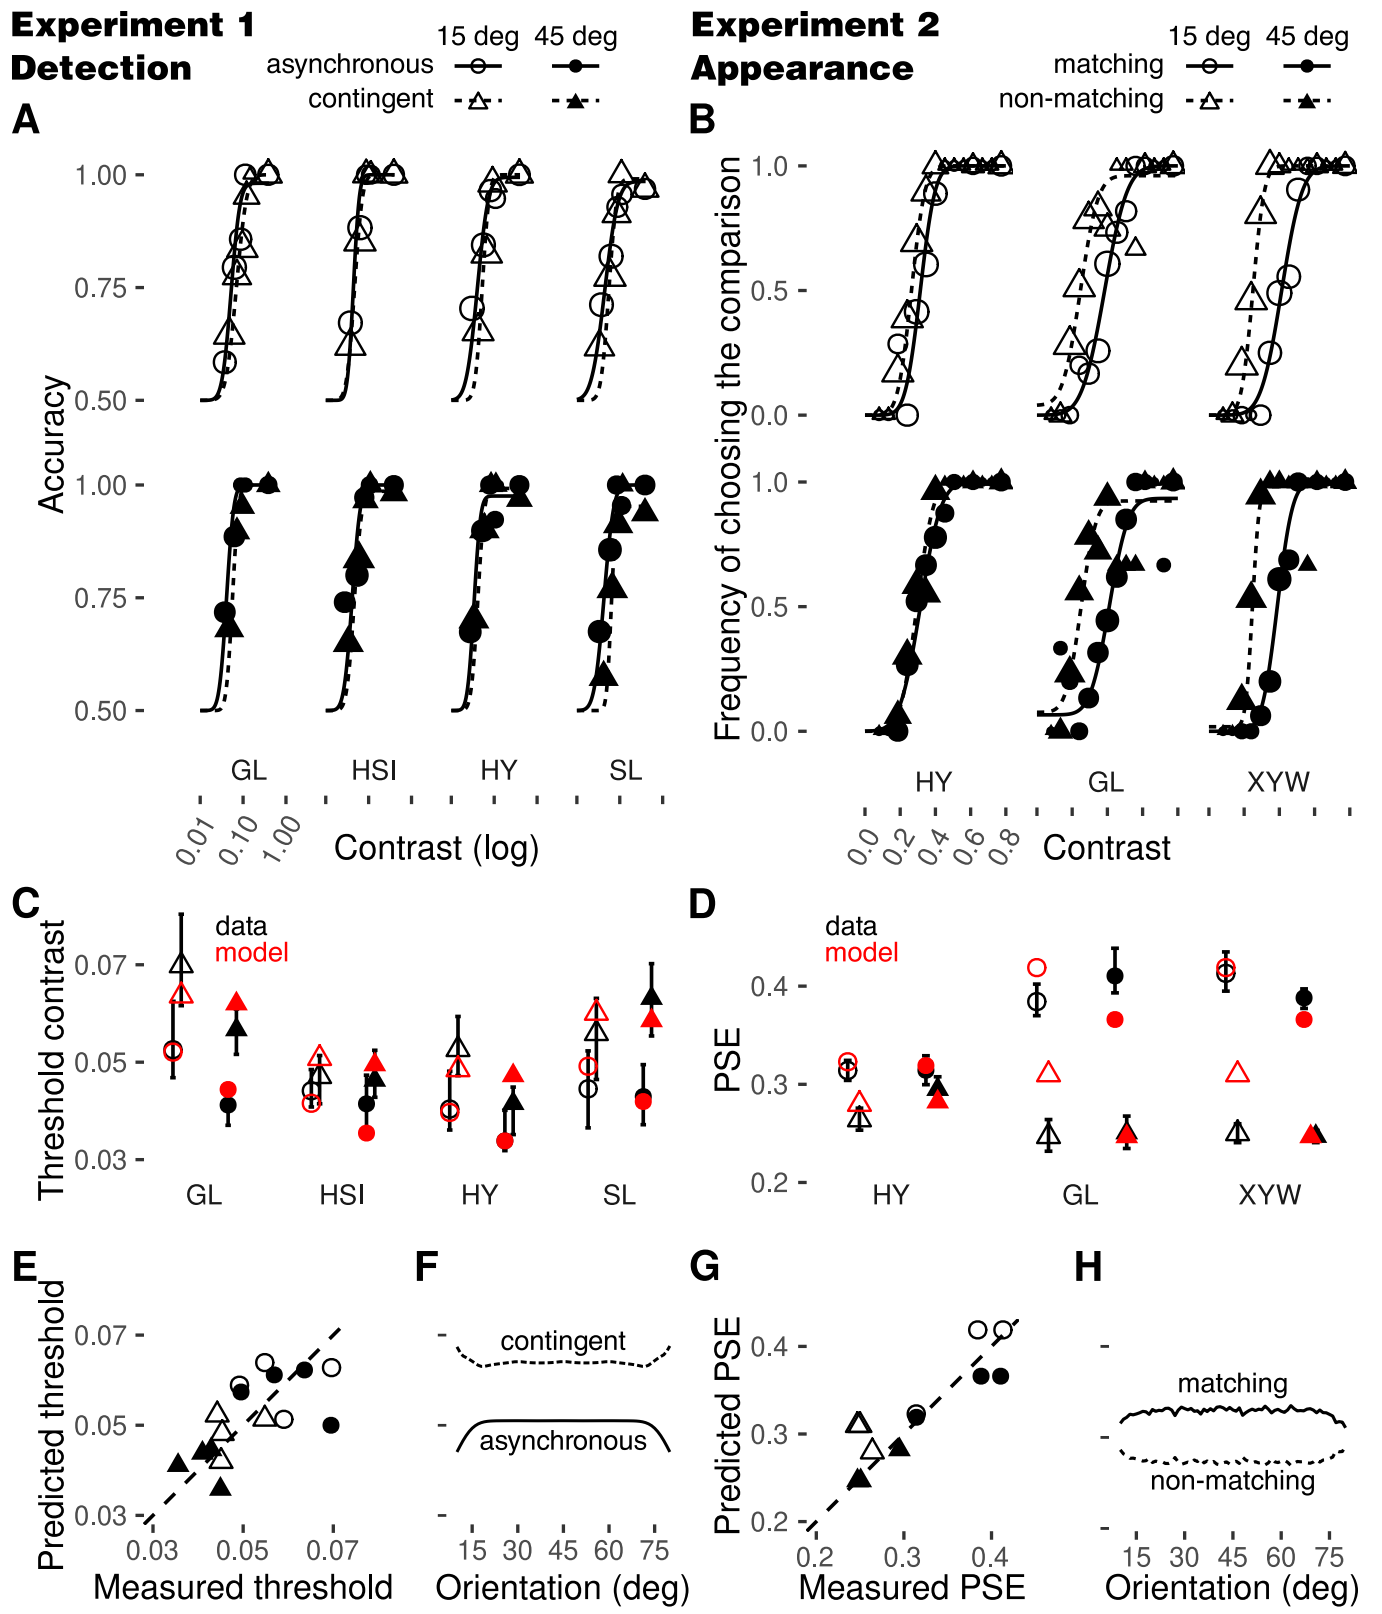

Fig. 7. Results for Experiment 1 – Detection (left, **A, C, E, F**) and Experiment 2 – Appearance (right, **B, D, G, H**). **A.** Proportion correct is plotted as a function of target contrast. Data were pooled into 6 log-contrast bins. Solid curves and circles: asynchronous adaptation; dashed curves and triangles: contingent adaptation. Curves: cumulative log-normal distributions fit to the data. **B.** Probability of perceiving the comparison stimulus as having higher contrast than the standard is plotted as a function of its physical contrast. Data were pooled into 15 contrast bins. Solid curves and circles: matching case; dashed curves and triangles: non-matching case. Curves: cumulative normal distributions fit to the data. **A, B.** Open symbols (upper panel): ± 15 deg orientations; filled symbols (lower panel): ± 45 deg orientations. Data point area is proportional to the log of the number of trials. **C, D.** Threshold estimates corresponding to 75% correct (left) and PSE estimates (right) of human observers (black) and model (red). Error bars: 16th to 84th percentiles of the parametric bootstrap threshold estimates with z-scores of ± 1. **E, G.** Model predictions versus data. **F, H.** Model predictions for a wider range of orientations for observer GL.

of subject and orientation condition tested separately (parametric bootstrap: ± 15 deg condition: all $p < .001$; ±45 deg condition: HY: $p = 0.16$, other two observers: $p < .001$). These results were based on pooling configurations in which both the standard and comparison stimuli were in the opposite-orientation configuration (Fig. 3, middle) and when they were in the same-orientation configuration (Fig. 3, bottom). The results did not change when examining each of these conditions separately (data not shown).

We simulated a neural network that followed the Hebbian normalization model with the stimulation paradigm of Experiment 2 (see Methods). The best-fit parameters for the network were $b = .16$, $B = 84$ and $\beta = 0$. The semi-saturation constants σ for our three observers (HY, GL and XYW) were .09, .03 and .03, respectively. The RMSEs were 0.011, 0.042 and 0.032. These parameter settings resulted in PSEs similar to those demonstrated experimentally (Fig. 7D, red points and Fig. 7G). Thus, our results are consistent with the predictions of the Hebbian normalization model.

With these model parameters, we simulated the stimulation paradigm for a wider range of orientations to extend our experimental findings. Fig. 7H shows predictions for observer GL (predictions for other observers were similar).

These results indicate that contingent adaptation affects appearance for stimuli at contrasts well above threshold. Again, since each orientation component occurred equally often in any of our conditions, a single-neuron gain-control model cannot explain our result.

4. Discussion

According to the classic single-neuron gain-control model, each neuron decreases its own gain in response to its own recent input history. Alternatively, the Hebbian normalization model predicts that a pair of neurons that co-fire will suppress each other more than they otherwise would have. Therefore, when two stimuli are presented simultaneously (contingent adaptation), pairs of neurons sensitive to each will co-fire and thus increase mutual suppression, leading to changes in the perception of subsequently presented stimuli. We designed two experiments to test this prediction. In both experiments, adapters consisted of pairs of gratings with orientations symmetric around vertical, hence balanced for oblique effects, and appeared equally often, hence balancing typical models of gain adaptation. However, we found that contingent adaptation increased masking effectiveness at detection threshold (overlap masking) and reduced apparent contrast in a center-surround configuration (i.e., increased surround suppression). It has been argued that masking and surround suppression are two distinct mechanisms in the early visual system (Petrov, Carandini, & McKee, 2005; Petrov & McKee, 2009). But, it appears that both effects have a similar underlying computation, such as Hebbian normalization, resulting in contingent adaptation effects.

Contingent adaptation requires the observer to be sensitive to second-order statistics of the adapting stimuli. Thus, one might argue that contingent adaptation is merely the consequence of well-known second-order texture mechanisms in spatial vision. This is an unlikely explanation for several reasons. Experiment 1 is a masking experiment in which one of the two plaid components has threshold contrast. At such contrast levels, first-order mechanisms dominate; second-order thresholds are generally much higher (Landy & Oruç, 2002). In Experiment 2, the stimulus pairs were non-overlapping and thus included no second-order pattern components, and yet we still found significant contingent-adaptation effects. Thus, contingent adaptation is distinct from second-order pattern vision.

It is possible that contingent adaptation emerges naturally as a result of connectivity between neurons within and across functionally similar areas, where connected neurons interact suppressively following the Hebbian rule: neurons that fire together wire together. However, this learning results in strengthened normalization, that is, suppression. As a result, this learning rule is effectively anti-Hebbian.

Threshold elevation for detection due to masking decreases as the orientation difference between masker and test increases (Meese & Holmes, 2010; Phillips & Wilson, 1984). Similarly, threshold elevation due to adaptation decreases as the orientation difference between adapter and test increases (Movshon & Blakemore, 1973). Physiological recordings from macaque V1 neurons show that adaptation to a plaid (the cell's preferred orientation and the orthogonal orientation) results in increased cross-orientation suppression (Carandini, Movshon, & Ferster, 1998).

Analogously, we found that threshold elevation due to contingent adaptation was significant for two individual subjects when adapted and tested with ± 45 deg gratings but not significant for ± 15 deg. Since the orientation tuning bandwidth of a V1 neuron is typically 15–30 deg, two input orientations 90 deg apart (in our ± 45 deg condition) excite distinct sub-populations of neurons, whereas there is a large overlap between the populations excited by the two grating components in our ± 15 deg condition. A greater difference in joint activation between contingent and asynchronous adaptation in the ± 45 deg condition predicts a correspondingly greater change in normalization weights and thus a greater threshold elevation, consistent with our results.

Neurons typically show peak surround suppression when center and extra-classical surround orientations both match the neuron's preferred orientation (Cavanaugh, Bair, & Movshon, 2002). This implies that surround suppression was stronger and could involve a greater number of neuron pairs for the ± 15 deg than the ± 45 deg opposite-orientation adapters. As a result, the model would predict a greater increase in surround suppression when adapting to ± 15 deg, resulting in reduced post-adaptation apparent contrast when the test configuration matched the adapter as shown in Fig. 7F and H. Our experimental results show a hint of stronger effects of contingent adaptation leading to reduced apparent contrast in the ± 15 deg than the ± 45 deg condition, but this was not statistically significant.

In our models, most of the model parameters were shared across all four observers, and we used a per-observer late-noise parameter to account for individual differences in Experiment 1. For Experiment 2, there was no need for a late-noise scaling parameter, hence all model parameters except the semi-saturation constant were shared across observers. Allowing the semi-saturation constant to differ across observers in Experiment 1 led to better fits, but here we reported the simpler model. Although model parameters were not shared across experiments, note that the best-fitting temporal leak parameter for both experiments was 0, and the best-fitting half-height full-width bandwidths were both 28 deg. This orientation tuning bandwidth is at the narrow end of the range for orientation bandwidth reported in the literature (Beaudot & Mullen, 2006; Blake & Holopigian, 1985; Campbell & Kulikowski, 1966; Cavanaugh et al., 2002; Movshon & Blakemore, 1973).

The existence of contingent adaptation suggests that adaptation in orientation-selective neurons in the early visual system depends on population activity rather than on the response history of individual neurons. In the Hebbian normalization model, the normalization pool is neuron-specific, and adaptation shapes the normalization pool based on network activity, in particular, in response to neural response products. This model, and the contingent-adaptation protocol are perfectly general, and might apply to other brain regions and other feature continua.

Note that earlier experiments have compared adaptation to gratings to adaptation to plaids (i.e., combinations of two grating components with different orientations) in V1 (Carandini, Barlow, O'Keefe, Poirson, & Movshon, 1997; Carandini et al., 1998), but were inconclusive due to the small number of cells examined. However, more recent data confirm an increase in normalization strength in response to contingent adaptation to plaids (Aschner et al., 2018). The perceptual effect of contingent adaptation is superficially similar to the McCollough effect (McCollough, 1965), a contingent adaptation after-effect of color and orientation. However, the McCollough effect acts at much longer time-

scales. Other cross-feature aftereffects combining luminance and orientation (McCollough, 1965), or color and orientation (Movshon & Lennie, 1979) have also been found. However, it is unclear whether and how functionally distinct neural populations (e.g., tuned to color vs. orientation) contribute to a shared normalization pool.

Our model assumes that contingent adaptation is explained by mechanisms with simple-cell-like tuning curves with a single preferred orientation. Could our results be explained by adaptation in a higher-level mechanism (e.g., in V2 or later) that is tuned to the plaid pattern used in the contingent adapter of Experiment 1? If this mechanism requires *both* orientations to respond, it will respond to the contingent but not the asynchronous adapter. However, such a mechanism is unlikely to be useful for the subsequent masked-detection task in which one grating is at threshold levels. If such a plaid mechanism responds when *either* orientation is present, it will behave similarly to a mechanism with an extremely broad orientation tuning bandwidth, and should be adapted by similar amounts by the contingent and asynchronous adapters, and thus will also fail to predict our data.

The experiments reported here demonstrate the existence of contingent adaptation in the orientation domain as a static phenomenon. But, the model is dynamic, and further tests of the model should investigate its time-varying behavior. Two time constants in the model are the learning rate for modifying normalization weights based on firing rate covariance as well as the window for determining whether activity is correlated. It will be straightforward to modify the current experimental designs to estimate these time constants and further test the predictions of the Hebbian normalization model.

Acknowledgements

This project was funded by the National Institutes of Health (grant EY 08266). The authors declare no competing financial interests.

References

Aschner, A., Solomon, S. G., Landy, M. S., Heeger, D. J., & Kohn, A. (2018). Temporal contingencies determine whether adaptation strengthens or weakens normalization. *Journal of Neuroscience, 38*, 10129–10142.

Barlow, H. B., & Földiák, P. (1989). Adaptation and decorrelation in the cortex. In R. Durbin, C. Miall, & G. Mitchison (Eds.). *The Computing Neuron* (pp. 54–72). Wokingham, England: Addison-Wesley.

Beaudot, W. H., & Mullen, K. T. (2006). Orientation discrimination in human vision: Psychophysics and modeling. *Vision Research, 46*, 26–46.

Benucci, A., Saleem, A. B., & Carandini, M. (2013). Adaptation maintains population homeostasis in primary visual cortex. *Nature Neuroscience, 16*, 724–729.

Blake, R., & Holopigian, K. (1985). Orientation selectivity in cats and humans assessed by masking. *Vision Research, 25*, 1459–1467.

Blakemore, C., & Campbell, F. W. (1969). On the existence of neurones in the human visual system selectively sensitive to the orientation and size of retinal images. *Journal of Physiology, 203*, 237–260.

Blakemore, C., & Nachmias, J. (1971). The orientation specificity of two visual after-effects. *Journal of Physiology, 213*, 157–174.

Blakemore, C., & Sutton, P. (1969). Size adaptation: A new aftereffect. *Science, 166*, 245–247.

Calzolari, E., Azañón, E., Danvers, M., Vallar, G., & Longo, M. R. (2017). Adaptation aftereffects reveal that tactile distance is a basic somatosensory feature. *Proceedings of the National Academy of Sciences USA, 114*, 4555–4560.

Campbell, F. W., & Kulikowski, J. J. (1966). Orientational selectivity of the human visual system. *The Journal of Physiology, 187*, 437–445.

Carandini, M., Barlow, H. B., O'Keefe, L. P., Poirson, A. B., & Movshon, J. A. (1997). Adaptation to contingencies in macaque primary visual cortex. *Philosophical Transactions of the Royal Society London B, Biological Sciences, 352*, 1149–1154.

Carandini, M., Movshon, J. A., & Ferster, D. (1998). Pattern adaptation and cross-orientation interactions in the primary visual cortex. *Neuropharmacology, 37*, 501–511.

Cavanaugh, J. R., Bair, W., & Movshon, J. A. (2002). Selectivity and spatial distribution of signals from the receptive field surround in macaque V1 neurons. *Journal of Neurophysiology, 88*, 2547–2556.

Clifford, C. W. G., Webster, M. A., Stanley, G. B., Stocker, A. A., Kohn, A., Sharpee, T. O., & Schwartz, O. (2007). Visual adaptation: Neural, psychological and computational aspects. *Vision Research, 47*, 3125–3131.

Condon, C. D., & Weinberger, N. M. (1991). Habituation produces frequency-specific plasticity of receptive fields in the auditory cortex. *Behavioral Neuroscience, 105*, 416–430.

Cornelissen, F. W., Peters, E. M., & Palmer, J. (2002). The Eyelink toolbox: Eye tracking with MATLAB and the Psychophysics Toolbox. *Behavioral Research Methods, Instruments & Computers, 34*, 613–617.

Dragoi, V., Sharma, J., & Sur, M. (2000). Adaptation-induced plasticity of orientation tuning in adult visual cortex. *Neuron, 28*, 287–298.

Földiák, P. (1990). Forming sparse representations by local anti-Hebbian learning. *Biological Cybernetics, 64*, 165–170.

Gibson, J. J., & Radner, M. (1937). Adaptation after-effect and contrast in the perception of tilted lines. I. Quantitative studies. *Journal of Experimental Psychology, 20*, 453–467.

Hahn, J. F. (1968). Low-frequency vibrotactile adaptation. *Journal of Experimental Psychology, 78*, 655–659.

Kleiner, M., Brainard, D. H., Pelli, D. G., Ingling, A., Murray, R. F., & Broussard, C. (2007). What's new in psychtoolbox-3. *Perception, 36*, 1–16.

Landy, M. S., & Oruç, I. (2002). Properties of 2nd-order spatial frequency channels. *Vision Research, 42*, 2311–2329.

Levitt, H. (1971). Transformed up-down methods in psychoacoustics. *Journal of the Acoustical Society of America, 49*, 467–477.

Li, L., Chan, A., Iqbal, S. M., & Goldreich, D. (2017). An adaptation-induced repulsion illusion in tactile spatial perception. *Frontiers in Human Neuroscience, 11*, 331.

Louie, K., Grattan, L. E., & Glimcher, P. W. (2011). Reward value-based gain control: Divisive normalization in parietal cortex. *Journal of Neuroscience, 31*, 10627–10639.

Maffei, L., Fiorentini, A., & Bisti, S. (1973). Neural correlate of perceptual adaptation to gratings. *Science, 182*, 1036–1038.

Maloney, L. T. (1990). Confidence intervals for the parameters of psychometric functions. *Perception & Psychophysics, 47*, 127–134.

McCollough, C. (1965). Color adaptation of edge-detectors in the human visual system. *Science, 149*, 1115–1116.

Meese, T. S., & Holmes, D. J. (2010). Orientation masking and cross-orientation suppression (XOS): Implications for estimates of filter bandwidth. *Journal of Vision, 10*(12), 9.

Movshon, J. A., & Blakemore, C. (1973). Orientation specificity and spatial selectivity in human vision. *Perception, 2*, 53–60.

Movshon, J. A., & Lennie, P. (1979). Pattern-selective adaptation in visual cortical neurones. *Nature, 278*, 850–852.

Müller, J. R., Metha, A. B., Krauskopf, J., & Lennie, P. (1999). Rapid adaptation in visual cortex to the structure of images. *Science, 285*, 1405–1408.

Ohzawa, I., Sclar, G., & Freeman, R. D. (1982). Contrast gain control in the cat visual cortex. *Nature, 298*, 266–268.

Pantle, A., & Sekuler, R. (1968). Size detecting mechanisms in human vision. *Science, 162*, 1146–1148.

Petrov, Y., Carandini, M., & McKee, S. P. (2005). Surround masking comes after cross-orientation masking and is only found in the periphery. *Journal of Vision, 5*(8), 226.

Petrov, Y., & McKee, S. P. (2009). The time course of contrast masking reveals two distinct mechanisms of human surround suppression. *Journal of Vision, 9*(1), 21.

Phillips, G. C., & Wilson, H. R. (1984). Orientation bandwidths of spatial mechanisms measured by masking. *Journal of the Optical Society of America A, 1*, 226–232.

Sekuler, R., & Pantle, A. (1967). A model for after-effects of seen movement. *Vision Research, 7*, 427–439.

Shapley, R. M., & Enroth-Cugell, C. (1984). Visual adaptation and retinal gain controls. *Progress in Retinal Research, 3*, 263–346.

Simoncelli, E. P., Freeman, W. T., Adelson, E. H., & Heeger, D. J. (1992). Shiftable multiscale transforms. *IEEE Transactions on Information Theory, 38*, 587–607.

Solomon, S. G., & Kohn, A. (2014). Moving sensory adaptation beyond suppressive effects in single neurons. *Current Biology, 24*, R1012–R1022.

Westrick, Z. M., Heeger, D. J., & Landy, M. S. (2016). Pattern adaptation and normalization reweighting. *Journal of Neuroscience, 36*, 9805–9816.

Wickens, T. D. (2001). *Elementary Signal Detection Theory*. New York: Oxford.