Visual motion aftereffects arise from a cascade of two isomorphic adaptation mechanisms

Alan A. Stocker   Department of Psychology, University of Pennsylvania, Philadelphia, PA, USA

Eero P. Simoncelli   Howard Hughes Medical Institute, Center for Neural Science and Courant Institute of Mathematical Sciences, New York University, New York, NY, USA

Prolonged exposure to a moving stimulus can substantially alter the perceived velocity (both speed and direction) of subsequently presented stimuli. Here, we show that these changes can be parsimoniously explained with a model that combines the effects of two isomorphic adaptation mechanisms, one nondirectional and one directional. Each produces a pattern of velocity biases that serves as an observable "signature" of the corresponding mechanism. The net effect on perceived velocity is a superposition of these two signatures. By examining human velocity judgments in the context of different adaptor velocities, we are able to separate these two signatures. The model fits the data well, successfully predicts subjects' behavior in an additional experiment using a nondirectional adaptor, and is in agreement with a variety of previous experimental results. As such, the model provides a unifying explanation for the diversity of motion aftereffects.

Introduction

Our visual system continually adapts to the recent history of visual input. Although we are generally not aware of it, this adaptation causes substantial changes in our perception. A prominent effect is that of *repulsion*: Adaptation exaggerates the perceived differences between stimuli and the adaptor. For example, in the well-known tilt aftereffect, adaptation to a vertical bar causes the orientation of a subsequently viewed near-vertical bar to appear further from vertical (Gibson, 1937). These changes, extensively documented in the experimental literature, provide a window into the properties of the underlying neural representations. The simplest and most accepted linking hypothesis is based on tuned populations of cells and gain changes (Blakemore, Nachmias, & Sutton, 1970; Sekuler & Pantle, 1967). Specifically, those neurons (or, more abstractly, "channels") that are activated by the adaptor reduce their gain, which leads to a characteristic repulsive bias in the perception of subsequently viewed stimuli (Clifford, Wenderoth, & Spehkar, 2000; Maffei, Fiorentini, & Bisti, 1973; Schwartz, Hsu, & Dayan, 2007). Additional physiological changes have been observed (e.g., changes in the shape of tuning curves), which, by themselves, may lead to different effects, but altogether the net effect seems to be consistent with repulsive perceptual aftereffects (Jin, Dragoi, Sur, & Seung, 2005; Kohn, 2007; Schwartz et al., 2007; Seriès, Stocker, & Simoncelli, in press). As such, regardless of the physiological details, we can think of the repulsive pattern of adaptation-induced biases as providing an externally observable perceptual "signature" for the underlying adaptation mechanism.

A modality in which adaptation has been particularly well studied is that of motion. Prolonged exposure to a moving stimulus leads to perceived illusory motion of static stimuli (Nishida & Johnston, 1999; Wright & Johnston, 1985; often referred to as the "waterfall illusion", Thompson, 1880), illusory motion of dynamic random stimuli (see, e.g., Hiris & Blake, 1992), biases in the perceived velocity of moving stimuli (Clifford & Wenderoth, 1999; Ledgeway & Smith, 1997; Mather, Verstraten, & Anstis, 1998; Schrater & Simoncelli, 1998; Smith, 1985; Smith & Edgar, 1994; Thompson, 1981), and changes in speed or direction discriminability (Clifford & Wenderoth, 1999; Kristjansson, 2001; Phinney, Bowd, & Patterson, 1997). Although these results all arise from adaptation to moving stimuli, the individual studies use stimuli with different spatial structures (e.g., dots, bars, gratings), different choices

for the velocities of adapting and test stimuli (especially with regard to moving versus stationary stimuli), and typically manipulate either speed or direction individually, rather than the joint two-dimensional quantity of velocity (with the exception of Schrater & Simoncelli, 1998). Furthermore, unlike adaptation-induced biases seen in other visual stimulus parameters, speed biases are asymmetric around the adaptor. As such, the published results are difficult to reconcile, and to date, there is no consistent and unified explanation for the observed motion aftereffects.

Here, we provide a characterization of the perceptual biases arising from motion adaptation over a broad range of test and adaptor velocities. We find that the resulting perceptual changes can be parsimoniously and accurately fit using not one but a *superposition of two signatures*, corresponding to two distinct underlying mechanisms—one directional and one nondirectional—thus providing a unified explanation for previously noted differences in the effects of adaptation on speed and direction.

Methods

Psychophysical experiments

Seven subjects with normal or corrected-to-normal vision participated in the psychophysical experiments. Protocols for selection of subjects and experimental procedures were approved by the human subjects committee of New York University and all subjects signed an approved consent form. All but one of the subjects
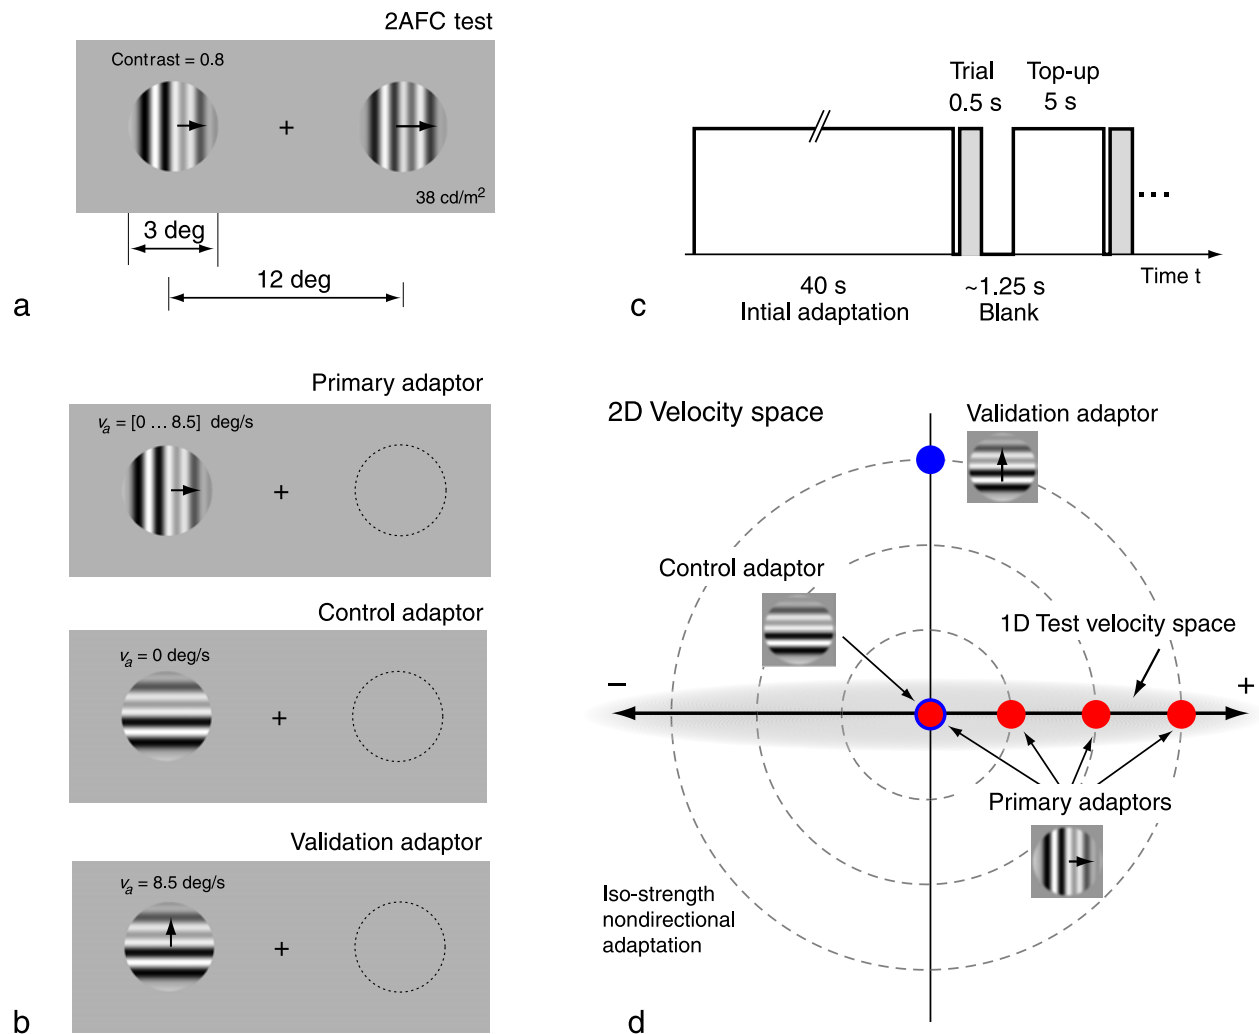

Figure 1. Methods. (a) Test stimuli and experimental setup. (b) Different adaptation conditions. (c) Timing diagram. (d) Two-dimensional velocity space and the locations of the different adaptors therein.

(number 1) were unfamiliar with the purpose of the study. Stimuli were high contrast (0.8) spatial broadband drifting gratings (frequency range from 1/3 cycles deg^{-1} to 2 cycles deg^{-1}) with phases randomized on each trial and spatial power spectrum falling as f^{-2} (see examples in Figure 1), identical to those used in one of our previous studies (Stocker & Simoncelli, 2006a). The mean luminance of gratings and background was held constant at 38 cd m^{-2}. Stimuli were presented on a CRT monitor running at 120-Hz refresh rate. We used the Psychophysics Toolbox for the controlled timing of the display (Brainard, 1997; Pelli, 1997). Stimulus size and eccentricity were as indicated in Figure 1a. Subjects' heads were stabilized using a chin rest. Each subject performed a primary, a control, and a validation adaptation experiment in a two-alternative-forced-choice (2AFC) paradigm (except subject 4, who did not complete the validation experiment). They were asked to fixate a central fixation mark while an adaptor stimulus was presented to the left of the fixation mark for 40 s. After this initial adaptation phase, a stimulus pair consisting of test (to the left of fixation, in the same location as the adaptor) and reference (to the right of fixation) was presented for 0.5 s, followed by a blank period of ~1.25 s in which the subject had to indicate which stimulus was moving faster by pushing an appropriate key. After a top-up period in which the original adaptation stimuli was presented again for 5 s, the next reference/test pair was presented. This trial sequence was identical for all experiments and is summarized in Figure 1c. The data were fit with a psychometric function (cumulative Gaussian) that characterizes the probability of a subject deciding that the reference is faster than the test, as a function of the reference velocity.

Primary experiment

The primary adaptation experiment was comprised of four separate temporal blocks, corresponding to adaptor velocities $v_a = [0.0, 0.8, 4.1, 8.5]$ deg s^{-1}. A period of at least 24 h was provided between blocks to allow complete recovery from the adaptation effects of previous blocks. Each block contained trials for 12 conditions, corresponding to test stimulus velocities of $v_t = \pm[0.8, 2.4, 4.1, 6.1, 8.5, 12]$ where negative velocity indicates leftward and positive indicates rightward motion (note that subject 1 was not tested at the highest speed). The reference stimulus had a variable speed v_r that was adjusted according to two interleaved adaptive staircase procedures, each starting from one end of the adaptive speed range of each condition. Staircases were of the type "one-up/one-down". Each test condition contained 60 trials, which were used to estimate a psychometric function for that condition. Individual trials for different test conditions were randomly interleaved.

Control experiment

The control adaptation experiment contained a single block with an adaptor stimulus that was identical in spatial frequency spectrum and contrast to the other adaptation stimuli but spatially rotated by 90 degrees and stationary (see Figures 1b and 1d). This control adaptor is completely ambiguous with regard to velocity along its grating orientation (the "aperture problem") and, as such, does not produce motion aftereffects in that direction. However, it retains all the other characteristics (e.g., contrast and temporal and spatial frequencies) of the adaptation stimuli used in the primary experiment. Thus it serves as a control to account for all possible nonmotion adaptation effects, such as apparent loss of contrast (Bex, Bedingham, & Hammett, 1999), as well as any potential spatial asymmetry between perception at the reference and test locations. The procedural sequence of this block was identical to that of the main experiment. All perceptual biases in the primary experiment were computed relative to the biases measured in this control experiment.

Validation experiment

The validation adaptation experiment was identical to the control experiment, with the exception that the rotated adaptor stimulus was not static but drifting upward with speed $v_a = 8.5$ deg s^{-1} (see Figures 1b and 1d).

Perceptual space

Throughout this paper we are using a transformed representation of stimulus motion. We (Stocker & Simoncelli, 2006a) and others (Nover, Anderson, & DeAngelis, 2005) have shown that a compressive nonlinear transformation, such as a normalized logarithm, allows a more homogeneous interpretation of both physiological and psychophysical data. Here, we used a circular representation of speed $\tilde{v} = 2 \arctan(v/r_0)$. This representation is very similar to the normalized logarithm with the difference that it is a closed (compact) space. For the findings presented here, the choice is not critical and is also relatively insensitive to the exact value of the parameter r_0 (here, $r_0 = 5$). However, we believe that a closed velocity space has some interesting implications for perception, which are beyond the scope of this paper.

Results

We adopt the traditional hypothesis that prolonged exposure to a single stimulus leads to a gain reduction in sensitivity of those sensory channels that are selectively responsive to the stimulus, and the strength of this gain
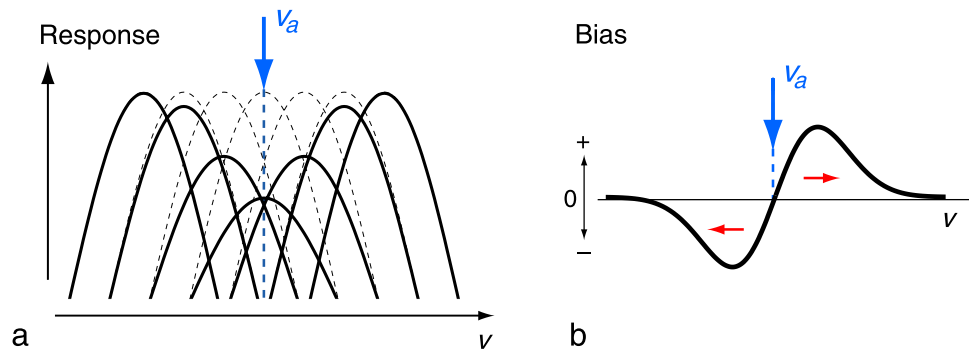
Figure 2. Signature of adaptation. (a) A generic mechanism for sensory adaptation: the response gain of channels that are selectively activated by the adaptor is reduced. The reduction is largest for channels that respond most. (b) Decoding such channel activity by taking the population average leads to a stereotypical *signature* in subsequent perception: the perceived value of stimulus v is repulsively biased away from the value of the adaptor, v_a. We define bias as the difference between the mean of perceived stimulus value v, before and after adaptation.

reduction is largest for channels that respond most strongly (Blakemore et al., 1970; see Figure 2a). A simple population average scheme for decoding channel activity leads to repulsive biases as shown in Figure 2b (see Blakemore et al., 1970; Clifford et al., 2000; Schwartz et al., 2007). The precise form of the bias curve depends on many details: the number of channels, their tuning, their noise properties, the relationship between gain reduction and response level to the adaptor, and the decoding method. However, these details are not of critical importance in the present context: we assume only that the bias curve takes on a characteristic repulsive form, which is anti-symmetric about the adaptor, and of limited extent. Aftereffects in many visual modalities (e.g., orientation, Gibson, 1937; spatial frequency, Blakemore et al., 1970) exhibit this behavior. As such, we interpret this repulsive bias as an externally visible manifestation of a single underlying physiological mechanism, a so-called *signature* of adaptation (see Supplementary material for an extension of this signature that includes also changes in the variability of the percept).

Psychophysical measurements

In order to examine the signature of motion adaptation, we measured the effect of adaptation on subsequent perception of visual velocity as described above. From the fit psychometric functions, we extracted perceptual biases, defined as the difference between mean perceived speed (point-of-subjective equality) under adapted and control conditions, which revealed a consistent pattern across all seven subjects. Three examples are shown in Figure 3, corresponding to those subjects that had complete data sets with the least trial variability.

For test stimuli moving in the same direction as the adaptor, we find that perceived speeds are typically reduced for test speeds equal to or smaller than the adaptor and are increased for higher speeds. This is in agreement with previously published motion adaptation results (see, e.g., Bex et al., 1999; Smith, 1985). However, we also find significant effects for test stimuli moving in the direction opposite to that of the adaptor, which have not been systematically examined in previous literature. For slow adaptor speeds, the perceived speed of test stimuli moving in the opposite direction is typically unchanged or slightly increased. However, as the adaptor speed increases, the perceived speed of opposite-moving tests is reduced. Bias amplitudes for test stimuli moving in either direction are seen to increase with adaptor speed. Finally, some subjects also show small but significant biases around a stationary adaptor, as has been previously reported (Ascher, Welch, & Festa, 1996).

Two isomorphic mechanisms account for motion aftereffects

The perceptual biases shown in Figure 3 clearly do not match the signature of Figure 2b. For all nonzero adaptor speeds, the bias is typically not zero at the location of the adaptor, and more generally, it is not symmetric around the adaptor. Perhaps this is not surprising. Visual motion is computed through a somewhat elaborate hierarchical spatio-temporal processing cascade in visual cortex. It seems plausible that adaptation could occur at several levels along this processing cascade,[1] and thus the observed perceptual changes may reflect the joint effects of multiple mechanisms (Langley & Anderson, 2007; Mather & Harris, 1998). Here, we propose that the observed adaptation effects represent a *superposition* of two signatures (as shown in Figure 2b) arising from a *directional* and a *nondirectional* mechanism.

To examine our hypothesis, we fit the full data set (i.e., the changes in mean perceived velocity for all test speeds, and under all four primary adaptation conditions) with a
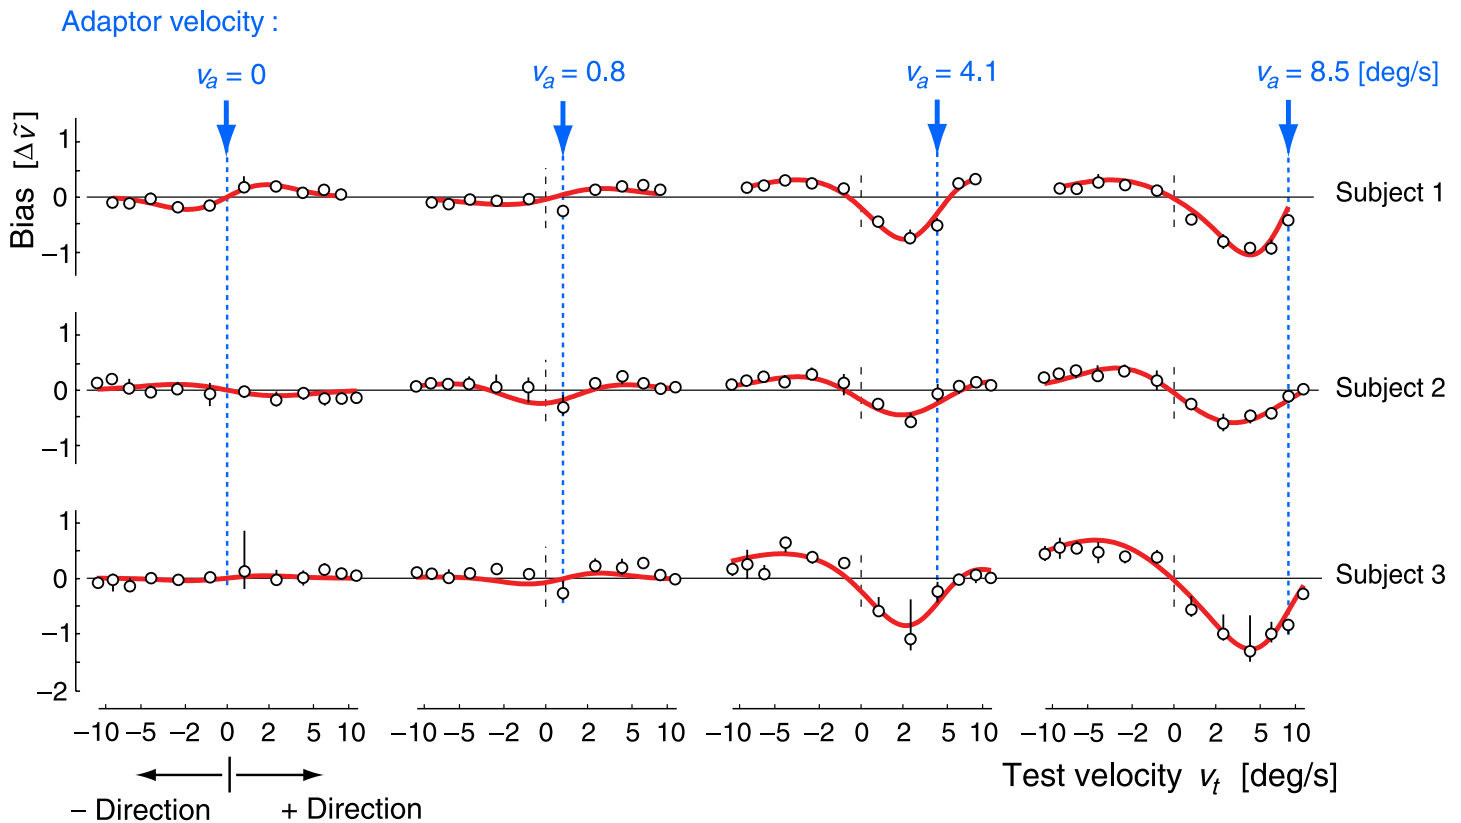
Figure 3. Perceptual estimation bias due to adaptation. Each row shows the perceptual biases of one subject, for each of four adaptation conditions, with adaptor velocities indicated by arrows and dashed vertical lines. Hollow points represent measured shifts in the point of subjective equality with respect to the control condition. Error bars represent the 5–95% quantiles of extracted biases, based on 1000 bootstrap samples of the data. Red bold lines indicate the bias of the two-mechanism model, fitted separately for each subject, using a total of 9 free parameters (see also Figure 4). Both axes represent nonlinearly transformed velocities (see Methods section), with horizontal and vertical distances scaled identically. Note that bias is a signed velocity difference; thus positive bias for negative test velocities indicates reduced perceived speed, while positive bias for positive test velocities indicates increased perceived speed. All subjects show similar behavior, with adaptation affecting subsequent perception over the whole range of test velocities, including test stimuli moving in the opposite direction.

single model consisting of a linear combination of two adaptation signatures, as illustrated in Figure 4a. We chose single-parameter functions (first derivatives of Gaussians of width s_D and s_{ND}, respectively) for modeling each signature of perceptual bias (Figure 2). We have also examined other function choices (with more parameters), which, however, did not significantly alter the main results presented here. An implicit feature of a nondirectional mechanism is its symmetry about zero motion. As a consequence, we assumed that the nondirectional signature was always centered at zero motion while the directional signature was centered at the adaptor. In addition, we constrained the widths of both signatures to be fixed across all adaptation conditions while the amplitudes were allowed to vary. Fits were computed by minimizing the total squared error.

As shown in Figure 3 the model provides a good account of the data. Figures 4b and 4c show the estimated model parameters for the three sample subjects and the average subject computed over the data of all seven subjects. The width of the nondirectional signature (s_{ND}) is approximately twice as large as the directional one (s_D). The amplitude of the directional signature (w_D) increases monotonically with adaptor speed, saturating for some subjects. The amplitude of the nondirectional signature (w_{ND}) becomes increasingly negative at higher adaptor speeds, indicating an increasingly attractive bias toward zero motion. In the Discussion section, we show that the fit superposition of signatures is consistent with a more mechanistic model based on a sequential cascade of two different populations of motion processing channels. Note that for a zero motion adaptor, the signatures are both centered at zero, and aside from their different widths, there is little to distinguish their relative contributions. Because in this case both the control adaptor and the zero-motion adaptor presumably produce equal activation (and adaptation) of the nondirectional mechanism, we assume that the measured net bias (adapted minus control) is
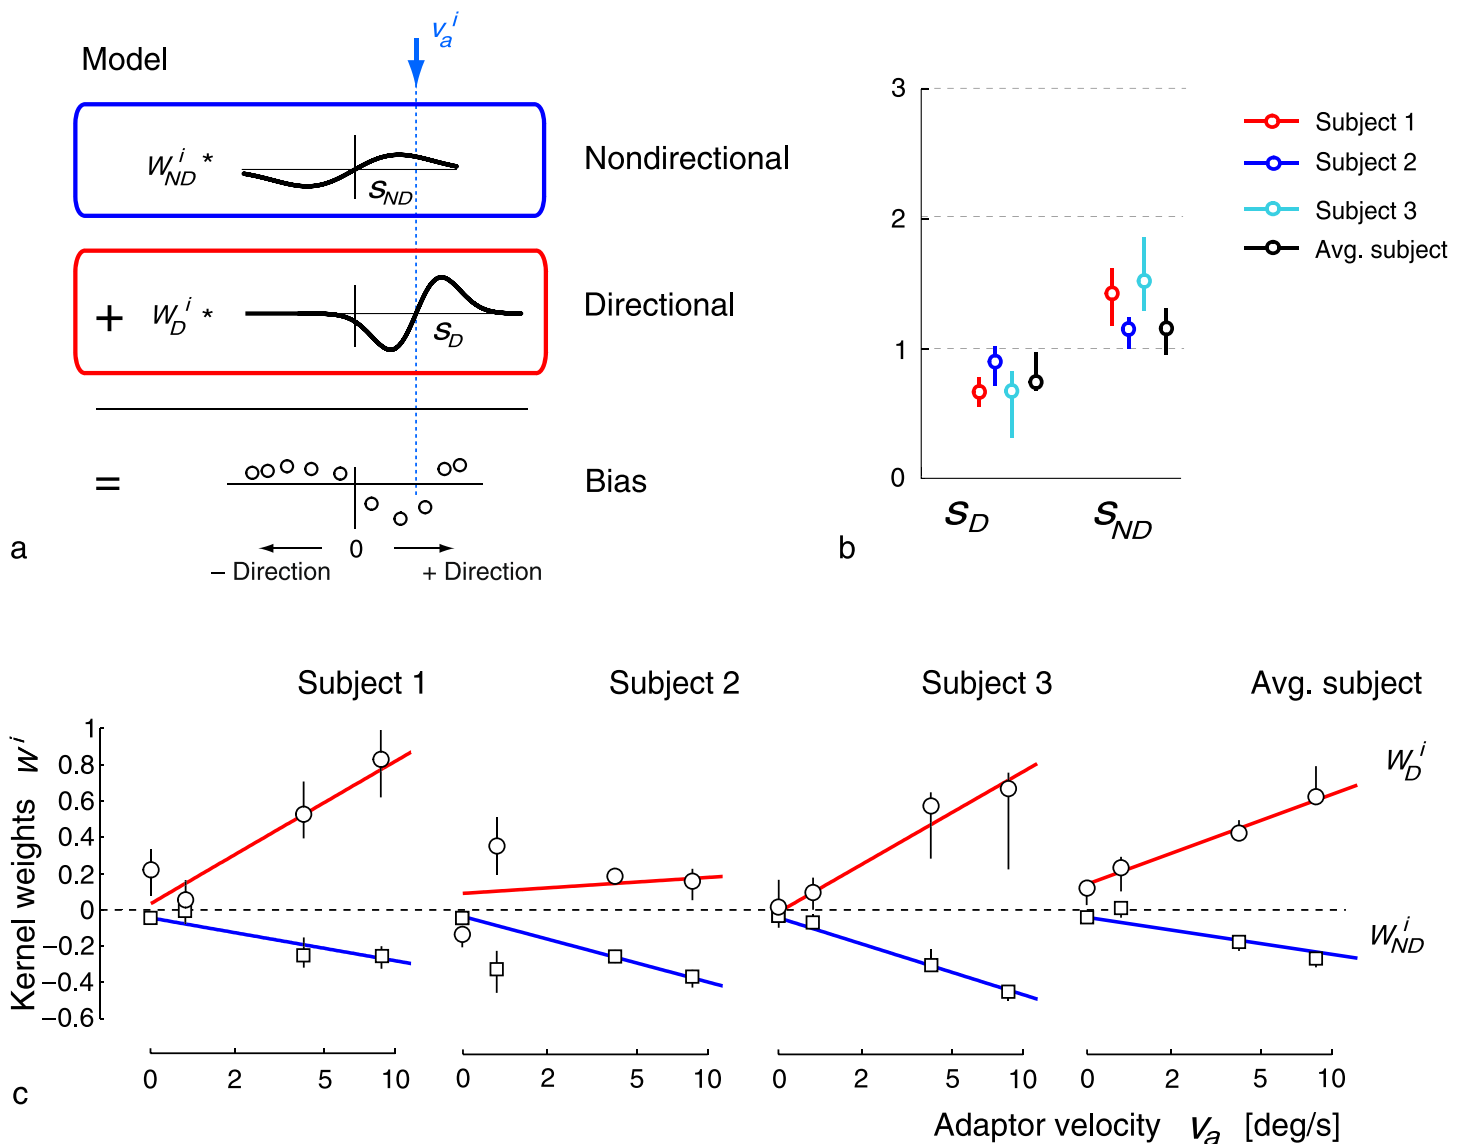
Figure 4. Superposition of a directional and a nondirectional signature. (a) Perceptual bias is modeled as a weighted sum of two signatures (Figure 2b): a nondirectional signature centered at zero motion and a directional signature centered at the adaptor. For each subject, we jointly fit the measured estimation bias gathered under all adaptation conditions. We constrained the signature widths s_D and s_{ND} (standard deviation of the Gaussians) to be constant across all adaptation conditions, while the fitted amplitudes w_D^i and w_{ND}^i were allowed to change with adaptor velocity v_a^i. Fits are shown as bold red lines in Figure 3. Note that for the nondirectional signature, the fitted weights were predominantly negative, which flips the polarity of the signature. (b) The estimated signature widths for each of the three sample subjects and the averaged subject (data of all seven subjects) suggest that the directional mechanism is constructed from more narrowly tuned channels than the nondirectional mechanism. (c) For the tested velocity range, the amplitude of the directional signature (circles) is typically positive and monotonically increasing with adaptor velocity v_a, while the amplitudes of the nondirectional signature (squares) become increasingly negative with increasing adaptation speed. For $v_a = 0$, the amplitude of the nondirectional signature was set to zero because the control and the stationary adaptor stimuli are activating the nondirectional mechanism equally. Bold lines (red and blue) represent linear fits to the estimated amplitude weights that are used to compute the predictions shown in Figure 7. Error bars indicate the 5–95% quantiles over 1000 bootstrap samples of the data.

entirely due to the directional signature, and thus the weight of the nondirectional signature is zero. Note that this assumption is specific to this particular control condition: adaptation effects in the nondirectional mechanism could be revealed under different control conditions.

Model comparison

We compare the fits of our two-mechanism model to those obtained for each of these mechanisms in isolation. The signatures of the two-mechanism model were constrained to

have a fixed shape across all four adaptors (resulting in a total of 9 free parameters), but we allowed each of the partial models the flexibility to fit each condition separately, in order to give them full opportunity to explain the data. We also compared to a model proposed by Hammett, Champion, Morland, and Thompson (2005), which accounts for adaptation biases in speed estimates based on a comparison of two physiologically inspired channels. The four free

parameters of this model were also optimized to fit the data of each subject.

Figure 5a shows the performance of the different models for all seven subjects, indicated as mean squared error relative to the error of the two-mechanism model. The two-mechanism model performs best for each individual subject, which is also true when we account for the different number of degrees of freedom by
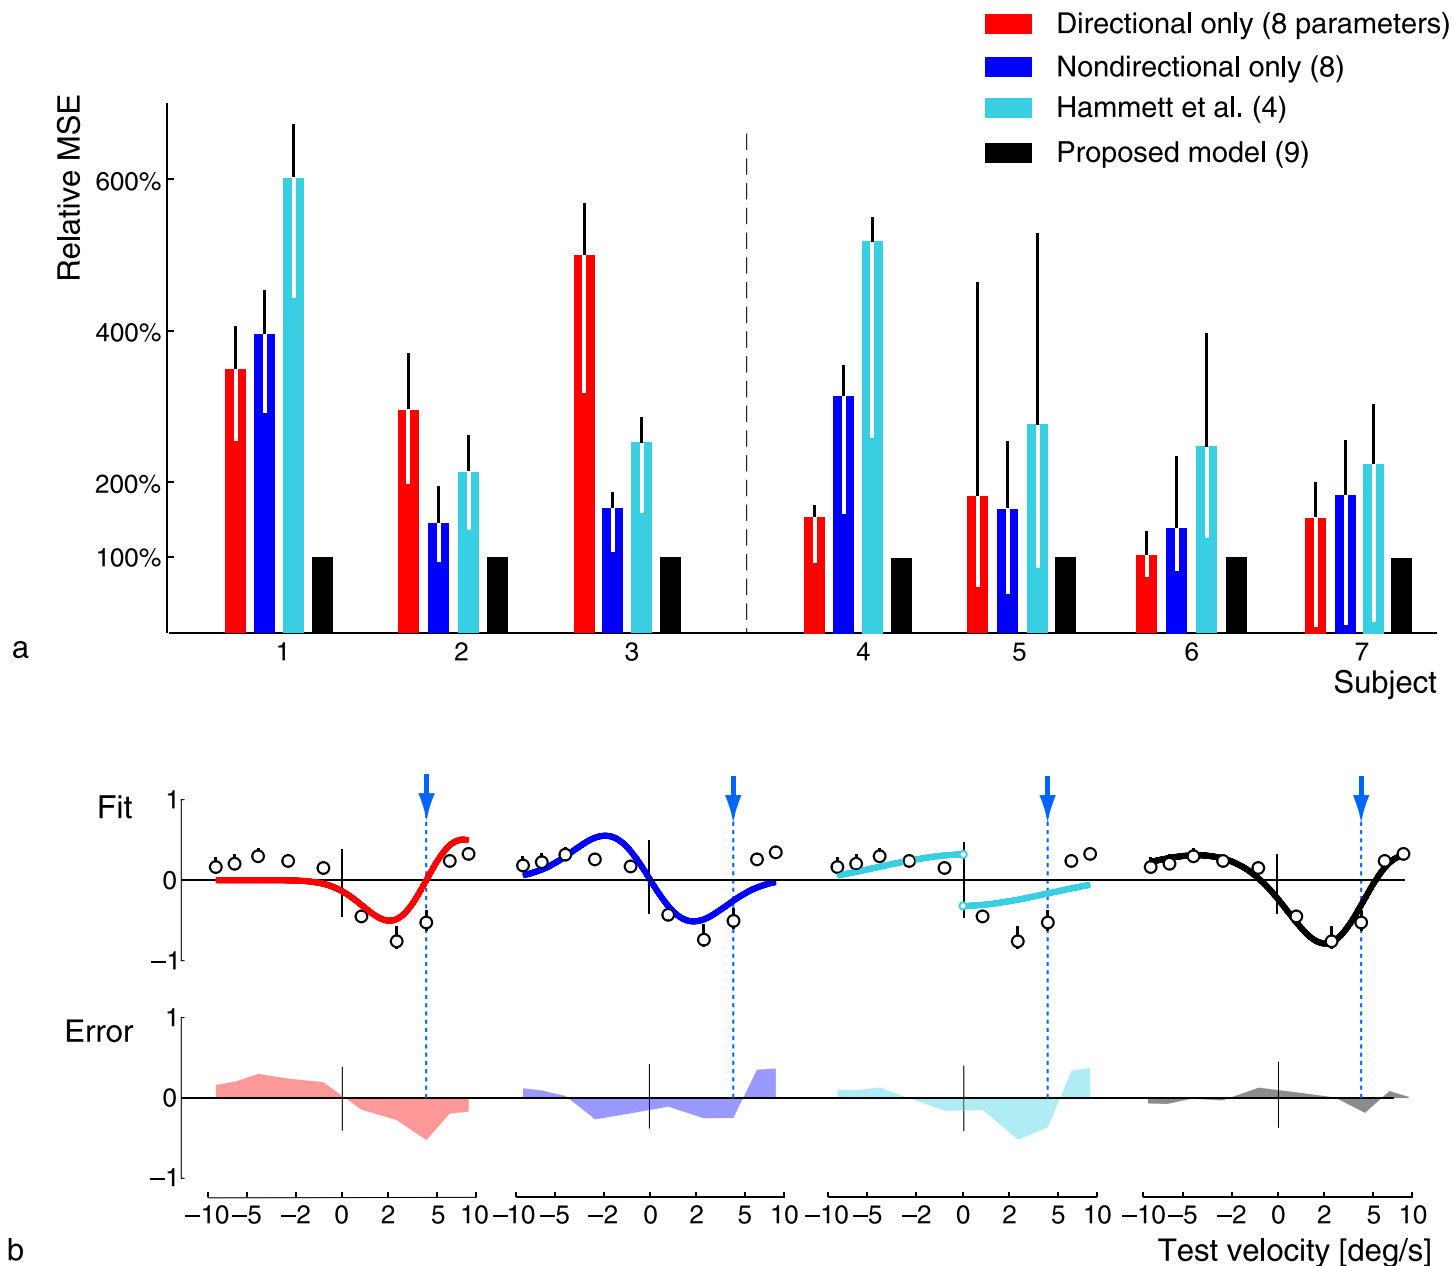

Figure 5. Goodness of fit and comparison to alternative, single-mechanism models. (a) Mean squared error of fitted alternative models, relative to the error of our proposed two-mechanism model. Error bars indicate the 5–95% quantiles for 1000 bootstrap samples of the data. Across all subjects, the two-mechanism model consistently outperforms the alternative models. (b) Fitted models and their fitting error, illustrated for a single subject and adaptation condition (subject 1, v_a = 4.1 deg s^{-1}). Note that the model by Hammett et al. (2005) is not defined at zero velocity.

performing an information criterion analysis (not shown). While each of the two partial models (directional/non-directional) might provide acceptable fits for some of the subjects, they each fail to account for the whole range of subjects' behavior, a fact that can be readily understood by examining the fits for a single adaptation condition, as shown in Figure 5b. The figure illustrates that the directional mechanism cannot account for bias effects at the adaptor and for velocities in opposite directions, while the nondirectional mechanism alone is not capable of accounting for the asymmetry across directions. The same drawback is seen in the model proposed by Hammett et al., which is also nondirectional. Moreover, the Hammett model produces a bias curve that is discontinuous at the origin (the model output is not defined for stationary stimuli).

Model validation

As a means of validating the model, we developed and ran an experiment with an adaptor designed to isolate the nondirectional mechanism. We used the same adaptor stimulus as in the control condition (grating rotated by 90 degrees relative to the original control adaptor) but drifting upward with speed equal to the fastest adaptor used in the primary adaptation experiment (see Figures 1b and 1d). According to our model hypothesis, this stimulus should drive the nondirectional mechanism as strongly as its 90-degree rotated counterpart from the original experiment. However, it should not engage the directional mechanism in the (horizontal) test direction, since its motion is completely ambiguous along this direction (aperture problem).

We tested each subject using this nondirectional adaptor, and then compared the resulting biases to those predicted by the nondirectional signature whose parameters were extracted from the primary experiment (Figure 4). As Figure 6 shows, the predictions are reasonably well matched to the measured biases of all three subjects shown in Figure 3, supporting the two-mechanism model.[2]

Predictions

The fitted model can be used to make a number of predictions. First, we can predict the classical static motion aftereffect (sMAE) results for our subjects. We performed a linear fit to the estimated model parameters (shown in Figure 4c) in order to continuously predict the perceived speed (in the opposite direction) of a static test stimulus after adaptation, i.e., the static motion aftereffect, as well as the speed of a "nulling" test stimulus (at which the test stimulus appears stationary after adaptation). The results are shown in Figure 7a. Both the shape and the peak speed of the predicted sMAE are similar for our subjects and remarkably consistent with previously reported data, despite substantial differences in stimulus details (adaptation duration, eccentricity, spatial frequency, etc.; Nishida & Johnston, 1999; Wright & Johnston, 1985). A key property of the model that allows prediction of the sMAE is the fact that the signatures are continuously and smoothly defined over the entire velocity space, including zero motion. Perhaps more importantly, the good qualitative match to the experimental data demonstrates that effects observed for zero motion (stationary) adapting stimuli can be viewed as consistent with those
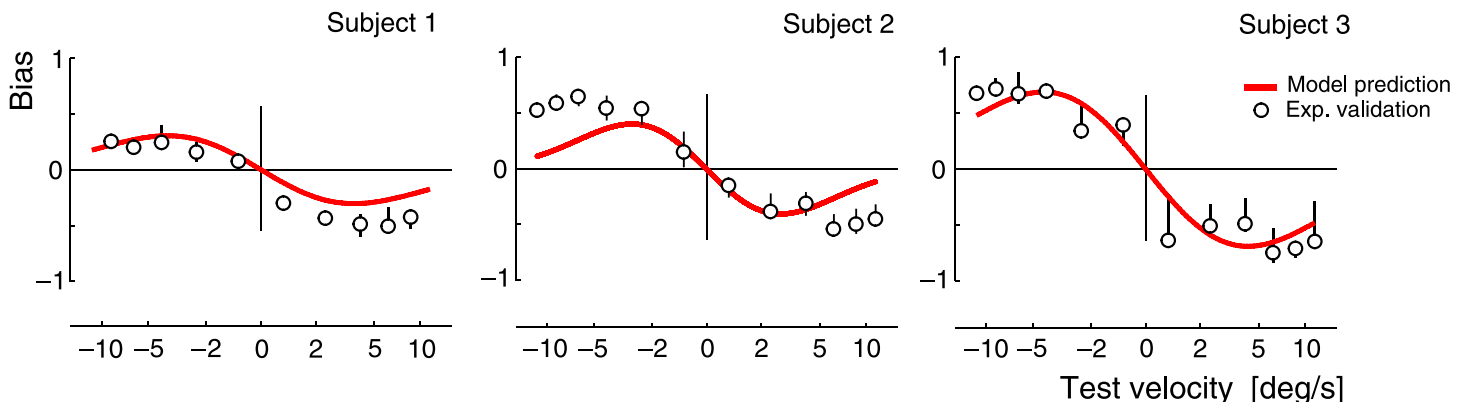

Figure 6. Predictions of the model for a nondirectional adaptor and experimental validation. Circles indicate measured biases after adaptation to the vertically drifting adaptor. The model predicts that this adaptor only engages the nondirectional mechanism. Bold lines represent the model prediction for the three subjects based on their individually estimated model parameters from the primary adaptation condition. The measured biases confirm the model prediction. In particular, they are all anti-symmetric about zero motion, as expected for the signature of a nondirectional mechanism. For subject 2, the prediction slightly underestimates the effect at high speeds.
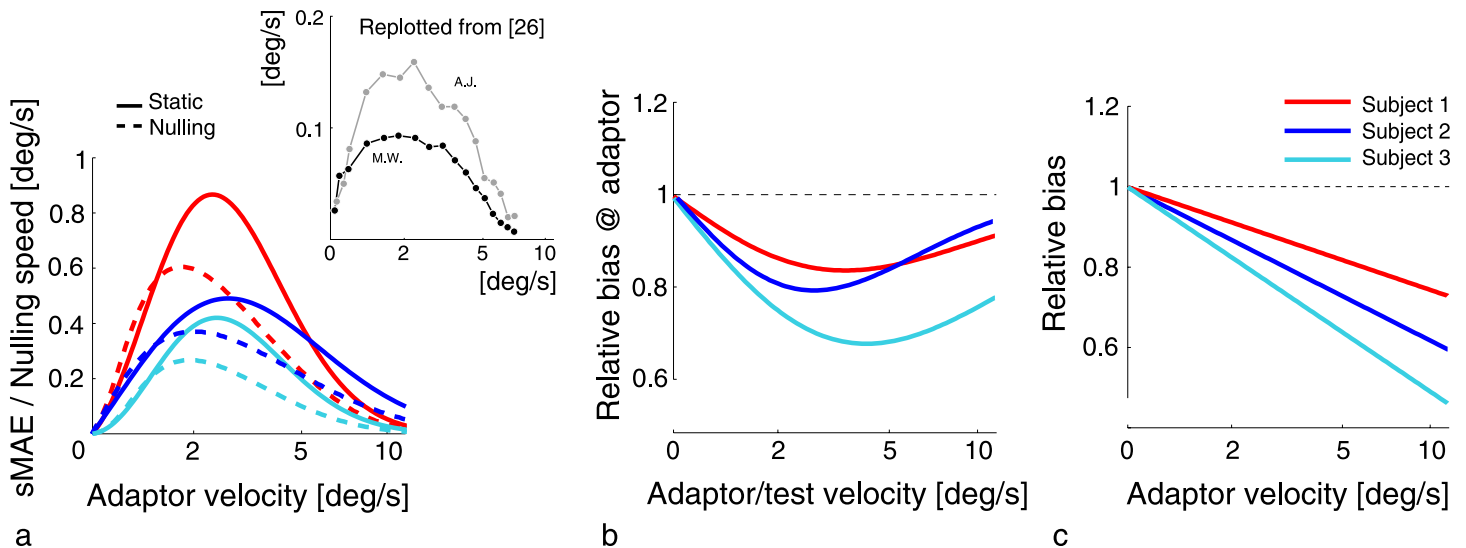
Figure 7. Predictions derived from estimated model parameters of individual subjects. (a) Classical motion aftereffects as a function of the adaptor speed. Solid lines indicate predicted perceived speed of a stationary test stimulus (in the direction opposite to the adaptor), as a function of adaptor speed. Dashed lines indicate the speed of a test stimulus (moving in the same direction as the adaptor) that "nulls" the aftereffect (i.e., makes the test stimulus appear to have no net motion). Inset shows measured nulling speeds for two subjects as reported by Wright and Johnston (1985; data replotted from their Figure 5). (b) Predictions of relative bias at the adaptor, as a function of adaptor speed. (c) Predictions of bias for a 6 deg s^{-1} test stimulus after exposure to a nondirectional adaptation stimulus. Note that all of the above bias predictions are plotted relative to the control adaptation condition.

observed for moving adapting stimuli, thus suggesting a continuous perceptual representation of velocity that includes zero.

Next, we can predict biases for test stimuli at the adaptor, as a function of adaptor velocity (Figure 7b). Because the directional signature is zero at the adaptor, biases for test stimuli moving at the same velocity as the adaptor are determined solely by the nondirectional signature. Thus, the model predicts that perceived speed at the adaptor is always reduced (Thompson, 1981). However, as predicted by the shape of the nondirectional signature the magnitude of this decrease is largest for intermediate adaptor speeds and fades for very low or high adaptor speeds.

The fact that perceived bias at the adaptor reflects only the nondirectional mechanism leads to a related set of interesting predictions. In particular, the aftereffects observed for test stimuli with the same velocity as the adaptor should be observable for a variety of other adapting stimuli, including some that are nondirectional. For example, when using sine-wave gratings of a fixed spatial frequency, a counter-phase grating adaptor should produce the same aftereffect at the adaptor as the grating adaptor. This has been verified by Clifford and Wenderoth (1999).

Last, and more generally, we can predict the bias for any fixed directional test stimulus velocity (here, 6 deg/s) after adaptation to a nondirectional adaptor, as a function of equivalent adaptor speed (see Figure 7c). We do not rule out the possibility that the biases in this case can be positive at very low adaptor speeds, when compared against a different control condition, i.e., when not constraining the amplitude of the nondirectional signature to be zero for a zero velocity adaptor.

Discussion

We have shown that the perceptual effects of adaptation to stimuli moving at a wide range of velocities (including zero motion) can be explained as a superposition of effects arising from two isomorphic adaptation mechanisms, one of which is directional and the other nondirectional. We posit that each mechanism leads to a characteristic change in perceived velocity that we refer to as the *signature of adaptation*. We have demonstrated that the superimposed signatures of the two mechanisms can account for perceptual biases measured over a broad range of speeds and across two opposite directions of motion. This two-mechanism model is significantly better in explaining the measured data than alternative models that assume only a single adaptation mechanism. We have validated our model by making an accurate quantitative prediction of individual subjects' velocity biases measured in a separate experiment using an adaptor designed to isolate the nondirectional mechanism.

We have also used the model to generate a set of predictions for aftereffects beyond those tested here. We find that predictions of subjects' static and nulling motion aftereffects for a continuous range of adaptor speeds are in agreement

with previously reported measurements of the static MAE (Nishida & Johnston, 1999; Wright & Johnston, 1985). In addition, we find that predicted biases of a test stimulus moving at the speed of the adaptor are always toward slower speeds, as is generally assumed (see, e.g., Thompson, 1981), but that they vanish for both low and high adaptor speeds. Last, because the model distinguishes directional and nondirectional effects, we can predict the form of after-effects arising from nondirectional adaptors. In total, this two-mechanism model serves to unify a broad and disjoint set of results from the literature on motion adaptation.

The nondirectional portion of our model is consistent with a previously developed mechanistic model for motion perception known as the "ratio model" (Hammett et al., 2005; Harris, 1980; Perrone & Thiele, 2002; Tolhurst, Sharpe, & Hart, 1973). In its simplest form, this model consists of a single mechanism with two temporal frequency channels, one low-pass and one high-pass,

whose relative responses may be used to estimate stimulus speed but not direction. Previous studies have proposed that differential gain changes in the two channels would be sufficient to account for motion aftereffects. Smith (1985), for example, shows that biases due to flicker adaptation can be qualitatively explained this way. This explanation is consistent with our model (see Figure 7c), since a flicker adaptor would affect the nondirectional mechanism only. In addition, Hammett et al. (2005) collected bias data for directional adaptors but used only a single test direction. They showed that their data could be approximately fit with the ratio model. However, as we have demonstrated (see Figure 5), this model cannot account for the asymmetry in the observed estimation biases across directions, and thus provides a relatively poor fit to data over the whole velocity range.

Our analysis does not specify the physiological details nor loci of the two mechanisms, but the basic response
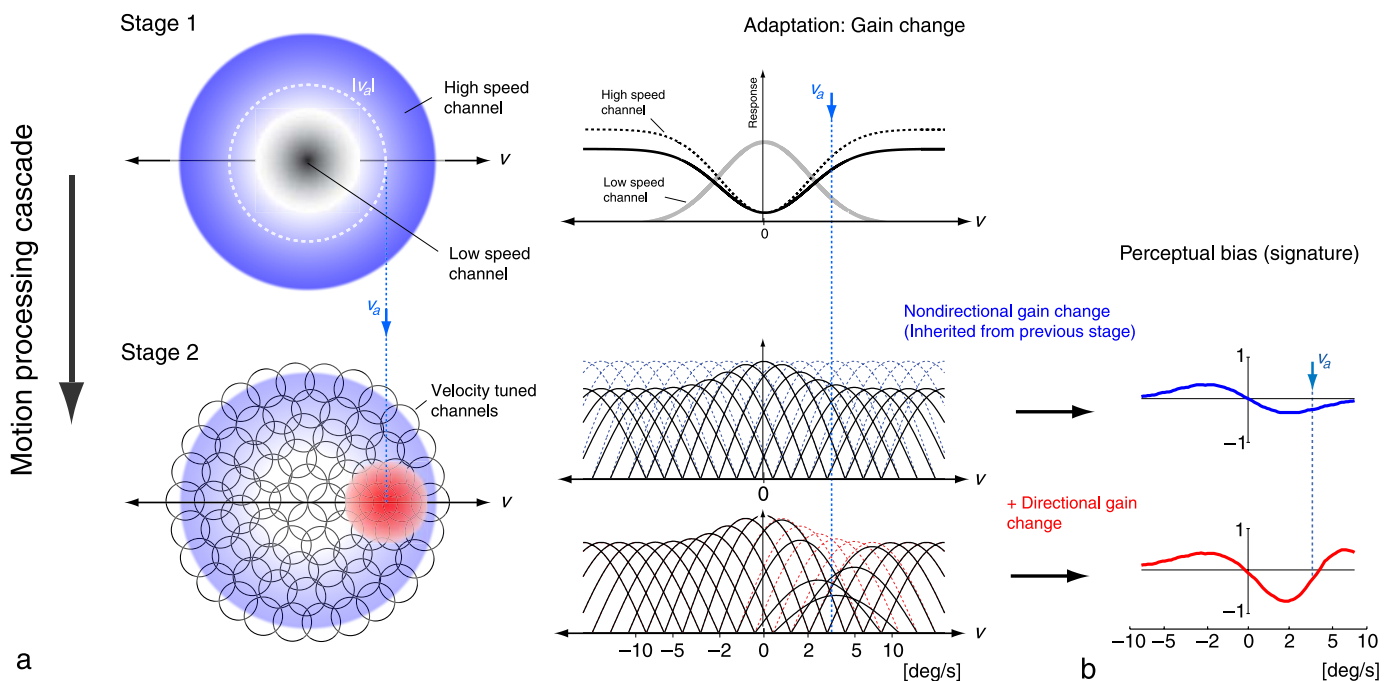

Figure 8. Simulation of adaptation effects arising from gain reduction at two stages of a velocity processing cascade. (a) The first stage consists of two (classes of) nondirectional channels, tuned for low and high speeds. In the illustrated example, adaptation changes the relative response gain of the high speed channel (blue region). Channels of the first stage then provide input to a second processing stage containing more narrowly tuned direction-selective velocity channels, which inherit the adaptation induced gain changes from the first stage but also undergo additional (directional) gain changes themselves (red region). The percept is decoded from the second stage. The middle column shows one-dimensional slices along the horizontal velocity axis. (b) Simulation results for the case of inherited gain reduction from the first stage only (blue), as well as the full model with gain changes in both stages (red). The population consisted of 21 equally spaced channels with rectified cosine tuning curves and equal tuning width when plotted in the chosen perceptual space. The mean response of each channel is determined by its tuning curve (shown in (a)) and response variability was assumed to follow a Poisson process. The perceptual estimate was assumed to be the population average, computed as $\hat{v}(v_s) = \sum_i v_0^i r^i v_s / \sum_i r^i v_s$, where v_0^i is the preferred velocity of channel i and r^i is its response to a stimulus with velocity v_s. The model qualitatively predicts the subjects' perceptual bias as shown in Figure 3. Results represent the average bias measured over 30,000 samples of the population response for each test velocity (see Supplementary material for simulation results for estimation variability). The gain reduction profile was Gaussian, with width twice and amplitude half as large for the nondirectional than for the directional mechanism, roughly matched to the fitted parameters from our subjects' data (see Figure 4b).

properties of neurons in the dorsal stream of the primate visual pathway suggest a plausible instantiation in terms of a sequential processing cascade. Figure 8a illustrates an implementation of such a cascade. The first stage consists of a population of broadly tuned nondirectional speed-tuned channels,[3] similar to those posited for the ratio model. These might correspond to early visual neurons, e.g., those in the lateral geniculate nucleus (LGN). Recent studies show that, in contrast to earlier claims, magnocellular cells in the LGN adapt to stimuli of rather low temporal frequency (Solomon, Peirce, Dhruv, & Lennie, 2004). The output of this first stage provides the input to a second stage of directional velocity-tuned channels. This second stage encodes velocity using more narrowly tuned direction-selective neurons (e.g., direction-selective neurons in cortical areas V1 or MT). The perceived velocity is decoded from this second stage. During adaptation, the neurons in each stage undergo a reduction in response gain according to their own response level. After adaptation, the net gain change seen in the second stage is the superposition of (nondirectional) changes inherited from the first stage, along with the (directional) gain changes that occur within the second stage.

To illustrate more precisely how this can lead to perceptual aftereffects, we simulated a simple implementation of this cascade model using a population average read-out rule (i.e., the sum of the peak tuning speed of each channel, weighted by the response of that channel, normalized by the total population response) to link channel responses to percept. The simulated perceptual biases arising from the nondirectional mechanism alone are in close agreement with the proposed signature of the nondirectional mechanism (blue curves in Figure 8b). Incorporating adaptation-induced gain changes in the second (directionally tuned) stage leads to a superposition of two signatures, as exhibited by our subjects. The simulated biases (red curve in Figure 8b) are in good qualitative agreement with the data shown in Figure 3. The cascade model also serves to explain some of the characteristics of the fitted parameters (see Figure 4). For example, the increasingly negative weights for the non-directional signature with increasing adaptor speed directly follow from the gain change in the first stage of the processing cascade: as adaptor speed increases, the gain of the high-speed channel is increasingly reduced relative to that of the low-speed channel.[4]

In addition to explaining adaptation-induced biases, we can also make a prediction of the adaptation-induced changes in the variability of the perceived stimulus velocity by assuming an appropriate noise model for the channels. If we assume that each channel's response corresponds to a sample from a Poisson distribution with rate determined by the associated tuning curve, the underlying adaptation mechanism will not only generate a signature in estimation bias (mean of the estimate) but an additional signature in the changes of estimation variability. We show simulations of this in the Supplementary material and demonstrate that this extended signature is consistent with the measured discrimination thresholds of our subjects.

Although the main experiment and analysis is based on data gathered for horizontal motion only, the results of the validation experiment suggest a generalization to the full two-dimensional motion (velocity) space and provide a natural explanation for the seemingly different motion aftereffects observed with respect to direction and speed. We assume that the adaptation signature will only reveal itself perceptually if the underlying adapted neural population is directly involved in solving the perceptual task on the test stimuli. Thus, in experiments for which subjects report only on the direction of a test stimulus, we would expect that one should see only the signature of the directional mechanism (i.e., repulsive biases in perceived direction away from that of the adaptor; Levinson & Sekuler, 1976) and not that of the nondirectional mechanism. If subjects are asked to report on both speed and direction of a test stimulus, we expect the nondirectional mechanism to generate additional effects, and previously reported joint aftereffects in direction and speed are qualitatively consistent with this (Schrater & Simoncelli, 1998). Thus, the model provides further evidence that direction and speed are not separately encoded in the brain (as has been suggested; Curran & Benton, 2006; Matthews & Qian, 1999) but are jointly encoded as a vector entity (Gardner, Tokiyama, & Lisberger, 2004; Heeger, 1987; Schrater & Simoncelli, 1998; Simoncelli & Heeger, 1998).

Although our model simulations are based on the parsimonious assumption that adaptation results from response-driven gain changes, our notion of a perceptual signature is agnostic about the detailed nature of the underlying mechanism. We note that a variety of physiological effects other than gain changes have been measured and modeled (Jin et al., 2005; Kohn, 2007; Schwartz et al., 2007; Seriès et al., in press). It has been suggested that some of these observed effects may arise from changes that occur at stages of processing other than the one being measured (Kohn & Movshon, 2004), analogous to the changes in the second stage of our model that are partly inherited from the first stage. Taking this concept to its logical extreme, we can speculate that all physiologically observed adaptation effects (including changes in tuning curve shape or position and changes in noise properties) might be explained as originating from adaptation-induced gain changes that have occurred at different stages of the system.

Finally, if adaptation is governed by a single, isomorphic mechanism that is repeated at different levels of a processing cascade, it is worth considering how that single mechanism can be understood from a normative perspective. Several functional roles for adaptation (which are not mutually exclusive) have been suggested: the minimization of correlations in sensory encoding (Barlow & Földiak, 1989; Wainwright, Schwartz, & Simoncelli, 2002), optimization of information transmission (Brenner,

Bialek, & de Ruyter van Steveninck, 2000; Fairhall, Lewen, Bialek, & de Ruyter van Steveninck, 2001; Wainwright, 1999), or optimization of perceptual discriminability and constancy (Abrams, Hillis, & Brainard, 2007; see also Clifford et al., 2007; Grzywacz & Balboa, 2002; Stocker & Simoncelli, 2006b). While it remains unclear which of these roles best explains both the physiological and psychophysical data, we believe that such an overarching principle is essential for understanding how the brain can perform reliable and consistent computation in the face of the constantly fluctuating adaptation state of its processing elements.

Acknowledgments

Footnotes

[1]Xu, Dayan, Lipkin, and Qian (2008) have recently presented an example of multi-level adaptation in face processing.

[2]Note that for subjects 1 and 3, there was almost a two-year interval between the original experiment (from which the model parameters are estimated and the prediction computed) and the validation experiment.

[3]Whether this first stage is tuned for speed or temporal frequency cannot be conclusively determined based on our experiments.

[4]Note that the weights of the nondirectional signature are not positive at low adaptor speeds because the perceptual changes are measured relative to the control stimulus condition (see Figure 1a), which, presumably, also adapts the low-speed channel.

References

Abrams, A., Hillis, J., & Brainard, D. (2007). The relation between color discriminability and color constancy: When is optimal adaptation task dependent? *Neural Computation, 19,* 2610–2637. [PubMed] [Article]

Ascher, D., Welch, L., & Festa, E. (1996). Adaptation to stationary gratings results in an increase in apparent speed of moving gratings. *Investigative Ophthalmology and Visual Science, 37,* S917.

Barlow, H., & Földiak, P. (1989). The computing neuron. In R. Durbin, C. Miall, & G. Mitchinson (Eds.), *Adaptation and decorrelation in the cortex* (pp. 54–72). New York: Addison-Wesley.

Bex, P. J., Bedingham, S., & Hammett, S. (1999). Apparent speed and speed sensitivity during adaptation to motion. *Journal of Optical Society of America, 16,* 2817.

Blakemore, C., Nachmias, J., & Sutton, P. (1970). The perceived spatial frequency shift: Evidence for frequency-selective neurons in the human brain. *The Journal of Physiology, 210,* 727–750. [PubMed] [Article]

Brainard, D. (1997). The psychophysics toolbox. *Spatial Vision, 10,* 433–436. [PubMed]

Brenner, N., Bialek, W., & de Ruyter van Steveninck, R. (2000). Adaptive rescaling maximizes information transmission. *Neuron, 26,* 695–702. [PubMed]

Clifford, C. W., 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,* 3125–3131. [PubMed]

Clifford, C. W., & Wenderoth, P. (1999). Adaptation to temporal modulation can enhance differential speed sensitivity. *Vision Research, 39,* 4324–4332. [PubMed]

Clifford, C. W., Wenderoth, P., & Spehkar, B. (2000). A functional angle on some after-effects in cortical vision. *Proceedings of the Royal Society of London B: Biological Sciences, 267,* 1705–1710. [PubMed] [Article]

Curran, W. W., & Benton, C. (2006). Test stimulus characteristics determine the perceived speed of the dynamic motion aftereffect. *Vision Research, 46,* 3284–3290. [PubMed]

Fairhall, A., Lewen, G., Bialek, W., & de Ruyter van Steveninck, R. (2001). Efficiency and ambiguity in an adaptive neural code. *Nature, 412,* 787–792. [PubMed]

Gardner, J., Tokiyama, S., & Lisberger, S. (2004). A population decoding framework for motion aftereffects on smooth pursuit eye movements. *Journal of Neuroscience, 24,* 9035–9048. [PubMed] [Article]

Gibson, J. (1937). Adaptation, after-effect, and contrast in the perception of tilted lines. *Journal of Experimental Psychology, 20,* 553–569.

Grzywacz, N., & Balboa, R. (2002). A Bayesian framework for sensory adaptation. *Neural Computation, 14,* 543–559. [PubMed]

Hammett, S. T., Champion, R. A., Morland, A. B., & Thompson, P. G. (2005). A ratio model of perceived speed in the human visual system. *Proceedings of the Royal Society of London B: Biological Sciences, 272,* 2351–2356. [PubMed] [Article]

Harris, M. (1980). Velocity specificity of the flicker to pattern sensitivity ratio in human vision. *Vision Research, 20,* 687–691.

Heeger, D. J. (1987). Model for the extraction of image flow. *Journal of the Optical Society of America A, Optics and Image Science, 4,* 1455–1471. [PubMed]

Hiris, E., & Blake, R. (1992). Another perspective on the visual motion aftereffect. *Proceedings of the National Academy of Sciences, 89,* 9025–9028. [PubMed] [Article]

Jin, D., Dragoi, V., Sur, M., & Seung, H. (2005). Tilt aftereffect and adaptation-induced changes in orientation tuning in visual cortex. *Journal of Neurophysiology, 94,* 4038–4050. [PubMed]

Kohn, A. (2007). Visual adaptation: Physiology, mechanisms, and functional benefits. *Journal of Neurophysiology, 97,* 3155–3164. [PubMed] [Article]

Kohn, A., & Movshon, J. (2004). Adaptation changes the direction tuning of macaque MT neurons. *Nature Neuroscience, 7,* 764–772. [PubMed]

Kristjansson, A. (2001). Increased sensitivity to speed changes during adaptation to first order, but not to second-order motion. *Vision Research, 41,* 1825–1832. [PubMed]

Langley, K., & Anderson, S. (2007). Subtractive and divisive adaptation in visual motion computations. *Vision Research, 47,* 673–686. [PubMed]

Ledgeway, T., & Smith, A. (1997). Changes in perceived speed following adaptation to first-order and second-order motion. *Vision Research, 37,* 215–224. [PubMed]

Levinson, E., & Sekuler, R. (1976). Adaptation alters perceived direction of motion. *Vision Research, 16,* 779–781. [PubMed]

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

Mather, G., & Harris, J. (1998). The motion aftereffect: A modern perspective. In G. Mather, F. Verstraten, & S. Anstis (Eds.), *Theoretical models of the motion aftereffect* (pp. 157–186). Cambridge, MA: MIT Press.

Mather, G., Verstraten, F., & Anstis, S. (Eds.) (1998). *The motion aftereffect.* Cambridge, MA: MIT Press.

Matthews, N., & Qian, N. (1999). Axis-of-motion affects direction discrimination, not speed discrimination. *Vision Research, 39,* 2205–2211. [PubMed]

Nishida, S., & Johnston, A. (1999). Influence of motion signals on the perceived position of spatial pattern. *Nature, 397,* 610–612. [PubMed]

Nover, H., Anderson, C., & DeAngelis, G. (2005). A logarithmic, scale-invariant representation of speed in macaque middle temporal area accounts for speed discrimination performance. *Journal of Neuroscience, 25,* 10049–10060. [PubMed] [Article]

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

Perrone, J., & Thiele, A. (2002). A model of speed tuning in MT neurons. *Vision Research, 42,* 1035–1051. [PubMed]

Phinney, R., Bowd, C., & Patterson, R. (1997). Direction-selective coding of stereoscopic (cyclopean) motion. *Vision Research, 37,* 865–869. [PubMed]

Schrater, P., & Simoncelli, E. (1998). Local velocity representation: Evidence from motion adaptation. *Vision Research, 38,* 3899–3912.

Schwartz, O., Hsu, A., & Dayan, P. (2007). Space and time in visual context. *Nature Reviews Neuroscience, 8,* 522–535. [PubMed]

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

Seriès, P., Stocker, A., & Simoncelli, E. (in press). Is the homunculus 'aware' of sensory adaptation?. *Neural Computation.*

Simoncelli, E., & Heeger, D. (1998). A model of neuronal responses in visual area MT. *Vision Research, 38,* 743–761. [PubMed]

Smith, A. (1985). Velocity coding: Evidence from perceived velocity shifts. *Vision Research, 25,* 1969–1976. [PubMed]

Smith, A., & Edgar, G. (1994). Antagonistic comparison of temporal frequency filter outputs as a basis for speed perception. *Vision Research, 34,* 253–365. [PubMed]

Solomon, S., Peirce, J., Dhruv, N., & Lennie, P. (2004). Profound contrast adaptation early in the visual pathway. *Neuron, 42,* 155–162. [PubMed]

Stocker, A., & Simoncelli, E. (2006a). Noise characteristics and prior expectations in human visual speed perception. *Nature Neuroscience, 9,* 578–585. [PubMed]

Stocker, A., & Simoncelli, E. (2006b). Sensory adaptation within a Bayesian framework for perception. In Y. Weiss, B. Schölkopf, & J. Platt (Eds.), *Advances in neural information processing systems* (vol. 18, pp. 1291–1298). Cambridge, MA: MIT Press.

Thompson, P. (1981). Velocity after-effects: The effects of adaptation to moving stimuli on the perception of subsequently seen moving stimuli. *Vision Research, 21,* 337–345.

Thompson, S. (1880). Optical illusions of motion. *Brain, 3,* 289–298.

Tolhurst, D., Sharpe, C., & Hart, G. (1973). The analysis of the drift rate of moving sinusoidal gratings. *Vision Research, 13,* 2545–2555. [PubMed]

Wainwright, M. (1999). Visual adaptation as optimal information transmission. *Vision Research, 39,* 3960–3974. [PubMed]

Wainwright, M. J., Schwartz, O., & Simoncelli, E. P. (2002). Natural image statistics and divisive normalization: Modeling nonlinearity and adaptation in cortical neurons. In R. Rao, B. Olshausen, & M. Lewicki (Eds.), *Probabilistic models of the brain: Perception and neural function* (chap. 10, pp. 203–222). Cambridge, MA: MIT Press.

Wright, M., & Johnston, A. (1985). Invariant tuning of motion aftereffect. *Vision Research, 25,* 1947–1955. [PubMed]

Xu, H., Dayan, P., Lipkin, R., & Qian, N. (2008). Adaptation across the cortical hierarchy: Low-level curve adaptation affects high-level facial-expression judgments. *Journal of Neuroscience, 28,* 3374–3383. [PubMed] [Article]