Vision Research 117 (2015) 91–99

Contents lists available at ScienceDirect

Vision Research

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

Tilt aftereffect due to adaptation to natural stimuli

Ron Dekel, Dov Sagi *

Department of Neurobiology, The Weizmann Institute of Science, Rehovot 76100, Israel

ARTICLE INFO

Article history:
Received 20 August 2014
Received in revised form 17 August 2015
Accepted 26 October 2015

Keywords:
Adaptation
Natural scenes
Natural image statistics
Orientation processing
Psychophysics
Serial dependence

ABSTRACT

The human visual system continuously adjusts to the current environment. To investigate these adjustments, biases in observers' perceptions owing to changes in the visual environment are measured (visual aftereffects). Typically, the stimuli used are synthetic and are composed of oriented patterns such as lines or gratings. These patterns are known to activate individual neurons in the visual cortex, but cover only a small subset of actual visual stimulations. To overcome this drawback, recent research has focused on synthetic patterns that mimic several aspects of natural stimulation. However, the aftereffects of natural stimulation per-se remain largely unexplored. Here, we interleaved presentations of unmodified natural image adaptors, selected according to criteria favoring content at a particular orientation, with presentations of targets that test a perceived orientation. This allowed us to measure the change in the perceived orientation, namely the tilt aftereffect (TAE), which resulted from repeated image presentations. Results show a close to standard TAE with adaptor durations around 500 ms, which is reduced with longer presentations. Importantly, our method can be generalized to investigate other aftereffects by selecting images differently.

1. Introduction

Visual adaptation can be defined as the adjustment of visual processing that occurs in response to changes in visual input (Clifford et al., 2007). Typically, research about visual adaptation is performed with oriented stimuli. This is motivated by the selectivity of visual neurons (Hubel & Wiesel, 1968), and by early psychophysical results (Blakemore & Campbell, 1969; Gibson & Radner, 1937). Adaptation to such stimuli leads to a change in the response properties of early visual neurons, as measured by electrophysiology (Kohn, 2007), and to visual aftereffects that are measured psychophysically (Webster, 2011). Such studies are typically performed using synthetic patterns. For example, research about the tilt aftereffect (TAE), which is the change in the perceived orientation for stimuli near the adaptor orientation, is typically performed using synthetic lines (Gibson & Radner, 1937), gratings (Campbell & Maffei, 1971; Mitchell & Muir, 1976), or Gabor patches (Knapen, Rolfs, Wexler, & Cavanagh, 2010). However, natural vision is more complex, and the visual system operates differently for synthetic and natural input (Alam, Vilankar, Field, & Chandler, 2014; Carandini et al., 2005; Olshausen & Field, 2005). Even for synthetic patterns with natural Fourier power-spectra (Geisler, 2008; Van der Schaaf & van Hateren, 1996), obtained by randomizing the phase of the Fourier transform of natural movies or images, this holds (Froudarakis et al., 2014; Goddard, Clifford, & Solomon, 2008).

Therefore, several studies have investigated adaptation with natural stimuli, for example, the change in contrast sensitivity resulting from exposure to natural movies or images (Bex, Solomon, & Dakin, 2009; Webster & Miyahara, 1997). In such experiments, the statistics of the stimuli are approximately the statistics in natural vision. It is therefore interesting how the visual system adapts to different statistics, for example different second-order statistics. Indeed, some studies investigated this by exposing observers to distorted natural stimuli (Bao & Engel, 2012; Haak, Fast, Bao, Lee, & Engel, 2014; Zhang, Bao, Kwon, He, & Engel, 2009), but no study has investigated this with natural stimuli that were not modified.

Here, we interleaved presentations of unmodified natural images with presentations of synthetic targets that test a perceived orientation. Images were either random (unbiased), or selected according to criteria favoring content at a particular orientation (biased). We show that exposure to biased images changes the perceived orientation, compared with a reference perceived orientation obtained while being exposure to unbiased images. The obtained TAE is compared with TAE resulting from synthetic noise images having similar oriented frequency content.

* Corresponding author.
E-mail address: Dov.Sagi@Weizmann.ac.il (D. Sagi).

http://dx.doi.org/10.1016/j.visres.2015.10.014

2. Materials and methods

2.1. Observers

Fourteen observers (aged 20–30 years, 5 male, 9 female) with normal or corrected-to-normal vision participated. All were naïve to the purpose of the experiment, were paid for their participation, and provided informed consent in accordance with the Declaration of Helsinki.

2.2. Apparatus

The stimuli were presented on a linearized Philips 201B4 21" CRT monitor (resolution: 1280 × 1024 pixels; refresh rate: 100 Hz), which was controlled by dedicated software. Observers were seated 100 cm from the display (occupying 23° × 18.5° of visual field) in an otherwise dark environment. The mean luminance of the display was 57.8 cd/m^2.

2.3. Stimuli and tasks

We used adaptors to affect the perceived orientation (Fig. 1A), and targets to test the perceived orientation (Fig. 1B).

2.3.1. Adaptor stimuli

Three types of adaptor stimuli were used: oriented noise patterns, biased images, and unbiased images (Fig. 1C).

Noise adaptors were random $1/f^{\alpha}$ ($\alpha = 2.5$) noise patterns (Geisler, 2008; Van der Schaaf & van Hateren, 1996) filtered in order to depict the orientation content at a particular orientation ('*noise*', Fig. 1C). A value of $\alpha = 2.5$ was obtained by fitting $1/f^{\alpha}$ to the Fourier spectrum of the biased images (described below). The filter for oriented content was the 'oriented band-pass filter' used to select the biased images (as described below, but with a Butterworth filter of order 4 instead of 2). This procedure was used to generate a pool of 100 oriented noise images that were then scaled to have a fixed RMS contrast of 23%. These images were presented in a circular window subtending 4.5° of the visual angle, whose surrounding edge was averaged smoothly with the background (linearly over 0.28°).

Image adaptors were unmodified natural images that were either selected to maximize the orientation content at a particular orientation ('*biased*'), or selected randomly ('*unbiased*') (Fig. 1C). Images had the same mean luminance as the background, had a mean RMS contrast of 23%, and were presented in a circular window subtending 4.5° of the visual angle, whose surrounding edge was averaged smoothly with the background (linearly over 0.28°).

Table 1
Session types.

Session	Adaptor trials				Test trials
	Stimuli	Presentation duration	Dummy task	Count	Count
noise	Oriented noise	Predefined varying	Random button	45	135
biased	Biased images	Until response	Image categorization	45	135
unbiased	Unbiased images	Until response	Image categorization	45	135

Test trials were identical across conditions. For noise, the dummy task was to press either the left or the right mouse button, randomized across days.

Biased images were obtained by the following method. First, images labeled as plants, fungi, or animals in the public ImageNet database (Deng et al., 2009) were downloaded ($N = 60{,}000$) and were converted to grayscale. Then, from each image, the sub-images of size 256 × 256 pixels were extracted (displaced by 20 pixels vertically or horizontally in the original image), resulting in ~10,000,000 sub-images. Each sub-image was padded with 0's on the sides to a size of 511 × 511 pixels, and its two-dimensional Fourier transform was calculated (using Matlab© function "*fft2*"). The power spectrum of this transform was then used to calculate the response of two filters: a band-pass filter, and an oriented band-pass filter. The band-pass filter was a second-order Butterworth spatial frequency filter with half-responses at 1.5 and 7.5 cycles/deg. The oriented filter was a Gaussian filter with a SD of 15° and a maximal response at 115° (i.e. the maximal response for edges oriented 25° clockwise to vertical). The response of the oriented filter was then divided by the response of the not oriented filter, and the 412 sub-images with the highest ratio were selected (all from different images). Of those, 285 were manually pruned, to remove sub-images depicting content that is unidentified, blurred, or artificial (the large number of sub-images with undesired content is an artifact of the biased selection; some images depicting a strong secondary orientation in content were also discarded). The final 129 sub-images were used in the experiment (Fig. A.1).

Unbiased images were obtained by randomly selecting 513 of the 60,000 images described above, and cropping their 256 × 256 pixel center (Fig. A.2).

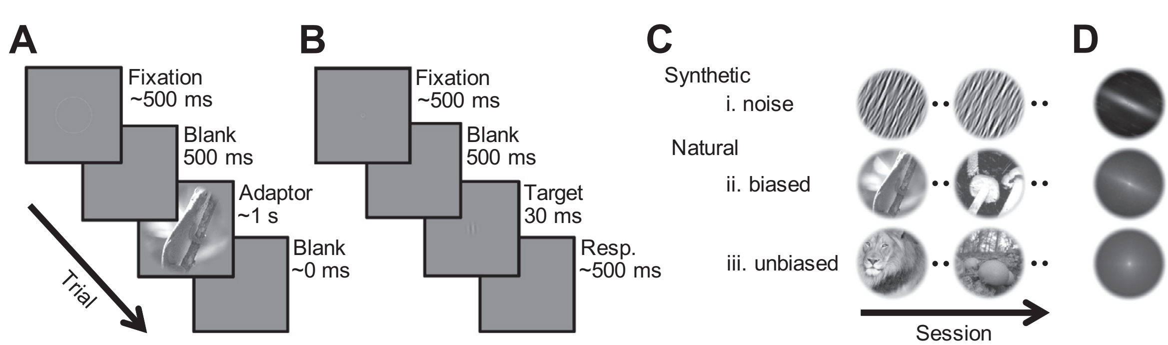

Fig. 1. Experimental design. Within session, adaptation trials with adaptors of a single type were randomly interleaved with target trials. (A) Adaptation trials were used to affect the perceived orientation by exposing observers to adaptors, either synthetic patterns or natural images (experimental differences summarized in Table 1). (B) Target trials were used to determine a perceived vertical orientation, by presenting observers with a near-vertical Gabor patch target, to which they reported whether the patch is oriented CW (clockwise) or CCW to vertical. (C) Example adaptors. (i) Oriented noise adaptors ('noise') are synthetic random $1/f^{\alpha}$ noise patterns that were filtered in order to depict the orientation content at a particular orientation. (ii) Biased image adaptors ('biased') are unmodified natural images selected according to criteria favoring content oriented at a particular orientation. (iii) Unbiased image adaptors ('unbiased') are unmodified natural images selected randomly. (D) Average of the two-dimensional Fourier power spectrum of adaptors. Because of selection, biased images had on average more Fourier power at the biased orientation (ii), similar to the Fourier power distribution of synthetic adaptors (i), whereas unbiased images had a natural distribution of Fourier-power (iii).

Table 2
Trial timing.

Session	Stimuli duration (ms)		Trial duration (ms)			Time between adaptors (ms)
	Adaptor	Target	Adaptor	Target	Average	
noise	946 ± 481	30	1863 ± 565	1471 ± 415	1568 ± 487	5329 ± 4661
biased	1030 ± 790	30	2025 ± 989	1536 ± 530	1658 ± 707	5601 ± 5277
unbiased	982 ± 871	30	1973 ± 1022	1557 ± 620	1663 ± 765	5671 ± 5342

Shown values are timing median (data pooled across observers) ± standard deviation (total across and within observers). Noise pattern adaptors were presented for a predefined duration, randomly one of: 500, 630, 750, 1000, 1850 ms. For images (biased or unbiased), the observers were required to categorize image content; therefore, the images were presented until there was a response. Time between adaptors is measured from the end of presentation of one adaptor stimuli to the start of the next.

2.3.2. Adaptor trials

Adaptor trials started with the presentation of a large fixation cue (a circle with a diameter of 1.8° and a luminance of 30% above the background) that was maintained until the observer continued the trial (~500 ms), followed by a blank screen (500 ms) and the presentation of the adaptor (Fig. 1A). Natural images were presented with no time limit (varying duration with average of ~1000 ms, see Table 2), during which observers categorized the content of the image (animal or not animal, 50% chance each, reported by the left or right mouse button, respectively), with auditory feedback indicating mistaken reports (accuracy: biased, 82 ± 7%; unbiased, 92 ± 4%; mean ± standard deviation across observers). Synthetic patterns were presented for a predefined random duration of 500, 630, 750, 1000, or 1850 ms, which were the 16.6, 33.3, 50, 66.6, and 83.3 percentiles of the distribution of presentation durations of biased images across observers. During pattern presentation, observers were instructed to fixate on the pattern and perform a dummy task of pressing either the left or the right mouse button, randomized across days. If the observers did not respond during the allotted time, a blank screen was presented until they responded.

2.3.3. Target trials

The stimuli used to test the perceived orientation consisted of Gabor patch targets oriented randomly at one of 9 near-vertical orientations (0°, ±1°, ±2°, ±3°, or ±5°) presented at fixation (Fig. 1B). Targets were low-contrast (carrier amplitude of 11.7%), small (Gaussian envelope of 0.25°), with medium spatial frequency (4 cycles/deg), and a random spatial phase. Each trial started by presenting a small fixation cue (a circle with a diameter of 0.3° and a luminance of 30% above the background) that was maintained until the observer continued the trial (~500 ms), followed by a blank screen (500 ms), a Gabor target presentation (30 ms), and then a response period with no time limit (~500 ms), during which observers reported whether the target is tilted CCW or CW relative to vertical (by using the left or right mouse button, respectively).

2.3.4. Adaptor strength

To compare the strength of the different adaptors with respect to the targets, we consider two measures: (a) adaptor overall contrast, quantified as the RSS (root sum of squares) of adaptor amplitude spectrum, and (b) spectral similarity of adaptors and targets, quantified as the vector inner product of the normalized amplitude spectrums of an adaptor and a target having the same orientation as the adaptor (i.e. 25°). Both measures were calculated for a square stimuli region (of size 256 × 256 pixels corresponding to 4.6° × 4.6° of visual field), where pixel values are between −127 to 128 and 0 is background luminance. The measured overall contrast was 8993 ± 1685 for oriented noise adaptors, 8696 ± 4678 for biased image adaptors, and 8598 ± 3736 for unbiased image adaptors (mean ± standard deviation). The measured spectral similarity to targets was 73 ± 2% for oriented noise adaptors, 62 ± 7% for biased image adaptors, and 54 ± 6% for unbiased image adaptors.

2.4. Procedure

The experimental sessions consisted of 46 adaptation trials of the same type, and 135 randomly interleaved target trials (with 15 targets for each of the 9 possible orientations). The typical daily sequence of sessions with oriented noise adaptors (4 observers, Table 3) was U-N-U-N-U-N were U stands for a session with unbiased image adaptors, and N stands for a session with oriented noise adaptors. The typical sequence for biased image adaptors (9 observers) was U-B-U-B-U-B-U-B, where B stands for a session with biased image adaptors. Sessions with unbiased images were also performed separately (14 observers in total, Table 3). A session with oriented adaptors (*noise* or *biased*) was always followed by a break of >1 min, to clear off aftereffects. New observers were gradually acquainted with the tasks: first, they performed a practice session of up to 300 target trials with auditory feedback indicating mistaken reports (no adaptors), then a similar session without feedback, and finally several regular sessions with unbiased image adaptors, until a satisfactory performance was reached. The variability in number of session repetitions of the different observers (Table 3) is not a result of attrition. After the experiments were concluded, several observers ($N = 5$) were asked whether they had noticed that some sessions contained similarly oriented image content. All answered in the negative. One observer was disqualified in the oriented noise condition, having highly atypical measured results likely due to her low orientation discrimination.

Table 3
Number of sessions performed by each observer under each condition.

Session	Observers														Total
noise	–	–	–	–	–	–	–	–	–	–	21	21	12	9	63
biased	19	16	10	7	3	3	2	2	2	–	–	–	–	–	64
unbiased	60	56	9	22	3	3	17	13	2	25	29	34	18	17	308

2.5. Analysis

2.5.1. Fitting a perceived orientation

We used reports about a perceived orientation of near-vertical targets (is the target CW or CCW relative to vertical?) to measure the orientation that is perceived as vertical. Specifically, the percentage of CW reports as a function of the target orientation was interpolated to find the orientation at which CW and CCW reports are equal (slightly adjusted to account for the measured lapse rates) (Fig. 2). The interpolation was performed by fitting this function with a cumulative Gaussian using the MATLAB interface of the Psignifit 3.0 software (lapse rate: up to 0.2, separately for lambda and gamma, uniform priors; bootstrap iterations: 1000; no sensitivity analysis; no BC_a) (Fründ, Haenel, & Wichmann, 2011).

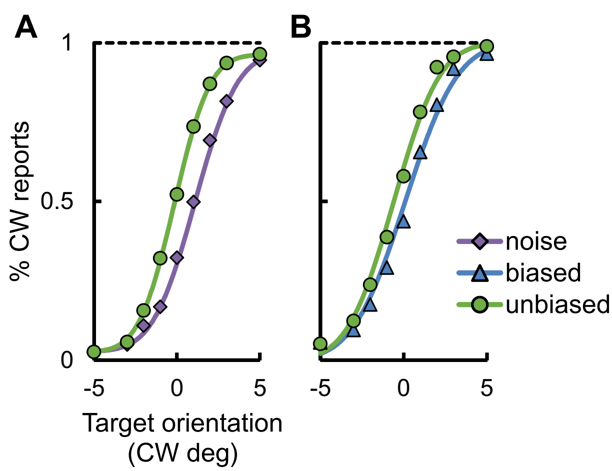

Fig. 2. Measuring the TAE. The change in perceived orientation due to adaptor exposure (TAE, tilt aftereffect) was quantified using the perceived vertical orientation (the orientation with equal probability for a CW and a CCW report considering lapse rates, see Section 2.5.1), as measured from behavioral CW reports by interpolation (continuous lines). (A and B) The percentage of CW reports as a function of target orientation, for sessions with the different oriented adaptors, and sessions with unbiased image adaptors performed on the same day (green circles). TAE is the difference in perceived vertical orientation between sessions with unbiased image adaptors to sessions with (A) orientated noise adaptors (purple rhombuses), or (B) biased image adaptors (blue rectangles). Data pooled across observers, and 95% CI at 50% CW reports are smaller than line width (obtained by bootstrapping).

2.5.2. Target dependencies

The effect of the target trials on the perceived orientation in subsequent target trials was measured as follows (Fig. 4). For a particular lag, N, and each of the 9 target orientations, θ, we selected all target trials that are located N trials after a target oriented θ, and measured the perceived orientation of these trials (see Section 2.5.1). The target dependency is then quantified as the linear slope of the perceived orientation as a function of θ, measured by least-squares fitting. This assumes that there is linearity of the effect with orientation (Campbell & Maffei, 1971), at least within the used range of orientations (−5°:5°).

3. Results

Fourteen participants underwent sessions in which they were repeatedly exposed to adaptors (Fig. 1A), interleaved randomly with orientation judgment targets which determined perceived orientation (Fig. 1B). Adaptors were either synthetic oriented $1/f^{\alpha}$ noise patterns, natural images selected according to criteria favoring content at a particular orientation (biased images), or natural images selected randomly (unbiased images) (Fig. 1C). Because of selection, biased images had on average more Fourier-power at the biased orientation, similar to the Fourier-power distribution of the synthetic adaptors, whereas unbiased images had a natural distribution of Fourier power (Fig. 1D).

3.1. Adaptor exposures

We first measured the change in perceived vertical orientation due to repeated exposure to adaptors (Figs. 2 and 3). This change is the TAE, and the sign of the TAE is defined here such that a positive TAE stands for a repulsive change of the perceived orientation (the change of the perceived orientation in a direction that is away from the adaptor orientation). Reported TAE is the change in perceived orientation relative to the perceived orientation in sessions with unbiased image adaptors (performed on the same day) (Fig. 2). Averaged across observers, TAE showed 1.08° ± 0.11° for

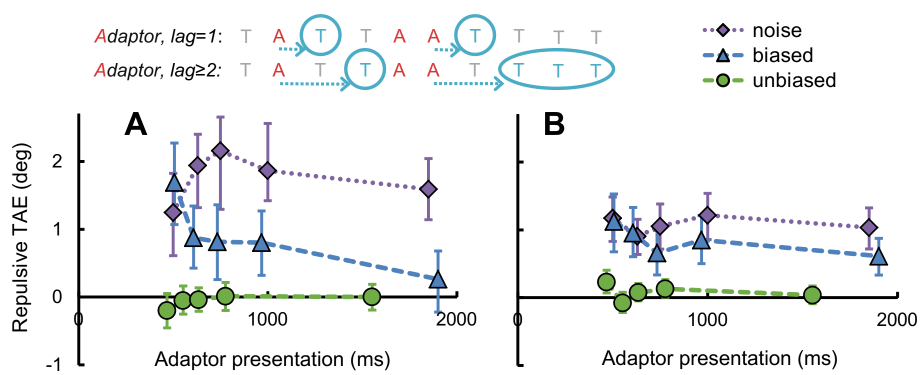

Fig. 3. TAE as a function of adaptor presentation duration, for different lags from adaptor, after > 10 adaptor exposures in the sequence. Targets ('T') were separated based on lag from adaptor ('A'), to targets that immediately follow the adaptor (lag = 1), and subsequent targets before the next adaptor (lag ⩾ 2). Then, separately for (A) lag = 1, and (B) lag ⩾ 2, targets were analyzed based on the adaptors' presentation duration (Section 2.5.1). For synthetic patterns (noise: purple rhombuses), presentation duration was predefined. For images (biased: blue triangles, unbiased: green circles), trials were grouped in five equally-sized bins based on the reaction-time-dependent presentation duration (excluding the upper and lower 1%), and x-axis value is the average duration. Because the observer composition varied between bins, the reference of each bin was obtained by weighting the individual references according to the bin composition. There are possible confound factors in the analyses of the biased condition (see Section 4.2). Data pooled across observers, and error bars are 95% CI obtained by bootstrapping.

oriented *noise* adaptors (mean ± SEM) with $p = 0.003$ (two-tail *t*-test, $t(3) = 9.2$), and 0.65° ± 0.17° for *biased* image adaptors with $p = 0.007$ ($t(8) = 3.6$). The TAE when pooling data across observers showed similar results (*noise*: 1.11°, 95% CI [0.99°, 1.27°]; *biased*: 0.65°, 95% CI [0.54°, 0.78°], confidence intervals obtained by bootstrapping) (Fig. 2). Therefore, exposure to both natural and synthetic adaptors has affected the perceived orientation. Across observers, the difference between the conditions (average of 0.43°), was nearly statistically significant ($p = 0.07$, two-tailed t-test for unequal variance).

3.2. Adaptor presentation duration and lag from adaptor

In the sequence of repeated exposures, adaptors differed in presentation duration, and each adaptor was typically followed by multiple targets. Therefore, we analyzed the TAE conditioned on the lag of a target from the most recent adaptor, further conditioned on the adaptors' presentation duration. For synthetic adaptors having predefined presentation durations (Table 1), this is straightforward; for image adaptors which were presented until response (*biased* and *unbiased*, Tables 1 and 2), target trials were binned by presentation duration (bins of equal size, irrespective of observer, and excluding the upper and lower 1%; because the observer composition varied between bins, the reference of each bin was obtained by weighting the individual references according to the bin composition, Table 3). This analysis was separately applied for targets that immediately follow the adaptor ('lag = 1', Fig. 3A), and targets that do not ('lag ⩾ 2', Fig. 3B). The first 10 adaptor exposures in the sequence were eliminated.

For *noise* adaptors (Fig. 3, purple rhombuses), TAE was significantly stronger at lag = 1 compared with lag ⩾ 2 (average diff of 0.75; $p = 0.03$, $F_{(1,12)} = 15.2$, two-factor repeated measures ANOVA for lag × duration, based on individual data), exhibiting a non-significant modulation by presentation duration at lag = 1 ($p = 0.13$, $F_{(4,12)} = 2.1$, one-way repeated measures ANOVA) and no modulation at lag ⩾ 2 ($p = 0.7$, $F_{(4,12)} = 0.5$).

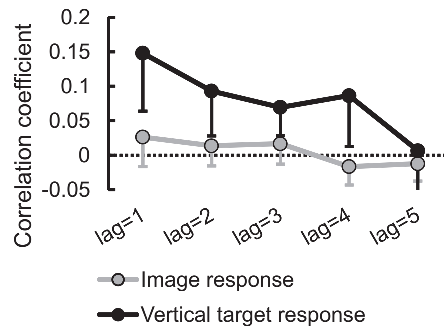

Fig. 5. Response correlations. For lag = *N*, the correlation between the current and *N*-back response, where the current response is to a target, and the *N*-back response is to either a vertical target (same task response correlation, black line) or an unbiased image (across tasks response correlation, gray line). Reported values are Pearson correlation coefficient averaged across observers, and error bars are 95% CI. Only sessions with unbiased image adaptors were considered.

For *biased* adaptors (Fig. 3, blue triangles), at lag = 1 the TAE consistently decreased for increasing presentation durations (linear regression showing negative slope with $p < 0.002$, two-tails, obtained by bootstrapping), with lag ⩾ 2 showing a similar but weaker trend ($p < 0.08$). A similar analysis which maintains a fixed amount of trials per bin per session (by separating trials of each session to five equal size bins then pooling across sessions), and so a fixed amount of trials per bin per observer, shows that this effect has an individual component (at lag = 1, $p < 0.01$). These analyses are not applicable as-is for individual observers due to limited number of repetitions (Table 3). Also, note that reaction-time is a possible confound of this analysis (among others, see Section 4.2). For *unbiased* adaptors (Fig. 3, green circles), performing the same analysis did not show systematic changes of perceived orientation.

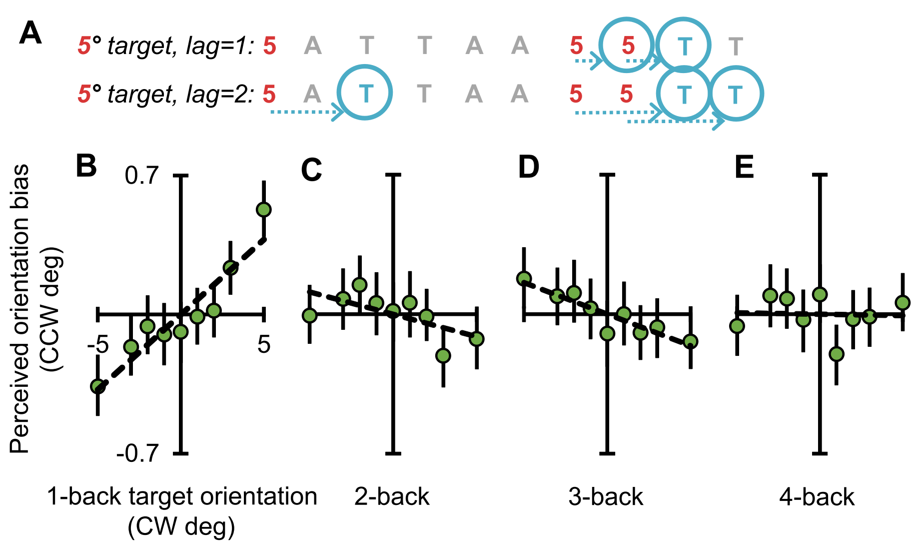

Fig. 4. Target dependencies in sessions with unbiased images. (A) For lag = *N*, targets ('T' or '5') of different orientations were grouped (across observers) according to the *N*-back target orientation (e.g. +5°). Current-target dependent psychometric functions were constructed for the different *N*-back orientations, from which perceived vertical was estimated. (B–E) Perceived orientation of targets in each group compared to average, as a function of the *N*-back orientation (abscissa) by which the targets were grouped (green symbols; dashed line – linear regression), for lag = 1 (B), lag = 2 (C), lag = 3 (D), and lag = 4 (E). Positive slope indicates repulsion (CCW bias due to a CW *N*-back target orientation). Error bars are 95% CI obtained by bootstrapping.

3.3. Target dependencies

We next checked whether the target trials affect the perceived orientation in subsequent target trials (Fig. 4). Only sessions with unbiased image adaptors were considered. To this end, we computed separate psychometric functions (CW reports as a function of orientation) for the different target orientations presented *N* trials back (Fig. 4A), allowing us to compute the dependency of the perceived orientation on the *N*-back target orientation (Fig. 4B–E, green circles, data pooled across observers). A nearly linear dependency was found within the tested range ($-5°$:$5°$), which permitted a convenient description in terms of the function slope (Fig. 4B–E, dashed line).

At lag = 1, the slope of target dependencies was positive, showing 0.076 with $p = 0.0003$ ($F(1,7) = 45.6$, linear regression slope test, data pooled across observers, Fig. 4B). At lag = 2 and 3, the dependency showed negative slopes (of -0.022 and -0.032 with $p = 0.058$ and 0.0005, respectively, Fig. 4C and D). Individual data showed a similar modulation, with a robust transition from a repulsive (CCW bias due to a CW *N*-back target orientation, lag = 1) to an attractive (lag = 2) effect (average slope difference of -0.096 with $p = 0.0002$, two-tailed *t*-test).

3.4. Response correlations

The dependency of the current response on a previous targets' orientation may arise from the orientation-dependent response on that previous trial. To address this possibility, we measured the correlation of current and *N*-back responses, conditioning that the current response is to a target, and the *N*-back response is either (a) to a vertical target, or (b) to an unbiased image (Fig. 5). This analysis isolates the effect of the *N*-back response, as exposure to these N-back stimuli did not change the perceived vertical orientation (Figs. 3 and 4B–E). Because the response to images is not orientation-dependent, its response correlation is a measure for the tendency to repeatedly use the same mouse key across tasks.

For a vertical target response, the correlation showed positive values at lag = 1:4 (averaged across observers, rho > 0.08, $p < 0.05$, two-tailed *t*-test). Therefore, at lag = 1, this attractive response correlation cannot account for the repulsive target dependency described above (Fig. 4B). For an unbiased image response, the correlation was not significantly different from zero (at all lags, $p > 0.25$), and was weaker than the vertical target response correlation (at lag = 1:4, $p < 0.06$ for a difference).

4. Discussion

4.1. Adapting to the natural environment

Our experiments show that exposure to a biased ensemble of natural images, selected according to criteria favoring image content oriented at a particular orientation, results in a repulsive TAE. Previous works have reported that TAE is induced by synthetic or modified natural stimuli (Bao & Engel, 2012), but this is the first report of TAE induced by natural stimuli that are not modified in any way. This result provides the strongest psychophysical evidence to date that the visual system continuously adapts to the distribution of orientations (edges) in the natural environment. Our methodology can be generalized to induce other aftereffects, or to investigate any statistical manipulation of the environment, simply by using different criteria in the biased selection of images. Also, because we used a public database with millions of images (Deng et al., 2009), our method can be scaled to select tens of thousands of sub-images (instead of ~100 here).

4.2. Presentation duration of natural adaptors

Image adaptors were presented with no time limit, during which observers identified image content. The intra- and inter-observer reaction time (and so presentation duration) variability in performing this task was strongly influential on TAE magnitude, showing a surprising negative correlation (*biased*, Fig. 3). Specifically, for ~500 ms the TAE was broadly standard considering experimental design differences (Magnussen & Johnsen, 1986), and for ~2000 ms was much reduced (with lag = 1, but less with lag > 1). This is at odds with previous reports about increasing or saturating TAE magnitude for increasing presentation durations around 300–5000 ms (Dickinson, Mighall, Almeida, Bell, & Badcock, 2012; Harris & Calvert, 1989), but see (Haak et al., 2014). We note that presentation duration is confounded with reaction time in experiments with natural adaptors. While observers with slower RTs may have reduced TAE, our results show reduced TAE with increasing stimulus duration for individual observers. This perhaps suggests that an image identification task can reduce TAE, e.g. due to a calibration mechanism applicable during recognition which reduces selective low-level adaptation effects (possibly detrimental for the current image identification task). Interestingly, a recent study found a reduced TAE when oriented adaptors were grouped into an object (He, Kersten, & Fang, 2012).

4.3. Natural compared with synthetic adaptors

Overall (irrespective of presentation duration), the TAE we find for natural images was somewhat reduced relative to the effects found with synthetic stimuli, which may be accounted for by differences in adaptor image statistics (Fig. 1D) adaptor variability (in both image statistics, see Section 2.3.4, and timing, see Table 2), or other experimental differences (see previous section). This trend is in accordance with earlier reports showing that the tilt illusion, which is the shift in the perceived orientation due to the simultaneously presented oriented stimuli, is weaker for natural compared with synthetic stimuli having identical power spectra (Goddard et al., 2008).

4.4. Target dependencies

We have shown that target dependencies result in a repulsive effect at lag = 1 and an attractive effect at lag = 2. Also, response correlations were shown to be attractive at both lag = 1 and lag = 2, implying that response dependencies (Fründ, Wichmann, & Macke, 2014; Treisman & Williams, 1984) can at most account for the attractive effect.

Acknowledgments

This research was supported by the Basic Research Foundation, administered by the Israel Academy of Science. We thank Misha Katkov and Noga Pinchuk-Yacobi for their suggestions and comments.

Appendix A

(See Figs. A.1 and A.2).

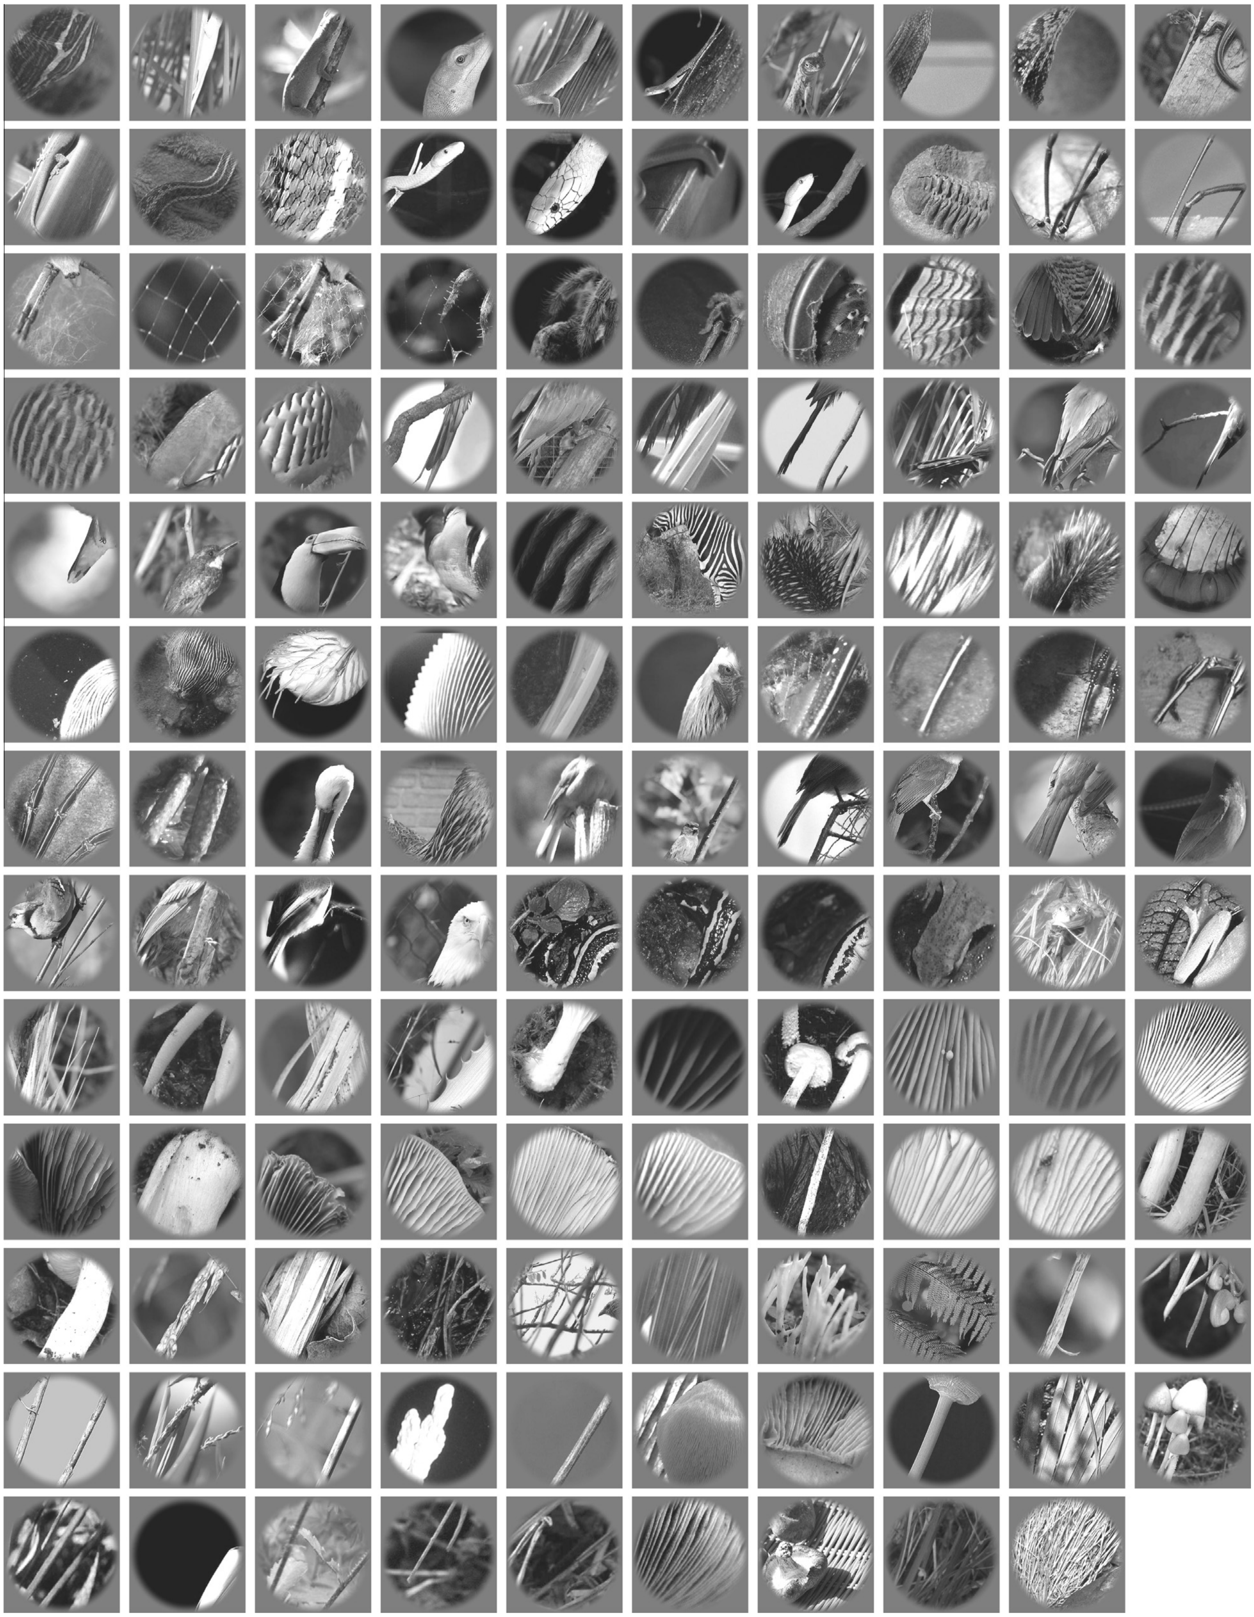

Fig. A.1. All 129 biased images.

Fig. A.2. 130 random unbiased images (out of 513).

References

Alam, M. M., Vilankar, K. P., Field, D. J., & Chandler, D. M. (2014). Local masking in natural images: A database and analysis. *Journal of Vision, 14*(8), 22. http://dx.doi.org/10.1167/14.8.22.

Bao, M., & Engel, S. A. (2012). Distinct mechanism for long-term contrast adaptation. *Proceedings of the National Academy of Sciences, 109*(15), 5898–5903.

Bex, P. J., Solomon, S. G., & Dakin, S. C. (2009). Contrast sensitivity in natural scenes depends on edge as well as spatial frequency structure. *Journal of Vision, 9*(10).

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. *The Journal of Physiology, 203*(1), 237–260.

Campbell, F. W., & Maffei, L. (1971). The tilt after-effect: A fresh look. *Vision Research, 11*(8), 833–840.

Carandini, M., Demb, J. B., Mante, V., Tolhurst, D. J., Dan, Y., Olshausen, B. A., et al. (2005). Do we know what the early visual system does? *The Journal of Neuroscience, 25*(46), 10577–10597.

Clifford, C. W. G., Webster, M. A., Stanley, G. B., Stocker, A. A., Kohn, A., Sharpee, T. O., et al. (2007). Visual adaptation: neural, psychological and computational aspects. *Vision Research, 47*(25), 3125–3131. http://dx.doi.org/10.1016/j.visres.2007.08.023.

Deng, J., Dong, W., Socher, R., Li, L. -J., Li, K., & Fei-Fei, L. (2009). Imagenet: A large-scale hierarchical image database. In *Computer vision and pattern recognition, 2009. CVPR 2009. IEEE Conference on* (pp. 248–255).

Dickinson, J. E., Mighall, H. K., Almeida, R. A., Bell, J., & Badcock, D. R. (2012). Rapidly acquired shape and face aftereffects are retinotopic and local in origin. *Vision Research, 65*, 1–11. http://dx.doi.org/10.1016/j.visres.2012.05.012.

Froudarakis, E., Berens, P., Ecker, A. S., Cotton, R. J., Sinz, F. H., Yatsenko, D., et al. (2014). Population code in mouse V1 facilitates readout of natural scenes through increased sparseness. *Nature Neuroscience, 17*(6), 851–857. http://dx.doi.org/10.1038/nn.3707.

Fründ, I., Haenel, N. V., & Wichmann, F. A. (2011). Inference for psychometric functions in the presence of nonstationary behavior. *Journal of Vision, 11*(6).

Fründ, I., Wichmann, F. A., & Macke, J. H. (2014). Quantifying the effect of intertrial dependence on perceptual decisions. *Journal of Vision, 14*(7), 9. http://dx.doi.org/10.1167/14.7.9.

Geisler, W. S. (2008). Visual perception and the statistical properties of natural scenes. *Annual Review of Psychology, 59*, 167–192. http://dx.doi.org/10.1146/annurev.psych.58.110405.085632.

Gibson, J. J., & Radner, M. (1937). Adaptation, after-effect and contrast in the perception of tilted lines. I. Quantitative studies. *Journal of Experimental Psychology, 20*(5), 453.

Goddard, E., Clifford, C. W. G., & Solomon, S. G. (2008). Centre-surround effects on perceived orientation in complex images. *Vision Research, 48*(12), 1374–1382. http://dx.doi.org/10.1016/j.visres.2008.02.023.

Haak, K. V., Fast, E., Bao, M., Lee, M., & Engel, S. A. (2014). Four days of visual contrast deprivation reveals limits of neuronal adaptation. *Current Biology, 24*(21), 2575–2579. http://dx.doi.org/10.1016/j.cub.2014.09.027.

Harris, J. P., & Calvert, J. E. (1989). Contrast, spatial frequency and test duration effects on the tilt aftereffect: Implications for underlying mechanisms. *Vision Research, 29*(1), 129–135.

He, D., Kersten, D., & Fang, F. (2012). Opposite modulation of high- and low-level visual aftereffects by perceptual grouping. *Current Biology, 22*(11), 1040–1045.

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

Knapen, T., Rolfs, M., Wexler, M., & Cavanagh, P. (2010). The reference frame of the tilt aftereffect. *Journal of Vision, 10*(1).

Kohn, A. (2007). Visual adaptation: Physiology, mechanisms, and functional benefits. *Journal of Neurophysiology, 97*(5), 3155–3164. http://dx.doi.org/10.1152/jn.00086.2007.

Magnussen, S., & Johnsen, T. (1986). Temporal aspects of spatial adaptation. A study of the tilt aftereffect. *Vision Research, 26*(4), 661–672.

Mitchell, D. E., & Muir, D. W. (1976). Does the tilt after-effect occur in the oblique meridian? *Vision Research, 16*(6), 609–613.

Olshausen, B. A., & Field, D. J. (2005). How close are we to understanding V1? *Neural Computation, 17*(8), 1665–1699.

Treisman, M., & Williams, T. C. (1984). A theory of criterion setting with an application to sequential dependencies. *Psychological Review, 91*(1), 68–111. http://dx.doi.org/10.1037//0033-295X.91.1.68.

Van der Schaaf, A., & van Hateren, H. J. (1996). Modelling the power spectra of natural images: Statistics and information. *Vision Research, 36*(17), 2759–2770.

Webster, M. A. (2011). Adaptation and visual coding. *Journal of Vision, 11*(5).

Webster, M. A., & Miyahara, E. (1997). Contrast adaptation and the spatial structure of natural images. *Journal of the Optical Society of America. A, 14*(9), 2355–2366.

Zhang, P., Bao, M., Kwon, M., He, S., & Engel, S. A. (2009). Effects of orientation-specific visual deprivation induced with altered reality. *Current Biology, 19*(22), 1956–1960.