Title:

Sensitivity to the slope (α) of the amplitude spectrum is dependent on the α of recently viewed environments: A visual adaptation study in modified reality.

Short Title: Sensitivity to α in Modified Reality

Authors:

Bruno Richard[1]

Patrick Shafto[1]

Author Affiliation:

[1]Department of Mathematics and Computer Science, Rutgers University - Newark, Newark, New Jersey, USA, 07102

Corresponding Author: Bruno Richard, Department of Mathematics and Computer Science, Rutgers University - Newark, 101 Warren Street, Rm 216, Newark, New Jersey, USA, 07102, Email: bruno.richard@rutgers.edu

Keywords: Natural Scenes, Amplitude Spectrum Slope, Head-Mounted Display, Modified Environment, Bayesian Observer Model, Multimodal Prior

Abstract

Scenes contain many statistical regularities that could benefit visual processing if accounted for by the visual system. One such statistic is the orientation-averaged slope (α) of the amplitude spectrum of natural scenes. Human observers show different discrimination sensitivity to α: sensitivity is highest for α values between 1.0 and 1.2 and decreases as α is steepened or shallowed. The range of α for peak discrimination sensitivity is concordant with the average α of natural scenes, which may indicate that visual mechanisms are optimized to process information at α values commonly encountered in the environment. Here we explore the association between peak discrimination sensitivity and the most viewed αs in natural environments. Specifically, we verified whether discrimination sensitivity depends on the recently viewed environments. Observers were immersed, using a Head-Mounted Display, in an environment that was either unaltered or had its average α steepened or shallowed by 0.4. Discrimination thresholds were affected by the average shift in α, but this effect was most prominent following adaptation to a shallowed environment. We modeled these data with a Bayesian observer and explored whether a change in the prior or a change in the likelihood best explained the psychophysical effects. Change in discrimination thresholds following adaptation could be explained by a shift in the central tendency of the prior concordant with the shift of the environment, in addition to a change in the likelihood. Our findings suggest that expectations on the occurrence of α that result from a lifetime of exposure remain plastic and able to accommodate for the statistical structure of recently viewed environments.

Our visual world is diverse but contains spatial regularities that span across the various environments we encounter. There is evidence that accounting for these spatial regularities in encoding can improve the visual system's efficiency in processing complex scene properties [1–4]. One such regularity is the slope of the amplitude spectrum (α) of natural scenes, which defines the association of amplitude to spatial frequency (f),

$$amplitude \propto 1/f^{\alpha}. \tag{1}$$

On average, the α value of natural scenes is 1.08 [5–11]. Observers are sensitive to and influenced by α: they exhibit different discrimination and recognition abilities for different α values [1, 12–14] and can adapt to α [15–17]. Biases in the discrimination ability of observers for α may be associated with the recent visual experience of observers [15]. Here, we explore how the discrimination sensitivity for α is related to prior expectations of the distribution of αs gathered from recent perceptual experience. We used a Head-Mounted Display (HMD) system paired with a camera to display and modify the natural visual environment of observers and measure discrimination thresholds to α. Finally, we interpret our findings according to a Bayesian observer model [18–20] and derive the shape and malleability of the prior for the αs of natural scenes.

Measuring the discriminability of α has been a popular method for determining the shape of the internal representation for α [9, 21–23]. For example, a common finding when measured with noise stimuli, modified to have a particular α value, is that the discriminability of α is best (thresholds are lowest) when the reference α is between 1.0 and 1.3, and rises as the α is steepened or shallowed [13, 24, 25]. Peak discriminability matches that of the average α of natural scenes. Thus it appears that the human visual system is best able to discriminate the spatial characteristics of scenes it encounters often. However, the peak in discriminability can depend on the original α value of the image. When observers are asked to discriminate against the original α of a natural image, thresholds are significantly higher than when observers complete the discrimination task against an image with a modified α[1] [26]. Tolerance for small deviations

[1]Modifying the α of a natural image will alter its perceived focus: steepened images will appear blurred while

in the natural α of an image [22, 23, 26] may assist in optimizing visual processing by facilitating the discrimination of various objects in the natural world [12, 14]. Thus, while the human visual system peak sensitivity to α is for values regularly encountered ($\alpha \approx 1.0$), it is simultaneously tolerant for minor deviations from the natural α of an image, which suggests some process of adaptation to the spatial characteristics of the environment.

Visual adaptation is the process by which exposure to a stimulus affects visual processing for subsequent stimuli [27–30]. When the adapting stimuli are spatially complex — noise stimuli modified to have a natural α value — adaptation acts to renormalize the percept of observers towards an internal norm; a neutral point in the responses of the detecting mechanisms [15]. Adaptation to artificially steepened or shallowed images generated a shift in the perceived α of a new image in the opposite direction to that of the adaptor α. No shift in the perceived α of images was measured when the adapter had an α of 1.0, which is coincidentally the average value of α across all natural scenes [15]. It appears that through visual experience, the human visual system has developed an internal representation or an expectation for α values commonly encountered.

The statistical structure of natural scenes is known to influence biases in sensitivity to visual features [19, 20, 31, 32]. For example, there is an orientation anisotropy in human vision whereby sensitivity to horizontal content embedded in natural scenes is worse than that of obliques (i.e., the horizontal effect; [33–35]). This bias in orientation sensitivity is believed to stem from the over-representation of horizontal content in natural scenes [11, 19, 33], which is perceptually normalized by reducing sensitivity to horizontal content. Indeed, if observers are immersed in an environment that does not contain the natural over-representation of horizontal content, the strength of the horizontal effect diminishes severely [31]. These effects occur relatively rapidly (within an hour or two of immersion), which is a strong indication that the expected structure of natural scenes can be relearned [31, 32]. The change in sensitivity to orientation contrast was well-explained by a change in the prior of a Bayesian observer model which matched the new orientation distribution of the modified environment [31]. If the internal representation for α is

shallowed images appear sharpened, which may aid in the discrimination task.

determined by the distribution of commonly viewed αs in the natural environment, then it may be able to adjust given a novel environment statistical structure, as has been demonstrated with orientation sensitivity in natural scenes.

Here, we explore whether sensitivity to α is associated with the statistical structure, or distribution, of α in natural scenes by measuring α discrimination thresholds prior to and following immersion in an environment that no longer contains the typical α distribution of scenes. Observers were immersed in the environment with a Head-Mounted Displays (HMDs) and a camera that recorded and displayed the visual world in near real-time (a display system that we have named Modified Reality; [31]). All experimental conditions, the measurement of α discrimination thresholds and the adaptation to the modified environment were conducted in the Head-Mounted Display. This method greatly facilitates the measurement of sensitivity to features in the natural environment of observers while maintaining sufficient experimental control on the presented stimuli. The environment was modified by steepening or shallowing the distribution of αs. Finally we develop a Bayesian observer model to explore the possible shape of a prior for α and verify if the prior can adjusts to match the novel environmental distribution of α [18, 19, 31].

Materials and Methods

Ethics Statement

Procedures were approved by the Rutgers Arts and Sciences Institutional Review Board (Study ID: Pro2020002981). The experiment was completed by three male participants who provided written informed consent.

Participants

Three male observers (age 29, 33, and 34) took part in this study including one author (BR). All participants were experienced psychophysical observers with prior training on the amplitude spectrum slope discrimination task and had normal or corrected to normal visual acuity.

Apparatus

The experiment was conducted in an *Oculus Rift Development Kit 2* (DK2). The DK2 has two 5.7 inch OLED displays, each having a monocular resolution of 960x1080 pixels and a refresh rate of 75Hz. The Field of View of the Oculus Rift DK2 is 100°, resulting in a visual resolution of approximately 11 pixels per degree of visual angle [2]. The observer's environment was captured by a *See3Cam_130* USB 3.0 camera mounted on the Oculus Rift. The camera continuously collected images of the observers' environment at a rate of 60 fps during the experiment. The images recorded by the camera were sent to the laptop computer (Dell Latitude 3470), using the MATLAB Image Acquisition Toolbox, filtered in MATLAB and displayed back to the Oculus Rift with the the Oculus VR library of *Psychtoolbox* [36, 37].

Stimuli

The environment images collected by the mounted camera were converted to grayscale and cropped to a size of 720x1080 pixels. The 120 pixels of the Oculus display above and below the presented environment were set to mean luminance (RGB [128,128,128]). At baseline condition, only the RMS contrast of the environment images was adjusted to match that of the target stimuli (RMS = 0.15). For the adaptation condition, images collected from the environment had their global α changed by either +0.4 or -0.4 (see **Figure 1**). The environment α was changed in an identical manner to the creation of the synthetic noise images used to measure discrimination thresholds (**Figure 1**). The amplitude of each spatial frequency coefficient was divided by the average amplitude (calculated across orientation for that spatial frequency) to create a flat ($\alpha = 0.0$) amplitude spectrum [22]. In this form, the α of the environment image was adjusted by multiplying each spatial frequency's coefficient by $f^{-\alpha_{adjust}}$ (the original image α + 0.4 for the steeper condition or - 0.4 for the shallower condition). The RMS contrast of α_{adjust} images was then adjusted to 0.15.

[2]This low visual resolution is a consequence of having displays with a large field of views very close to the observers' eyes, $\sim$ 5cm.

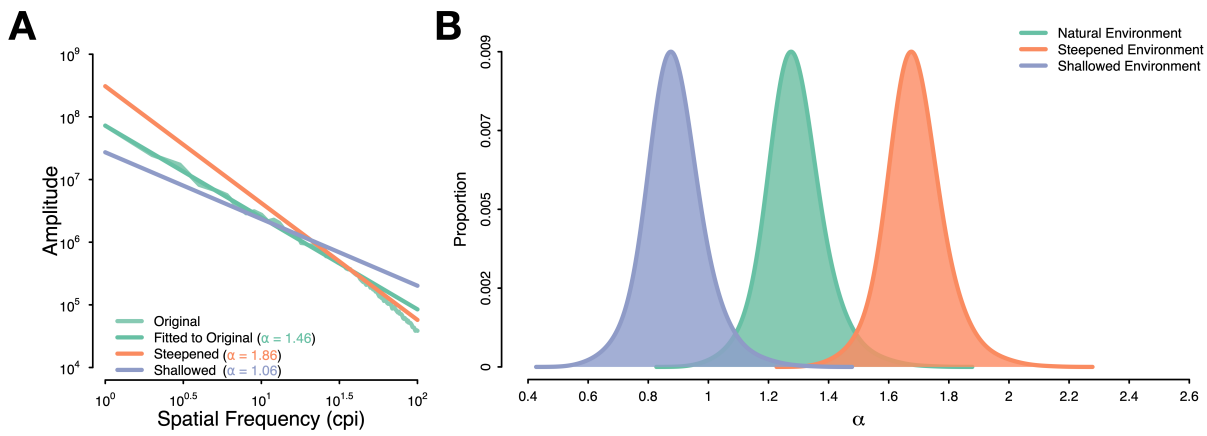

Figure 1: **A** α modification procedures: We fit a linear regression to the log orientation averaged amplitude spectrum of the original image (Original) to measure its $\alpha_{original}$. The new amplitude spectrum slope was created by first dividing the spatial frequency coefficients by the average amplitude to create a flat amplitude spectrum. The α of the original image could then be adjusted by multiplying each spatial frequency coefficient by $\alpha_{original} + 0.4$. The modified spectrum was then scaled vertically to ensure that the total contrast energy in the modified image was identical to that of the original. **B** The distribution of α values in the environment of our observers. The α distribution of our laboratory was similar to that reported by [6] for fully carpented environments with a mean α = 1.29 and standard deviation of 0.09. The steep environment distribution has a mean of 1.69 while the shallow environment had a mean of 0.89.

All test stimuli consisted of synthetic visual noise patterns constructed in the Fourier domain using MATLAB (Mathworks, Natick, MA) and corresponding Signal Processing and Image Processing toolboxes. We opted for noise patterns with different α values as test stimuli, and not natural images, as these were overlaid onto the modified environment when we measured discrimination thresholds and this greatly simplifies the identification of the test stimuli for observers. The visual noise stimuli were created by constructing a polar matrix for the amplitude spectrum and assigning all coordinates the same arbitrary amplitude coefficient (except at the location of the DC component, which was assigned a value of 0). The result is a flat isotropic broadband spectrum (i.e., α = 0.0), referred to as the template amplitude spectrum [22, 24]. In this form, the α of the template spectrum can be adjusted by multiplying each spatial frequency's amplitude coefficient by $f^{-\alpha}$. The phase spectra of test stimuli were constructed by assigning random values from $-\pi$ to π to the different coordinates of a polar matrix while maintaining an odd-symmetric phase relationship to maintain conjugate symmetry. The noise patterns were rendered into the spatial domain by taking the inverse Fourier transform of an α-altered template

amplitude spectrum and a given random phase spectrum. The phase spectrum for all stimuli presented within a trial was identical but randomized from trial-to-trial. The $1/f^{\alpha}$ stimuli were 20° in size (220 pixels). RMS contrast (the standard deviation of all pixel luminance values divided by the mean of all pixel luminance values) was fixed to 0.15.

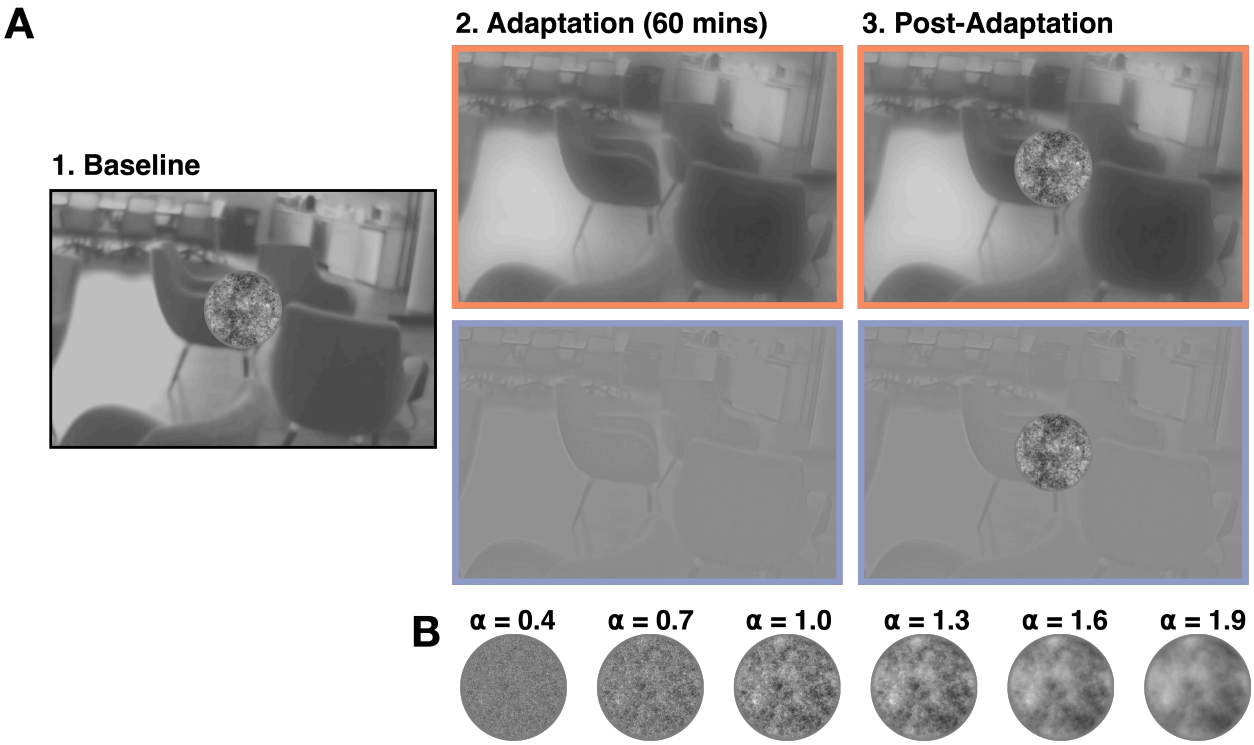

Figure 2: **A** Prior to adaptation, participants completed a baseline α discrimination task in their unaltered (gray-scaled and RMS contrast adjusted to 0.15). Participants did 2 staircases per reference α [0.4, 0.7, 1.0, 1.3, 1.6] (see **B** for example reference α). Participants then began the adaptation procedure, where they viewed the modified world (global α of the environment were steepened by +0.4) for a period of 10 minutes. Following adaptation, participants completed the same α discrimination task as baseline, but with the modified environment in the background. Note that the contrast of the test pattern has been increased for visibility.

Procedures

The experimental procedures measured sensitivity to α prior to and following adaptation to a modified environment. Participants first completed an α discrimination task with $1/f^{\alpha}$ noise stimuli placed atop the unaltered environment (RMS adjusted only; see **Figure 2**). Discrimination thresholds were estimated by a temporal three-interval, two-alternative "Odd-Man-Out" forced choice task [13, 24, 25]. The trial-to-trial change in the image's α was controlled by a 1-up, 2-down staircase procedure targeting 70.71% correct [38, 39]. The staircase began with a

difference between the odd and reference stimulus of 0.5, which decreased in linear steps (step size = 0.02) towards the reference α when the observer made two consecutive correct responses. The difference was increased by 0.02 when observers made an incorrect response (1-Up / 2-Down Rule).

Each trial began with a black (RGB [0, 0, 0]) fixation point (0.5^o) presented for 1000ms at the center of the display, followed by three stimulus presentation intervals (250ms each) interlaced by a blank screen (environment only) for 500ms. The second interval always contained the reference α while either the first or third contained the odd stimulus (the other being identical to the reference α interval). Participants indicated which interval, either the first or the third, they perceived as being the "Odd-Man-Out" via keyboard press. The staircase continued until 12 reversals occurred and thresholds were estimated by averaging the α values of the odd stimulus for the last 5 reversals. Observers completed two staircases per reference α value [0.4, 0.7, 1.0, 1.3, 1.6, 1.9]; the discrimination thresholds of each were averaged for data analysis.

Immediately following baseline, participants completed a 60 minute adaptation period immersed in the modified environment (either α shifted by +0.4 or -0.4) in the HMD. Participants were instructed to look around their environment and interact with objects and individuals present around them but to remain seated or standing in place. Observers then repeated the α discrimination task with the modified environment presented in the background. Participants repeated the experiment with the other modified environment (either α shifted by +0.4 or -0.4) no less than 24 hours following the first experimental session. Each experimental session took approximately 2.5 hours to complete (45 minutes per α discrimination task and 60 minutes of adaptation).

Results

Average α discrimination thresholds measured at baseline and following the adaptation conditions (+0.4 and -0.4). and resulting model fits (see below) are shown in **Figure 3A**. Discrimination thresholds measured in the Oculus Rift DK2 follow the typical peak and trough

commonly observed with traditional displays (i.e., CRTs) [13, 21, 22, 24, 25]. The magnitude of

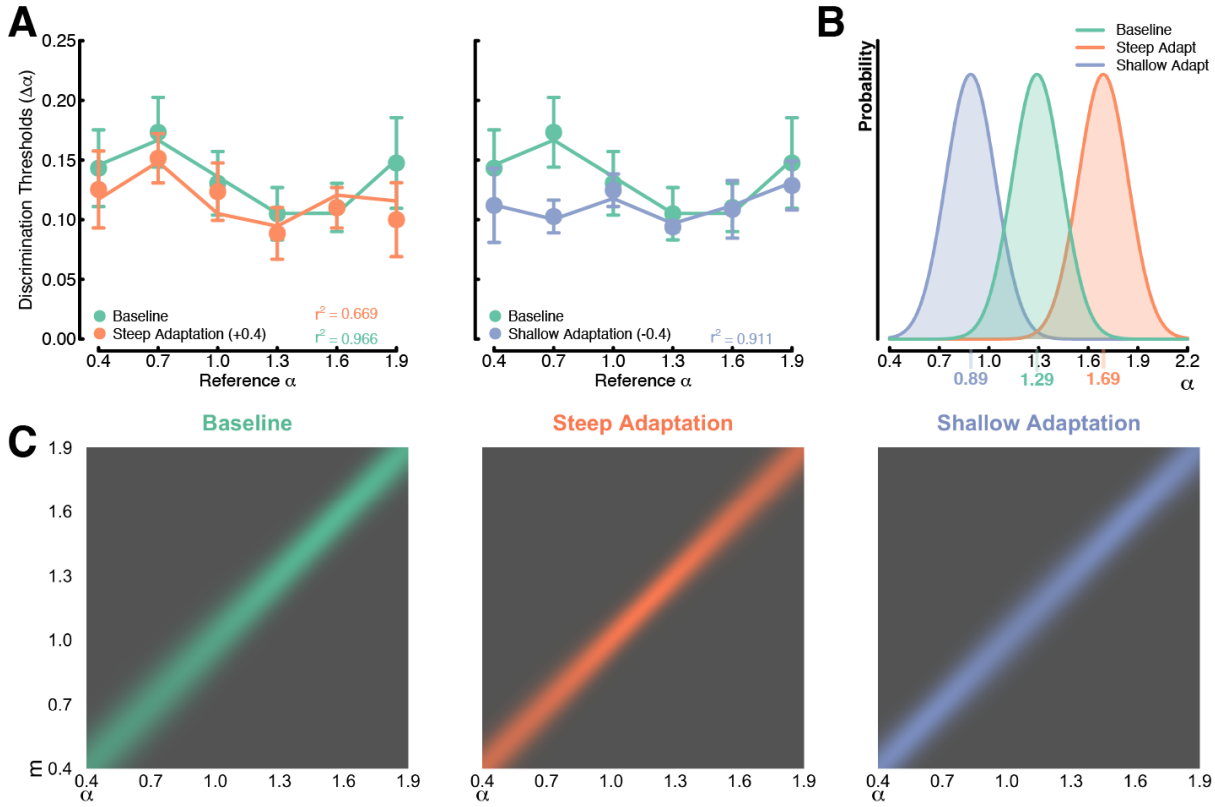

Figure 3: **A.** Average discrimination thresholds (points) and environmental priors model fit (lines) for the six reference αs used in this study. Green markers represent baseline thresholds. Discrimination thresholds measured after the positive (orange) and negative (blue) adaptation conditions are drawn on separate charts for clarity. Error bars represent $\pm$ 1 standard error of the mean. The environmental priors model captured observer discrimination thresholds very well at baseline (r^2 = 0.966) and following adaptation to a shallow environment (r^2 = 0.911) while it struggled with the steep adaptation condition (r^2 = 0.669). **B.** The environmental prior probability distributions used in this model variant. **C.** The resulting likelihoods from the model fitting procedure shown as a two-dimensional conditional density, $p(m|\alpha)$. At baseline, the likelihood was narrowest for steep α values. Following steep adaptation, the likelihoods were narrowest for α values between 1.0 and 1.3 while shallow adaptation narrowed the likelihood distributions at shallow reference α ($\sim$ 0.7).

change between the baseline and adaptation conditions, quantified with the effect size Hedge's g and exact 95% confidence intervals [40] are in **Table 1**. While the overall effects of adaptation on discrimination thresholds are small, we observe some meaningful changes in discrimination according to the adaptation condition. Adaptation to a steepened environment lead to a reduction in discrimination thresholds to very steep reference α (1.9), t(2) = -3.48, p = .037, g = -0.57 [-1.17, 0.05]. Adaptation to a shallowed environment reduced discrimination thresholds for a reference α of 0.7, although this decrease was large, it was not statistically significant, t = -1.84, p

= .104, g = -1.58 [-3.08, 0.72]. Our adaptation effects are moderate in magnitude, yet do follow the general effects of adaptation in the literature [15, 41–43]. Thus it appears that the environmental αs influence human sensitivity to α in a discrimination task.

	Positive Adaptation		Negative Adaptation	
Reference α	g	95% CI	g	95% CI
0.4	-0.23	[-0.56 0.20]	-0.30	[-0.66, 0.01]
0.7	-0.20	[-0.80, 0.53]	-1.58	[-3.08, 0.72]
1.0	-0.01	[-0.84, 0.98]	-0.29	[-0.67, 0.26]
1.3	-0.62	[-1.69, 1.02]	-1.07	[-1.52, 0.92]
1.6	0.34	[-0.73, 1.46]	-0.12	[-1.06, 0.93]
1.9	-0.57	[-1.17, 0.05]	-0.49	[-0.88, 0.17]

Table 1: The effect size (Hedge's g) and exact 95% confidence intervals [40] for differences in the average discrimination following either the positive (+0.4) or negative (-0.4) adaptation conditions.

In both adaptation conditions, we find an attraction effect towards the average α of the environment. Under a Bayesian observer framework, this change in discrimination thresholds could be attributed to an improvement in measurement reliability: a change in the likelihood [42–44] or in the prior [31] to match the novel environment distribution. In this context, adaption is thought to generate an increase in the signal-to-noise ratio near the adapter, leading to an increase in sensitivity for stimulus values near those of the adapter: a change in the likelihood. Our results can, however, also be attributed to a change in the expectation on the occurrence of αs resulting from the modified environment: a change in the prior for α. A change in the prior for α may not be far-fetched. An interesting factor to consider is that the mode of α distributions varies according to the type of environment they represent (e.g., human-made versus natural) [6]. Different α values will thus vary in occurrence when observers are in entirely man-made, natural, or hybrid environments. Previous work has shown that the shape of the prior for statistical regularities in natural scenes (i.e., orientation contrast) matches the distribution of the statistical regularity [19] and if the distribution is changed [31], a similar change in the prior can be observed. Thus, we built a Bayesian observer model to determine whether a change in experience (i.e., prior) and/or improvement in measurement (i.e., likelihood) accounts for the change in discrimination thresholds we observed here.

Under a Bayesian framework, the probability of observing a particular stimulus, α, given a measurement, m, is defined as the product of the likelihood, $p(m|\alpha)$, and prior probability distribution $p(\alpha)$

$$p(\alpha|m) = \frac{p(m|\alpha)p(\alpha)}{p(m)}. \tag{2}$$

We explore the influence of a change in the prior and/or a change in the likelihood of the Bayesian observer best explaining the observed change in discrimination thresholds by using different environmental priors. The priors were all defined as Gaussian distributions,

$$p(\alpha) = \mathcal{N}(\mu_\alpha, \sigma_\alpha^2) \tag{3}$$

with their means μ_α set to one of 5 possible conditions. In the first condition, the mean value of the priors matched those of the adaptation environment (the envrionmental priors model; baseline: μ_α = 1.29, steep adaptation: μ_α = 1.69, shallow adaptation: μ_α = 0.89), the second to the fourth alternatives were models with a prior fixed at one of the environmental distributions, (baseline, steep or shallow models) while the fifth alternative had a prior mean that matched that the distribution of αs recorded from a variety of environments (the natural environment prior; α = 1.08). The standard deviation (σ_α) of all priors was fixed to 0.15, slightly broader than the recorded environmental prior but a better approximation of the width of α distributions recorded in various environments [5].

The observer likelihood $p(m|\alpha_i)$ was defined as a Gaussian centered on the stimulus α for trial i,

$$p(m|\alpha_i) = \mathcal{N}(\alpha_i, \sigma_{\alpha_i}^2) \tag{4}$$

with a variance term, $\sigma_{\alpha_i}^2$ defined as a sine wave,

$$\sigma_\alpha = A\sin(2\pi f\alpha + \phi) + Z \tag{5}$$

scaled to range between 0 and 1. The free parameters A, f, and ϕ to control the Amplitude,

frequency, and phase, respectively. The value Z prevented the sine wave from reaching a value of 0. A sine wave was chosen to operationalize the variance of the likelihoods for α as it is a simple model with few parameters and approximated observer changes in variance for α well.

The expected α ($\hat{\alpha}$) for any given trial was taken as the mean of the posterior distribution with variance $\hat{\sigma}^2$. Model accuracy (percent correct) between the trial $\hat{\alpha}_i$ and the reference ($\hat{\alpha}_{ref}$) was defined as the probability returned by a cumulative normal distribution [45],

$$p_{correct} = \Phi\Big(\frac{d'}{2}\Big) \tag{6}$$

where

$$d' = \frac{\hat{\alpha}_i - \hat{\alpha}_{ref}}{\sqrt{\frac{\hat{\sigma}^2_{\alpha_i} + \hat{\sigma}^2_{\alpha_{ref}}}{2}}} \tag{7}$$

Discrimination thresholds were the $\Delta\alpha$ value (difference between the test α and the reference α) that elicited 70.71% correct performance, as in our behavioral task. Model parameters (A, f, and phi) were optimized for each prior condition with *fminsearch* in MATLAB, minimizing the sums of square error between the model output and observer α discrimination thresholds, prior to and following adaptation.

Of the prior variants explored here, the best performing model had prior probability distributions that matched the adapting environment of observers (**Figure 3**). The other model variants, with a prior fixed to either the baseline, steep, shallow or natural environment captured observer performance for their respective adaptation condition well, but were unable fit discrimination thresholds for other experimental conditions (see table 2 for the model RMS_{errors} and Appendix A for the resulting model fits). Given these results, it appears that adaptation in modified environments, where the distribution of α has been shifted towards steeper or shallower α can exert a change in the observer's prior for α and reduce discrimination thresholds (**Figure 3B**). As would be expected from an adaptation paradigm, we also observe a systematic change in the likelihood variance across adaptation condition (**Figure 3C**). To simplify visualization, **Figure 4** draws the envrionmental prior model standard deviation of the likelihoods for each

Model	R^2	RMS_e	AIC
Environmental Priors	0.88	0.007	-32.59
Baseline Prior	-	0.051	-1.55
Steep Prior	0.59	0.014	-32.62
Shallow Prior	-	0.065	0.43
Natural Prior	-	0.056	-24.27

Table 2: The R^2, Root Mean Square error (RMS_e) and the Akaike Information Criterion (AIC) calculated across all experimental conditions for each model investigated in this study. Only the environmental priors model and Steep prior had average deviations small enough to calculate R^2 values for the model fits. Both models have similar AIC values, but the environmental priors model captured more of the variability in our data and returned a smaller average deviation than the steep prior model and is elected the best performing model.

environments. At baseline, the standard deviation of the likelihood peaks at α of approximately 0.7 and decreases as α steepens, with a trough around α of 1.6. Adaptation to a shallow environment reduced the width of the likelihoods for shallow α values while seemingly leaving steeper α values unaffected. Adaptation to a steep environment reduced the width of the likelihoods for α values between 1.0 and 1.3, mirroring the standard deviation of the likelihoods in the shallow adaptation condition.

Discussion

We measured discrimination thresholds to the slope of the amplitude spectrum of noise patches before and following immersion in a modified environment, all within an HMD. We found baseline α discrimination thresholds to have the same pattern across reference αs as is typically measured in more traditional laboratory equipment. Adaptation generated a reduction in α discrimination thresholds for reference α values near the average α of the modified environment. Specifically, immersion in a steepened world generated a decrease in thresholds for very steep reference α (i.e., 1.9) while immersion in the shallowed environment decreased thresholds for a reference α of 0.7. Behaviorally, our findings are similar to those of other visual paradigms (e.g., speed and color perception) that found discriminability to improve at values near those of the adapter [27, 42, 43]. These adaptation effects are generally interpreted as an improvement in sensory encoding (an increase in the signal to noise ratio or a narrowing of the likelihood [42–44].

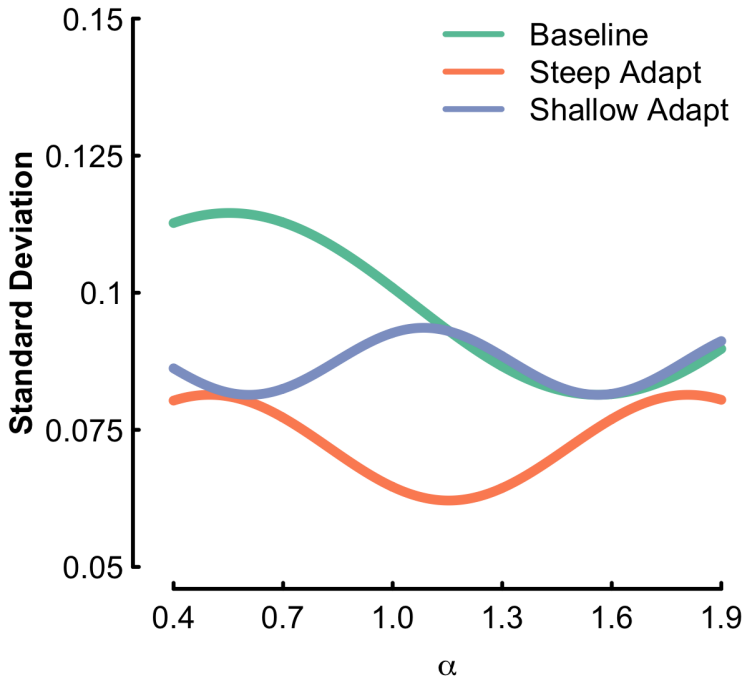

Figure 4: The standard deviation of the environmental priors model likelihoods across α. The standard deviation of the likelihoods was defined as a sine wave with amplitude, frequency and phase as free parameters.

However, there is accumulating evidence that prolonged or repeated exposure to particular stimuli (i.e., adaptation) can have a meaningful effect on the prior probability distribution of observers, shifting them towards the values of the adapting stimuli [31, 32, 46, 47]. In order to determine whether improvement in sensory encoding or a change in the expectation of the stimulus properties best explain our findings, we modeled our psychophysical results with a Bayesian observer similar to previous studies investigating the statistics of natural scenes [19, 31]. We built five model variants that differed in their prior probability distribution and found the environmental prior model, prior probability distributions that represented the current distribution of αs, best captured observer thresholds. In addition, observer likelihoods were also affect by the adaptation paradigm, where a shallow adaptation narrowed the width of the likelihoods near α values of 0.7 and a steep adaptation narrowed the likelihood for α values between 1.0 and 1.3. We find both a

change in the likelihood and in the prior explain the change psychophysical responses. Our findings support that sensitivity to α is associated with the statistical structure or distribution of α in natural scenes and sufficiently malleable to adapt in environments with novel α distributions.

Sensitivity to visual features is dependent on the probability of occurrence of these features in the environment of observers [32]. For example, the typical anisotropic sensitivity to orientation (i.e., the oblique effect) is not solely attributed to inhomogeneity in sensory noise for oblique orientations but also from their relatively lower occurrence compared to cardinal orientations in natural environments [19]. Our modeling results suggest a similar phenomenon concerning α discrimination, which can—in part—be attributed to the occurrence of α values in natural environments. Following exposure to an environment with a modified α distribution increased sensitivity to α at reference values near the peak of the modified α distribution. To capture this effect in observer discrimination thresholds the peak probability of the prior distributions had to match those of the distribution of αs in the environment. Changes in the prior probability distribution in the direction of the environment distributions have been reported for other visual paradigms [31, 32, 46, 47]. Exposure to higher speed stimuli over multiple days changed the speed prior of observers from one with high probabilities for low speeds (the low speed prior) to a prior probability distribution with a peak at higher speeds [47]. This change in the prior is also similar to effects shown by [31] where immersion in an isotropic world (in an HMD) flattened the prior for orientation [19].

Changes to the prior alone, however, do not appear to be sufficient to fully characterize the change in discrimination thresholds following adaptation. This was particularly evident in the shallow adaptation condition, where the model variance, or widths of the model likelihood narrowed for α values near the mean α of the adapting environment. This effect is concordant with previous Bayesian observer models that have found the increase in sensitivity following adaptation to be attributed to a narrowing of the likelihood distributions at the adapter and not a change in the prior [42, 43]. It is plausible that the change in the likelihood following shallow adaptation contributed to the reduction in discrimination thresholds at a reference α of 0.7. Still, a

change in the likelihood alone is unable to adequately capture the discrimination thresholds across all three adaptation conditions. Model variants that were clamped to a set prior (i.e, baseline, steep, shallow or the natural [5] priors), while able to fit their respective experimental condition were unable to capture all experimental conditions as well as the model with flexible priors. Thus, while we do find evidence that a change in the likelihood does contribute to the change in discrimination thresholds, this change alone does not explain the adaptation effects we measured following immersion in a modified environment. Our modified reality procedure subjects observers to the adaptation environment for an extended period (1 hour + 45 minutes of subsequent testing), which is significantly longer than traditional adaptation studies that adapt for periods of seconds or minutes. More extended adaptation periods are have been shown to target different visual mechanisms, which operate on longer-term visual representations than short-term adaptation, which alters mechanisms responsible for moment-to-moment fluctuations [48, 49]. Short adaptations periods are therefore unlikely to generate any meaningful change in the prior of observers. At the same time, a prolonged immersion in modified reality [31] or repeated exposure to certain stimuli over days [47] will stimulate the learning of novel stimulus properties / environmental statistics and effect a change in the prior probability distribution.

Finally, it is worth noting that psychophysical measurements of α discrimination thresholds are traditionally completed in laboratory environments with monitors (e.g., CRTs) that are large and of a sufficiently high resolution to display a wide range of spatial frequencies [9, 13, 24, 25]. Traditional CRTs appear not to be necessary when measuring sensitivity to the slope of the amplitude spectrum of natural scenes. While implementing psychophysical paradigms in HMDs is not novel [31, 50], the measurement of α discrimination threshold, which is dependent on the spatial frequency resolution of the display, is novel. We have demonstrated that discrimination thresholds to α can be measured accurately with the smaller screens of HMDs limited in the range of spatial frequencies they can display. The spatial frequency resolution of the Oculus Rift DK2 only reaches 5.5 cycles/^{o} of visual angle. Still, α discrimination thresholds measured in the Oculus Rift DK2 showed the expected peak and trough for reference αs of 0.7 and 1-1.3,

respectively. The peak and trough were also of similar magnitude to thresholds measured in other studies (using the same staircasing methodology) [24, 25]. Human observers do not need a wide range of spatial frequencies to discriminate α and complete the task with relatively limited spatial frequency content.

Conclusion

Discrimination thresholds to the slope of the amplitude spectrum (α) are subject to the recently viewed environment of observers. Spending a prolonged amount of time in a modified environment where the typical α values are steeper or shallower than the normal visual world of the observer will decrease discrimination thresholds to steep or shallow α, respectively. These effects were well-described by a Bayesian observer model, which showed a systematic shift in its prior probability distribution as a function of adaptation. The priors were shifted in the same direction as the adaptation condition. Our findings demonstrate that expectations on the occurrence of α values that result from a lifetime of exposure remain plastic, accommodating for the statistical structure of the recently viewed environment.

Acknowledgements

We would like to thank Bruce C. Hansen for his insightful comments on this manuscript.

References

[1] David J. Field and Nuala Brady. "Visual sensitivity, blur and the sources of variability in the amplitude spectra of natural scenes". In: *Vision Research* 37.23 (1997), pp. 3367–3383. DOI: 10.1016/S0042-6989(97)00181-8.

[2] Horace B. Barlow. "Possible principles underlying the transformations of sensory messages". In: *Sensory Communication*. Ed. by Walter A Rosenblith. Vol. 6. 2. Cambridge, MA: MIT Press, 1961. Chap. 13, pp. 217–234. ISBN: 9780262518420. DOI: 10.7551/mitpress/9780262518420.003.0013. URL: http://www.trin.cam.ac.uk/horacebarlow/21.pdf.

[3] Bruce C Hansen and Edward A Essock. "A horizontal bias in human visual processing of orientation and its correspondence to the structural components of natural scenes". In: *Journal of vision* 4.12 (2004), pp. 5–5.

[4] Joseph J Atick. "Could information theory provide an ecological theory of sensory processing?" In: *Network: Computation in Neural Systems* 3.2 (May 1992), pp. 213–251. ISSN: 0954-898X. DOI: 10.1088/0954-898X/3/2/009.

[5] V Billock. "Neural acclimation to 1/f spatial frequency spectra in natural images transduced by the human visual system". In: *Physica D: Nonlinear Phenomena* 137.3-4 (2000), pp. 379–391. ISSN: 01672789. DOI: 10.1016/S0167-2789(99)00197-9.

[6] Bruce C Hansen and Edward A Essock. "Influence of scale and orientation on the visual perception of natural scenes". In: *Visual Cognition* 12.6 (2005), pp. 1199–1234. ISSN: 1350-6285. DOI: 10.1080/13506280444000715.

[7] David J Tolhurst, Yoav Tadmor, and T. Chao. "Amplitude spectra of natural images." In: *Ophthalmic and Physiological Optics* 12.2 (1992), pp. 229–232. ISSN: 0275-5408. DOI: 10.1111/j.1475-1313.1992.tb00296.x.

[8] A. van der Schaaf and J.H. van Hateren. “Modelling the Power Spectra of Natural Images: Statistics and Information”. In: *Vision Research* 36.17 (1996), pp. 2759–2770. ISSN: 00426989. DOI: 10.1016/0042-6989(96)00002-8.

[9] David J. Field. “Relations between the statistics of natural images and the response properties of cortical cells”. In: *Journal of the Optical Society of America A* 4.12 (1987), pp. 2379–94. ISSN: 0740-3232. DOI: 10.1364/JOSAA.4.002379.

[10] G J Burton and Ian R Moorhead. “Color and spatial structure in natural scenes”. In: *Applied Optics* 26.1 (1987), pp. 157–170. ISSN: 0003-6935. DOI: 10.1364/AO.26.000157.

[11] Aude Oliva and Antonio Torralba. *Modelling the shape of the scene: A holistic representation of the spatial envelope*. 2001. DOI: 10.1023/A:1011139631724. URL: http://people.csail.mit.edu/torralba/code/spatialenvelope/%20http://www.springerlink.com/index/K62TG81W8352G71H.pdf (visited on 10/22/2011).

[12] Carlos Alejandro Párraga and David J. Tolhurst. “The effect of contrast randomisation on the discrimination of changes in the slopes of the amplitude spectra of natural scenes”. In: *Perception* 29.9 (2000), pp. 1101–1116. ISSN: 0301-0066. DOI: 10.1068/p2904.

[13] Aaron P Johnson et al. “The magnitude of center – surround facilitation in the discrimination of amplitude spectrum is dependent on the amplitude of the surround”. In: *Journal of Vision* 11.7 (2011), pp. 1–10. DOI: 10.1167/11.7.14.

[14] Carlos Alejandro Párraga, Tom Troscianko, and David J. Tolhurst. “The effects of amplitude-spectrum statistics on foveal and peripheral discrimination of changes in natural images, and a multi-resolution model.” In: *Vision research* 45.25-26 (2005), pp. 3145–68. ISSN: 0042-6989. DOI: 10.1016/j.visres.2005.08.006.

[15] Sarah L Elliott, Mark A Georgeson, and Michael A Webster. “Response normalization and blur adaptation: Data and multi-scale model”. In: *Journal of Vision* 11.2 (2011), pp. 1–18. DOI: 10.1167/11.2.7.

[16] Michael A. Webster, Mark A Georgeson, and Shernaaz M Webster. “Neural adjustments to image blur.” In: *Nature neuroscience* 5.9 (Sept. 2002), pp. 839–840. ISSN: 10976256. DOI: 10.1038/nn906. URL: http://www.ncbi.nlm.nih.gov/pubmed/12195427%20http://www.nature.com/doifinder/10.1038/nn906.

[17] Michael A. Webster and Eriko Miyahara. “Contrast adaptation and the spatial structure of natural images”. In: *Journal of the Optical Society of America A* 14.9 (1997), pp. 2355–2366. DOI: 10.1364/JOSAA.14.002355.

[18] Alan A. Stocker and Eero P. Simoncelli. “Noise characteristics and prior expectations in human visual speed perception”. In: *Nature Neuroscience* 9.4 (2006), pp. 578–585. ISSN: 10976256. DOI: 10.1038/nn1669.

[19] Ahna R Girshick, Michael S Landy, and Eero P Simoncelli. “Cardinal rules: visual orientation perception reflects knowledge of environmental statistics”. In: *Nature neuroscience* 14.7 (2011), pp. 926–932.

[20] Ruyuan Zhang, Oh-Sang Kwon, and Duje Tadin. “Illusory Movement of Stationary Stimuli in the Visual Periphery: Evidence for a Strong Centrifugal Prior in Motion Processing”. In: *Journal of Neuroscience* 33.10 (2013), pp. 4415–4423. DOI: 10.1523/JNEUROSCI.4744-12.2013.

[21] D C Knill, David J. Field, and D Kersten. “Human discrimination of fractal images.” In: *Journal of the Optical Society of America. A, Optics and image science* 7.6 (1990), pp. 1113–1123. ISSN: 0740-3232. DOI: 10.1364/JOSAA.7.001113.

[22] Yoav Tadmor and David J Tolhurst. “Discrimination of changes in the second-order statistics of natural and synthetic images”. In: *Vision Research* 34.4 (Feb. 1994), pp. 541–554. ISSN: 00426989. DOI: 10.1016/0042-6989(94)90167-8. URL: http://www.ncbi.nlm.nih.gov/pubmed/8303837%20http://linkinghub.elsevier.com/retrieve/pii/0042698994901678.

[23] David J Tolhurst and Yoav Tadmor. “Band-limited contrast in natural images explains the detectability of changes in the amplitude spectra”. In: *Vision research* 37.23 (Dec. 1997), pp. 3203–15. ISSN: 0042-6989. DOI: 10.1016/S0042-6989(97)00119-3.

[24] Bruce C. Hansen and Robert F. Hess. “Discrimination of amplitude spectrum slope in the fovea and parafovea and the local amplitude distributions of natural scene imagery”. In: *Journal of Vision* 6.7 (2006), pp. 696–711. ISSN: 1534-7362. DOI: 10.1167/6.7.3.

[25] Bruno Richard et al. “Spatial summation of broadband contrast”. In: *Journal of Vision* 19.5 (May 2019), p. 16. ISSN: 1534-7362. DOI: 10.1167/19.5.16.

[26] M. G. A. Thomson and David H Foster. “Role of second- and third-order statistics in the discriminability of natural images”. In: *Journal of the Optical Society of America A* 14.9 (Sept. 1997), pp. 2081–2092. ISSN: 1084-7529. DOI: 10.1364/JOSAA.14.002081. URL: http://www.opticsinfobase.org/abstract.cfm?URI=josaa-14-9-2081.

[27] Michael A. Webster. “Visual Adaptation”. In: *Annual Review of Vision Science* 1.1 (2015), pp. 547–567. DOI: 10.1146/annurev-vision-082114-035509.

[28] Michael A Webster. “Adaptation and visual coding.” In: *Journal of vision* 11.5 (2011), pp. 1–23. DOI: 10.1167/11.5.3.

[29] C. Blakemore and F W. Campbell. “On the existence of neurones in the human visual system selectively sensitive to the orientation and size of retinal images”. In: *The Journal of physiology* 203.1 (1969), pp. 237–260. ISSN: 00223751. DOI: 10.1113/jphysiol.1969.sp008862.

[30] Hugh R. Wilson and Richard Humanski. “Spatial frequency adaptation and contrast gain control”. In: *Vision Research* 33.8 (1993), pp. 1133–1149. DOI: 10.1016/0042-6989(93)90248-U.

[31] April Schweinhart, Patrick Shafto, and Edward A. Essock. “Distribution of content in recently-viewed scenes whitens perception”. In: *Journal of Vision* 17.3 (2017), p. 8. DOI: 10.1167/17.3.8.

[32] Peggy Series and Aaron Seitz. “Learning what to expect (in visual perception)”. In: *Frontiers in Human Neuroscience* 7 (2013), p. 668. ISSN: 1662-5161. DOI: 10.3389/fnhum.2013.00668. URL: https://www.frontiersin.org/article/10.3389/fnhum.2013.00668.

[33] Edward A Essock et al. “Oblique stimuli are seen best (not worst!) in naturalistic broad-band stimuli: A horizontal effect”. In: *Vision Research* 43.12 (2003), pp. 1329–1335. ISSN: 00426989. DOI: 10.1016/S0042-6989(03)00142-1.

[34] Bruce C Hansen et al. “A cortical locus for anisotropic overlay suppression of stimuli presented at fixation”. In: *Visual Neuroscience* 32 (2015), E023. ISSN: 0952-5238. DOI: 10.1017/S0952523815000255.

[35] Andrew M Haun and Edward A Essock. “Contrast sensitivity for oriented patterns in 1/f noise: contrast response and the horizontal effect.” In: *Journal of vision* 10 (2010), p. 1. ISSN: 1534-7362. DOI: 10.1167/10.10.1.

[36] David H Brainard. “The psychophysics toolbox”. In: *Spatial Vision* 10.4 (1997), pp. 433–436. ISSN: 0169-1015. DOI: 10.1163/156856897X00357.

[37] Denis G. Pelli. “The VideoToolbox software for visual psychophysics: Transforming numbers into movies”. In: *Spatial Vision* 10.4 (1997), pp. 437–442. ISSN: 0169-1015. DOI: 10.1163/156856897X00366.

[38] Nicolaas Prins and Frederick A. A. Kingdom. *Palamedes: Matlab routines for analyzing psychophysical data.* 2009. URL: http://www.palamedestoolbox.org/.

[39] Christian Kaernbach. “Simple adaptive testing with the weighted up-down method”. In: *Perception & psychophysics* 49.3 (Mar. 1991), pp. 227–229. ISSN: 0031-5117. DOI: 10.3758/BF03214307. URL:

http://www.ncbi.nlm.nih.gov/pubmed/2011460%20http://www.springerlink.com/index/10.3758/BF03214307.

[40] Bruno Richard et al. “The effects of tDCS across the spatial frequencies and orientations that comprise the Contrast Sensitivity Function”. In: *Frontiers in Psychology: Perception Science* 6 (2015), p. 1784. DOI: 10.3389/fpsyg.2015.01784.

[41] Michael A. Webster, John S. Werner, and David J. Field. “Adaptation and the phenomenology of perception”. In: *Fitting the Mind to the World*. Ed. by Colin W. G. Clifford and G. Rhodes. Oxford: Oxford University Press, 2005, pp. 241–277. ISBN: 9780198529699.

[42] Alan A. Stocker and Eero P. Simoncelli. “Sensory Adaptation within a Bayesian Framework for Perception”. In: *Proceedings of the 18th International Conference on Neural Information Processing Systems*. NIPS’05. Vancouver, British Columbia, Canada: MIT Press, 2006, pp. 1289–1296.

[43] Colin W G Clifford et al. “Visual adaptation: Neural, psychological and computational aspects”. In: *Vision Research* 47.25 (2007), pp. 3125–3131. ISSN: 00426989. DOI: 10.1016/j.visres.2007.08.023.

[44] Peggy Seriès, Alan A Stocker, and Eero P Simoncelli. “Is the Homunculus “Aware” of Sensory Adaptation?” In: *Neural Computation* 21.12 (2009), pp. 3271–3304. DOI: 10.1162/neco.2009.09-08-869. URL: https://doi.org/10.1162/neco.2009.09-08-869.

[45] Wilson S Geisler. “Ideal observer analysis”. In: *The visual neurosciences* 10.7 (2003), pp. 12–12.

[46] Wendy J Adams, Erich W Graf, and Marc O Ernst. “Experience can change the ’light-from-above’ prior.” In: *Nature neuroscience* 7.10 (2004), pp. 1057–1058. ISSN: 1097-6256. DOI: 10.1038/nn1312.

[47] Grigorios Sotiropoulos, Aaron R. Seitz, and Peggy Seriès. "Changing expectations about speed alters perceived motion direction". In: *Current Biology* 21.21 (2011), R883–R884. DOI: https://doi.org/10.1016/j.cub.2011.09.013.

[48] Min Bao and Stephen A Engel. "Distinct mechanism for long-term contrast adaptation". In: *Proceedings of the National Academy of Sciences* (2012). ISSN: 0027-8424. DOI: 10.1073/pnas.1113503109.

[49] Carlyn A. Patterson, Stephanie C. Wissig, and Adam Kohn. "Distinct Effects of Brief and Prolonged Adaptation on Orientation Tuning in Primary Visual Cortex". In: *Journal of Neuroscience* 33.2 (2013), pp. 532–543. ISSN: 0270-6474. DOI: 10.1523/JNEUROSCI.3345-12.2013. eprint: https://www.jneurosci.org/content/33/2/532.full.pdf. URL: https://www.jneurosci.org/content/33/2/532.

[50] Peng Zhang et al. "Effects of Orientation-Specific Visual Deprivation Induced with Altered Reality". In: *Current Biology* 19.22 (2009), pp. 1956–1960. DOI: https://doi.org/10.1016/j.cub.2009.10.018.

A Fits of the other models investigated in this study

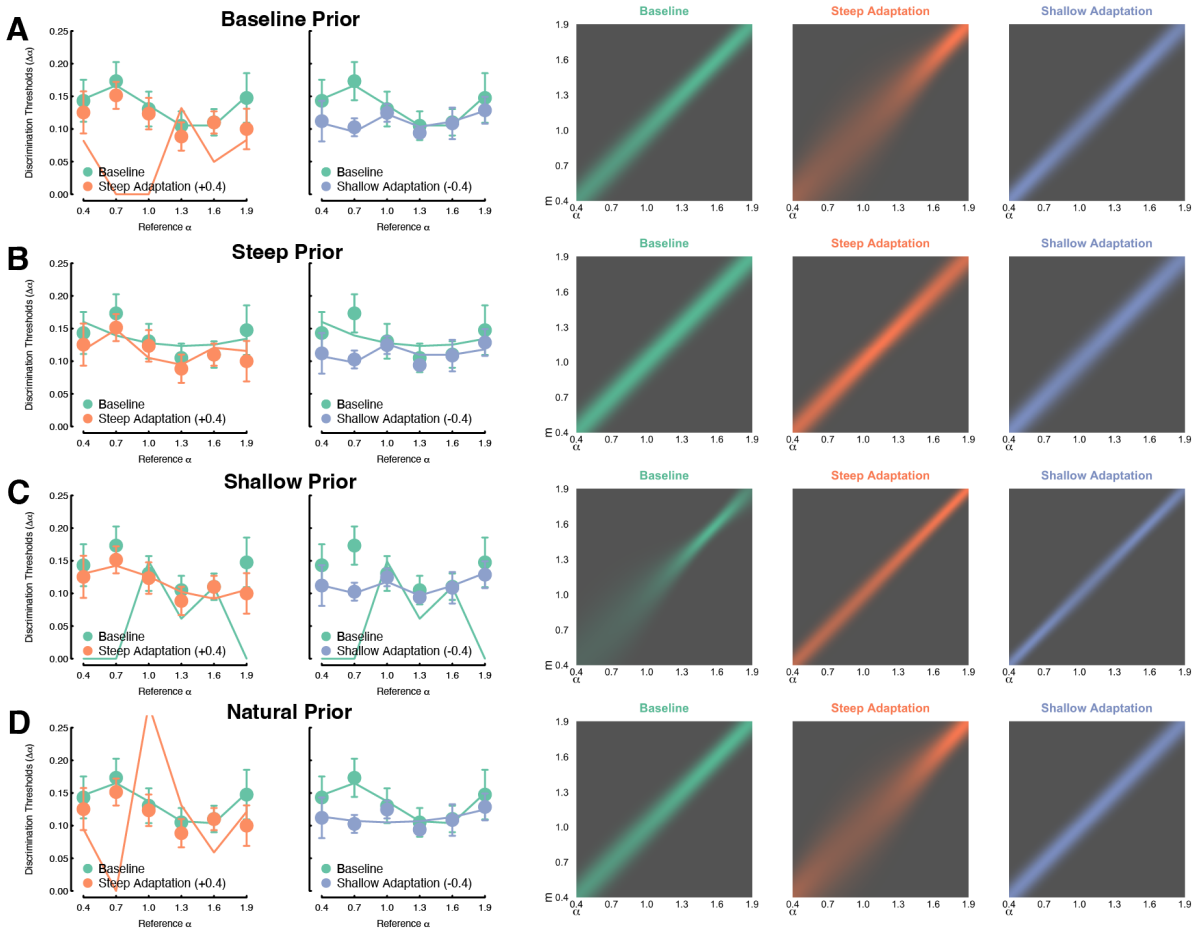

Figure 5: **A** The resulting fits and likelihoods of a Bayesian observer model with a prior probability distribution fixed to the baseline environmental distribution of α (α = 1.29). **B** The resulting fits and likelihoods of a Bayesian observer model with a prior probability distribution fixed to the steepened environmental distribution of α (α = 1.69). textbfC The resulting fits and likelihoods of a Bayesian observer model with a prior probability distribution fixed to the shallowed environmental distribution of α (α = 0.89). The resulting fits and likelihoods of a Bayesian observer model with a prior probability distribution fixed to the natural environmental distribution [5] of α (α = 1.08). Models A - C captured the discrimination thresholds of their respective experimental condition well, but failed when fitting other experimental conditions. The natural prior performed well with baseline discrimination thresholds but could not fit the other experimental conditions.

	Baseline		Steep Adaptation		Shallow Adaptation	
Model	r^2	RMS_e	r^2	RMS_e	r^2	RMS_e
Environmental Priors	0.966	0.004	0.669	0.012	0.911	0.004
Baseline Prior	0.966	0.004	-	0.087	0.817	0.005
Steep Prior	0.321	0.019	0.669	0.012	0.530	0.008
Shallow Prior	-	0.111	0. 729	0.011	0.911	0.004
Natural Prior	0.947	0.005	-	0.096	0.291	0.010

Table 3: Goodness of fit measures for each model explored in this study for all experimental conditions presented separately. Some models generated errors too large to calculate an $^r2$, in these scenarios a - is used to mark the missing value.