Supplementary Information

for *A functional and perceptual signature of the second visual area in primates*

Jeremy Freeman, Corey M. Ziemba, David J. Heeger, Eero P. Simoncelli, & J. Anthony Movshon

Contents

Nature Neuroscience: doi:10.1038/nn.3402

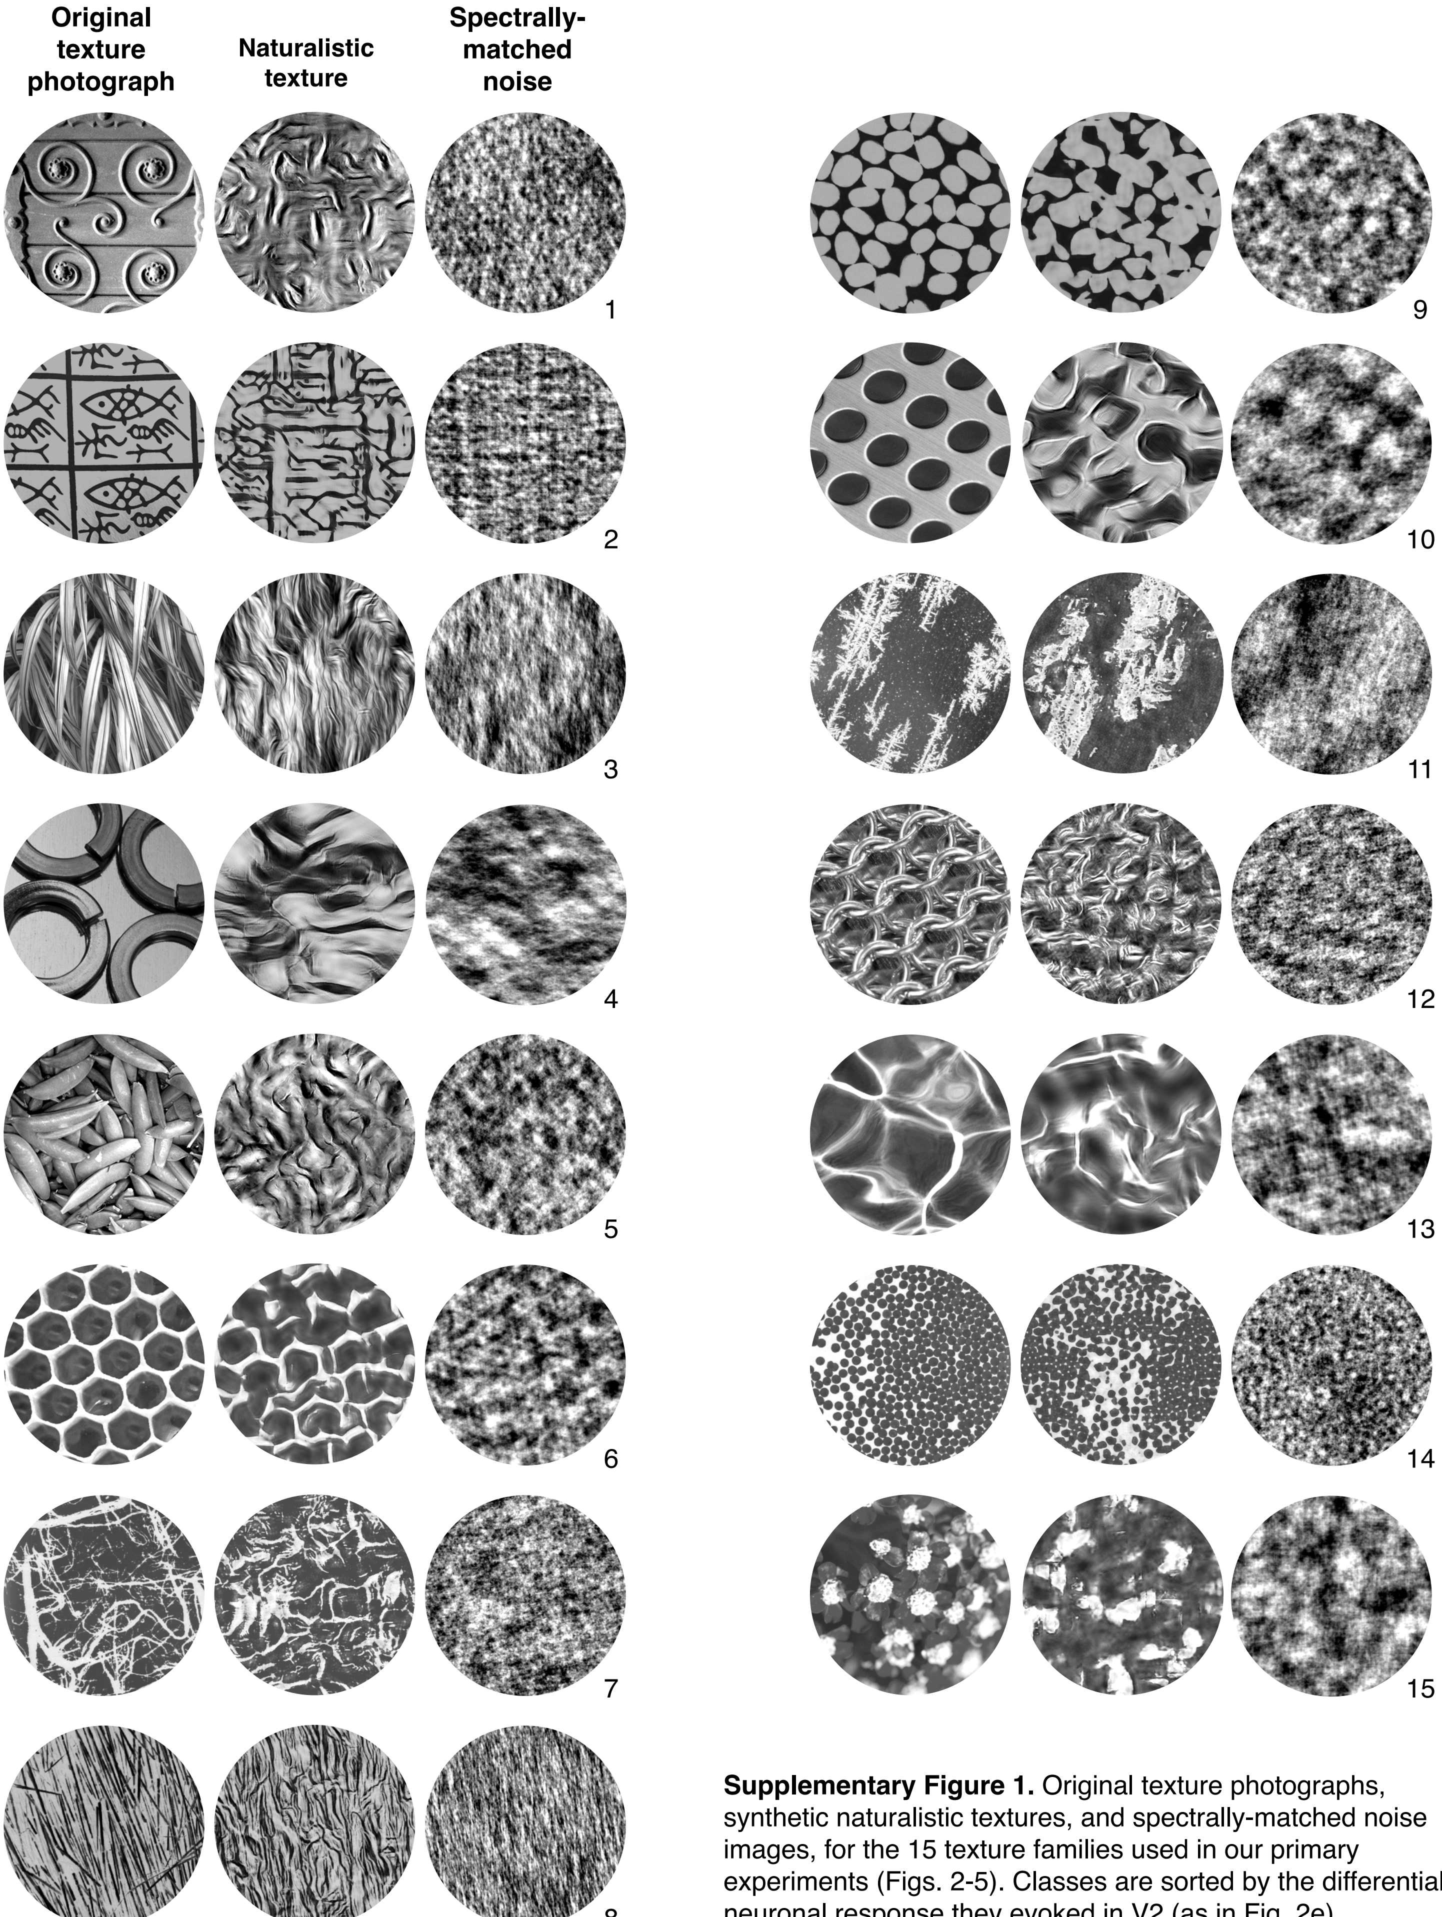

Nature Neuroscience: doi:10.1038/nn.3402

Supplementary Figure 1. Original texture photographs, synthetic naturalistic textures, and spectrally-matched noise images, for the 15 texture families used in our primary experiments (Figs. 2-5). Classes are sorted by the differential neuronal response they evoked in V2 (as in Fig. 2e).

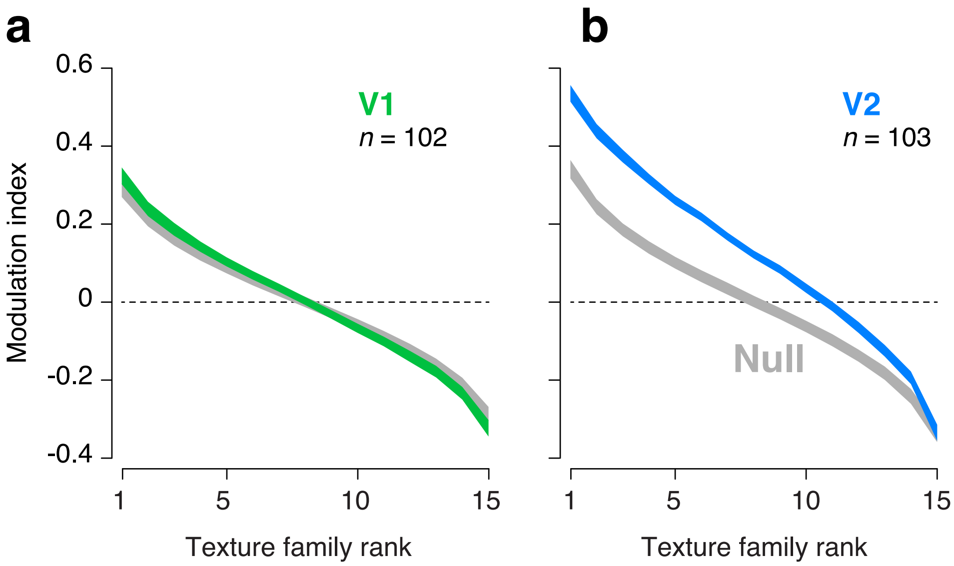

Supplementary Figure 2. Average texture family ranking by modulation. For each cell, the 15 texture families were ranked and sorted according to modulation index. The sorted modulation indices were then averaged across all cells in V1 (green, panel a) and V2 (blue, panel b). A null distribution was obtained for both areas by permuting the naturalistic and noise labels and iterating 1000 times. Gray areas in both panels indicate the 2.5th and 97.5th percentiles of the null distribution. On average, any differences between naturalistic and noise stimuli exhibited by V1 cells, either positive or negative, were not distinguishable from those expected by chance, but this was not the case for V2 cells.

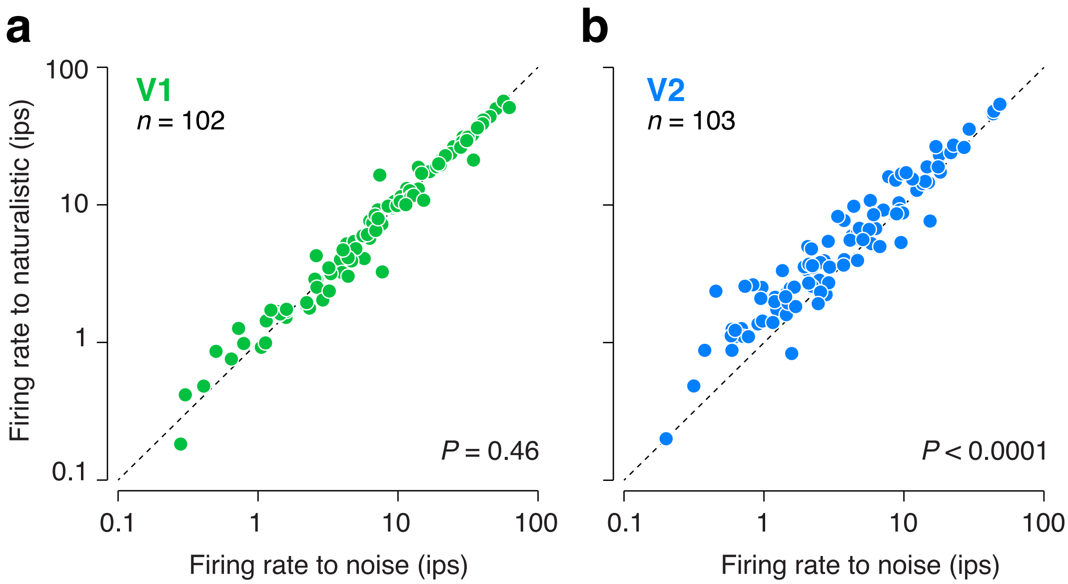

Supplementary Figure 3. Difference between V1 and V2 in terms of firing rate. Responses to naturalistic and spectrally-matched noise images in V1 (green, panel a) and V2 (blue, panel b). Diagonal dashed line is the line of equality. The difference between naturalistic and noise was statistically significant in V2 but not in V1 (paired *t*-tests), and the difference of differences between the two areas was statistically significant ($P < 0.0001$, unpaired *t*-test). Neurons with comparable responsiveness (~5-10 ips) show a differential response to naturalistic images in V2 but not in V1.

Supplementary Modeling

For each texture family, we fit data from all Turkers simultaneously with a mixture model. The model consists of a psychometric function common to all Turkers, parameterized with a slope and threshold, and a parameter that controls the quality of each Turker. Specifically, for each texture family, we assume that N Turkers performed the psychophysical task. The task contained C conditions (the different levels of naturalness), and there were T trials for each condition. On every trial, the Turker provided a response x_{nct} that was either correct ($x_{nct} = 1$) or incorrect ($x_{nct} = 0$). The probability of a response being correct is governed by an Turker-independent function $p_c = F(c, \theta)$ which relates the conditions to a probability of correct response via the parameters θ of a cumulative Wiebull function. The probability of correct response is also determined by an Turker-dependent lapse parameter λ_n which gives the probability that an observer will lapse on any trial, that is, respond randomly rather than according to p_c. Let Θ represent all parameters (those governing the psychometric function, and the lapse rates for all observers). We introduce the latent variable z_{nct} to represent whether or not an observer lapsed on a particular condition/trial combination. We make use of the indicator variable z_{nctk} : if $z_{nct} = 1$, then $z_{nct1} = 1$ and $z_{nct0} = 0$; if $z_{nct} = 0$, then $z_{nct1} = 0$ and $z_{nct0} = 1$.

Consider a particular Turker, trial, and condition. If the Turker lapses, she will respond correctly at chance, so the distribution of her response is given by a Bernoulli random variable

$$P(x_{nct} \mid z_{nct} = 1, \theta) = \gamma^{x_{nct}} (1 - \gamma)^{1 - x_{nct}} \text{ ,}$$

where γ is $1/3$ for the 3AFC task. And if she does not lapse, her response will be governed by the psychometric function

$$P(x_{nct} \mid z_{nct}=0,\theta)=p_c^{x_{nct}}\left(1-p_c\right)^{1-x_{nct}}.$$

We can use the indicator variables to write the joint distribution over the data and the latent variables as

$$P(x_{nct},z_{nct} \mid \theta,\lambda_n)=\left[\left(\gamma^{x_{nct}}(1-\gamma)^{1-x_{nct}}\right)\lambda_n\right]^{z_{nct1}}\left[\left(p_c^{x_{nct}}\left(1-p_c\right)^{1-x_{nct}}\right)(1-\lambda_n)\right]^{z_{nct0}}. \qquad \text{Eq. 1}$$

Note the dependence on the marginal probability of a lapse, λ_n. When $z_{nct}=1$, the above reduces to the first term alone, which is

$$P\left(x_{nct} \mid z_{nct}=1\right)P\left(z_{nct}=1\right),$$

and likewise for the second term. Thus, the expression for the joint uses the indicator variable to capture what is essentially a piecewise combination of Bernoulli distributions.

The complete log likelihood of the data under this model is

$$\ln P(X \mid \Theta) \ =\ln\prod_{nct}\left(P(x_{nct} \mid \theta,\lambda_n)\right)=\ln\prod_{nct}\left(\sum_z P(x_{nct},z_{nct} \mid \theta,\lambda_n)\right).$$

Directly maximizing this function with respect to θ and λ_n would be difficult. However, note that if the true values of the latent variables were known, maximizing the log likelihood of the data would become linear in the parameters (by taking the log of Eq. 1). Thus, this problem is naturally suited to the EM (expectation-maximization) algorithm. Given a current setting of the parameters $\lambda_n^{(t)}$ and $\theta^{(t)}$, we can write the expected log likelihood of the data with respect to the conditional distribution of the latent variables,

$$Q(\Theta \mid \Theta^{(t)})=\mathbb{E}_{Z \mid X,\Theta^{(t)}}\left[\ln P(X,Z \mid \Theta)\right]. \qquad \text{Eq. 2}$$

We alternate between computing this expected value, and then estimating the parameters that maximize Eq. 2. By the linearity of expectation, it will suffice to compute a point estimate of the expected value of each z_{nct}. To compute that expected value, we need the probability of the latent variables given a known set of parameter values, which we obtain using Bayes' rule,

$$P(Z \mid X, \Theta^{(t)}) = \frac{P(X \mid Z, \Theta^{(t)}) P(Z \mid \Theta^{(t)})}{\sum_{Z'} P(X \mid Z', \Theta^{(t)}) P(Z' \mid \Theta^{(t)})}.$$

Because $\tilde{z}_{nct} = \mathbb{E}(z_{nct} \mid x_{nct}, \Theta^{(t)}) = P(z_{nct} = 1 \mid x_{nct}, \Theta^{(t)})$, we need only compute

$$P(z_{nct} = 1 \mid x_{nct}, \Theta^{(t)}) = \frac{(\gamma^{x_{nct}} (1-\gamma)^{1-x_{nct}}) \lambda_n}{\left(p_c^{x_{nct}} \left(1-p_c\right)^{1-x_{nct}}\right)(1-\lambda_n) + (\gamma^{x_{nct}} (1-\gamma)^{1-x_{nct}}) \lambda_n}, \qquad \text{Eq. 3}$$

where p_c and λ_n depend on $\Theta^{(t)}$. We now consider the quantity to be maximized,

$$\begin{aligned} \mathbb{E}_{Z|X,\Theta^{(t)}}[\ln P(X,Z \mid \Theta)] = \sum_{nct} & \tilde{z}_{nct}[x_{nct} \ln(\gamma) + (1-x_{nct}) \ln(1-\gamma) + \ln(\lambda_n)] \\ & + (1-\tilde{z}_{nct})[x_{nct} \ln(p_c) + (1-x_{nct}) + \ln(1-p_c) + \ln(1-\lambda_n)] \end{aligned}$$

where we have used the linearity of expectation to replace z_{nct} with $\tilde{z}_{nct}$ from Eq. 3. Differentiating with respect to the parameters of interest yields maximum likelihood estimates. In practice, we want to differentiate with respect to the parameters θ that control p_c, or find maximum likelihood estimates through numerical optimization if the derivatives are non-trivial. But for simplicity and intuition, here we differentiate with respect to p_c directly to obtain an estimate of the fraction of correct responses:

$$\hat{p}_c = \frac{\sum_{nt} (1-\tilde{z}_{nct}) x_{nct}}{\sum_{nt} (1-\tilde{z}_{nct})} \qquad \text{Eq. 4}$$

Nature Neuroscience: doi:10.1038/nn.3402

$$\hat{\lambda}_n = \frac{\sum_{ct} \tilde{z}_{nct}}{CT}. \qquad \text{Eq. 5}$$

The expression for $\hat{p}_c$ is simply the fraction of correct responses weighted by the lapse occurrence; if there were no lapses, the denominator would contain only 1s, and Eq. 4 would reduce to the number of correct responses divided by the total number of trials. The expression for $\hat{\lambda}_n$ (Eq. 5) is similarly intuitive: the number of lapse trials divided by the total number of trials (across all conditions). Having obtained these estimates on the M step, they are used on the E step to compute the expectation in Eq. 3.

We confirmed that the parameter estimates obtained from this algorithm reliably converged from multiple random initializations. As expected, the analysis estimated high lapse rates for Turkers with outlier behavior (e.g. near chance performance in all conditions), and the analysis ensured that these Turkers contributed minimally to estimates of slope and threshold.