Psychonomic Bulletin & Review
2010, 17 (5), 624-629
doi:10.3758/PBR.17.5.624

Brief Reports

Subjective randomness and natural scene statistics

Anne S. Hsu
University College London, London, England

Thomas L. Griffiths
University of California, Berkeley, California

and

Ethan Schreiber
University of California, Los Angeles, California

Accounts of subjective randomness suggest that people consider a stimulus random when they cannot detect any regularities characterizing the structure of that stimulus. We explored the possibility that the regularities people detect are shaped by the statistics of their natural environment. We did this by testing the hypothesis that people's perception of randomness in two-dimensional binary arrays (images with two levels of intensity) is inversely related to the probability with which the array's pattern would be encountered in nature. We estimated natural scene probabilities for small binary arrays by tabulating the frequencies with which each pattern of cell values appears. We then conducted an experiment in which we collected human randomness judgments. The results show an inverse relationship between people's perceived randomness of an array pattern and the probability of the pattern appearing in nature.

People are very sensitive to deviations from their expectations about randomness. For example, the game Yahtzee involves repeatedly rolling 5 six-sided dice. If you were to roll all sixes 6 times in a row, you would probably be quite surprised. The probability of such a sequence arising by chance is $1/6^{30}$. However, the low probability of such an event is not sufficient to explain its apparent nonrandomness, since any other ordered sequence of the same number of dice rolls has the same probability. Consequently, recent accounts of human subjective randomness (our sense of the extent to which an event seems random) have focused on the regularities in an event. These regularities suggest that a process other than chance might be at work (Falk & Konold, 1997; Feldman, 1996, 1997; Griffiths & Tenenbaum, 2003, 2004). The basic idea behind these accounts is that stimuli will appear random when they do not express any regularities.

An important challenge for any account of subjective randomness based on the presence of regularities is to explain why people should be sensitive to a particular set of regularities. In the example given above, systematic runs of the same number may suggest loaded dice or some other nonrandom process influencing the outcomes. However, for other kinds of stimuli, such as the one- or two-dimensional binary arrays used in many subjective randomness experiments, explanations are more difficult to come by. A common finding in these experiments is that people consider arrays in which cells take different values from their neighbors (such as the one-dimensional array 0010101101) more random than arrays in which cells take the same values as their neighbors (such as 0000011111) (Falk & Konold, 1997). This result makes it clear that people are sensitive to certain regularities, such as cells having the same values as their neighbors. However, it is difficult to explain why these regularities should be more important than others that seem a priori plausible, such as neighboring cells differing in their values.

In this article, we explore a possible explanation for the origins of the regularities that influence subjective randomness judgments for one class of stimuli: two-dimensional binary arrays. These stimuli are essentially images, with the cells in the array having the appearance of a grid of black and white pixels (see Figure 1). We might thus expect that the kinds of regularities detected by the visual system should play an important role in determining their perceived randomness. A great deal of recent research suggests that the human visual cortex efficiently codes for the structure of natural scenes—scenes containing natural ele-

T. L. Griffiths, tom_griffiths@berkeley.edu

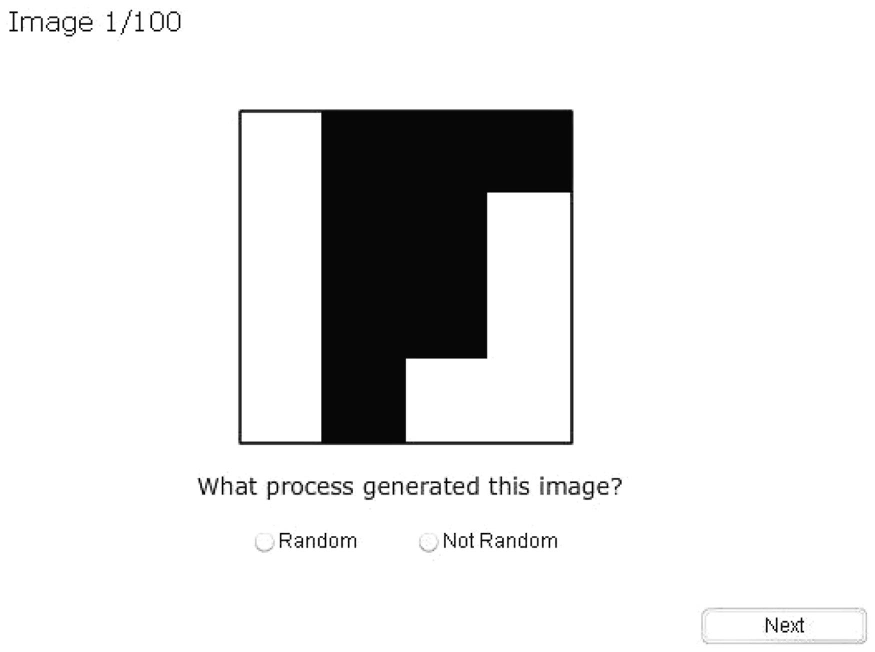

Figure 1. An example of a two-dimensional binary array. This is a screenshot from the experiment, indicating how the participants made their responses.

ments, such as trees, flowers, and shrubs, that represent the visual environment in which humans evolved (Olshausen & Field, 2000; Simoncelli & Olshausen, 2001). We consider the possibility that the kinds of regularities that people detect in two-dimensional binary arrays are those that are characteristic of natural scenes.

Preliminary support for the idea that the statistics of natural scenes may explain subjective randomness was provided by a study conducted by Schreiber and Griffiths (2007). In this study, human randomness judgments were found to correspond to the predictions of a simple probabilistic model estimated from images of natural scenes. This model examined the frequency with which neighboring regions of an image had the same intensity value. It was found that neighboring regions tend to have similar intensity values, providing a potential explanation for why people consider binary arrays in which neighboring cells differ in their values more random. However, a full characterization of natural scene statistics is not feasible for large binary arrays, due to the exponentially large number of patterns of cell values that can be expressed in such arrays. In our present work, we use small binary arrays, which allow us to directly estimate a probability distribution over all possible patterns of cell values from images of natural scenes. We then conduct an experiment with human participants to examine the relationship between this probability distribution and the subjective randomness of the image.

Subjective Randomness As Bayesian Inference

One explanation for human randomness judgments is to view them as the result of an inference as to whether an observed stimulus, X, was generated by chance or by some other, more regular process (Feldman, 1996, 1997; Griffiths & Tenenbaum, 2003, 2004). If we let $P(X \mid \text{random})$ denote the probability that X will be generated by chance and $P(X \mid \text{regular})$ be the probability of X under the regular generating process, Bayes's rule gives the posterior odds in favor of random generation as

$$\frac{P(\text{random} \mid X)}{P(\text{regular} \mid X)} = \frac{P(X \mid \text{random})\, P(\text{random})}{P(X \mid \text{regular})\, P(\text{regular})}, \tag{1}$$

where $P(\text{random})$ and $P(\text{regular})$ are the prior probabilities assigned to the random and regular processes, respectively. Only the first term on the right-hand side of this expression, the likelihood ratio, changes as a function of X, making it a natural measure of the amount of evidence X provides in favor of a random generating process. Hence, we can define the randomness of a stimulus X as

$$\text{random}(X) = \log \frac{P(X \mid \text{random})}{P(X \mid \text{regular})}, \tag{2}$$

where the logarithm simply places the result on a linear scale.

The measure of subjective randomness defined in Equation 2 has been used to model human randomness judgments for single-digit numbers and one-dimensional binary arrays (Griffiths & Tenenbaum, 2001, 2003, 2004). Following Schreiber and Griffiths (2007), we examine how subjective randomness might be applied to two-dimensional binary arrays of the kind shown in Figure 1. A reasonable choice of $P(X \mid \text{random})$ is to assume that each cell in the array takes on a value of 1 or 0 with equal probability, making $P(X \mid \text{random}) = 1/2^m$, where m is the number of cells in the array. However, defining $P(X \mid \text{regular})$ is more challenging. Direct estimation of the probability of all 2^m binary arrays becomes intractable as m becomes large. Thus, we use 4×4 binary arrays for which we can exhaustively tabulate the frequencies of all

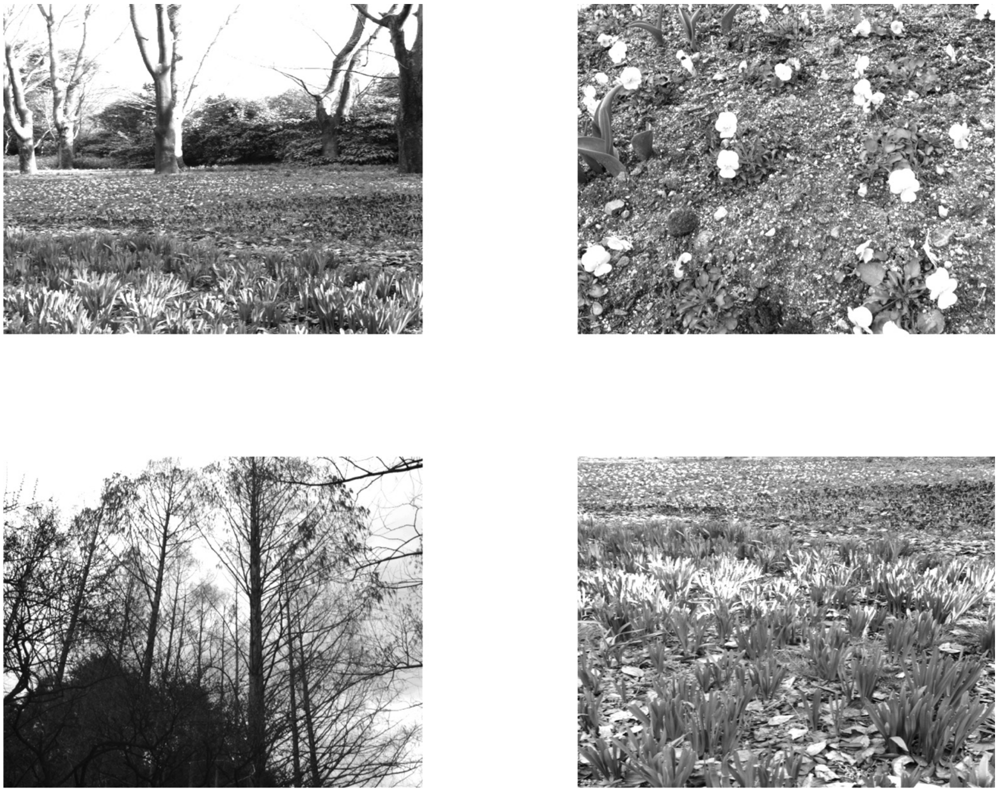

Figure 2. Examples of natural scene images used in our study. The images are from a data set used in the study "Spatiochromatic Receptive Field Properties Derived From Information-Theoretic Analyses of Cone Mosaic Responses to Natural Scenes," by E. Doi, T. Inui, T. W. Lee, T. Wachtler, and T. J. Sejnowski, 2003, *Neural Computation*, *15*, published by MIT Press. Printed with the authors' permission.

patterns that can appear, providing a nonparametric estimate of $P(X \mid \text{regular})$ that allows for a comprehensive test of our hypothesis.

We estimated the values of $P(X \mid \text{regular})$ for a set of stimuli X corresponding to 4×4 binary arrays. Since there are only 2^{16} possible patterns that can be expressed in such an array, we can count the frequency with which each pattern appears in natural images. These stimuli were extracted from a set of images of natural scenes that have been used in previous research (Doi, Inui, Lee, Wachtler, & Sejnowski, 2003). This set consisted of 62 still nature shots containing trees, flowers, and shrubs, as shown in Figure 2. There were no images of humans, animals, or cityscapes. Image patches of varying sizes were extracted from each natural image to measure statistics at a range of scales. A total of 700,000 patches were sampled at random from among the 62 images with dimensions $n \times n$, for $n = 4, 8, 16, 32, 64, 128$, and 256 pixels. All the patches were then reduced through averaging down to 4×4 arrays, binarized by setting pixels with intensity greater than 0 (the overall mean intensity) to 1 and with all other pixels being set to 0. The resulting 4,900,000 binary arrays were then divided into the 2^{16} possible patterns, and the frequency of each pattern was recorded. Normalizing these frequencies gives us an estimate for the probability distribution $P(X \mid \text{regular})$, which we can use to compare the Bayesian measure of randomness given in Equation 2 with human judgments.

METHOD

Participants

The participants were 77 members of the University of California, Berkeley community, who participated for course credit.

Stimuli

The participants were shown one hundred 4×4 binary arrays. The stimuli were restricted to evenly balanced 4×4 patches, containing an equal number of black and white pixels, in order to avoid any effects resulting from the ratio of black to white pixels in each patch. All stimuli with this property were rank-ordered by the probability with which they appeared in natural scenes. We determined the stimulus for which this probability equaled $P(X \mid \text{random}) = 1/2^{16}$, and labeled this the *neutral stimulus* because its probability is equal under the regular and random generating processes and its randomness score is thus zero. Test stimuli were selected from 50 evenly spaced quantiles on either side of the neutral stimulus for a total of 100 images and are shown in Figure 3. Five practice stimuli were selected from the 10%, 30%, 50%, 70%, and 90% quantiles of the ordered images.

Procedure

The participants were told that the stimuli were created using either a random process or another undefined process. The proportion of each was unspecified. They were asked to decide which process generated each stimulus. Below the stimulus were two buttons, as shown in Figure 1. The first button was labeled "Random," and the second was labeled "Not Random." The participants were instructed to press the button corresponding to their intuition about which process had generated the stimulus. In order to familiarize the participants with

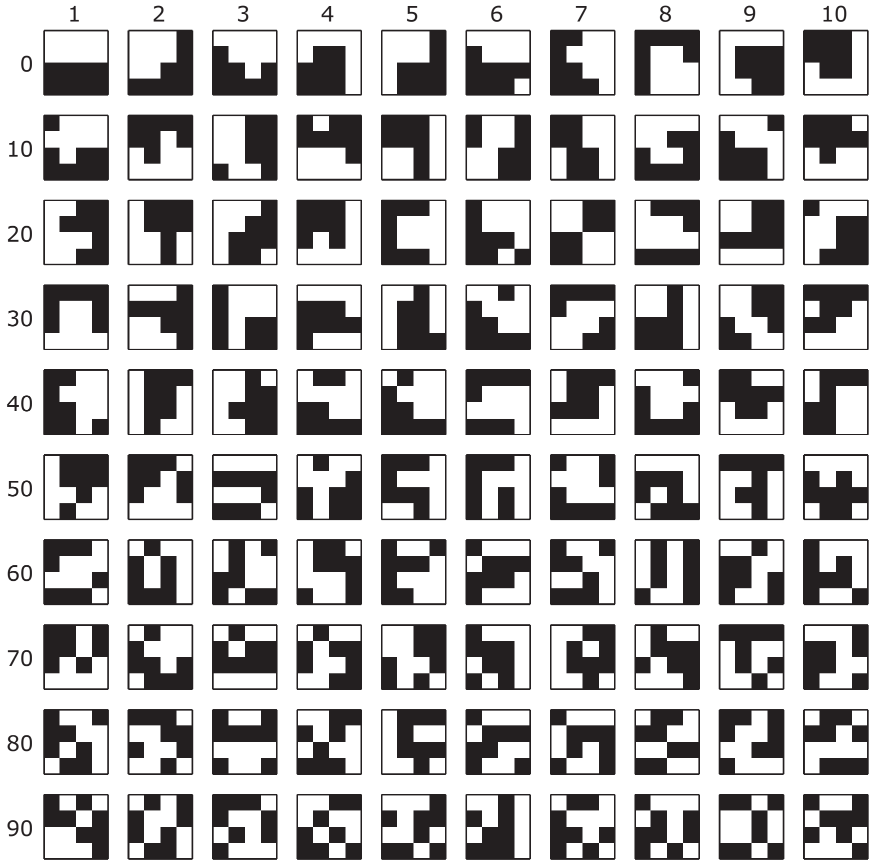

Figure 3. Stimuli for the experiment: Binary images for which a randomness score is estimated from a nonparametric histogram of natural images. The stimuli are ordered from the lowest to the highest value of random(X), with numbers increasing from left to right and then from top to bottom.

this procedure, the first 5 stimuli were presented as practice trials, followed by the 100 stimuli that constituted the main experiment.

RESULTS

The results are shown in Figure 4. A one-way ANOVA showed a statistically significant effect of image [$F(99,7524) = 30.26$, $MS_e = 0.17$, $p < .001$]. The linear correlation between random(X) and the probability that the stimulus would be classified as random was $r(98) = .75$, $p < .001$, and the rank–order correlation (taking into account only the relative ordering of these different measures) was $\rho = .75$ ($p < .01$). These results bear out the predictions of the Bayesian model and indicate that, for the most part, the distribution estimated using a nonparametric histogram of natural scene images provides a reasonable candidate for $P(X \mid \text{regular})$.

A few stimuli deviated noticeably from the model predictions. For example, Image 26 (3rd row, 6th column, of Figure 3) was rated as significantly more random, on average, than were other images with similar random(X) values. Here, the patch seemed to capture an intersection of a curve with a line (a semicircle on the bottom half with a line on the left side). This is a common occurrence in natural images—hence, its relatively high value of $P(X \mid \text{regular})$. However, the pixelation resulting from reducing all images to 4 × 4 binary arrays caused this stimulus to appear more random than it may appear at full resolution. Other deviations from model predictions include stimuli with symmetry, such as Image 83 (9th row, 3rd column) and Image 99 (10th row, 9th column). This gridlike design is an example of a pattern that is not probable among natural scenes but still appears highly structured. Such examples illustrate that properties other than those of natural scenes also give rise to perceived structure.

DISCUSSION

Accounts of subjective randomness that appeal to people's ability to detect regularities in stimuli need to be able to characterize those regularities and their origins. For example, if subjective randomness is viewed as the statistical evidence that a stimulus provides for having been produced from a random generating process, rather than from one with a more regular structure (Griffiths & Tenenbaum, 2003, 2004), we need to know what distribution over stimuli is induced by the more regular process, $P(X \mid \text{regular})$. Studying the subjective randomness of two-dimensional binary arrays provides one way to approach this problem, making it possible to explore the hypothesis

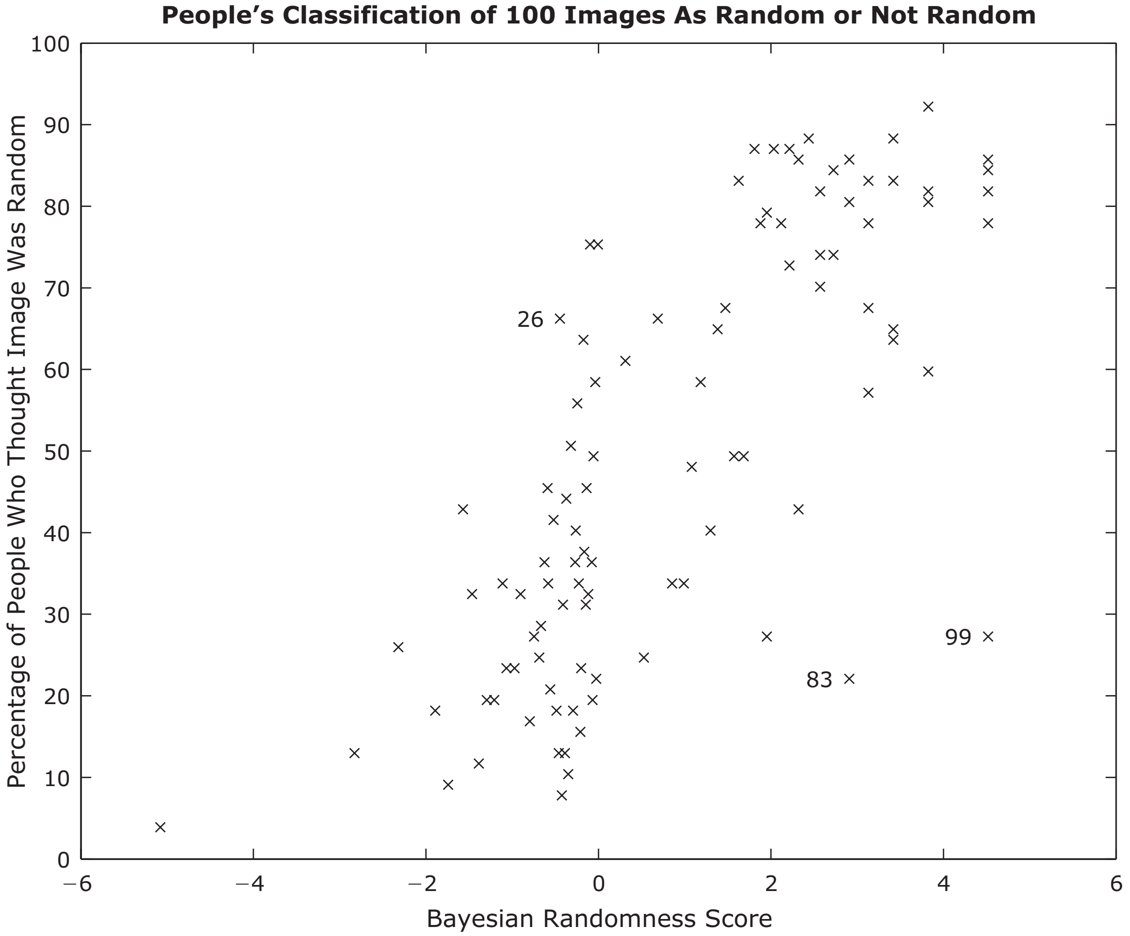

Figure 4. Results of the experiment. The proportion of the participants who classified each stimulus as having come from a random source as a function of random(X), computed using the distribution estimated from a nonparametric histogram of natural scenes as $P(X|\text{regular})$. Randomness judgments generally followed model predictions, with a few exceptions. Images 26, 83, and 99 (marked on the figure) are examples of patches that did not follow model predictions, for reasons that we discuss in the text.

that the regularities that people detect in these arrays and the resulting evaluation of their randomness are influenced by the statistical structure of our visual world. This statistical structure can be estimated from images of natural scenes, providing an objective method for estimating $P(X|\text{regular})$. In this study, we explored the consequences of estimating $P(X|\text{regular})$ for binary arrays extracted from natural scenes through an exhaustive enumeration of patterns that can appear in small arrays. We found that a Bayesian model of randomness perception using this distribution as the alternative to random generation provides a reasonably good model of human judgments, with the subjective randomness of a pattern decreasing as its probability of appearing in natural scenes increases.

These results demonstrate that it is possible to define good models of subjective randomness using objective sources of regularities—in our case, the statistics of natural scenes. However, there is room for future work toward a more complete analysis of the hypothesis that the statistics of these images can explain human randomness judgments. The method for estimating $P(X|\text{regular})$ we used in this study is suitable only for small binary arrays, thus capturing only limited aspects of the structure of natural scenes. Schreiber and Griffiths (2007) used a simple probabilistic model of images to explore randomness perception for larger binary arrays, but this analysis was limited to the extent to which neighboring cells shared the same value. With more sophisticated image models (e.g., Freeman, Pasztor, & Carmichael, 2000; Gimel'farb, 1996; Roth & Black, 2005; Zhu, Wu, & Mumford, 1998), we may be able to capture more of the nuances of subjective randomness judgments for larger binary arrays, enabling us to provide a more exhaustive exploration of how these judgments correspond to the statistics of our natural visual environment.

A deeper understanding of human randomness judgments may also be achieved by considering the structure of the human visual system in more detail. High-level properties of our stimuli, such as symmetry, seem important to human randomness judgments but are not captured by our models. Other phenomena, such as translation invariance, might fall out of the neural processes supporting vision.

For example, images could be initially passed through oriented filters, simulating receptive fields for V1 neurons. If this results in a significant transformation of the stimulus space, arrays that we considered distinct in our analysis could be considered the same by the visual system, and their probabilities should be combined accordingly. Other, more sophisticated models of visual processing could also be used, with the potential to further improve the fit between predicted and actual human randomness judgments and to help us understand how we come to form our judgments about visual structure.

AUTHOR NOTE

This work was supported by Grant FA9550-07-1-0351 from the Air Force Office of Scientific Research and Grant RES-000-22-3275 from the Economics and Social Research Council. The authors thank Evan Moss and Joseph Vuong for assistance in data collection. Correspondence concerning this article should be addressed to T. L. Griffiths, Department of Psychology, University of California, 3210 Tolman Hall #1650, Berkeley, CA 94720-1650 (e-mail: tom_griffiths@berkeley.edu).

REFERENCES

DOI, E., INUI, T., LEE, T. W., WACHTLER, T., & SEJNOWSKI, T. J. (2003). Spatiochromatic receptive field properties derived from information-theoretic analyses of cone mosaic responses to natural scenes. *Neural Computation*, **15**, 397-417.

FALK, R., & KONOLD, C. (1997). Making sense of randomness: Implicit encoding as a basis for judgement. *Psychological Review*, **104**, 301-318.

FELDMAN, J. (1996). Regularity vs genericity in the perception of collinearity. *Perception*, **25**, 335-342.

FELDMAN, J. (1997). Curvilinearity, covariance, and regularity in perceptual groups. *Vision Research*, **37**, 2835-2848.

FREEMAN, W., PASZTOR, E., & CARMICHAEL, O. (2000). Learning low-level vision. *International Journal of Computer Vision*, **40**, 24-47.

GIMEL'FARB, G. (1996). Texture modeling by multiple pairwise pixel interactions. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, **18**, 1110-1114.

GRIFFITHS, T. L., & TENENBAUM, J. B. (2001). Randomness and coincidences: Reconciling intuition and probability theory. In J. D. Moore & K. Stenning (Eds.), *Proceedings of the 23rd Annual Conference of the Cognitive Science Society* (pp. 370-375). Mahwah, NJ: Erlbaum.

GRIFFITHS, T. L., & TENENBAUM, J. B. (2003). Probability, algorithmic complexity, and subjective randomness. In R. Alterman & D. Kirsh (Eds.), *Proceedings of the 25th Annual Conference of the Cognitive Science Society* (pp. 480-485). Mahwah, NJ: Erlbaum.

GRIFFITHS, T. L., & TENENBAUM, J. B. (2004). From algorithmic to subjective randomness. In S. Thrun, L. K. Saul, & B. Schölkopf (Eds.), *Advances in neural information processing systems* (Vol. 16, pp. 953-960). Cambridge, MA: MIT Press.

OLSHAUSEN, B. A., & FIELD, D. J. (2000). Vision and the coding of natural images. *American Scientist*, **88**, 238-245.

ROTH, S., & BLACK, M. J. (2005). Fields of experts: A framework for learning image priors. In *Proceedings of the 2005 IEEE Computer Science Society Conference on Computer Vision and Pattern Recognition* (Vol. 2, pp. 860-867). Washington, DC: IEEE Computer Society.

SCHREIBER, E., & GRIFFITHS, T. L. (2007). Subjective randomness and natural scene statistics. In D. S. McNamara & J. G. Trafton (Eds.), *Proceedings of the 29th Annual Conference of the Cognitive Science Society* (pp. 1449-1454). Mahwah, NJ: Erlbaum.

SIMONCELLI, E. P., & OLSHAUSEN, B. A. (2001). Natural image statistics and neural representation. *Annual Review of Neuroscience*, **24**, 1193-1216.

ZHU, S., WU, Y., & MUMFORD, D. (1998). Filters, random fields and maximum entropy (frame): Towards a unified theory for texture modeling. *International Journal of Computer Vision*, **27**, 107-126.

(Manuscript received September 21, 2009;
revision accepted for publication March 9, 2010.)