

Subjective Randomness and Natural Scene Statistics

Ethan Schreiber (els@cs.brown.edu)

Department of Computer Science, 115 Waterman St
Providence, RI 02912 USA

Thomas L. Griffiths (tom_griffiths@berkeley.edu)

Department of Psychology, 3210 Tolman Hall #1650
Berkeley, CA 94720 USA

Abstract

We examine factors influencing whether people perceive two-dimensional binary arrays as random. Accounts of subjective randomness often appeal to the idea that people consider stimuli random when they cannot detect any regularities characterizing the structure of those stimuli. We explore the possibility that the kinds of regularities people can detect are shaped by the statistics of their natural environment. For two-dimensional binary arrays – grids of black and white pixels – we can measure these statistics directly from images of natural scenes. By estimating the parameters of a Markov random field from these images, we obtain a potential explanation for why people find arrays in which cells differ in value from their neighbors more random, with the reverse being more characteristic of structured images. We present an experiment in which we find a correlation between the extent to which people consider an array random, and the evidence it provides for having been generated by chance rather than nature.

Keywords: subjective randomness; Markov random fields; Boltzmann machine; Gibbs sampling; natural scenes

Introduction

People are very sensitive to deviations from their expectations about randomness. For example, the game Yahtzee involves repeatedly rolling five six-sided dice. If you were to roll all sixes six 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 non-randomness, as any other 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 that suggest a process other than chance might be at work (Griffiths & Tenenbaum, 2003, 2004; Falk & Konold, 1997). 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 explaining why people should be sensitive to a particular set of regularities. In the example given above, we might be able to appeal to the fact that systematic runs of the same number suggest loaded dice, or some other non-random 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). These results make 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 paper, 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, consisting of a grid of black and white pixels (see Figure 1), and 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 (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. We show that the statistics of images of natural scenes provide an explanation for why arrays in which neighboring cells differ in value should be considered more random, and describe an experiment illustrating that a Bayesian model of subjective randomness that evaluates whether a particular array is more likely to be generated by chance or be part of a natural scene provides good predictions of human randomness judgments.

The plan of the paper is as follows. The next section presents a formal definition of subjective randomness from the perspective of Bayesian statistics. This establishes the need to define a distribution over stimuli based on a “regular” process, which we approach by introducing a simple statistical model for images. We then describe how the parameters of this model can be estimated from natural scenes and outline the predictions that result from using this model as part of the Bayesian account of subjective randomness, before going on to present an experiment testing those predictions. We conclude by discussing the implications and limitations of this experiment.

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

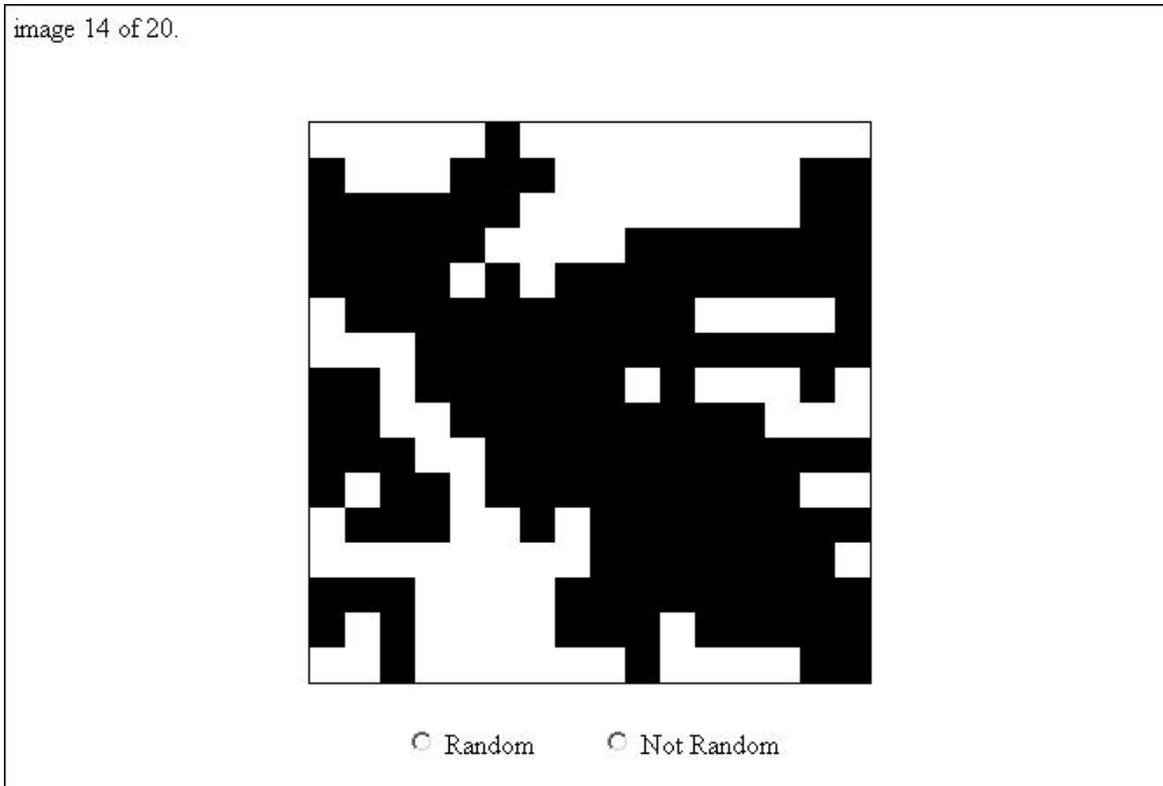


Figure 1: An example of a two-dimensional binary array. This is a screenshot from the experiment described in the main text, indicating how participants made their responses.

regular process (Griffiths & Tenenbaum, 2003, 2004). If we let $P(X|\text{random})$ denote the probability of X being generated by chance, and $P(X|\text{regular})$ be the probability of X under the regular generating process, then Bayes' rule gives the posterior odds in favor of random generation as

$$\frac{P(\text{random}|X)}{P(\text{regular}|X)} = \frac{P(X|\text{random}) P(\text{random})}{P(X|\text{regular}) P(\text{regular})} \quad (1)$$

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

$$\text{random}(X) = \log \frac{P(X|\text{random})}{P(X|\text{regular})} \quad (2)$$

where the logarithm simply scales the result.

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). Here, we consider how it might be applied to two-dimensional binary arrays, of the kind shown in Figure 1. A reasonable

choice of $P(X|\text{random})$ is to assume that each cell in the array takes on a value of 1 or 0 with equal probability, making $P(X|\text{random}) = 1/2^n$, where n is the number of cells in the array. However, defining $P(X|\text{regular})$ is more challenging. In the remainder of the paper, we develop and test a model in which $P(X|\text{regular})$ is chosen to reflect the statistics of natural scenes.

A statistical model for images

In order to define $P(X|\text{regular})$, we need to be able to specify a statistical model for two-dimensional binary arrays that captures the structure of natural images. A standard model for images is a Markov random field (MRF) (Geman & Geman, 1984; Pérez, 1998). We can view a digital image as a set of m pixel values and model each pixel value as a random variable X_i , where X is the set of possible values of each X_i . Using X to denote X_1, X_2, \dots, X_m , the joint distribution over all of the pixels is $P(X)$. To define this distribution, we lay a two-dimensional lattice over the pixels (as shown in Figure 2), and assume that the value of each pixel (a node in the lattice) is independent of the value of all other pixels when conditioned on its neighbors in the lattice.

If we constrain X to contain just two values – 0 and 1 – then we can draw on work on Ising models in statistical physics (Newman & Barkema, 1999) and Boltzmann machines in neural network research (Hinton & Sejnowski, 1986) to de-

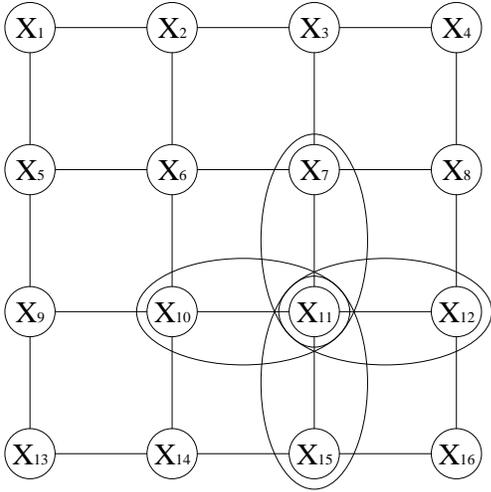


Figure 2: The dependency structure of a Markov random field for a distribution over images. Each variable corresponds to one pixel in the image, and the variables are connected in a lattice reflecting their spatial layout. Each variable is independent of all others when conditioned on its neighbors in the lattice. However, variables are dependent on their neighbors, making it possible to capture regularities involving relationships between neighboring pixels. The probability of a particular image depends on a set of functions defined over neighboring pairs of variables, illustrated with ellipses.

fine a distribution over images based on a global energy function $E(X)$. Specifically, we take

$$E(X) = \lambda_1 \sum_i \psi(X_i) + \lambda_2 \sum_{i < j} \psi(X_i, X_j) \quad (3)$$

where $\psi(X_i) = 2X_i - 1$, $\psi(X_i, X_j)$ is 1 if $X_i = X_j$ and -1 otherwise, and the final sum is constrained such that X_i and X_j are neighbors in the lattice. We can then define the distribution

$$P(X|\lambda_1, \lambda_2) = \frac{\exp\{-E(X)\}}{Z(\lambda_1, \lambda_2)} \quad (4)$$

where $Z(\lambda_1, \lambda_2) = \sum_X \exp\{-E(X)\}$ is the *partition function* of the distribution. This is a *pairwise* MRF, since the distribution on X depends only on functions computed over neighboring pairs of variables. The properties of the resulting distribution depend on the values of the parameters λ_1 and λ_2 . If $\lambda_1 > 0$, it will favor images in which more pixels take the value 0, and if $\lambda_2 > 0$, it will favor images in which pixels tend to differ from their neighbors.

Estimating the statistics of natural scenes

The statistical model introduced in the previous section provides a coarse-grained means of defining a distribution over two-dimensional binary arrays that will reflect the statistics of our visual world: take $P(X|\text{regular}) = P(X|\lambda_1, \lambda_2)$, and estimate λ_1 and λ_2 from images of natural scenes. However, in order to do this we need to define an algorithm for estimating

the parameters of this model. In this section, we outline such an algorithm and describe its application to images of natural scenes.

Finding λ_1 and λ_2

The simplest way to estimate λ_1 and λ_2 is to use the learning algorithm for Boltzmann machines (Hinton & Sejnowski, 1986), which is gradient descent on the log-likelihood of a collection of n images $\mathbf{X} = \{X^{(1)}, X^{(2)}, \dots, X^{(n)}\}$. If we assume that each of the images is generated independently from the distribution $P(X|\lambda_1, \lambda_2)$ defined by the MRF, we have the log-likelihood

$$\begin{aligned} \log P(\mathbf{X}|\lambda_1, \lambda_2) &= \sum_{i=1}^n \log P(X^{(i)}|\lambda_1, \lambda_2) \\ &= - \left(\lambda_1 \left[\sum_{i=1}^n (m_i^1 - m_i^0) \right] + \lambda_2 \left[\sum_{i=1}^n (m_i^- - m_i^\neq) \right] \right) - \sum_{i=1}^n \log Z(\lambda_1, \lambda_2) \\ &= - \left(\lambda_1 \left[\sum_{i=1}^n (m_i^1 - m_i^0) \right] + \lambda_2 \left[\sum_{i=1}^n (m_i^- - m_i^\neq) \right] \right) - n \log Z(\lambda_1, \lambda_2) \end{aligned}$$

where m_i^1 and m_i^0 are the number of pixels in image i taking values 1 and 0 respectively, and m_i^- and m_i^\neq are the number of neighboring pixels in image i of equal and different value respectively. Differentiating the log-likelihood with respect to λ_1 gives

$$\begin{aligned} \frac{\partial}{\partial \lambda_1} &= - \left(\sum_{i=1}^n (m_i^1 - m_i^0) \right) - n \cdot \frac{\frac{\partial Z(\lambda_1, \lambda_2)}{\partial \lambda_1}}{Z(\lambda_1, \lambda_2)} \\ &= - \left(\sum_{i=1}^n (m_i^1 - m_i^0) \right) + n \cdot E_{P(X|\lambda_1, \lambda_2)}[(m^1 - m^0)] \quad (5) \end{aligned}$$

where $E_{P(X)}[f(X)]$ is the expectation of the function $f(X)$ with respect to the distribution $P(X)$. The last line follows from the fact that

$$\begin{aligned} \frac{\frac{\partial Z(\lambda_1, \lambda_2)}{\partial \lambda_1}}{Z(\lambda_1, \lambda_2)} &= \frac{\frac{\partial}{\partial \lambda_1} \sum_X \exp\{-\lambda_1(m^1 - m^0) - \lambda_2(m^- - m^\neq)\}}{Z(\lambda_1, \lambda_2)} \\ &= - \sum_X (m^1 - m^0) \frac{\exp\{-E(X)\}}{Z(\lambda_1, \lambda_2)} \\ &= - \sum_X (m^1 - m^0) P(X|\lambda_1, \lambda_2). \quad (6) \end{aligned}$$

A similar calculation for λ_2 gives

$$\frac{\partial}{\partial \lambda_2} = - \left(\sum_{i=1}^n (m_i^- - m_i^\neq) \right) + n \cdot E_{P(X|\lambda_1, \lambda_2)}[(m^- - m^\neq)] \quad (7)$$

providing us with all we need in order to apply gradient descent.

The expectations in Equations 5 and 7 cannot be computed analytically, and thus need to be approximated. We follow the procedure used in (Hinton & Sejnowski, 1986), approximating these expectations using a Markov chain Monte Carlo algorithm (Newman & Barkema, 1999; Andrieu, Freitas, Doucet, & Jordan, 2003). The basic idea behind these algorithms is to repeatedly sample from a Markov chain until

it converges to its stationary distribution. If we construct the chain carefully, the stationary distribution will be the distribution from which we want to sample. Gibbs sampling (Geman & Geman, 1984) provides one way to construct such a chain. To perform one sweep of Gibbs sampling, we iterate over all pixels in the image, drawing that pixel from its conditional distribution given the current values of all other pixels. This procedure defines a Markov chain which will, after many sweeps, converge to the distribution $P(X|\lambda_1, \lambda_2)$, providing us with samples from that distribution. We simply take the average of the quantities that appear in the expectations over the samples produced by the Markov chain.

Fitting the model to image data

We estimated the parameters of a 16×16 pixel Markov random field model from a dataset of 62 images of natural scenes, each 1000×1280 pixels in size (Doi, Inui, Lee, Wachtler, & Sejnowski, 2003). From each of these images, we sampled 10 sub-images at each of 16×16 , 32×32 , 64×64 , 128×128 and 256×256 pixels in size uniformly at random. We then resized each of these images down to 16×16 pixels, giving us a set of images at different resolutions. Finally, the images were converted to binary by finding the median pixel value of all 3100 images, and mapping all pixel values less than the median to black (0) and all other pixel values to white (1). This procedure guaranteed that white and black pixels occurred equally often across the entire dataset, but preserved some of the variation in intensity across different images.

Gradient descent was used to find the values of λ_1 and λ_2 that maximized the log-likelihood $\log P(\mathbf{X}|\lambda_1, \lambda_2)$. Inspection of Equations 5 and 7 reveals that the derivative is zero if the expected number of white pixels and of matching neighbors under $P(X|\lambda_1, \lambda_2)$ equals the empirical average observed in the images. The gradient descent algorithm thus consisted of using Gibbs sampling to compute the expectations that appear in the gradients, and then modifying λ_1 and λ_2 in the direction indicated by those gradients. We ran the Gibbs sampler for 250 sweeps in each gradient descent step, and halted the algorithm after 100 iterations, when the log-likelihood had stabilized.

The estimates of λ_1 and λ_2 for the binary arrays produced from the natural scenes were -0.001 and -0.189 . The value of λ_1 is approximately 0 because we ensured an equal number of black and white pixels through the thresholding procedure used to construct the binary images, and thus simply indicates the validity of our method of estimating these parameters. The value of λ_2 reflects this propensity for neighboring pixels in our natural scenes to take on the same intensity value: in our training set of 3100 image samples, each pixel had an average of 2.99 neighbors with the same value and 1.01 neighbors with a differing pixel value.

Testing the model predictions

If we return to the definition of subjective randomness introduced in above, we can now examine the predictions that re-

sult from using our simple model of images of natural scenes as $P(X|\text{regular})$. Using the definition of randomness given in Equation 2 and taking $P(X|\text{random})$ to be the probability resulting from choosing pixel values with equal probability, we obtain

$$\text{random}(X) = \lambda_1 \sum_i \psi(X_i) + \lambda_2 \sum_{i < j} \psi(X_i, X_j) + C \quad (8)$$

where C is a constant that absorbs the log partition function and $\log P(X|\text{random})$. Thus, the randomness of an array depends on the number of white pixels and the number of matching neighbors, with the nature of this dependency being determined by the values of λ_1 and λ_2 .

Based on the estimates of λ_1 and λ_2 obtained from natural images, we should expect that people will find images with pixels that tend to be the same as their neighbors less random than images with pixels that tend to differ from their neighbors. This is consistent with previous subjective randomness experiments, as mentioned in the introduction, and potentially provides an explanation for why a match between the values of neighboring pixels should constitute a regularity: it is something that is more characteristic of structured natural scenes than of images produced by chance. Beyond these qualitative predictions, the model makes quantitative predictions about which images will be considered random. We conducted an experiment to test those predictions.

Method

Participants Participants were 25 members of the Brown university community.

Stimuli Participants were shown 16×16 pixel black and white binary arrays that were generated from one of two distributions: $P(X|\text{random})$ and $P(X|\text{regular})$. Two thousand images were generated from each distribution, with the images from $P(X|\text{regular})$ being produced by 300 sweeps of Gibbs sampling. In order to ensure a range of randomness values under the model, $\text{random}(X)$ was computed for each image using Equation 8. The 4000 images were sorted according to this score, and 20 images that evenly spanned the range of values of $\text{random}(X)$ were selected, together with another 5 practice images. These 25 images were used as stimuli in the experiment.

Procedure Participants were told that the stimuli were created using either a random process or another undefined process. 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.” They were instructed to press the button corresponding to their intuition about which process generated the stimulus. In order to familiarize participants with this procedure, the first 5 stimuli were presented as practice trials, followed by the 20 stimuli that constituted the main experiment.

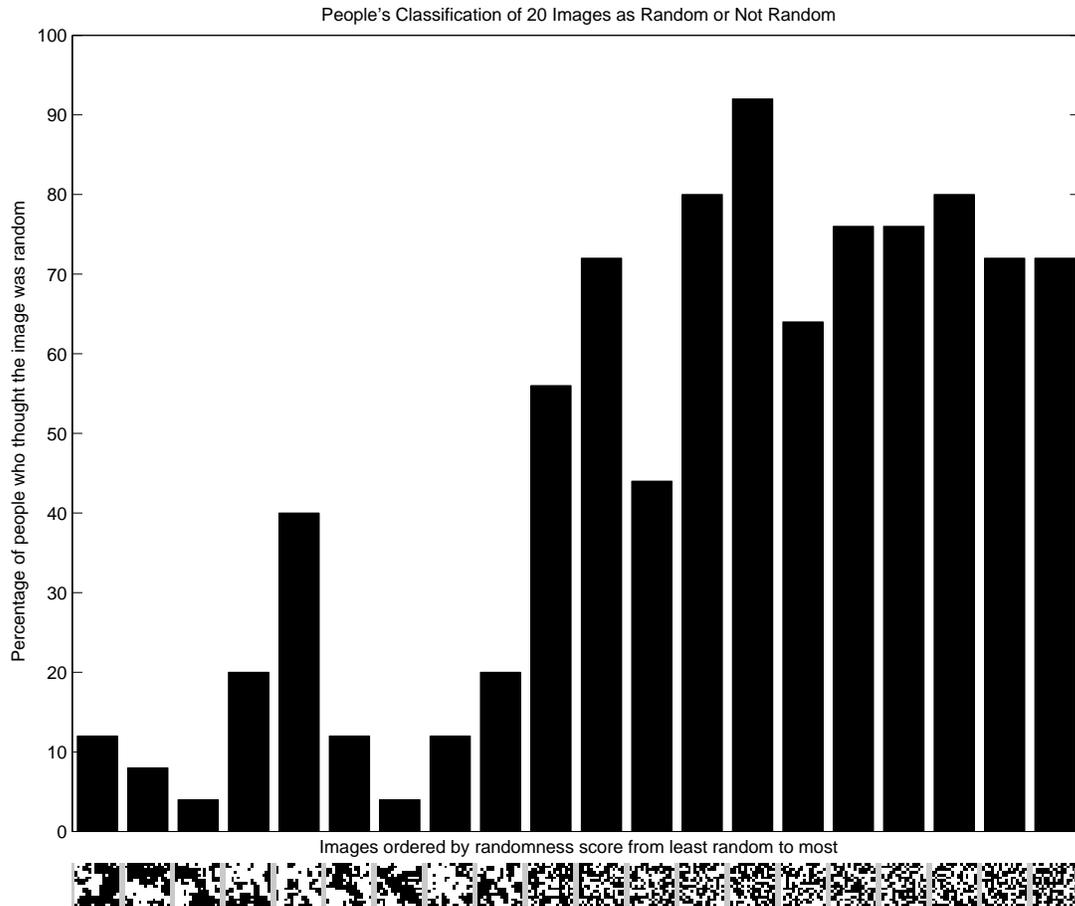


Figure 3: Results of the experiment. The actual stimuli used in the experiment appear along the bottom of the figure, ordered from lowest to highest value of $\text{random}(X)$, computed using the distribution estimated from natural scenes as $P(X|\text{regular})$. Bars show the proportion of participants who classified each stimulus as having come from a random source. There is a close correlation between the predictions of the Bayesian model and the subjective randomness of the stimuli.

Results

The results of the experiment are shown in Figure 3. A one-way ANOVA showed a statistically significant effect of image ($F(19, 456) = 15.18, MSE = 0.16, p < .001$). If we order the stimuli by $\text{random}(X)$, as was done in the figure, we can see the locus of this effect: as $\text{random}(X)$ increases, so does the percentage of participants who classified a stimulus as having been generated by the random process. The linear correlation between $\text{random}(X)$ and the probability of the stimulus being classified as random was $r = 0.87$, and the rank-order correlation (taking into account only the relative ordering of these different measures) was $\rho = .82$. Notably, the model has no free parameters: the randomness of a stimulus is determined solely by $P(X|\text{random})$, which is fixed a priori, and $P(X|\text{regular})$, which is estimated from the images of natural

scenes.

The results show a strong correspondence between a stimulus having higher probability under the distribution estimated from natural images than under the chance distribution and people deciding that the stimulus was not random. Very few people thought that the images that were more likely to be from the “regular” distribution were random. The images whose statistics were on the border between regular and random split people’s opinions. The stimuli that were least like the images were almost unanimously declared to be random. On the whole, these results bear out the predictions of the Bayesian model, and indicate that the distribution estimated from images of natural scenes provides a reasonable candidate for $P(X|\text{regular})$.

Summary and conclusions

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 one with more regular structure (Griffiths & Tenenbaum, 2003, 2004), we need to know what distribution over stimuli is induced by the more regular process, $P(X|\text{regular})$. Studying the subjective randomness of two-dimensional binary arrays provides one way to approach this problem, making it possible to explore the hypothesis that the regularities that people detect in these arrays, and the resulting evaluation of their randomness, is 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 paper, we explored the consequences of using a very simple image model to estimate $P(X|\text{regular})$: a pairwise Markov random field. This model captures an important aspect of images of natural scenes – that neighboring pixels are more likely to have the same value than would be expected by chance – which provides a potential explanation for the common finding that people give higher randomness ratings to arrays in which neighboring cells take different values (Falk & Konold, 1997). Our experiment indicates that the quantitative predictions of a model using this choice of $P(X|\text{regular})$ (and with no free parameters) correspond quite well to human randomness judgments.

These results illustrate 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, a more complete analysis of the hypothesis that these statistics can explain human randomness judgments would require using a more sophisticated model of images. The simple MRF that we used in this paper can only capture very limited aspects of the structure of natural scenes. In particular, with the current parameter values, it predicts that the scene with the lowest probability (and hence the one that is most random) is one in which the pixels alternate perfectly, as on a chess board. Intuitively, such a stimulus would seem to be quite non-random, suggesting that we might need to investigate more sophisticated image models (e.g., Freeman, Pasztor, & Carmichael, 2000; Gimel'farb, 1996; Roth & Black, 2005; Zhu, Wu, & Mumford, 1998) in order to be able to capture all of the nuances of subjective randomness judgments for two-dimensional binary arrays, and to provide a more exhaustive exploration of how these judgments correspond to the between natural statistics of our visual environment.

Acknowledgments

The authors thank John Donaldson for assistance in writing the web application used for the experiment.

References

- Andrieu, C., Freitas, N. de, Doucet, A., & Jordan, M. I. (2003). An introduction to MCMC for machine learning. *Machine Learning*, 50, 5-43.
- Doi, E., Inui, T., Lee, T., 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.
- Freeman, W., Pasztor, E., & Carmichael, O. (2000). Learning lowlevel vision. *International Journal of Computer Vision*, 40, 24-47.
- Geman, S., & Geman, D. (1984). Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 721-741.
- Gimel'farb, G. (1996). Texture modeling by multiple pairwise pixel interactions. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 18, 1110-1114.
- Griffiths, T. L., & Tenenbaum, J. B. (2001). Randomness and coincidences: Reconciling intuition and probability theory. In *Proceedings of the 23rd annual conference of the cognitive science society*.
- Griffiths, T. L., & Tenenbaum, J. B. (2003). Probability, algorithmic complexity, and subjective randomness. In *Proceedings of the 25th annual conference of the cognitive science society*.
- Griffiths, T. L., & Tenenbaum, J. B. (2004). From algorithmic to subjective randomness. In *Advances in neural information processing systems 16*.
- Hinton, G. E., & Sejnowski, T. J. (1986). Learning and relearning in Boltzmann machines. In *Parallel distributed processing: Explorations in the microstructure of cognition, volume 1: Foundations* (p. 282-317). Cambridge, MA: MIT Press.
- Newman, M. E. J., & Barkema, G. T. (1999). *Monte Carlo methods in statistical physics*. Oxford: Clarendon Press.
- Olshausen, B. A., & Field, D. J. (2000). Vision and the coding of natural images. *American Scientist*, 88, 238-245.
- Pérez, P. (1998). Markov random fields and images. *CWI Quarterly*, 11, 413-437.
- Roth, S., & Black, M. J. (2005). Fields of experts: A framework for learning image priors. In *Proceedings of the conference on computer vision and pattern recognition* (p. 860-867).
- 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.