

Running head: ABSTRACT STRUCTURE

Selectivity for Abstract Visual Structure in Ventral Temporal Cortex

John Myles White^{1,2}, Jiaying Zhao¹, Ray Lee², Nicholas B. Turk-Browne^{1,2}, & Daniel
Osherson^{1,2}

1. Department of Psychology, Princeton University, Princeton, NJ, USA

2. Princeton Neuroscience Institute, Princeton University, Princeton, NJ, USA

Corresponding Author:

John Myles White

Department of Psychology

Princeton University

Princeton, NJ 08540

jmw1729@princeton.edu

Number of figures and tables: 6

Number of pages: 14

Number of words: 228 (abstract), 3651 (main text)

Abstract

Uncovering and explaining the organization and selectivity of visual cortex is a chief goal of cognitive and systems neuroscience. In low-level visual cortex, this endeavor has been facilitated by the use of stimuli that vary along single feature dimensions (e.g., spatial frequency, orientation) that can be adjusted parametrically. In high-level visual cortex, the canonical stimuli used to study selectivity are complex objects (e.g., faces and scenes) that cannot easily be decomposed into basic dimensions. As a result, it has proven difficult to explain why a given region of high-level visual cortex is selective for one kind of object but not another. To help study the response properties of high-level visual cortex, we created a novel class of abstract stimuli that can be modulated parametrically like simple visual features, but which also possess more complex structure. We found that these stimuli, which did not afford object recognition or conceptual categorization, produced graded activity in ventral temporal cortex. Specifically, a region of the collateral sulcus responded more strongly as the proportion of random pixels in the stimuli decreased. We therefore identified a class of stimuli that evoke activity in high-level visual cortex, yet bear little surface resemblance to the typically preferred stimuli for these regions. Abstractly structured stimuli may ultimately help us characterize the selectivity of ventral temporal cortex and understand why certain regions are selective for particular object categories.

Introduction

Natural images form a small subset of the space of all possible images (Field, 1987). Evolutionary theory suggests that the visual system should therefore be organized to represent the visual environment in a reduced fashion for reasons of efficiency, particularly when a class of stimuli has ecological relevance to the organism. Empirical evidence supports this hypothesis: V1 and other regions in low-level visual cortex are responsive to image chunks that are well-suited to encoding natural images efficiently (Olshausen & Field, 1996). Moreover, cognitive neuroscientists have discovered regions of high-level visual cortex, most notably in the fusiform gyrus and parahippocampal cortex, which are specialized for the representation of specific classes of ecologically-relevant objects, namely faces and scenes respectively (Kanwisher et al., 1997; McCarthy et al., 1997; Epstein & Kanwisher, 1998). However, the question of why the fusiform gyrus is selective for faces and the parahippocampal cortex is selective for scenes (among other examples of category-selectivity in ventral temporal cortex) remains unanswered. Indeed, why do we not observe the reverse organization? We believe that answering such questions may be simplified by studying the basic response properties of high-level visual cortex in the absence of object recognition or conceptual categorization.

Understanding these response properties has been complicated by the sheer complexity and within-category variability of the stimuli used to elicit activity in regions like the fusiform face area (FFA) and parahippocampal place area (PPA). Studies of low-level visual cortex, in contrast, have benefited from the use of stimuli that can be manipulated along one feature dimension at a time, such as the angular

orientation or spatial frequency of a Gabor patch. The use of Gabor patches has allowed a rigorous characterization of the response properties of V1, including a sophisticated understanding of its local organization and connectivity with higher-level regions of the visual system (Jones & Palmer, 1987; Ringach, 2002). Similarly, the use of pure color lights has allowed the characterization of response properties of V4 (Schein & Desimone, 1990), and videos of moving dots have been used to explore the response properties of MT (Britten et al., 1992).

In contrast, exemplars from complex object categories, such as faces, bodies, word strings, and scenes, have been used to study high-level visual cortex (Kanwisher, 2010). These stimuli have proven more difficult to decompose into orthogonal dimensions because of their holistic and semantic nature. Indeed, the complexity of such images may be partially responsible for past and ongoing debates about the response properties of higher-level visual regions, including the debates about whether or not the PPA processes general contextual associations or places/scenes (Bar et al., 2008; Epstein & Ward, 2010), and whether or not the FFA is specialized for objects of expertise or for faces per se (Gauthier et al., 2000; McKone et al., 2007).

To make progress in understanding why high-level visual cortex is organized the way it is, here we develop a more general and basic language of visual structure, freed from the complexity, variability, and familiarity of real-world object categories. Specifically, we generated a novel stimulus class with varying levels of abstract structure embedded in a visual matrix. These stimuli are more complex than those used to study low-level vision, but are also simpler than the natural

images used previously to study high-level vision. Moreover, these stimuli do not allow for object recognition or conceptual categorization. Perhaps most importantly, the stimuli are defined parametrically, allowing for precise experimental control over the amount of structure in each image. In what follows, we introduce the stimuli and then describe the results of behavioral and neuroimaging studies concerning their perception.

Materials and Methods

The p -Randomness Algorithm

We designed a novel stimulus generation algorithm, which we call the p -Randomness algorithm. This algorithm generates bit matrices that can be rendered as images using any binary color scheme. We have used both a green/blue color scheme and a black/white color scheme in our work so far. The p -Randomness algorithm generates an image by traversing a matrix entry-by-entry, moving either horizontally or vertically at each step (in a constant direction for an entire matrix). At the first entry, a bit is generated using a random number generator. At each subsequent entry, the next bit is either generated by the same random number generator with probability p , or set to a value dependent on the previous bit with probability $1 - p$. In the latter case, the bit currently being set is either a repetition of the previous bit or an inversion of the previous bit (using a constant rule for an entire matrix). In this formulation, the p -Randomness algorithm has three free parameters: (1) p , the probability of generating the next bit entirely at random; (2) the traversal direction of the matrix, which can be either horizontal or vertical; and

(3) the function of the previous bit used to set the next bit, which can be either a repetition function or an inversion function.

Sample stimuli are shown in Figure 1. As can be seen, stimuli generated with $p = 1$ are pure noise images in which all bits are drawn fully at random. Because the bits are random, the other two parameters of the p -Randomness algorithm have no distinguishable effects on the resulting images. As p decreases, the stimuli take on the types of structure induced by the other two parameters: vertical repetition stimuli exhibit vertical contours, while horizontal repetition stimuli exhibit horizontal contours; and both the vertical and horizontal inversion stimuli exhibit a checkerboard-like structure. Though individual types of structure vary in spatial frequency and orientation, by pooling over matrices with varying parameters, we seek to localize regions responsive to the presence of structure *per se*, and not to low-level features of the stimuli, such as contours. This approach is vindicated by the results we present, in which we do not find significant activity in any of the low-level visual cortical areas that we would expect to see if such features were responsible for the activity seen in high-level visual cortex. In addition, the parameter p in our algorithm correlates with other natural measures of structure, such as the 2-D entropy of the resulting bit matrices and their compressibility using standard text compression tools (Figures 2A and 2B).

Behavioral Experiment

To confirm that subjects were sensitive to the types of structure generated by the p -Randomness algorithm, we used a perceptual discrimination task. Thirty

adult subjects (18 female, mean age = 20.2) viewed 100 60 x 60 matrices in which one half of the matrix (divided along either the horizontal or vertical axis) was generated using the p -Randomness algorithm with $p < 1$, while the other half of the matrix consisted of entirely random bits, which is equivalent to the output of the p -Randomness algorithm when $p = 1$. Subjects were instructed to determine for each image whether the boundary line dividing the halves was horizontal or vertical (Figures 3A and 3B), which required discriminating the random half from the non-random half. For this reason, reliable discrimination implied that the amount of structure at the p level of the non-random half could be registered perceptually. Each image was visible for 4s and no feedback regarding performance was given.

As the value of the p parameter approached 1, the task became increasingly difficult. At $p = 1$, the task was theoretically impossible to perform because the two halves were generated by the same process. As the value of p approached 0, the task became trivial to perform because the structured half contained no random bits. We therefore sought to determine the highest value of p at which the task could be performed with above-chance accuracy. We interpreted this level as the perceptual threshold for the structure induced by the p -Randomness algorithm. To increase the efficiency of our experimental design, we used a standard psychometric staircasing procedure. For every subject, the initial value of p was 0.7. This value was selected based on data from pilot studies indicating near-perfect performance by a majority of subjects up to this value. We then staircased p based on subjects' performance: p was decreased by 0.01 for every incorrect answer (which provides easier stimuli to discriminate); p was increased by 0.01 after two successive correct answers; and, p

was left unchanged otherwise. We computed separate p thresholds for the repeat versus inversion functions. Stimuli generated from these two classes were interleaved and counterbalanced across subjects.

Main Neuroimaging Experiment

We tested how the visual system processes the structure generated by the p -Randomness algorithm using fMRI. There are several possible outcomes: First, the level of abstraction that defines this structure may be outside the realm of what object-selective cortex can discriminate. If so, activity might be expected throughout or nowhere in object-selective cortex. Second, the classes of structure generated by the p -Randomness algorithm may overlap with the perceptual building blocks used by object-selective cortex to discriminate between stimuli. If so, activity might be expected in a subset of high-level visual regions.

To test these hypotheses, a total of 520 images were generated using the p -Randomness algorithm for 13 different levels of p , ranging from 0.52 to 1.00 in increments of 0.04. For each level of p , we generated ten images for each of the two traversal orientations (vertical, horizontal), and each of the two bit functions (repeat, invert). Thirteen adult subjects (5 male, 8 female) viewed 300 different stimuli from our pool of images over the course of 5 runs while EPI images were acquired using a Siemens 3T Allegra scanner. In our final analyses, data from one subject were excluded because we could not localize her PPA. An MPRAGE anatomical image was acquired for each subject using 1 x 1 x 1mm voxels and 160 slices. We acquired EPI images using 3.5 x 3.5 x 3.5 voxels and 34 unspaced slices

with interleaved acquisition. The EPI images were all AC-PC aligned. The TR was 2s, the TE was 20ms, the matrix was 64 x 64 and the flip angle was 90 degrees.

Each run lasted six minutes and consisted of 60 trials during which an exemplar of each of the 52 types of images (2 traversals x 2 functions x 13 p levels) was presented once along with eight randomly selected images during the final 8 trials of a run. All images appeared for 2000 ms and subtended 12 x 12 degrees of visual angle. The intertrial interval was jittered and sampled uniformly between 1, 3, and 5s to allow for statistical separation of the hemodynamic responses. A fixation crosshair blinked off and on during the ITI at 500 ms before the onset of the next stimulus.

To minimize conceptual categorization and idiosyncratic strategies related to p , subjects performed an orthogonal cover task during each run. Specifically, they detected a small gray target patch that appeared asynchronously on a subset of the images. The gray targets occurred roughly once in every eight trials and were uniformly distributed within each group of eight trials.

The EPI data were preprocessed in AFNI before statistical analysis (Cox, 1996). All images were corrected for slice timing and head movement and were aligned to the subject's anatomical data. The data were then mean-scaled within each voxel separately and an 8 mm FWHM kernel was used for spatial smoothing to improve signal-to-noise. The first four images from each EPI run were discarded to remove outliers caused by equilibration of the scanner's magnetic field.

For each subject, trials were sorted into 13 different regressors — one for each level of p — and convolved with AFNI's standard gamma HRF. The GLM also

included regressors of no interest for signal drift, head movement, target onset and subject response events. After fitting coefficients for all regressors, a linear contrast test was performed to identify voxels in which there was a monotonic trend in the BOLD response as a function of p . The reliability of the linear contrast values was tested at the group level (in Talairach space) with a one-sample t-test against 0 for all voxels. Clusters showing a significant effect were identified at uncorrected $p < 0.01$ with 115-voxel extent correction. This threshold resulted in a family-wise cluster-corrected error rate of $p < 0.05$. We used separate regressors for each level of p so that we could detect non-linear monotonic trends without fitting separate parametric regressors for each possible non-linear model. To verify our results, we also conducted all analyses using linear parametric regressors (Büchel et al., 1998).

Follow-Up Neuroimaging Experiment

In light of the results from the first neuroimaging experiment, we attempted to recruit our original subjects back for an additional FFA/PPA localizer task. Seven subjects (3 male, 4 female) returned to participate in this second task. We acquired new MPRAGE anatomical images for each subject. We then acquired two runs of EPI data using the same parameters as in the first experiment. A new localizer task was used, involving two runs, each of which contained 180 images of faces or scenes split up into blocks of 30 images of one category. Subjects were asked to determine whether each face was male or female during face blocks and whether each scene occurred indoors or outdoors during the scene blocks.

The resulting EPI scans were preprocessed before GLM analysis using the

same approach as in the first experiment. The data were analyzed in AFNI using a GLM model that included regressors of no interest for signal drift, head movement and subject response events, and two separate regressors for face and scene blocks, each of which was convolved with the canonical AFNI gamma HRF. The coefficients for the face and scene regressors were used in a linear contrast test (face vs. scene) and the results were tested at the group level with a one-sample t-test against 0 for all voxels. Clusters showing a significant effect were identified at uncorrected $p < .01$ with 20-voxel extent correction. The data from one subject did not allow us to identify the PPA and so they were excluded from our analyses. To assess the relationship of our original findings to the PPA, we recomputed our original p -level linear trend analysis from the first experiment using only the six subjects who had returned and successfully completed the PPA localizer task.

Results

Behavioral Experiment

For every subject, we identified the highest value of p at which the subject could reliably discriminate between the random half and the p -random half of the image. Specifically, at a given value of p we counted the number of correct trials at or above the value of p and the number of incorrect trials at or above the value of p . We then tested whether the number of correct trials was more than 50% of the sum of the number of correct and incorrect trials via a binomial test ($p < .05$). Across subjects, the highest value of p for the repetition function was 0.85 ($SD = 0.05$), and the highest value of p for the inversion function was 0.78 ($SD = 0.05$). The difference

between these levels was reliable ($t(29) = 5.42, p < .01$). These results show that subjects can detect abstract structure imposed by the p -Randomness algorithm and provide a behavioral estimate of the average subjective threshold for discriminating randomness. The distribution of behavioral thresholds is shown in Figure 4A.

Main Neuroimaging Experiment

All subjects performed the cover task with high accuracy (85% mean accuracy). No subjects reported awareness of our randomness manipulation during debriefing and all reported believing that the cover task was the primary experimental manipulation. For that reason, we conclude that any effect of p reflects task-irrelevant incidental processing of the abstract structure in the images.

The GLM analysis revealed three regions for which the amplitude of the BOLD response had an inverse monotonic relationship to p , i.e. greater responses as randomness decreased (Figures 5A and 5B). The two most prominent regions were in bilateral collateral sulcus (center-of-mass Talairach coordinates: left = -32, 59, -7; right = 32, 59, -7). One additional region was located in left posterior parietal cortex (-32, 37, 26). These results were verified using a parametric analysis in which the amplitude of the predicted BOLD response was scaled as a linear function of p .

To explore the generality of these effects, we conducted additional conjunction analyses in terms of the direction in which the matrices were traversed (horizontal, vertical) and the function used to set non-random bits (repetition, inversion). Specifically, we examined whether an effect was present for *both* horizontal and vertical traversal and for *both* repeat and invert functions by

assigning each voxel the minimum of the two contrasts in question. In these conjunction analyses, we observed bilateral clusters in the same regions as in our earlier analysis at $p < .01$ with a 20-voxel cluster correction. This weaker cluster correction was used due to the conservativeness of conjunction analyses. Overall, the direction of the effect, combined with its selectivity to a small number of regions and the symmetry of these regions across hemispheres, strongly suggests that parahippocampal cortex is involved in processing abstract structure.

To compare the sensitivity of these brain regions with the discrimination sensitivity from the behavioral experiment, we fit a threshold model of BOLD activity as a function of p in the region identified by our whole brain analysis. In this model, we assume that BOLD activity decreases linearly as a function of p until a threshold t is reached; after this point, BOLD activity is constant. To estimate the parameters of this model, we used MCMC methods (Gelfand & Smith, 1990; Robert & Casella, 1999; Plummer, 2003) to produce an approximation to the posterior distribution of the threshold, t , conditional on the observed BOLD signal. The posterior of the threshold for BOLD signal is shown in Figure 4B. As can be seen, this posterior has a mode near the peak of the distribution of behavioral thresholds.

Follow-up Neuroimaging Experiment

The results of the PPA localizer task are shown in Figure 6A. The center of mass of the PPA in our group of subjects was $(-27, 42, -7)$ for the right PPA and $(24, 43, -7)$ for the left PPA. Additionally, the results of recomputing our original GLM analysis on the six subjects who successfully completed the PPA localizer are shown

in Figure 6B. There was clear overlap between the group PPA and the clusters modulated by p . Further research will be needed to fully characterize this relationship, due to the small sample of subjects who could return for a second session and also to the post-hoc nature of this analysis. Nevertheless, we report these findings as preliminary seed evidence that the PPA may be sensitive to the abstract structure generated by the p -Randomness algorithm.

Discussion

We developed a novel class of visual stimuli with parameterized structure. These stimuli may serve as a useful extension to traditional stimuli like Gabor patches, which have been used to study coding in low-level visual cortex for orientation and spatial frequency. Our stimuli better match the complexity of stimuli traditionally used to study high-level visual cortex and, once pooled across traversal direction and previous bit function, cannot be reduced to simpler visual features. Indeed, the abstract structure generated by the p -Randomness algorithm selectively activated a bilateral region of ventral temporal cortex and not earlier areas.

The parametric nature of our stimuli makes it possible to study the response properties of high-level visual cortex in a graded fashion, and one can use this to determine the thresholds of perception for minimal structure. More generally, stimuli that can be varied parametrically are important tools when formulating quantitative models of the response properties of the visual system. The p parameter of the p -Randomness algorithm is closely related to the entropy of the resulting images, which provides another formalization of the continuum from noise

to structure. The close relationship with entropy may make it easier to employ information theoretic analyses of the response properties of high-level visual cortex.

In light of the selectivity of our results, the construction of other classes of abstract structure (i.e., using alternative algorithms) may help neuroscientists to understand the organization of object-selective cortex. For example, one could examine how primitive forms of structure that activate a given category-selective region relate to typical stimuli from that category. Moreover, such an endeavor might allow categorically-different stimuli (e.g., faces and scenes) to be evaluated and compared along a small number of rudimentary and continuous dimensions. Finally, in addition to improving our understanding of how complex features are coded, we may also be able to answer questions about why certain categories tend to co-localize, such as faces and bodies (e.g., Schwarzlose et al., 2005).

The abstractness of the stimuli means that our data do not depend on object recognition or conceptual categorization. This is the most unexpected aspect of our findings: images that are not *prima facie* ecologically-relevant objects elicit graded activity in parts of visual cortex that are strongly associated with object and scene recognition. The particular brain regions that correlated with p are located close to the PPA — a region generally linked to the processing of spatial layout, but the subject of some controversy (Bar et al., 2008; Epstein & Ward, 2010). The existing literature did not suggest that we would observe selective activity in the PPA in response to our stimuli. Future work is needed to characterize the relationship between our stimuli and typical place stimuli that elicit activity in the PPA.

The response properties of high-level visual areas are typically studied by

finding a preferred stimulus category, and comparing the response for exemplar objects from that category that vary in one aspect: for example, facial identity in the FFA (Rotshtein et al., 2005) or scene viewpoint in the PPA (Epstein et al., 2003). In contrast, we started bottom-up by generating stimuli that bear little resemblance to typical objects, but have been defined via algorithm. Our findings provide initial evidence that this may ultimately be a fruitful and complementary approach for understanding selectivity in the visual system.

References

- Bar, M., Aminoff, E., & Schacter, D. L. (2008). Scenes unseen: The parahippocampal cortex intrinsically subserves contextual associations, not scenes or places per se. *The Journal of Neuroscience*, *28*, 8539-8544.
- Britten, K. H., Shadlen, M. N., Newsome, W. T., & Movshon, J. A. (1992). The analysis of visual motion: A comparison of neuronal and psychophysical performance. *The Journal of Neuroscience*, *12*, 4745-4765.
- Büchel, C., Holmes, A. P., Rees, H. G., & Friston, K. J. (1998). Characterizing stimulus-response functions using nonlinear parametric regressors in parametric fMRI experiments. *NeuroImage*, *8*, 140-148.
- Cox, R. (1996). AFNI: Software for analysis and visualization of functional magnetic resonance neuroimages. *Computers and Biomedical Research*, *29*, 162-173.
- Downing, P. E., Chan, A. W. Y., Peelen, M. V., Dodds, C. M., & Kanwisher, N. (2006). Domain specificity in visual cortex. *Cerebral Cortex*, *16*, 1453-1461.
- Epstein, R., & Kanwisher, N. (1998). A cortical representation of the local visual environment. *Nature*, *392*, 598-601.
- Epstein, R., & Ward, E. J. (2010). How reliable are visual context effects in the parahippocampal place area? *Cerebral Cortex*, *20*, 294-303.
- Gauthier, I., Skudlarski, P., Gore, J. C., & Anderson, A. W. (2000). Expertise for cars and birds recruits brain areas involved in face recognition. *Nature Neuroscience*, *3*, 191-197.
- Gelfand, A. E., & Smith, F. M. (1990). Sampling-based approaches to calculating marginal densities, *Journal of the American Statistical Association*, *85*, 398-

409.

- Gonsalves, B. D., Kahn, I., Curran, T., Norman, K. A., & Wagner, A. D. (2005). Memory strength and repetition suppression: Multimodal imaging of medial temporal cortical contributions to recognition. *Neuron*, *57*, 751-761.
- Jones, J. P., & Palmer, L. A. (1987). An evaluation of the two-dimensional gabor filter model of simple receptive fields in cat striate cortex. *Journal of Neurophysiology*, *58*, 1233-1258.
- Kanwisher, N., McDermott, J., & Chun, M. M. (1997). The fusiform face area: A module in human extrastriate cortex specialized for face perception. *The Journal of Neuroscience*, *17*, 4302-4311.
- Kanwisher, N. (2010). Functional specificity in the human brain: A window into the functional architecture of the mind. *Proceedings of the National Academy of Sciences*, *107*, 11163-11170.
- MacKay, D. (2003). *Information Theory, Inference and Learning Algorithms*. Cambridge, UK: Cambridge University Press.
- McCarthy, G., Puce, A., Gore, J. C., & Allison, T. (1997). Face-specific processing in the human fusiform gyrus. *Journal of Cognitive Neuroscience*, *9*, 605-610.
- McKone, E., Kanwisher, N., & Duchaine, B. C. (2007). Can generic expertise explain special processing for faces. *Trends in Cognitive Science*, *11*, 8-15.
- Miyashita, Y. (1993). Inferior temporal cortex: Where visual perception meets memory. *Annual Reviews of Neuroscience*, *16*, 245-263.
- Olshausen, B. A., & Field, D. J. (1996). Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, *381*, 607-

609.

- Plummer, M. (2003). JAGS: A program for analysis of Bayesian graphical models using Gibbs sampling. In K. Hornik, & F. Leisch (Eds), *Proceedings of the 3rd International Workshop on Distributed Statistical Computing*, Vienna, Austria.
- Ringach, D. (2002). Spatial structure and symmetry of simple-cell receptive fields in macaque primary visual cortex. *Journal of Neurophysiology*, *88*, 455-463.
- Robert, C., & Casella, G. (1999). Monte Carlo Statistical Methods. New York: Springer-Verlag.
- Schein, S. J., & Desimone, R. (1990). Spectral properties of V4 neurons in macaques. *The Journal of Neuroscience*, *10*, 3369-3389.
- Schwarzlose, R. F., Baker, C. I., & Kanwisher, N. (2005). Separate face and body selectivity on the fusiform gyrus. *The Journal of Neuroscience*, *25*, 11055-11059.
- Simoncelli, E. P., & Olshausen, B. A. (2001). Natural image statistics and neural representation. *Annual Reviews of Neuroscience*, *24*, 1193-1216.
- Turk-Browne, N. B., Scholl, B. J., & Chun, M. M., & Johnson, M. K. (2009). Neural evidence of statistical learning: Efficient detection of visual regularities without awareness. *Journal of Cognitive Neuroscience*, *21*, 1934-1945.

Figures

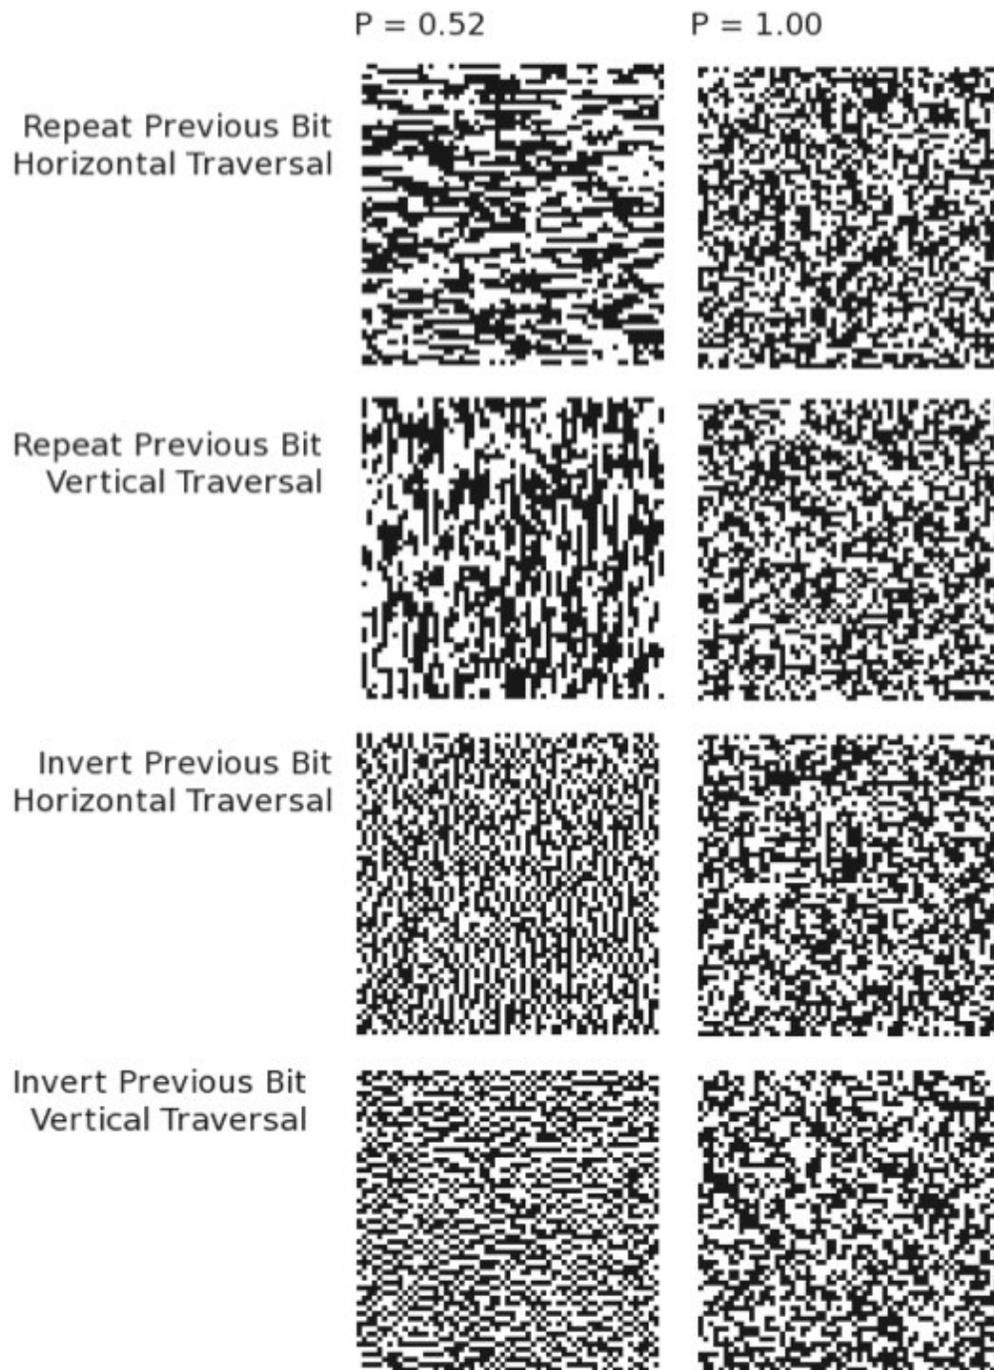


Figure 1. Sample matrices for the fMRI experiment. Each matrix is generated by the p -Randomness algorithm.

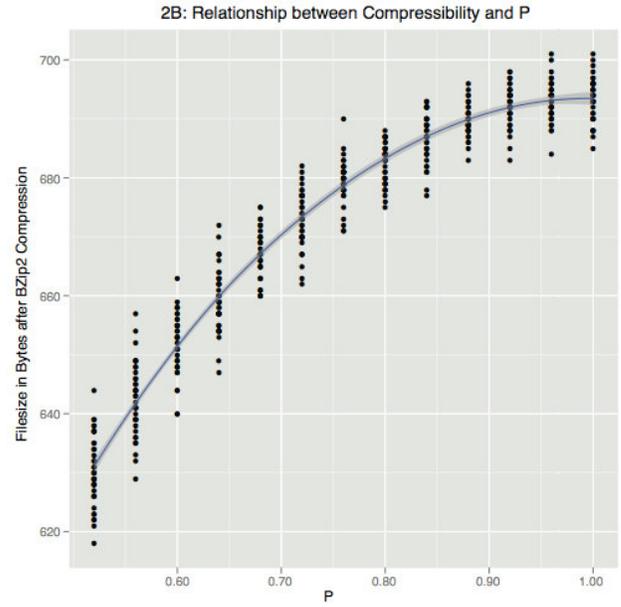
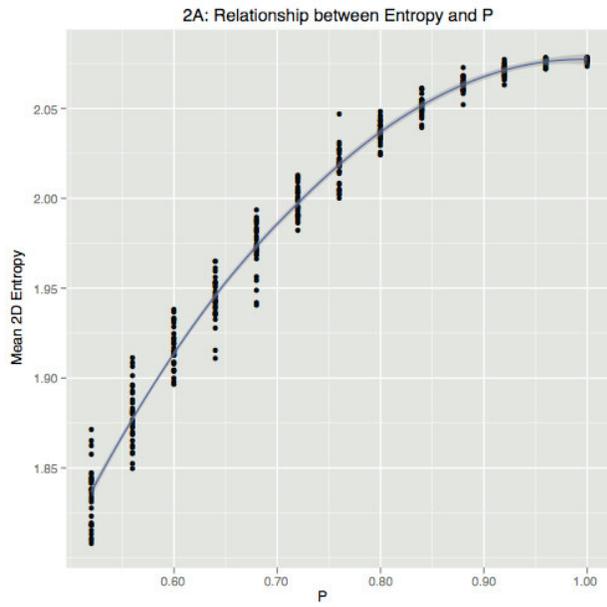


Figure 2. Correlations between p and 2D entropy (A) and between p and compressibility (B).

Figure 3A: Horizontally Divided Stimulus with $p = 0.60$

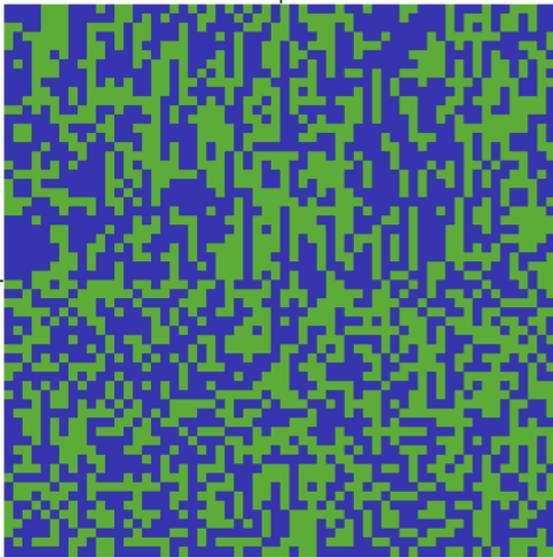


Figure 3B: Vertically Divided Stimulus with $p = 0.60$

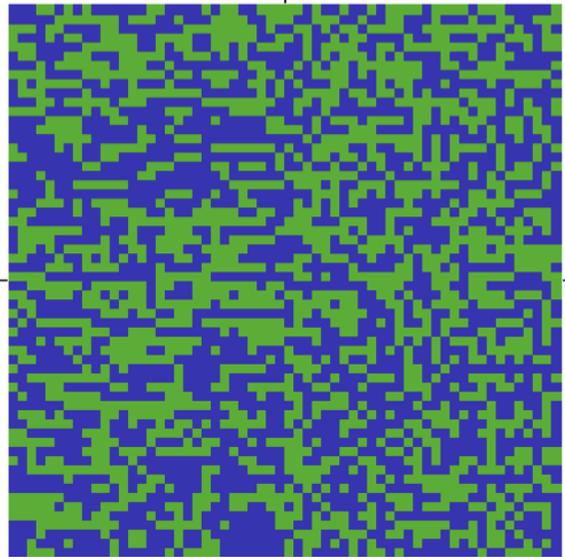


Figure 3. Sample matrices for the behavioral experiment. Each matrix consists of two halves, one half generated by the p -Randomness algorithm and the other half fully random. The task is to identify the boundary between the two halves.

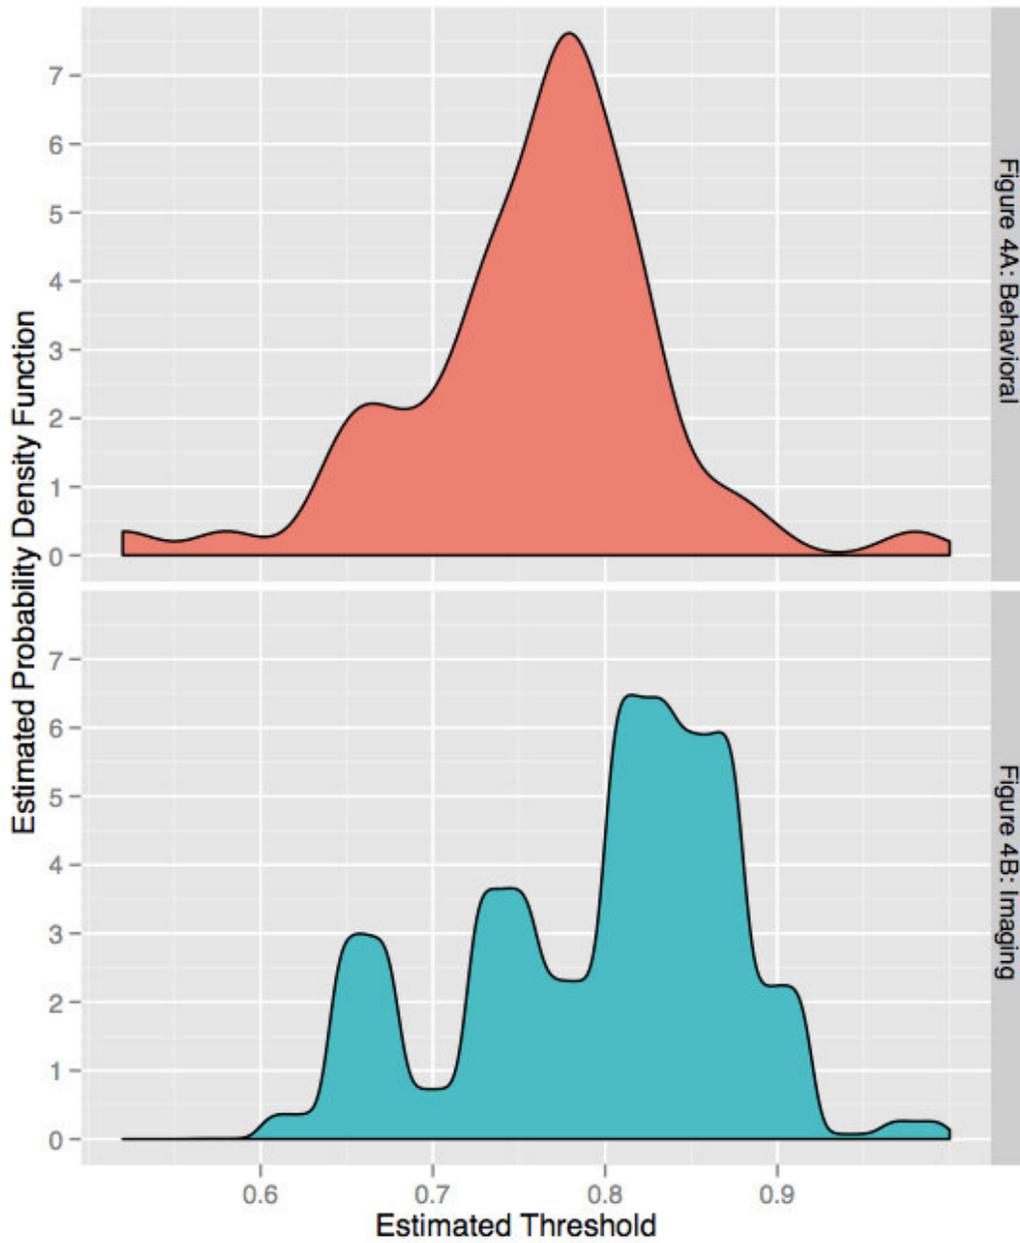


Figure 4. Perceptual thresholds estimated from behavioral data (A) and neuroimaging data (B).

Figure 5A: Axial Slice, Z = -7, $p < .01$, 115 Voxel Cluster Corrected



Figure 5B: Coronal Slice, Y = 46, $p < .01$, 115 Voxel Cluster Corrected

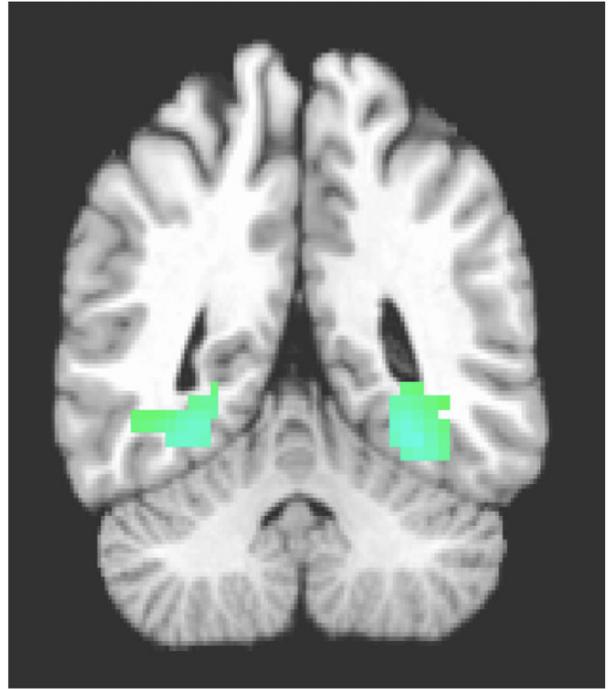


Figure 5. Regions that exhibit an inverse monotonic relationship to p (i.e., greater responses as randomness decreases)

Figure 6A: PPA

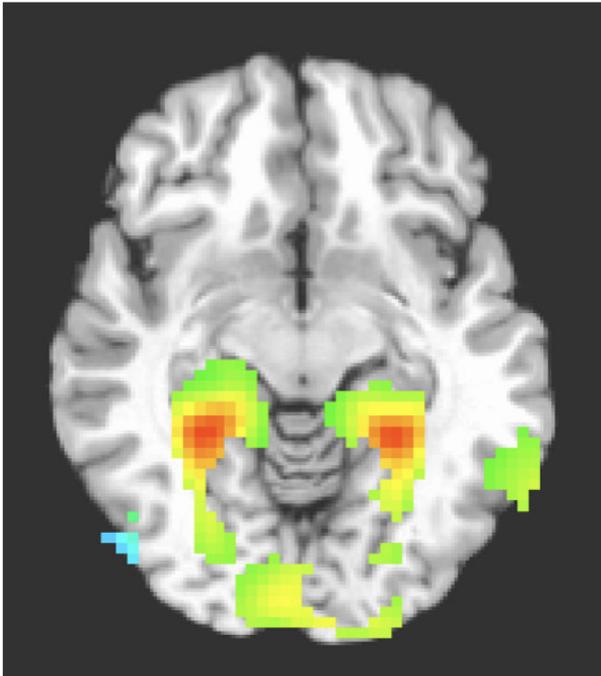


Figure 6B: Structure Sensitive Region

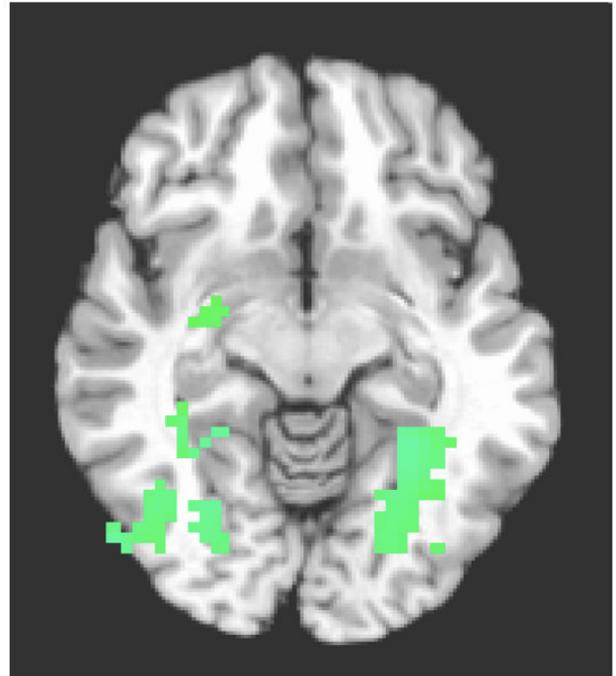


Figure 6. Comparison between PPA and regions that are responsive to structure.