

Multivariate Pattern Analysis of Functional Magnetic Resonance Imaging Data Reveals Deficits in Distributed Representations in Schizophrenia

Jong H. Yoon, Diana Tamir, Michael J. Minzenberg, J. Daniel Ragland, Stefan Ursu, and Cameron S. Carter

Background: Multivariate pattern analysis is an alternative method of analyzing functional magnetic resonance imaging (fMRI) data, which is capable of decoding distributed neural representations. We applied this method to test the hypothesis of the impairment in distributed representations in schizophrenia. We also compared the results of this method with traditional general linear model (GLM)-based univariate analysis.

Methods: Nineteen schizophrenia and 15 control subjects viewed two runs of stimuli—exemplars of faces, scenes, objects, and scrambled images. To verify engagement with stimuli, subjects completed a 1-back matching task. A multivoxel pattern classifier was trained to identify category-specific activity patterns on one run of fMRI data. Classification testing was conducted on the remaining run. Correlation of voxelwise activity across runs evaluated variance over time in activity patterns.

Results: Patients performed the task less accurately. This group difference was reflected in the pattern analysis results with diminished classification accuracy in patients compared with control subjects, 59% and 72%, respectively. In contrast, there was no group difference in GLM-based univariate measures. In both groups, classification accuracy was significantly correlated with behavioral measures. Both groups showed highly significant correlation between interrater correlations and classification accuracy.

Conclusions: Distributed representations of visual objects are impaired in schizophrenia. This impairment is correlated with diminished task performance, suggesting that decreased integrity of cortical activity patterns is reflected in impaired behavior. Comparisons with univariate results suggest greater sensitivity of pattern analysis in detecting group differences in neural activity and reduced likelihood of nonspecific factors driving these results.

Key Words: Distributed representation, fMRI, multivariate pattern analysis, schizophrenia, visual association cortex, visual processing

Modern theories on the nature of neural representation have emphasized a connectionist view in which representations are emergent properties of coordinated and distributed neural activity (1,2). Influenced by these insights and the availability of methods capable of measuring combinatorial patterns of activity, we examined the proposition that schizophrenia involves impairment in the stable and reliable activation of distributed representations.

Traditionally, functional neuroimaging studies have relied on univariate analysis, which examines each voxel in isolation and is primarily designed to identify regions showing differential activity between conditions and groups. These studies have been invaluable in identifying and focusing attention on specific brain regions in the pathophysiology of schizophrenia (3–6). However, univariate methods provide limited information regarding the distributed nature of neural processing. While there is a growing number of studies in schizophrenia that have highlighted the relevance of large-scale connectivity across distant regions (7–10), there have been few studies on another potentially important aspect of distributed processing, namely the coding of information in distributed patterns of activity within a local region.

From the Department of Psychiatry and Imaging Research Center, University of California Davis, Sacramento, California.

Address reprint requests to Jong H. Yoon, M.D., Ph.D., University of California Davis Imaging Research Center, 4701 X Street, Sacramento, CA 95817; E-mail: jhyoon@ucdavis.edu.

Received December 19, 2007; revised July 7, 2008; accepted July 24, 2008.

0006-3223/08/\$34.00
doi:10.1016/j.biopsych.2008.07.025

Basic cognitive neuroscience has begun to develop approaches capable of assessing, in a combinatorial or multivariate manner, the distribution of voxels providing optimal differentiation of distinct brain states. In a series of studies, Haxby *et al.* (11) and O'Toole *et al.* (12) demonstrated that the representation of specific categories of visual objects are coded by category-specific patterns of activity within the visual association cortex. Recently, it has been recognized that methods for automated recognition of spatial patterns, which have been widely employed in other fields, e.g., biometrics, can be fruitfully applied to functional neuroimaging (13,14). In studies with healthy subjects, it has been shown that pattern analysis algorithms can accurately identify activity patterns associated with specific perceptual states (11,12,15–17). The viewing of different categories of stimuli is represented as distinct patterns of activity in the visual cortex and, conversely, these patterns can be decoded to identify the category of objects being viewed. The development of these methods has implications for schizophrenia, as they may be more amenable than traditional methods to addressing hypotheses based on a distributed connectionist model of brain functioning (13,14).

In this study, we applied a multivariate pattern analysis algorithm to test the hypothesis that distributed perceptual representations are impaired in schizophrenia. We studied a cohort of subjects with schizophrenia and a healthy control group while they viewed exemplars from diverse categories of visual stimuli (Figure 1). A multivoxel pattern analysis program was trained to identify category-specific functional magnetic resonance imaging (fMRI) activity patterns and then to classify an independent set of fMRI data. The main dependent measure was the program's accuracy in this decoding. High classification accuracy denotes the presence of category-specific canonical

BIOL PSYCHIATRY 2008;64:1035–1041
© 2008 Society of Biological Psychiatry

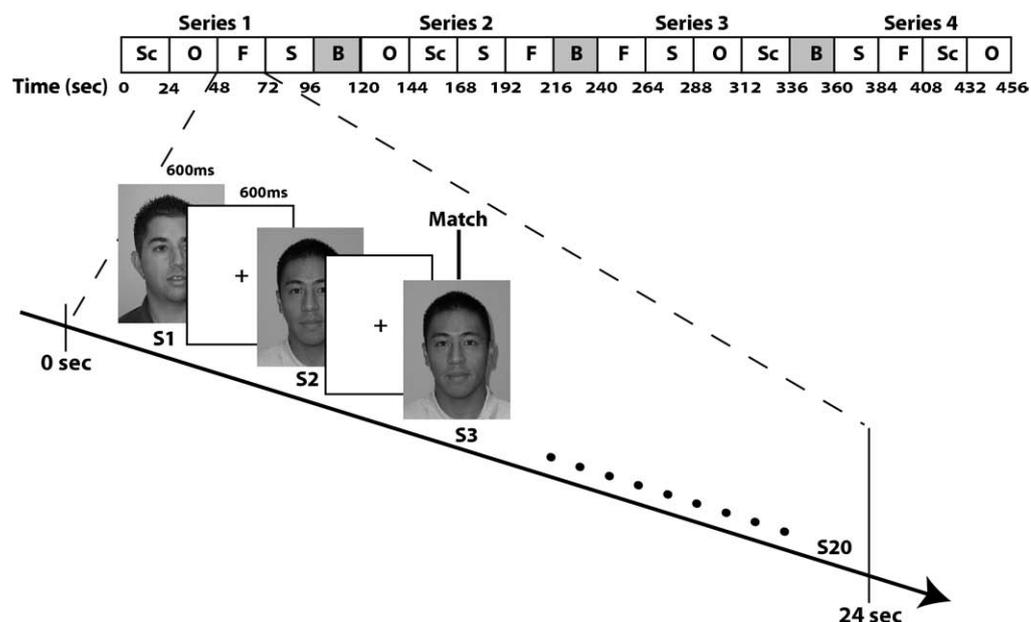


Figure 1. Visual processing task. The subjects were shown a series of exemplars of four categories of visual objects—faces (F), scenes (S), everyday objects (O), and scrambled images of everyday objects (Sc). Within each category block, 20 exemplars were shown. To verify task engagement, the subjects were required to conduct a 1-back match. A fixation baseline (B) was displayed between series of category blocks.

pattern of activity and the stable engagement of this pattern. We predicted high classification accuracy for control subjects. For patients, we predicted lower classification accuracy, reflecting impaired coordinated neural activity supporting distributed representations. We assessed the stability of the category-specific patterns by quantifying the voxelwise correlation in activity over time and predicted that patients would show diminished interrun correlations. And finally, we conducted a parallel set of univariate general linear model (GLM)-based analyses to assess the presence of nonspecific activation differences between groups and to provide insights into differential sensitivity between the two methods in detecting group differences.

Methods and Materials

Subjects

Nineteen individuals with schizophrenia (SZ) and 15 healthy control subjects (C) were studied. Demographic and clinical data are displayed in Table 1. Data from a subset of these subjects have been published in a separate study (18). Patients were clinically stable and outpatients. Diagnostic status was evaluated with the Structured Clinical Interview for DSM-IV-Text Revision (SCID-I) conducted by masters- or doctoral-level clinicians and confirmed by consensus conference. Symptoms were quantified with the Brief Psychiatric Rating Scale (BPRS), Scale for the Assessment of Negative Symptoms (SANS), Scale for the Assessment of Positive Symptoms (SAPS), Strauss Carpenter Outcome Scale (SCOS), and Global Assessment Scale (GAS). All patients were taking antipsychotics with all but one taking atypical neuroleptics. Exclusion criteria for all were IQ < 70, drug or alcohol dependence or abuse within 3 months of testing, major medical or neurological illness, significant head trauma, or contraindication to magnetic resonance imaging (MRI). Exclusion criteria for control subjects were lifetime diagnosis of Axis I disorder or first-degree relative with a psychotic disorder. Groups were well matched except in years of education, 13.0 ± 1.9 versus 15.9 ± 2.6 , SZ and C, respectively, $p < .05$; and IQ, 95.5 ± 13.7

versus 110.3 ± 6.7 , SZ and C, respectively, $p < .05$. After description of the study, written informed consent was obtained from all subjects. This study was approved by the Institutional Review Board (IRB) at the University of California Davis.

Activation Paradigm

Stimuli were presented and responses recorded with E-Prime (Psychology Software Tools, Inc., Pittsburgh, Pennsylvania; <http://www.pstnet.com>). Within each run, there were four series of stimulus blocks, and in each series, there were four blocks of stimuli (Figure 1). Within a block, 20 exemplars of faces, scenes,

Table 1. Demographics

	Patients (n = 19)		Control Subjects (n = 15)		p <.10
	Mean	SD	Mean	SD	
Age (years)	33.4	11.1	28.2	7.4	<.10
Gender (% male)	68		53		<.10
Handedness (% right)	95		100		<.10
WRAT	95.5	13.7	110.3	6.8	<.05 ^a
Education (years)	13.0	1.9	15.9	2.6	<.05 ^a
Parental Education (years)	12.9	3.0	14.9	3.6	<.10
GAS	59.7	9.1			
SCOS	9.9	3.7			
SANS	1.53	.81			
SAPS	1.56	1.04			
BPRS	44.5	11.2			
Mediations (atypical/ typical)	18/1				

BPRS, Brief Psychiatric Rating Scale; GAS, Global Assessment Scale; SANS, Scale for the Assessment of Negative Symptoms; SAPS, Scale for the Assessment of Positive Symptoms; SCOS, Carpenter Outcome Scale; WRAT, Wide Range Achievement Test.

^aStatistically significant difference.

everyday objects, and scrambled images of everyday objects were displayed. The series were interleaved with a baseline period consisting of a crosshair displayed on a blank background for 24 seconds. Each subject underwent two runs of this paradigm. To obtain online verification that subjects were attending to the stimuli, subjects were required to press two fingers on the response pad when the current image was the same as the image immediately preceding it.

MRI Data Acquisition and Preprocessing

Images were acquired with a GE Signa 1.5T MRI system (General Electric, Milwaukee, Wisconsin) using single-shot T2*-weighted echo-planar imaging (EPI) in the anterior commissure-posterior commissure (AC-PC) aligned axial plane with repetition time (TR) of 2.0 seconds, an echo time (TE) of 32 milliseconds, and a flip angle of 90°. Twenty-seven interleaved slices for whole-brain coverage were collected with a 22-cm field of view, slice thickness of 4.0 mm and a gap of .5 mm, and in-plane resolution of 3.5 × 3.5 mm. Functional MRI processing was conducted using SPM2 (Functional Imaging Laboratory, University College London, United Kingdom, <http://www.fil.ion.ucl.ac.uk/spm>) and included slice time correction and six-parameter rigid body spatial realignment. All subjects displayed less than 4 mm of movements and the groups did not differ either in cumulative [$F(6,27) = 1.36, p = .265$] or mean movements [$F(6,27) = 1.37, p = .262$]. To better preserve the spatial configuration of activations in individual subjects, images were not smoothed or spatially normalized into a common space. Coplanar, T1-weighted scans with .88 × .88 mm in-plane resolution were acquired just prior to functional images.

Behavioral Data Analysis

The “hit rate” performance (number of correct responses during the presentation of a 1-back divided by the total number of 1-back matches) and reaction time (RT) for hits were used to evaluate performance.

Multivariate Pattern Analysis

Data analysis was conducted using the Princeton Multi-Voxel Pattern Analysis (MVPA) toolbox (Princeton University, Princeton, New Jersey, <http://www.csmb.princeton.edu/mvpa>) following the methods described in Polyn *et al.* (17). Details of the method can be found on the MVPA website and in the supplement to the Polyn *et al.* (17) paper. In addition, two recent reviews on this method have recently been published (13,14). The analysis involved three main stages: feature selection and training and testing of the automated classifier. Prior to these steps, the images underwent voxelwise z-scoring of intensity values to account for baseline drift and shifting condition regressors by 6 seconds to account for hemodynamic delay in the blood oxygenation level-dependent (BOLD) signal.

Feature Selection. An anatomic mask was used to constrain the region in which the multivariate and univariate analyses were conducted. This mask, defined by the union of the parahippocampal, lingual, and fusiform gyri, limited the analysis to visual regions most active during visual object processing (18,19). The boundaries of these regions were defined according to Duvernoy (20) and Duvernoy and Bourgouin (21). The anatomical masks were drawn by a single rater blind to diagnosis on T1 coplanar images acquired right before EPI scans for each subject. We conducted test-retest reliability analysis on 10 randomly selected subjects. The Spearman-Brown coefficient was .90. Within this region, the feature space was further restricted by

eliminating voxels with low task-related activity assessed by a voxelwise analysis of variance (ANOVA) of BOLD intensity across all conditions (face, scene, objects, scrambled images, and baseline). These less informative voxels serve as sources of noise and thus impair detection of classifiable patterns. Those voxels active at a threshold of $p < .001$, uncorrected, were retained for inclusion in the subsequent classification analysis. This area contained an average of 346 voxels per subject (SD = 127).

Training and Testing. These stages were carried out using a linear pattern classifier following a full cross-validation design: data from run 1 were used to feature select and train the classifier, which was then tested on data from run 2, and then run 2 data were used to feature select and train the classifier, which then tested data from run 1. We employed a two-level (input and output) backpropagation neural network (NN) classifier. The input layer consisted of one unit for each voxel in the feature-selected image. The output layer consisted of units for each of the five stimulus categories.

During training, data from one run, with each TR labeled by condition, were fed into the NN. Using a conjugate gradient variant of the backpropagation classifier, the weights between input and output layers were adjusted using a minimum square error algorithm between target categorization and actual categorization. In this way, over many iterations of weight adjustment, the classifier “learned” the pattern of activation of each stimulus category. These weighted connections were then applied during testing to data from the remaining run. Each input of a TR’s activation pattern resulted in an output layer activation pattern, with the maximally active output unit being recorded as the classifier’s guess of the condition of that TR. Accuracy of the classifier’s guesses during this testing stage (1 for correct, 0 for incorrect) was averaged across all TRs of each stimulus type.

Interrun Correlations of Activity. A secondary analysis, what we refer to as the interrun correlations of category-specific activity patterns, evaluated the stability of activity for each stimulus type across the two runs. Activity patterns for each category were represented by a voxelwise statistical map of p values generated from an ANOVA of the BOLD signal. The ANOVA contrasted each scan’s signal during the condition of interest to all other time points. The ANOVA was restricted to the mask of the visual regions described above. Activity of each voxel was represented by its p value in a vector with a length defined by the number of voxels in the anatomic mask. A vector of this form was generated for each run and stimulus type. Bivariate Pearson product-moment correlations were then calculated across the stimulus-specific maps for the two runs. This correlation value represented the across-run functional stability, quantified for each stimulus type.

Univariate Analysis

We conducted a parallel set of analyses using univariate GLM-based methods so that we could compare these results with those from the multivariate pattern analysis. We chose to analyze faces because the neural correlate of face processing has been an increasingly common area of research in schizophrenia, and our prior work (18), which showed relatively normal univariate measures of fusiform face area (FFA) activity, suggested that comparisons between multivariate and univariate analyses of face processing may provide a good test of the former’s sensitivity. We analyzed the same images preprocessed for multivariate analysis, meaning after slice time correction and realignment, images were not spatially smoothed or transformed into a standard space. Instead, analysis proceeded in the subject’s

Table 2. Behavioral Results

	Accuracy—% (SD)			RT—msec (SD)		
	Control Subjects	SZ	<i>p</i>	Control Subjects	SZ	<i>p</i>
Face	96 (6)	86 (17)	.02 ^a	525 (109)	601 (137)	.08
Scene	93 (10)	78 (21)	.01 ^a	569 (109)	602 (74)	.32
Object	94 (6)	89 (19)	.26	567 (109)	616 (131)	.25
Scrambled	86 (11)	72 (27)	.05 ^a	601 (95)	651 (198)	.34
Total	93 (6)	81 (19)	<.01 ^a	565 (96)	617 (117)	.17

RT, reaction time; SZ, schizophrenia.

^aStatistically significant difference.

“native space” (18,19). Univariate analysis was conducted in VoxBo (University of Pennsylvania, Philadelphia, Pennsylvania, www.voxbo.org). The convolution matrix included a time-domain representation of the 1/f power structure (22,23), a low-pass filter to remove frequencies above .45 Hz, and nuisance covariates to model an intercept and global signal orthogonal to the covariates of interest. Activity was estimated using a four-covariate model, a covariate for each of the stimulus categories, convolved with a canonical hemodynamic response function. We performed linear contrasts of face-scenes and face-objects to identify face-specific voxels within the anatomically defined region described above, thresholded at a *t* value corresponding to $p < .05$, Bonferroni corrected for multiple comparisons (24). The dependent measures were the maximum and mean parameter estimates (i.e., beta values) and volume of the FFA (number of above-threshold voxels). These measures were first calculated for each subject and then they were group averaged for a random effects between-group analysis with two-sample *t* tests (18,19).

Results

Both groups performed the task well, but control subjects showed significantly higher accuracy than patients, 93% versus 81% ($t = -5.5$, $p < .001$). For all categories, control subject performance was higher, and for all but the everyday objects, the differences were significant. Control subjects responded faster, but RT differences were not significant ($M_C = 565$ msec, $SD_C = 96$ msec, $M_{SZ} = 617$ msec, $SD_{SZ} = 117$ msec, $p = .17$). For faces, there was a nonsignificant difference in reaction time ($t = 1.8$, $p = .08$). These and more detailed results are displayed in Table 2.

Accuracy of the pattern classifier was significantly different between groups (Figure 2). Overall, activity patterns in control subjects were more accurately classified than in patients, with averages of 71.7% and 58.6%, respectively. For all but one category (everyday objects), accuracy was significantly higher for

control subjects ($t > 2.1$, $p < .05$). An ANOVA showed a main effect of group ($F = 12.2$, $p < .001$). We reanalyzed the data with the baseline condition excluded and these results were virtually identical to the above results.

Results in Figure 3 show a significant correlation between classification accuracy across all conditions and behavioral performance. For patients, behavioral accuracy ($r[n = 19] = .48$, $r^2 = .23$, $p < .05$) and RT ($r[n = 19] = -.48$, $r^2 = .23$, $p < .05$) were significantly correlated. For control subjects, RT correlated significantly with classification accuracy ($r[n = 15] = -.52$, $r^2 = .27$, $p < .01$) but behavioral accuracy did not, possibly due to ceiling effects. The correlation between behavioral performance and pattern classification accuracy remained significant after controlling for possible confounds of age, subject and parental education, and Wide Range Achievement Test (WRAT) scores: classification accuracy and RT [$df = 25$] = $-.526$, $r^2 = .28$, $p < .05$]; classification accuracy and behavioral accuracy [$df = 25$] = $.505$, $r^2 = .26$, $p < .05$]. The pattern classification results did not correlate with any of the clinical measures.

To explore the possibility that increased temporal variability in activity contributed to impaired classification in patients, we conducted a secondary analysis measuring the correlation of voxelwise activity across runs (Figure 4). Results showed that interrater correlations were lower for patients for all stimulus categories, except objects. This difference was significant for faces ($t = -2.4$, $p < .05$) and scenes ($t = -2.3$, $p < .05$) but nonsignificant with baseline ($t = -2.0$, $p = .057$). An ANOVA of the interrater correlations showed a nonsignificant effect of group ($F = 3.7$, $p = .065$). There was very high correlation between interrater stability of engagement of activity and classification accuracy in both groups, patients ($r[n = 19] = .93$, $r^2 = .86$, $p < .01$) and control subjects ($r[n = 15] = .89$, $r^2 = .79$, $p < .01$).

A univariate, GLM-based analysis of the same data was conducted to examine the impact of nonspecific factors in the multivariate results and to allow comparisons between these methods. We focused on faces because we predicted, based on prior results (18), that traditional univariate measures would not yield significant differences between groups, thereby offering a point of contrast between the univariate and pattern classification approaches. In the contrasts of face versus nonface blocks, e.g., face-scenes and face-objects, there was no group difference in maximum beta [$F(1,32) = .77$, $p = .39$], mean beta [$F(1,32) = 1.4$, $p = .224$], or voxel count [$F(1,32) = 1.2$, $p = .25$] (Table 3) within regions showing above-threshold activity ($p > .05$, Bonferroni corrected for multiple comparisons). These results suggest that reduced classifier accuracy in patients is not simply due to a nonspecific factor, such as reduced signal-to-noise. These results also highlight a divergence between multivariate and a standard univariate analysis with the former showing group differences in the setting of undetectable differences in the latter.

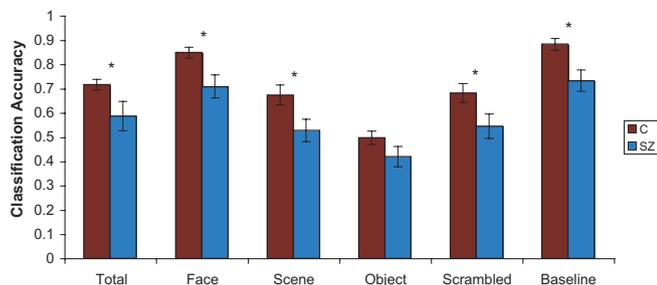


Figure 2. Multivoxel pattern analysis classification accuracy. The outcome, in terms of proportion of correct classification, for patients and control subjects across all stimulus types. *Significant difference between groups, two-tailed *t* test, $p < .05$. C, control subjects; SZ, schizophrenia.

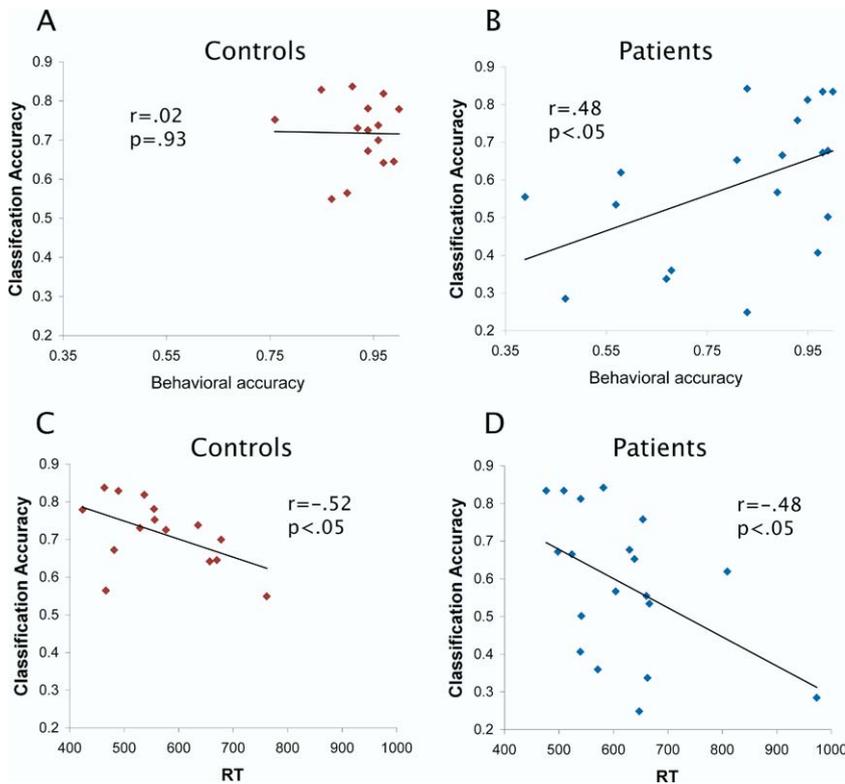


Figure 3. Correlations between classification accuracy and behavioral measures. (A) Control subjects' and (B) patients' behavioral accuracy in the 1-back task versus classification accuracy. (C) Control subjects' and (D) patients' reaction times (RT) in the 1-back task versus classification accuracy.

Discussion

The accuracy of a multivariate pattern classifier was high for healthy subjects and in close agreement with other studies (15,25). Classification accuracy was significantly lower in subjects with schizophrenia. In both groups, there was an inverse correlation between classification accuracy and reaction time, and in the schizophrenia group, there was a positive correlation between classification accuracy and behavioral accuracy. Diminished voxelwise correlation in activity across runs in patients and strong correlation between this measure and classification accuracy suggest that greater variability in engaging canonical activity patterns could be the cause of the lower classification accuracy in patients. There were no differences between groups in the standard GLM-based univariate measurements, demonstrating greater sensitivity of the pattern analysis method. The univariate results also suggest that nonspecific factors are not confounding the pattern analysis findings.

These results point to a possible neural mechanism of impaired visual information processing in schizophrenia—the instability of cortical networks and its failure to support coherent representations and their dependent cognitive processes. The relationship between category-specific representations, cognition, and behavior in healthy subjects was examined in a recent study by Polyn *et al.* (17). This study demonstrated that the time course of engagement of object-specific activity patterns, which had been shifted by 6 seconds to account for the lag in the BOLD signal, preceded the recall of an object. For example, when recalling a face, the face-specific pattern of activity is engaged prior to the actual response execution. The temporal order of events strongly suggests that the engagement of neural representations shapes the recall process. A logical extension of this model is to propose that the fidelity of representations influences the proper execution of dependent cognitive processes. Accord-

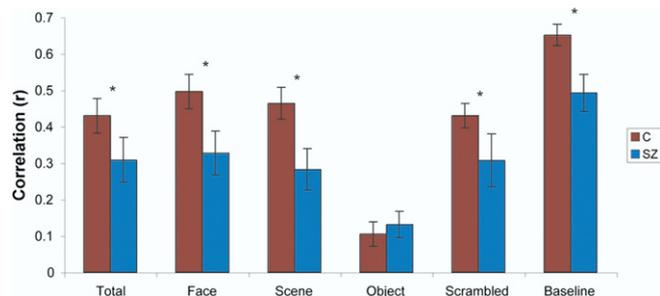


Figure 4. Interrun correlation in univariate activity. Interrun correlation values (Pearson's *r*) for patients and control subjects across all stimulus types. *Significant difference between groups, two-tailed *t* test, *p* < .05. C, control subjects; SZ, schizophrenia.

Table 3. Univariate Results

Contrast	Measure	Group	Mean	Standard Deviation
Face-Scene	No. voxels	SZ	25.2	30.7
		C	32.7	26.4
	Mean Beta	SZ	.019	.010
		C	.022	.008
	Maximum Beta	SZ	.066	.053
		C	.074	.051
Face-Object	No. voxels	SZ	12.5	14.6
		C	19.5	14.1
	Mean Beta	SZ	.017	.010
		C	.023	.011
	Maximum Beta	SZ	.045	.050
		C	.065	.043

C, control subjects; SZ, schizophrenia.

ingly, we interpret our findings as suggesting that the impairment in category-specific canonical representations is the reason that patients cannot perform the matching task and is not merely the reflection of poor performance. The decrement in stability of representations during baseline fixation blocks, where performance is not an issue, is consistent with this supposition. However, these statements about causality must be balanced by an acknowledgement that they are based on correlational results and more definitive empirical studies in the future are required.

An important issue related to the model of information processing deficits in schizophrenia is whether impairments reflect intrinsic neural abnormalities within the visual cortex or are the result of perturbations in top-down influences. Our findings could be accommodated by either model: a top-down account could be consistent with our findings if impaired coherence in activity patterns were due to a deficit in top-down signals and their failure to stabilize posterior networks (1,26). This potential mechanism is consistent with the task demands of this experiment, the 1-back, which imposes a low working memory load. According to contemporary theories and empirical findings related to prefrontal cortex (PFC) interaction with sensory cortex during working memory (2,19), a top-down signal from the PFC modulates activity and facilitates processing within the visual cortex during visual working memory. Consequently, altered top-down signals could contribute to impairments in visual networks. According to the bottom-up model, our findings would be further evidence of impaired neural function in sensory cortex, which then gives rise to higher-order processes in a feed-forward manner (27). This study was not designed to compare these competing models and the resolution of this issue will be the subject of future studies.

The comparison between univariate and pattern analysis results revealed some of the advantages of the latter approach, which is often suggested as a more sensitive method (14). The pattern analysis results were consistent with behavioral results and revealed group differences in activity, while traditional univariate analysis did not. The divergent findings between analyses cannot be attributed to differences in processing of the fMRI data, as they were processed identically across analytic streams. For pattern classification, spatial smoothing or transformation into a common template space is usually avoided to preserve as much spatial resolution as possible. Therefore, to control for possible effects of differences in image processing on our results, we did not follow the common practice of normalization into a standard space for group comparison in the univariate analysis.

Since we controlled for differences in image processing, the inherent differences in the analytic approaches likely account for the divergent results. While univariate analysis identifies strongly and specifically responsive voxels, pattern analysis is sensitive to the spatial combination of voxels. Consequently, a voxel that may be weakly active, i.e., subthreshold in the traditional univariate sense, may in the aggregate or in combination with others provide useful information in the pattern analysis approach (14).

The largely uniform difference between groups across conditions brings up the possibility of nonspecific, global factors such as neural excitability as the cause of the differential classification accuracy. This potential confound is addressed by two findings. First, there was an absence of group differences in the univariate measures of spatial extent and activation levels. Second, in the everyday objects category, in the context of very good and nearly equivalent behavioral performance across both groups, there was no significant group difference in classification accuracy and

interrun correlations. If nonspecific factors were driving the results, there should be a group difference in all categories. The greater interrun correlation among patients for this category suggests that the absence of difference is not just due to lack of power.

This study adds to a rapidly growing list of publications documenting impairments in what may be broadly categorized as neural connectivity in schizophrenia (7). Prior reports have identified deficits in long-range connectivity, either structural (28) or functional, with the latter being found during the resting state (8) and during cognitively activated states (9,10,29). We interpret the findings of disturbances in distributed representations revealed by pattern classification as a consequence of impaired local connectivity. While functional neuroimaging studies using traditional univariate analysis have provided very important information on the neural basis of schizophrenia, multivariate pattern analysis may be better suited to questions related to the distributed nature of neural representations and processes. As suggested by this study, the broad application of multivoxel pattern analysis to schizophrenia may yield novel insights into the underlying neural impairments leading to behavioral and cognitive deficits of this disorder.

Dr. Yoon, Ms. Tamir, Dr. Minzenberg, Dr. Ursu, and Dr. Ragland reported no biomedical financial interests or potential conflicts of interest.

Dr. Carter disclosed consultant fees for Eli Lilly and Roche.

1. Miller EK, Cohen JD (2001): An integrative theory of prefrontal cortex function. *Annu Rev Neurosci* 24:167–202.
2. Postle BR (2006): Working memory as an emergent property of the mind and brain. *Neuroscience* 139:23–38.
3. Barch DM, Carter CS, Braver TS, Sabb FW, MacDonald A 3rd, Noll DC, Cohen JD (2001): Selective deficits in prefrontal cortex function in medication-naïve patients with schizophrenia. *Arch Gen Psychiatry* 58:280–288.
4. Berman KF, Zec RF, Weinberger DR (1986): Physiologic dysfunction of dorsolateral prefrontal cortex in schizophrenia. II. Role of neuroleptic treatment, attention, and mental effort. *Arch Gen Psychiatry* 43:126–135.
5. Callicott JH, Bertolino A, Mattay VS, Langheim FJ, Duyn J, Coppola R, *et al.* (2000): Physiological dysfunction of the dorsolateral prefrontal cortex in schizophrenia revisited. *Cereb Cortex* 10:1078–1092.
6. Heckers S, Rauch SL, Goff D, Savage CR, Schacter DL, Fischman AJ, Alpert NM (1998): Impaired recruitment of the hippocampus during conscious recollection in schizophrenia. *Nat Neurosci* 1:318–323.
7. Friston KJ, Frith CD (1995): Schizophrenia: A disconnection syndrome? *Clin Neurosci* 3:89–97.
8. Garrity AG, Pearlson GD, McKiernan K, Lloyd D, Kiehl KA, Calhoun VD (2007): Aberrant “default mode” functional connectivity in schizophrenia. *Am J Psychiatry* 164:450–457.
9. McGuire PK, Frith CD (1996): Disordered functional connectivity in schizophrenia. *Psychol Med* 26:663–667.
10. Meyer-Lindenberg A, Poline JB, Kohn PD, Holt JL, Egan MF, Weinberger DR, Berman KF (2001): Evidence for abnormal cortical functional connectivity during working memory in schizophrenia. *Am J Psychiatry* 158:1809–1817.
11. Haxby JV, Gobbini MI, Furey ML, Ishai A, Schouten JL, Pietrini P (2001): Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science* 293:2425–2430.
12. O’Toole AJ, Jiang F, Abdi H, Haxby JV (2005): Partially distributed representations of objects and faces in ventral temporal cortex. *J Cogn Neurosci* 17:580–590.
13. Haynes JD, Rees G (2006): Decoding mental states from brain activity in humans. *Nat Rev Neurosci* 7:523–534.
14. Norman KA, Polyn SM, Detre GJ, Haxby JV (2006): Beyond mind-reading: Multi-voxel pattern analysis of fMRI data. *Trends Cogn Sci* 10:424–430.
15. Carlson TA, Schrater P, He S (2003): Patterns of activity in the categorical representations of objects. *J Cogn Neurosci* 15:704–717.

16. Haynes JD, Rees G (2005): Predicting the stream of consciousness from activity in human visual cortex. *Curr Biol* 15:1301–1307.
17. Polyn SM, Natu VS, Cohen JD, Norman KA (2005): Category-specific cortical activity precedes retrieval during memory search. *Science* 310:1963–1966.
18. Yoon JH, D'Esposito M, Carter CS (2006): Preserved function of the fusiform face area in schizophrenia as revealed by fMRI. *Psychiatry Res* 148:205–216.
19. Ranganath C, DeGutis J, D'Esposito M (2004): Category-specific modulation of inferior temporal activity during working memory encoding and maintenance. *Brain Res Cogn Brain Res* 20:37–45.
20. Duvernoy HM (1988): *The Human Hippocampus: An Atlas of Applied Anatomy*. Munchen, Germany: Bergmann.
21. Duvernoy HM, Bourgouin P (1999): *The Human Brain: Surface, Three-Dimensional Sectional Anatomy with MRI, and Blood Supply, 2nd ed.* New York: Springer.
22. Aguirre GK, Zarahn E, D'Esposito M (1998): The variability of human, BOLD hemodynamic responses. *Neuroimage* 8:360–369.
23. Zarahn E, Aguirre G, D'Esposito M (1997): A trial-based experimental design for fMRI. *Neuroimage* 6:122–138.
24. Postle BR, Zarahn E, D'Esposito M (2000): Using event-related fMRI to assess delay-period activity during performance of spatial and nonspatial working memory tasks. *Brain Res Brain Res Protoc* 5:57–66.
25. Hanson SJ, Matsuka T, Haxby JV (2004): Combinatorial codes in ventral temporal lobe for object recognition: Haxby (2001) revisited: Is there a "face" area? *Neuroimage* 23:156–166.
26. Silver MA, Ress D, Heeger DJ (2007): Neural correlates of sustained spatial attention in human early visual cortex. *J Neurophysiol* 97:229–237.
27. Butler PD, Zemon V, Schechter I, Saperstein AM, Hoptman MJ, Lim KO, *et al.* (2005): Early-stage visual processing and cortical amplification deficits in schizophrenia. *Arch Gen Psychiatry* 62:495–504.
28. Lim KO, Hedehus M, Moseley M, de Crespigny A, Sullivan EV, Pfefferbaum A (1999): Compromised white matter tract integrity in schizophrenia inferred from diffusion tensor imaging. *Arch Gen Psychiatry* 56:367–374.
29. Yoon JH, Minzenberg MJ, Ursu S, Walters R, Wendelken C, Ragland JD, Carter CS (2008): Prefrontal cortex dysfunction is associated with disrupted coordinated brain activity in schizophrenia: Relationship to impaired cognition, behavioral disorganization and global function [published online ahead of print June 2]. *Am J Psychiatry*.