



ACADEMIC
PRESS

Available online at www.sciencedirect.com

SCIENCE @ DIRECT®

NeuroImage

NeuroImage 20 (2003) 1096–1109

www.elsevier.com/locate/ynimg

Statistical properties of BOLD magnetic resonance activity in the human brain

Chien-Chung Chen,^{a,b} Christopher W. Tyler,^{a,*} and Heidi A. Baseler^a

^a *The Smith-Kettlewell Eye Research Institute, 2318 Fillmore Street, San Francisco, CA 94115, USA*

^b *Department of Psychology, National Taiwan University, 1, Sec. 4 Roosevelt Rd., Taipei 106, Taiwan*

Received 3 December 2002; revised 4 March 2003; accepted 10 June 2003

Abstract

We investigated the random variability of BOLD (blood oxygen level dependent) activation during rest, or null-hypothesis, conditions in which the observers were neither receiving controlled sensory stimuli nor performing cognitive tasks. The data indicate that the distributions for the BOLD variation across space are skewed, with non-Gaussian tails, while the distributions for the temporal variation within individual voxels are predominantly Gaussian. The proportion of voxels that show non-Gaussian properties is highly correlated with the magnitude of head movement of the observers. In all observers, the white matter showed less variability than the gray matter. The distributions for the spatial and the temporal variations are robust across observers despite differences in the data acquisition methods (EPI vs. spiral) and magnetic field strength (1.5 vs. 3 T). In most cases, the non-Gaussian tails of the spatial distribution can be eliminated by normalizing the amplitude in each voxel to its standard deviation before cumulating across voxels. We therefore recommend such a normalization procedure before any data manipulations are performed on fMRI data.

© 2003 Elsevier Inc. All rights reserved.

Keywords: Noise; Probability distribution; Functional magnetic resonance imaging; fMRI; EPI

Functional magnetic resonance imaging (fMRI) measures the regional change of blood oxygenation (Kwong et al., 1992; Ogawa et al., 1992; Bandettini et al., 1992) that signals neural activity in human brain when the observer is stimulated with sensorimotor or cognitive tasks. Since its inception, fMRI has evolved into a popular noninvasive tool to investigate human brain function.

One major limiting factor for fMRI is the low signal-to-noise ratio (SNR) of the blood-oxygen-level-dependent (BOLD; Ogawa et al., 1990) activity. Here, “signal” means the BOLD activity driven by the psychological tasks designated by the experimenters and ‘noise’ means the activity unrelated to such tasks. In a typical fMRI experiment, hemodynamic modulation induced by a psychological task is about 1~5% from the mean level in a block design and less in an event-related design. On the other hand, as shown below, random variation of the BOLD activation can easily

reach 2~4% in magnitude. Thus, fMRI data analysis is more susceptible to noise than many other neuroscience methodologies.

There are many studies on the MRI noise and artifacts resulting from the physics of the magnets (e.g., Henkelman, 1985; Gudbjartsson and Patz, 1995). However, since the physical equipment noise is small compared with the physiological noise (Krüger et al., 2001) these studies provide limited information about fMRI noise.

Studies on the noise of BOLD activation have focused on two issues: test-retest reliability and artifacts with well-defined origins. The test-retest reliability measures how well an experimental result can be reproduced in different sessions (Noll et al., 1997; Aguirre et al., 1998a; Tegeler et al., 1999; McGonigle, 2000; Waldvogel et al., 2000; Salli et al., 2001). This measurement provides an index of the magnitude of the noise, as high reliability is possible only when the noise is small. There have been studies on identifying distinct sources of artifacts in fMRI images, such as cardiovascular and respiratory activities (Hu et al., 1995; Noll,

* Corresponding author. Fax: +1-415-3458455.

E-mail address: cwt@ski.org (C.W. Tyler).

1998; Biswal et al., 1996), and correlated head or eye movements (Hajnal et al., 1994; Glover and Lee, 1995; Friston et al., 1996). The purpose of these studies was to design appropriate filters or algorithms to correct specific artifacts. Those analyses, of course, cannot cover all possible sources of fMRI noise. In addition, none of these analyses provides information about the probability distribution of the fMRI noise.

The latter, as it turns out, is crucial to current fMRI research. With a few exceptions that used nonparametric methods (Benson et al., 1996; Bullmore et al., 1996; Fernandez et al., 1999), most statistical fMRI data analyses assume a certain noise distribution. The most common approach is to apply a general linear model to the time series of fMRI data (Friston et al., 1995; Baudendistel et al., 1995; Worsley et al., 1997). It is assumed that the fMRI time series is a convolution of a certain hemodynamic response function with an experimental design sequence, corrupted by Gaussian noise. Therefore, the problem of fMRI data analysis is reduced to the general linear statistic problem. The Gaussian assumption is essential for these analysis using standard statistic tools such as the *t* test, the *F* test, or the ANOVA.

On the other hand, other methods make non-Gaussian assumptions. The most notable one is the independent component analysis (ICA; McKeown et al., 1998; McKeown and Sejnowski, 1998). ICA requires that the observed data are from independent non-Gaussian sources (Cardoso, 1998; Hyvärinen and Oja, 2000) because solutions for Gaussian sources can be found only up to an orthogonal transform.

While different data analysis methods are valid only under different noise distributions assumptions, there is little consensus about the empirical noise distribution. For instance, Hanson and Bly (2001) argued that the BOLD activation noise is non-Gaussian but Gamma-distributed, while Aguirre et al. (1998b) argued the noise distribution is essentially Gaussian. Our purpose for this study is to provide an empirical analysis of the noise distribution in BOLD activation within specified brain structures such as cerebral cortex (gray matter) or connecting fiber tracts (white matter). We recorded BOLD activation during rest scans, or null-hypothesis conditions, in which the human observers received no controlled sensory stimulation and were not instructed to perform any mental tasks. We then analyzed the variation of BOLD activation from a certain mean activation level.

Methods

Observers

Seven observers (5 males, 2 female, age 26–55 years) with normal or corrected-to-normal vision were scanned for this analysis. All observers provided informed consent to

participate in the project. Among them, observer R.K. was scanned on a GE 1.5-T Signa, and A.W., L.K., and C.T. were scanned on a GE 3-T Signa scanner at the Lucas Center at Stanford University. Observers C.L., S.C., and C.S. were scanned on a Bruker 3-T scanner at the MRI Laboratory of the National Taiwan University.

Scanning procedure

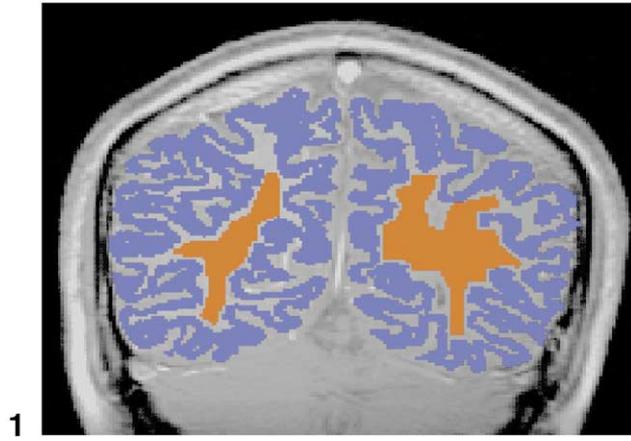
To evaluate the statistical properties of fMRI noise, data were collected for rest scans, in which observers were instructed to keep their eyes closed during the duration of the scan. The observer's head was immobilized using either a bite bar or padding within the head coil. The data were collected both at the Lucas Center in the Stanford University in Palo Alto and at the MRI center in the National Taiwan University in Taipei.

Palo Alto

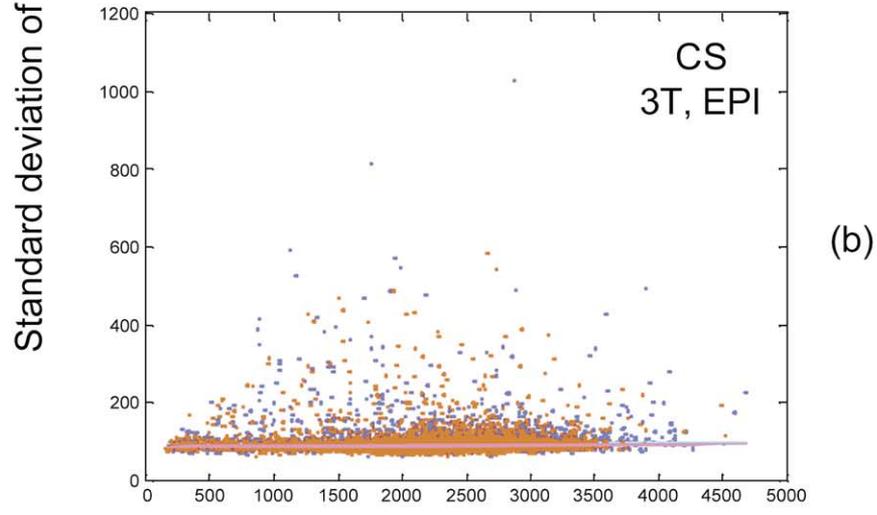
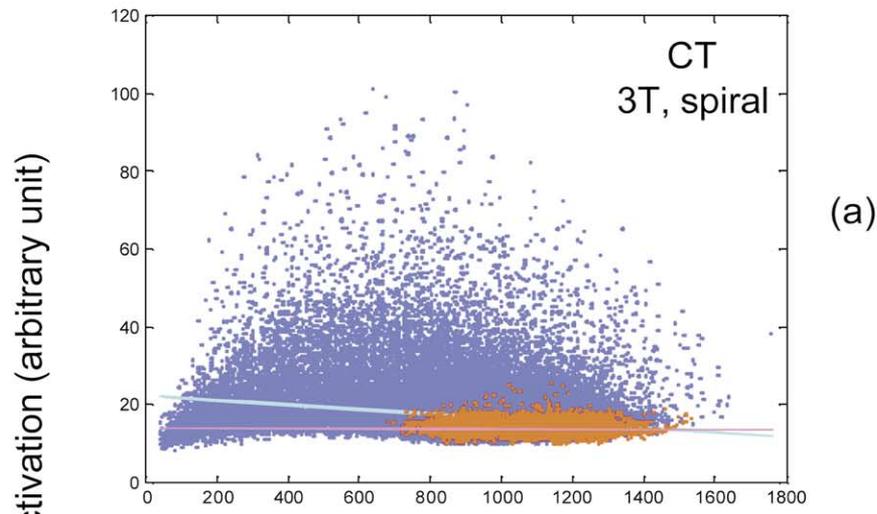
Observers were tested in a supine position in either a 1.5-T or 3-T GE Signa scanner equipped with a custom-made full or semicylindrical head coil. A high-resolution anatomical (T1-weighted) MRI volume scan of the entire head was run once on each observer (voxel size = $0.9 \times 0.9 \times 1.2$ mm). Within each scanning session, both functional (T2*-weighted BOLD) responses and anatomical (T1-weighted) images were acquired in identical planes. In the GE 1.5-T scanner, images were collected in eight planes perpendicular to the calcarine sulcus through the occipital lobes. A 2D spiral sequence was used for the functional images (Meyer et al., 1992) (two spiral interleaves, TR = 1500 ms, TE = 40 ms, flip angle = 90° , voxel resolution = $1.9 \times 1.9 \times 4$ mm). In the GE 3.0-T scanner, images were collected in 16 coronal planes spanning the posterior half of the brain. A 2D spiral sequence was again used for acquiring functional data (two spiral interleaves, TR = 1000 ms, TE = 30 ms, flip angle = 61° , voxel resolution = $2 \times 2 \times 4$ mm). Each scan lasted either 192 (96 images, 3 T), or 252 (84 images, 1.5 T) s. The first 16 s (8 images, 3 T) or 36 s (12 images, 1.5 T) was not included in the analyses, to remove any transient hardware-related artifacts at the beginning of the scan. Therefore, the data analyzed for each scan spanned 176 (88 images, 3 T) or 216 (72 images, 1.5 T) s.

Taipei

The images were collected with a Bruker 3-T scanner with a cylindrical head coil. A high-resolution anatomical (T1-weighted) MRI volume scan of the entire head was run once on each observer (voxel size = $1 \times 1 \times 0.75$ mm). Within each scanning session, both functional (T2*-weighted, BOLD) responses and anatomical (T1-weighted) images were acquired in identical planes. The images were collected in 18 transverse planes parallel to the AC-PC (anterior commissure–posterior commissure) line. An echoplanar imaging sequence (Stehling et al., 1991) was used to



- Gray
- White

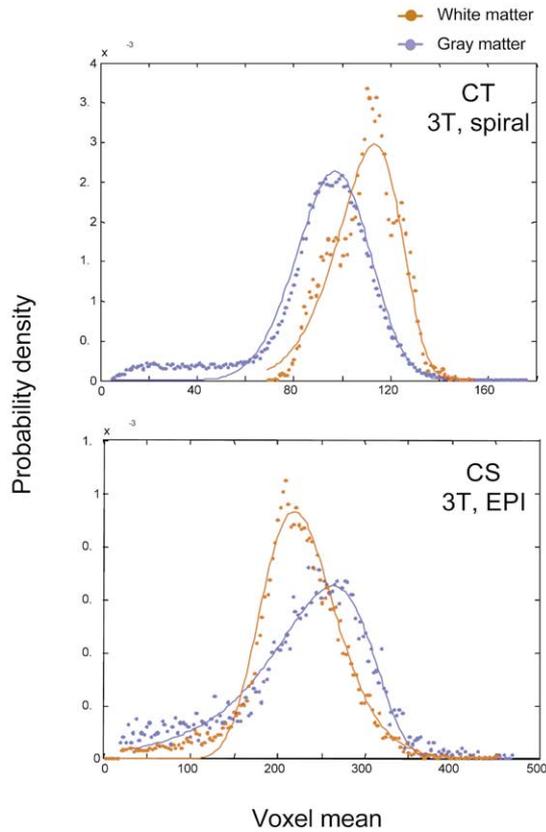


2

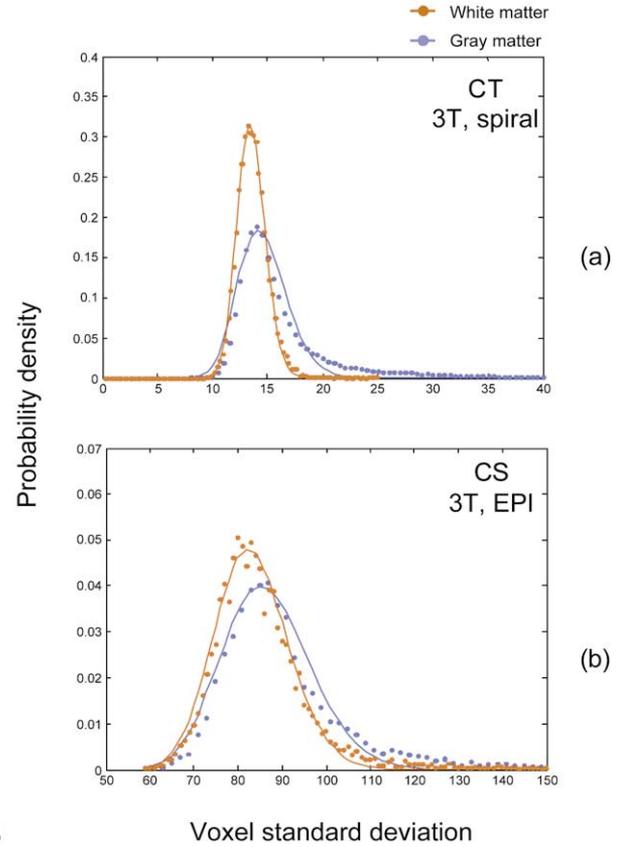
Mean activation (arbitrary unit)

Fig. 1. Segmentation of the gray (blue area) and the white matter (red area). Notice that the borders between the gray matter and the white matter, and between the gray matter and the CSF were not included in the analysis.

Fig. 2. Scatterplot for voxel means and standard deviations. Showing no correlation between the BOLD standard deviations and the mean levels. The blue dots denote voxels from the gray matter and the red dots from the white matter. The cyan and magenta lines are the best-fitting linear regression lines for the gray matter and the white matter respectively. Panel (a) shows data from one observer (C.T.) scanned with the spiral acquisition method while panel (b) shows data from one observer (C.S.) scanned with the EPI acquisition method.



3



4

Fig. 3. Probability density functions for spatial variation of the voxel means. The blue dots denote the empirical probability density for the gray matter and the red dots for the white matter. The blue and red smooth curves are the best-fit Gamma PDFs for the gray and white matter respectively. Panel (a) shows the data from observer C.T. and panel (b), C.S., as in Fig. 1.

Fig. 4. Probability density functions for spatial variation of the voxel standard deviations. The blue and red smooth curves are the best-fit Gamma PDFs for the white matter (red dots) and gray matter (blue dots), respectively. Panel (a) shows the data from observer C.T. and panel (b), C.S.

acquired the functional data (TR = 3500 ms, TE = 35 ms, flip angle = 90°, voxel resolution = 2.34 × 2.34 × 3 mm). Each scan lasted 262.5 s (75 images). The first 10.5 s (3 images) was excluded in the analyses. Thus, the data analyzed for each scan spanned 252 s (72 images).

Data analysis and visualization

Preprocessing

Voxels were selected in two categories: in the white matter and in the gray matter (cerebral cortex) of the brain. Voxels from the white matter and the gray matter were selected and segmented based on anatomical images in register with the functional images collected during each scanning session using available software (Teo et al., 1997; Wandell et al., 2000). The volume of the fMRI data from each scanning session was also aligned, projected, and interpolated to the high-resolution anatomical volume for each observer. Hence, the functional data voxels could be segmented into the gray matter and the white matter based on their correspondence with the anatomical segmentations. Fig. 1 shows an example of the segmentation. Notice that

we made an effort to avoid voxels near the border between gray and white matter.

The images acquired in the Stanford facility only covered the rear half the brain. Hence, for a comparison purpose, while the images acquired in the Taipei facility covered the whole brain, we only utilized the rear half of the brain for data processing.

Formulation

The variability of the voxel activations may be partitioned into two components: the within-voxel variability and the between-voxel variability. If x_{ij} is the activation in voxel i at time j , x_{ij} becomes

$$x_{ij} = \bar{X}_i + (x_{ij} - \bar{X}_i) \tag{1}$$

where \bar{X}_i is the mean activation of voxel i over time in a scan session. If we focus on a region of interest, then we can rewrite Eq. (1) as

$$x_{ij} = \bar{X}_r + (\bar{X}_i - \bar{X}_r) + (x_{ij} - \bar{X}_i) \tag{1'}$$

where \bar{X}_r is the overall mean activation of all voxels in a scan session in region of interest r . For instance, in this

study, we analyze activity for the voxels from the gray matter and the white matter separately. Hence, \overline{X}_r denotes the mean activation of either the gray matter ($r = g$) or the white matter ($r = w$). The term $(x_{ij} - \overline{X}_i)$ is the deviation of activation in a voxel i from a mean level, which gives us the within-voxel variation. Since it is a variation across time, the within-voxel variation also represents the temporal variation of the signal, denoted by the term \overline{X}_i while the comparative activity for different voxels gives us the between-voxel variation. Since each voxel occupies a different spatial location, the between-voxel variation also represents the spatial variation.

Probability density

When the sample size is large, it is possible to estimate the probability density function (PDF) directly. To do so, we first create the histogram of the quantity of interest (e.g., \overline{X}_i) with a reasonable number of bins. We count the frequency of occurrence in each bin, and scale the frequency by the total number of samples and the bin width. The result is taken as the probability density at the center of the bin. We then fit a theoretical PDF to the empirical densities with a least-squares algorithm.

Meta-analysis

There are situations where the sample size is too small to allow an estimation of PDF directly. For instance, with fewer than 100 data points in the time series for each voxel, the data are insufficient to allow reliable estimation of the single-voxel PDF. However, with a meta-analysis method, it is possible to decide if the activations are from a certain family of probability distributions. We may generate quantile-quantile plot (qq-plot; Johnson and Wichern, 1992) to test if the temporal variation in a voxel was Gaussian. For each voxel, we first sort the activations and calculate their corresponding probability values. We then calculate the quantiles corresponding to the probability values (e.g., “z-score” for the standard Gaussian distribution) and compare them with the Gaussian quantiles calculated from the original data. If the activations were Gaussian, the quantiles calculated from probability values should accurately match the quantiles calculated from the original data.

Head movement and motion index

One concern is the effect of head movement on the BOLD readout. One common approach to deal with this problem is to use a motion correction algorithm to align images collected at different time points. Such algorithms produce the voxel values in the transformed image by interpolating, transforming, or combining the voxel values of raw data. Hence, the statistical properties of the BOLD readout may be altered by these algorithms. Since different algorithms may alter the statistical properties in different ways, we chose to analyze the BOLD data without applying any motion correction. Instead, we designed a motion index

to study how head movement affect the statistic properties of the BOLD noise.

We use the rigid-body transform algorithm provided by the Statistic Parametric Mapping package (SPM; Friston et al., 1995) to determine the rotation parameters (θ_{x_i} , θ_{y_i} , θ_{z_i}) and translation parameters (δ_{x_i} , δ_{y_i} , δ_{z_i}) required to realign the i -th time periods to the first time period in a time series. From those parameters, one can compute which location $\lambda_0 = (x', y', z')$ in the first time period corresponds to a location $\lambda_i = (x_i, y_i, z_i)$ in the i -th time period. That is,

$$\begin{bmatrix} x' \\ y' \\ z' \\ 1 \end{bmatrix} = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & \cos \theta_{x_i} & -\sin \theta_{x_i} & 0 \\ 0 & \sin \theta_{x_i} & \cos \theta_{x_i} & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \times \begin{bmatrix} \cos \theta_{y_i} & 0 & -\sin \theta_{y_i} & 0 \\ 0 & 0 & 0 & 0 \\ \sin \theta_{y_i} & 0 & \cos \theta_{y_i} & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \times \begin{bmatrix} \cos \theta_{z_i} & -\sin \theta_{z_i} & 0 & 0 \\ \sin \theta_{z_i} & \cos \theta_{z_i} & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \times \begin{bmatrix} 1 & 0 & 0 & \delta_{x_i} \\ 0 & 1 & 0 & \delta_{y_i} \\ 0 & 0 & 1 & \delta_{z_i} \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_i \\ y_i \\ z_i \\ 1 \end{bmatrix}$$

With this relationship, we can also calculate the corresponding point of the same λ_0 in every image i , λ_i . We selected a point $\lambda_0 = (100, 100, 100)$ mm from the origin for the first volume and then computed the corresponding position in the following volumes. This point moves a distance $d_i = [(\lambda_i - \lambda_{i-1})^T (\lambda_i - \lambda_{i-1})]^{1/2}$, where T denotes the vector transpose, from image $i - 1$ to image i . The motion index is the average of d_i for the whole time series. Notice that λ_0 is chosen as a point 10 cm away from the origin, set at the anterior commissure. Given the size of most brains, λ_0 should be about the furthest point on the brain surface from the origin. Hence, the motion index is actually the mean motion of a hypothetic voxel near the surface of the brain and its magnitude, in turn, is about the same as the largest mean motion in brain voxels.

Results and discussion

Spatial variation

Fig. 2 shows examples of a scatter plot of the mean \overline{X}_i in Eq. (1) and the standard deviation for each voxel in the gray matter (blue dots) and in the white matter (red dots) from two of the seven observers (one scanned with spiral and the other with EPI) that we tested. The cyan and magenta lines are the least-squares linear regression lines for the

Table 1
Parameters for spatial variation

Observer			A.W.	C.T.	L.K.	R.K.	C.S.	S.C.	C.L.
Scanning method			3T, spiral	3T, spiral	3T, spiral	1.5T, spiral	3T, EPI	3T, EPI	3T, EPI
Relation between voxel mean and standard deviation	1. Correlation	g ^a	0.0592	−0.2064	−0.0107	−0.0455	0.0090	−0.0263	0.0184
		w	0.0956	−0.0273	0.1561	0.1314	0.0385	0.0173	0.1318
	2. Slope of linear regression line	g	0.0014	−0.0059	−0.0005	−0.0017	0.0005	−0.0048	0.0004
		w	0.0032	−0.0003	0.0109	0.0022	0.0018	0.0032	0.0037
Distribution of voxel means	3. Mean	g	877.4224	908.3182	1259.2353	1051.3474	2289.8989	12607.9745	13896.4382
		w	923.8768	1091.7140	1430.5659	1026.8614	2208.6519	12315.0192	14101.7513
	4. Standard deviation	g	195.6271	226.0435	423.4837	148.0756	751.6882	3220.1486	3647.7267
		w	117.1469	130.9665	205.9136	100.1934	531.0950	2453.1614	2544.6381
	5. Skewness	g	0.2419	−1.3543	−0.8760	−0.1137	−0.6873	−1.1904	−1.5658
		w	0.3427	−0.2028	−0.3884	0.5682	−0.5991	−0.8877	−1.0781
Distribution of voxel standard deviation	6. Mean	g	23.6032	16.7333	36.6020	16.7698	95.1060	833.5960	396.2135
		w	21.4714	13.6148	28.5594	12.5380	87.5916	663.4226	377.2543
	7. Standard deviation	g	4.4697	6.4736	18.2998	5.5768	38.2530	596.1706	76.7927
		w	3.9452	1.3949	14.3844	1.6962	24.5002	455.5924	71.7557
	8. Skewness	g	1.3196	3.8697	1.9108	4.1636	9.2518	2.8077	3.5056
		w	1.1280	0.8194	2.6861	2.2692	8.1744	3.3425	5.4680
Distribution of voxel standard deviation to mean ratio	9. Mean	g	0.02880	0.0223	0.0387	0.0164	0.0531	0.0778	0.0393
		w	0.02357	0.0127	0.0204	0.0123	0.0447	0.0588	0.0302
	10. Standard deviation	g	0.0130	0.0208	0.0412	0.0081	0.0507	0.0804	0.0569
		w	0.0051	0.0021	0.0123	0.0019	0.0329	0.0568	0.0317
	11. Skewness	g	6.7471	4.2700	3.5278	12.6100	4.1831	9.9918	6.0274
		w	1.0405	0.9005	6.5431	2.1456	6.5758	5.0995	9.9918

^a g: gray matter; w: white matter.

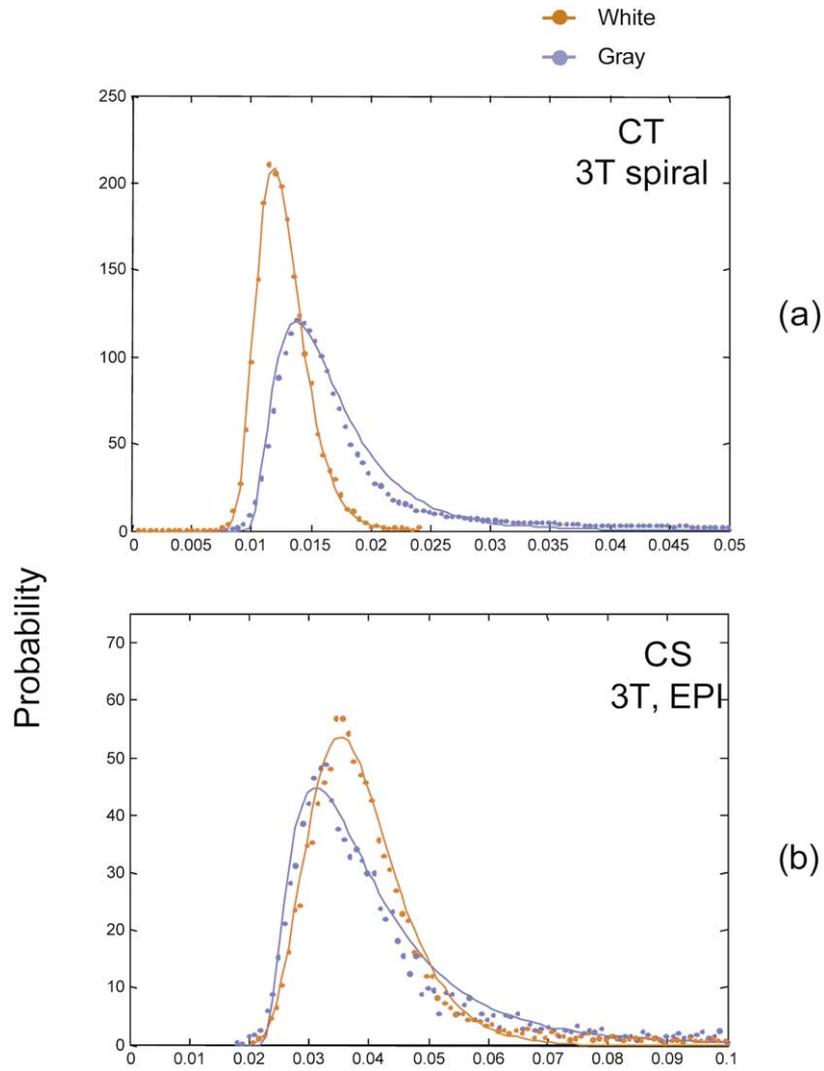
gray matter and the white matter regions, respectively. The slope of the linear regression lines and the correlation between the standard deviation and the mean for all seven observers are shown in Table 1 (rows 1 and 2). Some observers showed a slight correlation between the standard deviation and the mean of the voxels, but all slopes less than 0.01, giving the overall picture is that the standard deviation is statistically independent of the mean activation level over the full range from about 20 to 1000 (or 100 to 4000) in relative BOLD activation units.

Fig. 3 shows the distribution of voxel means for the same two observers. The blue dots represent the probability density of the distribution for the gray matter and red dots represent the probability density for the white matter. The smooth curves are best-fit Gamma density functions. In general (except for one observer, A.W.), the distributions of the voxel means are negatively (left) skewed, with the gray matter consistently showing greater skewness toward the left (Table 1, row 5). The voxel means also spread out more in the gray matter than in the white matter, with the standard deviation of the gray matter distributions being at least 25% greater than the white matter distribution. The three observers from the 3-T scanner with the spiral acquisition method show the greatest discrepancy between the gray matter and the white matter, with the gray-matter standard deviations being almost double those of the white matter. The white-matter distribution means in some observers are greater than those of the gray-matter distribution. These differences, however, are not statistically significant for any observer

and may just reflect the difference in the standard deviation and the skewness of the two distributions.

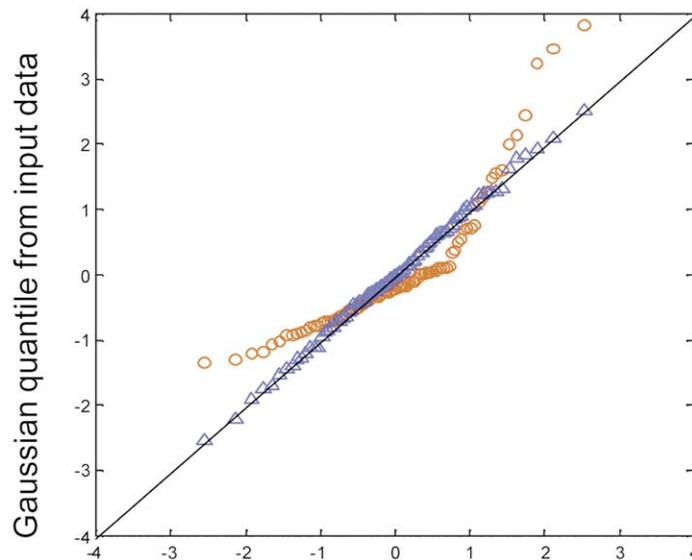
Next, we examine how the standard deviation for each voxel changes over space. The standard deviation for a voxel i , σ_i , is calculated as $[\sum_j (X_{ij} - \bar{X}_i)^2 / t]^{1/2}$, where t is the total number of samples in the time series of a voxel. Fig. 4 shows the distribution of the standard deviations. In both the gray matter and the white matter, the distributions are positively (right) skewed (as contrasted with the negative skew of the mean distributions). For all observers, the average voxel standard deviation is smaller in the white matter than in the gray matter. Hence, the white matter not only has smaller between-voxel variation over space, but also, on average, has smaller within-voxel variation over time.

Fig. 5 shows the distribution of the ratio between voxel standard deviation and the voxel mean, σ_i/\bar{X}_i (SMR) for the BOLD signal. The median voxel standard deviation-to-mean ratio (SMR; or coefficient of variation) is between 1.5% and, in one extreme case, 5.0%. Notice that BOLD signal strength is usually expressed as the percentage change of activation from a baseline that is usually the mean activation in a resting state (e.g., pre-stimulus activation). The signal strength is commonly seen at around 1~3% for event-related designs and 2~5% for block designs and, hence, is in the same ballpark as random variation. The distributions are positively skewed. For all observers, the distributions of this quantity are close to Gamma probability density functions (smooth curves in Fig. 5).



5

Voxel standard deviation to mean



6

Gaussian quantile from probability

Fig. 5. Probability density functions for spatial variation of the ratio between voxel standard deviation and the voxel mean, which are both negatively skewed. The blue and red smooth curves are the best-fit Gamma PDFs for white matter (red dots) and gray matter (blue dots), respectively. Panel (a) shows the data from observer C.T. and panel (b), C.S.

Fig. 6. Quantile-quantile plot for two voxels. The blue triangles denote a voxel with a Gaussian temporal variation. Hence, all the data points lie close to a straight line. The red circles denote a voxel with a non-Gaussian temporal variation. The data points markedly deviate from the straight line.

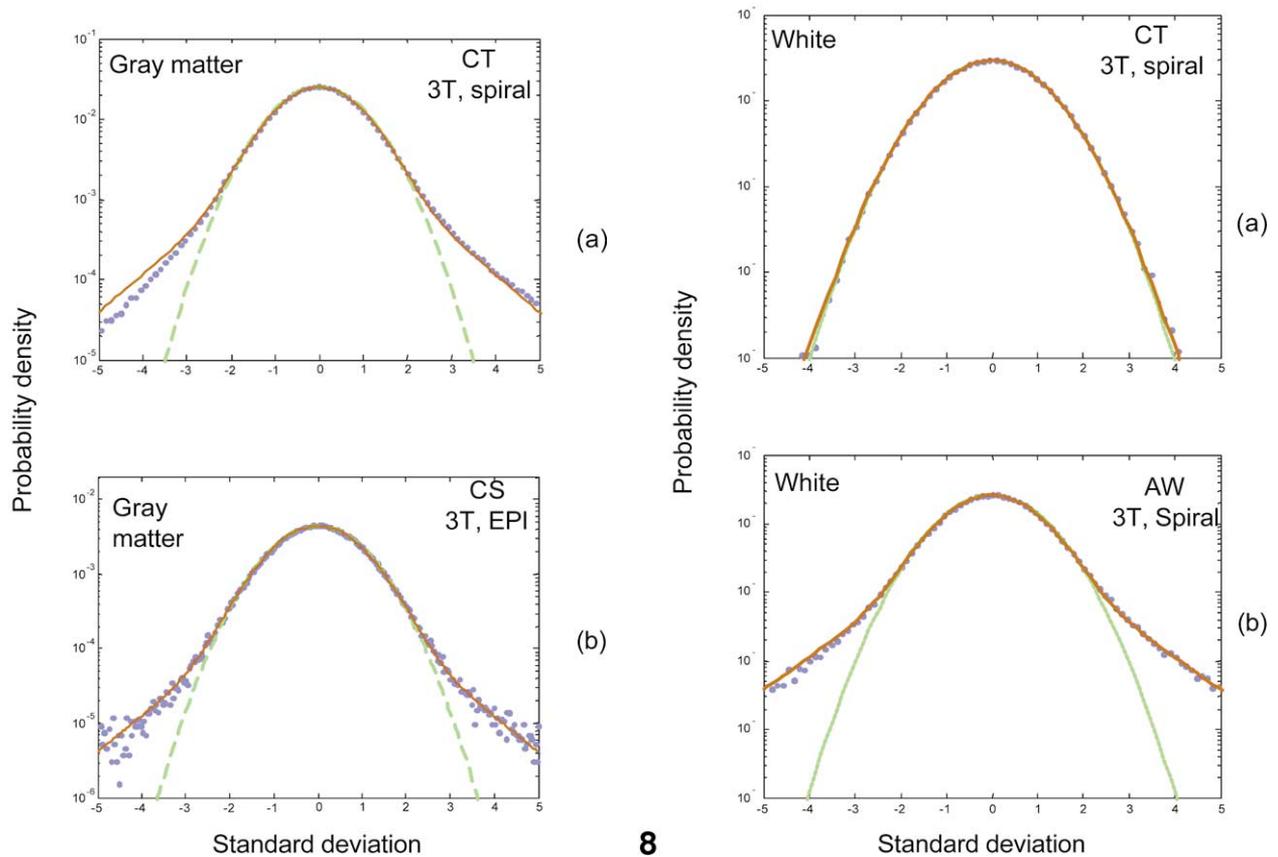


Fig. 7. Probability density functions for spatial-temporal variation in the gray matter in log-linear coordinates. The blue dots are the empirical probability densities. Note that, in this form, the Gaussian distribution appears as a parabola. The green dashed curves are the best Gaussian fit and the red solid curves are the fit from a hybrid of a Gaussian and a hyperbolic secant density functions. Panel (a) shows the data from observer C.T. and panel (b), C.S.

Fig. 8. The spatiotemporal variation in the white matter. For most observers, the Gaussian probability density function alone can fit the variation in the white matter as in panel (a) with observer C.T. In the worst case, (b) the white matter behaves like the gray matter and cannot be characterized by the hybrid distribution. The green curves are the Gaussian fits and the red curves are best-fit Gaussian-hyperbolic hybrid density functions.

The above analysis shows that, despite the wide differences in the data acquisition methods (spiral vs. EPI) and magnetic field strength (1.5 T vs. 3 T), the SMR are from the same family of distributions for all observers. In addition, we did not observe a better SMR for 3-T observers than the 1.5-T observer. We did notice that the smallest SMR observed in the EPI observers was greater than the largest SMR in the spiral observers. However, it is unclear if this difference is due the difference in the acquisition method or the quality of the magnets.

Temporal variation

Due to the limited number of data points in the series for individual voxels, we use the quantile-quantile analysis to test the normality of the temporal variation in a voxel $[(x_{ij} - \bar{X}_i)]$ in Eq. (1)]. Fig. 6 shows an example of such analysis on two voxels. The abscissa of the plot is the normal quantile computed from the raw data while the ordinate is the quantile computed from the percentiles of the Gaussian fit to the data points. If the variation in a voxel is Gaussian-distrib-

uted, the qq-plot of its activation will fall on a straight line and the correlation between the raw quantile and the percentile quantile will be close to 1.0. However, if the variation is non-Gaussian, the qq-plot of the voxel will deviate significantly from a straight line and result in a lower correlation. Thus, we can use the magnitude of the correlation to test the normality of the temporal (within voxel) distribution.

Table 2 lists the proportion of voxels that are classed as non-Gaussian distributed with the correlation test at the $\alpha = 0.01$ and 0.001 levels. In most observers, less than 5% of voxel activations are non-Gaussian at the $\alpha = 0.01$ level, and less than 1% at the $\alpha = 0.001$ level. However, in the worst case (A.W.), there were 35% of voxels with non-Gaussian activation. A possible cause of the non-Gaussian behavior will be discussed in later sections.

Spatiotemporal variation

Hanson and Bly (2001) reported that the noise of BOLD activations was Gamma distributed. However, in their anal-

Table 2
Percentage of non-Gaussian voxels in each observer at two type I error (α) level in the gray (g) and the white matter

Observer	ROI	Proportion of non-Gaussian voxels (%)	
		$\alpha = 0.01$	$\alpha = 0.001$
A.W.	g	34.7122	23.3678
	w	27.8418	18.1069
C.T.	g	3.0462	0.8892
	w	0.7997	0.1060
L.K.	g	14.9523	6.6947
	w	14.9303	5.3558
R.K.	g	2.3988	0.4525
	w	2.4017	0.2173
C.S.	g	1.5842	0.2821
	w	1.0992	0.1718
S.C.	g	12.7841	4.0031
	w	9.1330	2.2458
C.L.	g	1.0553	0.2499
	w	1.1693	0.0763

ysis, they pooled data points from both temporal samples and spatial locations. As shown above, separation of the set of variables reveals that the temporal variation of BOLD activation is essentially Gaussian but the spatial variation is close to Gamma. Hence, Hanson and Bly’s result is likely to be due to the confounding of the spatial with the temporal variation (2001).

Fig. 7 shows the distributions of the within-voxel variation $X_{ij} - \bar{X}_i$ for all voxels i and time points j in the data sets separated for the gray matter and the white matter. This analysis is similar to that of Hanson and Bly (2001), except that we first subtract the voxel mean from the data. The distributions are plotted in semilog coordinates to give a better look at the tails of the distributions down to densities of 10^{-4} . Notice that the distributions are symmetric and do not conform to the Gamma distribution. Hence Hanson and Bly’s (2001) conclusion applies only to the spatial variation of the BOLD noise.

Nevertheless, these spatiotemporal distributions are not Gaussian distributions either. In the gray matter, the distributions show two components: (1) in the center, the probability density functions have the form of a projectile on this logarithmic ordinate similar to the Gaussian; and (2) in the tails, the probability density functions decrease exponentially as the deviation from the voxel mean increases. The “projectile” portion can be fit by a Gaussian distribution, shown as the dashed curves in Fig. 7, while the linear tails imply an exponential probability density.¹ Therefore, we can fit the probability density functions with a linear com-

¹ The Gaussian probability density function has the form $y = (2\pi\sigma^2)^{-1/2} \exp(-x^2/2\sigma^2)$, where σ is the standard deviation. Hence, $\ln y = -0.5/\ln(2\pi\sigma^2) - x^2/2\sigma^2$, which appears as a projectile in a semilog coordinate system. The probability density function for an exponential distribution is $y = c \exp(-x/c)$ where c is the scale constant. Hence, $\ln y = \ln c - x/c$ implies a straight line in a semilog coordinate system.

ination of a Gaussian probability density function and a hyperbolic secant (*sech*) probability density function; that is, the probability density is defined by the function

$$p(x) = a * g(0, \sigma) + (1 - a) * s(0, \zeta)$$

$$= a \frac{1}{\sqrt{2\pi}\sigma} e^{-\left(\frac{x^2}{2\sigma^2}\right)} + (1 - a) \frac{2}{\pi\zeta} \frac{1}{e^{x/\zeta} + e^{-x/\zeta}} \tag{2}$$

where $g(0, \sigma)$ is a Gaussian density function with mean 0 and space constant σ , $s(0, \zeta)$ is a hyperbolic secant density function with mean 0 and space constant ζ , and a is the weight coefficient ranging between 0 and 1.

The advantage of the hyperbolic function is to have linear tails in the semilog plot while avoiding the exponential function’s discontinuity at the center. This compound distribution, as shown as the solid curves in Fig. 6, captures the characteristics of the data well. The parameters of this hybrid distribution for all observers are shown in Table 3. For most observers, the weight of the Gaussian [a in Eq. (2)] ranged from about 0.70 to about 0.90 in the gray matter. Three observers with a larger proportion of non-Gaussian voxels (Table 2) also show lower Gaussian weights. It is worth noting that many other candidate distributions, such as Student’s t , Cauchy, double exponential, log normal, and inverse-Gaussian do not provide good fits to the data. Some of these distributions can best fit either the central region or the tails, but not both. Other distributions, such as Cauchy and Student’s t , to minimize the sum of squared error, show considerable deviation from data in both the center and the tails.

In the white matter, the distributions for deviation from the voxel mean are closer to the Gaussian density function

Table 3
The parameters for the hybrid distribution^a

Observer	ROI	Empirical parameters for the distributions		Fitted parameters for the hybrid distributions		
		s.d.	skewness	a	σ	ζ
A.W.	g	19.6087	0.0696	0.7056	15.2279	17.4583
	w	16.3924	0.0636	0.8175	14.1287	15.3173
C.T.	g	17.8169	0.2139	0.7761	14.7029	16.5329
	w	13.5907	0.0041	0.9739	9.1285	13.5905
L.K.	g	40.6361	-0.2911	0.4148	24.9234	31.2743
	w	31.7541	-0.5839	0.5824	21.2275	26.8724
R.K.	g	15.2707	0.0047	0.8409	13.5139	13.9764
	w	12.5463	-0.0066	1.0000	12.5095	N/A
C.S.	g	101.7950	-0.1818	0.8521	86.8726	92.2935
	w	90.3194	-0.09739	0.9449	83.6917	101.3071
S.C.	g	1017.6578	-0.3361	0.4052	514.4062	731.8799
	w	799.1692	-0.0330	0.4499	442.2795	590.9691
C.L.	g	400.7730	-0.02956	0.9027	386.0048	348.5565
	w	381.3407	-0.03567	0.9461	368.3021	384.3380

^a The letter g and w in the second column denotes the gray and white matter respectively. Column 3 and 4 are the standard deviation and the skewness of raw data. The mean of the distributions is always zero. Columns 5 to 7 are parameters in Eq. (3).

than in the gray matter. For all observers, the non-Gaussian proportion [a in Eq. (2)] is less in white matter than in the gray matter. For some observers, the Gaussian density function alone provides a good description of the data (Fig. 8a) while in the worst case the white matter behaves like the gray matter (Fig. 8b).

In the previous section, we concluded that the activations in the vast majority of the voxels have accurately Gaussian distributions. In addition, we know that the standard deviation varies from voxel to voxel, suggesting that the non-normality of the overall distribution may be due to differences in the standard deviations of individual voxels. It follows that, if the activations are scaled by the standard deviations of the respective voxels, the resulting standard scores should be identically distributed for all the voxels. Thus, the overall distribution, normalized by the voxel standard deviations, should also be Gaussian.

Fig. 9 illustrates the result. In this figure, the same data shown in Fig. 7 are replotted after normalization in panels (a) and (b). Fig. 10 shows the qq-plots before and after normalization for observer CT. In panel (a), the plot deviates significantly from a straight line while in panel (b) the standardized data lie on virtually the same line. It is clear that the normalized distributions are much closer to Gaussian, with virtually no voxels generating non-Gaussian tails. For all observers, the sum-of-squared-error (SSE) for the Gaussian fit improves 3- to 6-fold in the gray matter when the distribution is normalized.

In summary, while the variability of voxels across space is not Gaussian and the standard deviation is inhomogeneous across space, the variability of individual voxels over time is uniformly Gaussian.

Head movement effects

It is known that head movements can produce artifacts and increase noise in BOLD images. We characterize this artifact by the head index described in the Methods section for each observer (Table 4). The proportion of the non-Gaussian voxels in the gray matter of an observer increases with the head index in a linear relationship between the two variables ($R^2 = 0.79$) in log-log coordinates, and a correlation between them of 0.89. Hence, we conclude that a great proportion of the non-normality in BOLD temporal activation is due to head movements.

It may be a concern that the head motion along one dimension may be greater than the others or has a greater influence on the statistic properties of voxels. A Tukey multiple-comparison test shows that three of seven observers (A.W., R.K., and L.K.) had the same size of head motion in all the three dimensions. Four observers had smaller head motion on the lateral-medial dimension than on the other two dimensions. However, none of the three dimensions of head motion provided substantially better prediction of the number of non-Gaussian voxels than the motion index: only the motion in the slice selection dimension has a correlation

with the number of non-Gaussian voxels slightly (but non-significantly) larger than the motion index (0.92 vs. 0.89), the other two dimensions having a smaller correlation (0.86 and 0.88). Hence, the motion index is a valid representation of the size of the head movement.

On the other hand, we found negligible correlation between the head motion and the spatial variation. The standard deviation of the voxel mean distribution (Fig. 2) characterizes the extent of the spatial variation. The correlation between the head motion index and the standard deviation scaled by the mean (i.e., dividing row 4 by row 3 in Table 1) is -0.19 in the gray matter and 0.18 in the white matter, while the critical point for a two-tailed correlation test at $\alpha = 0.05$ and $df = 5$ ($n = 7$) is 0.754 . Hence, the head motion contributes little, if any, spatial variation of BOLD activity compared with other possible sources such as the distribution of blood vessels.

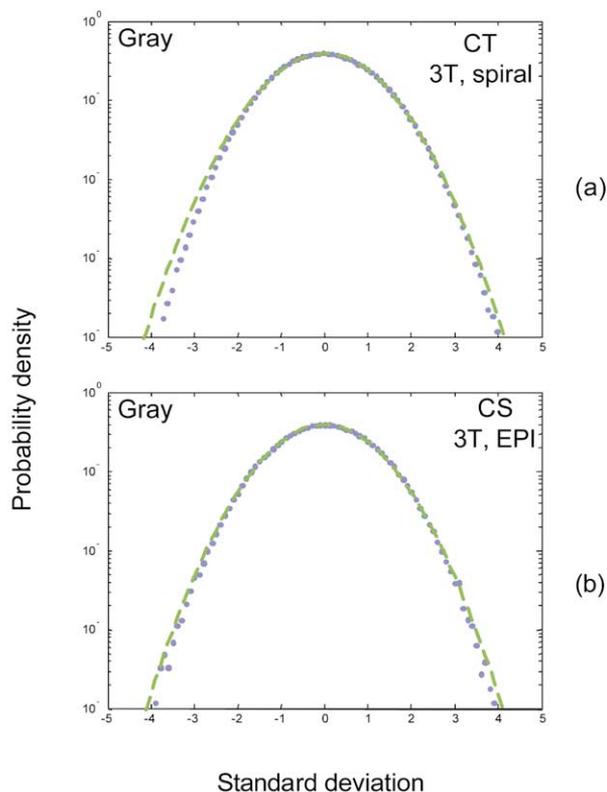
Spatial distribution of non-Gaussian voxels

Fig. 11 shows examples of the distribution of the temporal non-Gaussian voxels in the brain. The left column shows the non-Gaussian voxels from an observer with little head motion. The blue dots are the non-Gaussian voxels in the gray matter and the red dots, the white matter. The occurrence of non-Gaussian voxels is sparse. However, there are still clusters of non-Gaussian voxels. Those clusters tend to be at or near blood vessels (bright white spots or curves in the images). Hence, the larger blood vessels provide another source of non-normality.

The right column of Fig. 11 shows another observer with large head movements and in turn a greater proportion of non-Gaussian voxels. For comparison, we identify the Gaussian gray matter voxels in the image as cyan regions. It is clear that the non-Gaussian voxels are concentrated in the slices away from the center of the brain (top). The proportion of the non-Gaussian voxels in the slices near the center of the brain (bottom) is not much greater than that of the observer with little head motion. Since the voxels near the surface of the brain are more liable to the head motion (either from the size of head movements and the overlapping between brain and non-brain voxels in a time series), this result further demonstrates the contribution of non-normality from head movement.

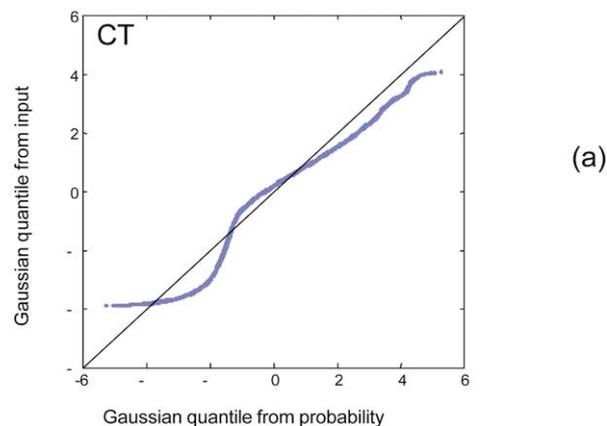
Table 4
Head movement index (in mm)

Observer	Head movement index
A.W.	3.2184
C.T.	0.1217
L.K.	2.4527
R.K.	0.0604
C.S.	0.1097
S.C.	0.1622
C.L.	0.1156

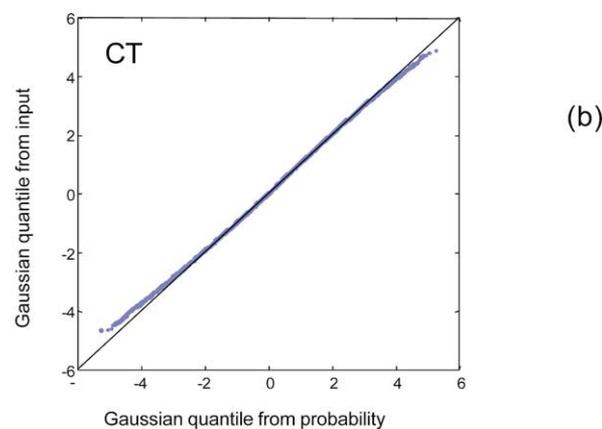


9

Standard deviation



(a)



(b)

10

Gaussian quantile from probability

Fig. 9. Probability density functions for spatiotemporal variation in the gray matter scaled by the standard deviations of the corresponding voxels. The blue dots are the empirical probability densities. The green dashed curves are the best Gaussian fits. Panel (a) shows the data from observer C.T. and panel (b), C.S.

Fig. 10. The quantile-quantile plot for the spatiotemporal variation observer C.T. before (a) and after (b) standardization of the individual voxel distributions.

Implications for data analysis

A common type of data analysis for BOLD activation is to apply a general linear model to an fMRI time series and to test whether the activation evoked by the cognitive/perceptual tasks is significantly greater than the activation at a base level (Friston et al., 1994, 1995; Baudendistel et al., 1995; Worsley et al., 1997). These analyses are based on t or F statistics to compare whether the mean activation for the task is sufficiently larger than that of the base level or whether the parameters for a hypothetical response structure (e.g., regression coefficients) sufficiently deviate from zero. Due to the definitions of the t and F distributions (Johnson and Kotz, 1973), these statistics rely on the assumption that the noise is Gaussian. We have shown that, when the head motion is small, the temporal variation in individual voxels is indeed Gaussian distributed. Hence, it is reasonable to apply such linear analysis for individual voxels. However, we also show that the activation pooled across locations yields a non-Gaussian distribution with broad tails. Thus, it is questionable to apply linear statistics after certain prepro-

cessing procedures such as spatial normalization and smoothing (Ashburner and Friston, 1999). These procedures pool activations from neighboring voxels together and distort the noise distribution when the voxel distributions are nonhomogeneous. Hence, they introduce substantial non-normality into the data, which invalidates the use of Gaussian-based statistics in principle. The net result of the non-normality is to increase the size of the tails from t and F distributions that are used to test the significance of mean differences or regression parameters. Hence, the experimenters have to either (1) increase the critical point to control alpha level (type I error) and reduce the power of statistics, or (2) keep the same critical point, and commit a lot of type I errors and in turn pinpoint false activations. We therefore recommend the normalization procedure of scaling the voxel amplitudes to the standard deviation of each voxel before any data manipulations are performed on fMRI data.

A second problem is the use of independent component analysis (ICA), which has been gaining popularity in fMRI recently (e.g., McKeon and Sejnowski, 1998; Calhoun et al.,

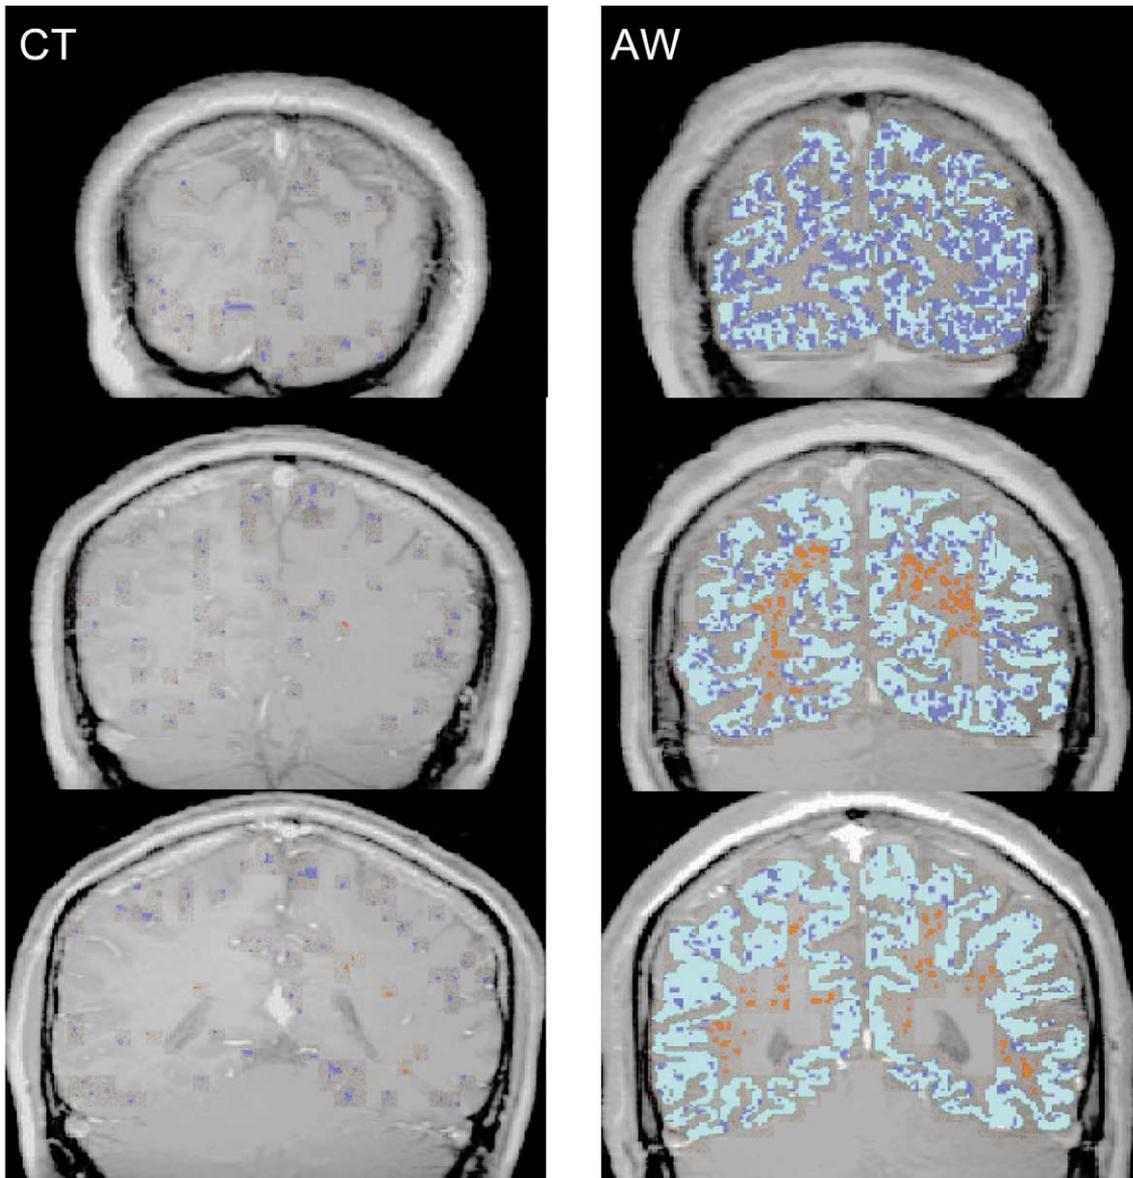


Fig. 11. The spatial distribution of non-Gaussian voxels (the correlation between the raw quantile and the Gaussian quantile smaller than the critical point of 0.01 α -level). The left column shows the distribution of observer C.T. and right column, observer A.W. In each column, images from top to bottom are slices from posterior to anterior. The blue dots are non-Gaussian voxels in the gray matter and the red dots, the white matter. The cyan regions in the right column are the Gaussian voxels in the gray matter.

2001; Stone et al., 2002; Svensen et al., 2002). The goal of the analysis is to derive independent response components from the overall distribution of voxel responses in a particular experimental paradigm. However, the analysis operates only on signals from non-Gaussian sources. While ICA is appropriate for spatial components, the temporal ICA may not be appropriate for application to spontaneous activity such as we have measured, because much of the non-normality derives from head motion artifacts. Quigley et al. (2002) reported that ICA performs better than conventional linear statistics only when the data set is corrupted by head motion or incorrect task performance. This conclusion is

consistent with our observation that head motion introduces non-normality in the temporal data. However, this approach is unlikely to provide any benefit in the component analysis of the neural signals because the non-normality introduced by head movements is extrinsic to the neural signals; the ICA analysis will focus largely on independent components of the motion artifact.

It is important to note, however, that these conclusions may not apply to stimulus-driven BOLD activity, which can readily be identified by ICA methods. Although the spontaneous activity is accurately Gaussian in each voxel, the response to a stimulus is likely to be non-Gaussian. Multiple

identifiable BOLD components with different spatial and temporal signatures can be revealed by appropriate analysis (Tyler et al., 2001, 2002).

Acknowledgments

We thank Dr. Chen Jyh-Hong of the MRI Laboratory at the National Taiwan University for allowing us to access the MRI facility at NTU; Dr. Wang Yao-Hong and Ms. Liu Chia-Li for helping us collect the data at NTU. This study was supported by NIH EY7890 and EY13025 to C.W.T. and an R.C. Atkinson Fellowship to C.C.C.

References

- Aguirre, G.K., Zarahn, E., D'Esposito, M., 1998a. The variability of human BOLD hemodynamic responses. *NeuroImage* 8, 360–369.
- Aguirre, G.K., Zarahn, E., D'Esposito, M., 1998b. A critique of the use of the Kolmogorov-Smirnov (KS) statistic for the analysis of BOLD fMRI data. *Magn. Reson. Med.* 39, 500–505.
- Ashburner, J., Friston, K.J., 1999. Nonlinear spatial normalization using basis functions. *Hum. Brain Mapp.* 7, 254–266.
- Bandettini, P.A., Wong, E.C., Hinks, R.S., Tikofsky, R.S., Hyde, J.S., 1992. Time course EPI of human brain function during task activation. *Magn. Reson. Med.* 25, 390–397.
- Baudendistel, K., Schad, L.R., Friedlinger, M., Wenz, F., Schroder, J., Lorenz, W., 1995. Postprocessing of functional MRI data of motor cortex stimulation measured with a standard 1.5 T imager. *Magn. Reson. Imag.* 13, 701–707.
- Benson, R.R., Logan, W.J., Cosgrove, G.R., Cole, A.J., Jiang, H., LeSueur, L.L., Buchbinder, B.R., Rosen, B.R., Caviness Jr., V.S., 1996. Functional MRI localization of language in a 9-year-old child. *Can. J. Neurol. Sci.* 23, 213–219.
- Biswal, B., DeYoe, E.A., Hyde, J.S., 1996. Reduction of physiological fluctuation in fMRI using digital filters. *Magn. Reson. Med.* 35, 107–113.
- Bullmore, E.T., Brammer, M.J., Williams, S.C.R., Rabe-Hesketh, S., Janot, N., David, A.S., Mellers, J.D.C., Howard, R., Sham, P., 1996. Statistical methods of estimation and inference for functional MR image analysis. *Magn. Reson. Med.* 35, 261–277.
- Calhoun, V.D., Adali, T., Pearlson, G.D., Pekar, J.J., 2001. Spatial and temporal independent component analysis of functional MRI data containing a pair of task-related waveforms. *Hum. Brain Mapp.* 13, 43–53.
- Cardoso, J.-S., 1998. Blind signal separation: statistical principles. *Proc. IEEE* 9, 2009–2025.
- Fernandez, G., Brewer, J.B., Zhao, Z., Glover, G.H., Gabrieli, J.D., 1999. Level of sustained entorhinal activity at study correlates with subsequent cued-recall performance: a functional magnetic resonance imaging study with high acquisition rate. *Hippocampus* 9, 35–44.
- Friston, K.J., Holmes, A.P., Worsley, K.J., Poline, J.P., Frith, C.D., Frackowiak, R.S.J., 1995. Statistical parametric maps in functional imaging: a general linear approach. *Hum. Brain Mapp.* 2, 189–210.
- Friston, K.J., Jezzard, P., Turner, R., 1994. Analysis of functional MRI time series. *Hum. Brain Mapp.* 1, 153–171.
- Friston, K.J., Williams, S.R., Howard, R., Frackowiak, R.S.J., Turner, R., 1996. Movement-related effects in fMRI time-series. *Magn. Reson. Med.* 35, 346–355.
- Glover, G.H., Lee, A.T., 1995. Motion artifacts in fMRI: comparison of 2DFT with PR and spiral scan methods. *Magn. Reson. Med.* 33, 624–635.
- Gudbjartsson, H., Patz, S., 1995. The Rician distribution of noisy MRI data. *Magn. Reson. Med.* 34, 910–914.
- Hajnal, J.V., Myers, R., Oatridge, A., Schwieso, J.E., Young, I.R., Bydder, G.M., 1994. Artifacts due to stimulus correlated motion in functional imaging of the brain. *Magn. Reson. Med.* 31, 283–291.
- Hanson, S.J., Bly, B.M., 2001. The distribution of BOLD susceptibility effects in the brain is non-Gaussian. *NeuroReport* 12, 1971–1977.
- Henkelman, R.M., 1985. Measurement of signal intensities in the presence of noise in MR image. *Med. Phys.* 12, 232–233.
- Hu, X., Le, T.H., Parrish, T., Erhard, P., 1995. Retrospective estimation and correction of physiological fluctuation in functional MRI. *Magn. Reson. Med.* 34, 201–212.
- Hyvärinen, A., Oja, E., 2000. Independent component analysis: a tutorial. *Neural Network* 13, 411–430.
- Johnson, N.L., Kotz, S.I., 1973. *Continuous Univariate Distributions*. Wiley, New York.
- Johnson, R.A., Wichern, D.W., 1992. *Applied Multivariate Statistical Analysis*, third ed. Prentice Hall, Englewood Cliff, NJ.
- Krüger, G., Kastrup, A., Glover, G.H., 2001. Neuroimaging at 1.5 and 3.0T: comparison of oxygenation-sensitive magnetic resonance imaging. *Magn. Reson. Med.* 45, 595–604.
- Kwong, K.K., Belliveau, J.W., Chesler, D.A., Goldberg, I.E., Weisskoff, R.M., Poncelet, B.P., Kennedy, D.N., Hoppel, B.E., Cohen, M.S., Turner, R., Cheng, H., Brady, T.J., Rosen, B.R., 1992. Dynamic magnetic resonance imaging of human brain activity during primary sensory stimulation. *Proc. Nat. Acad. Sci. USA* 89, 5675–5679.
- McGonigle, D.J., Howseman, A.M., Athwal, B.S., Friston, K.J., Frackowiak, R.S., Holmes, A.P., 2000. Variability in fMRI: an examination of intersession differences. *NeuroImage* 11, 708–734.
- McKeown, M.J., Makeig, S., Brown, G.G., Jung, T.P., Kindermann, S.S., Bell, A.J., Sejnowski, T.J., 1998. Analysis of fMRI data by blind separation into independent spatial components. *Hum. Brain Mapp.* 6, 160–188.
- McKeown, M.J., Sejnowski, T.J., 1998. Independent component analysis of fMRI data: examining the assumptions. *Hum. Brain Mapp.* 6, 368–372.
- Meyer, C., Hu, B., Nishimura, D., Macovski, A., 1992. Fast spiral coronary artery imaging. *Magn. Reson. Med.* 28, 202–213.
- Noll, D.C., Genovese, C.R., Nystrom, L.E., Vazquez, A.L., Forman, S.D., Eddy, W.F., Cohen, J.D., 1997. Estimating test-retest reliability in functional MR imaging. II: Application to motor and cognitive activation studies. *Magn. Reson. Med.* 38, 508–517.
- Noll, D.C., Genovese, C.R., Vazquez, A.L., O'Brien, J.L., Eddy, W.E., 1998. Evaluation of respiratory artifact correction techniques in multishot spiral functional MRI using receiver operator characteristic analyses. *Magn. Reson. Med.* 40, 633–639.
- Ogawa, S., Lee, T.M., Ray, A.R., Tank, D.W., 1990. Brain magnetic resonance imaging with contrast dependent on blood oxygenation. *Proc. Natl. Acad. Sci. USA* 87, 9868–9872.
- Ogawa, S., Tank, D.W., Menon, R., Ellermann, J.M., Kim, S.G., Merkle, H., Ugurbil, K., 1992. Intrinsic signal changes accompanying sensory stimulation: functional brain mapping with magnetic resonance imaging. *Proc. Natl. Acad. Sci. USA* 89, 5951–5955.
- Quigley, M.A., Haughton, V.M., Carew, J., Cordes, D., Moritz, C.H., Meyerand, M.E., 2002. Comparison of independent component analysis and conventional hypothesis-driven analysis for clinical functional MR image processing. *Am. J. Neuroradiol.* 23, 49–58.
- Salli, E., Korvenoja, A., Visa, A., Katila, T., Aronen, H.J., 2001. Reproducibility of fMRI: effect of the use of contextual information. *NeuroImage* 13, 459–471.
- Stehling, M.K., Turner, R., Mansfield, P., 1991. Echo-planar imaging: magnetic resonance imaging in a fraction of a second. *Science* 254, 43–50.
- Stone, J.V., Porrill, J., Porter, N.R., Wilkinson, I.D., 2002. Spatiotemporal independent component analysis of event-related fMRI data using skewed probability density functions. *NeuroImage* 15, 407–421.
- Svensen, M., Kruggel, F., Benali, H., 2002. ICA of fMRI group study data. *NeuroImage* 16, 551–563.

- Tegeler, C., Strother, S.C., Anderson, J.R., Kim, S.G., 1999. Reproducibility of BOLD-based functional MRI obtained at 4 T. *Hum. Brain Mapp.* 7, 267–283.
- Teo, P.C., Sapiro, G., Wandell, B.A., 1997. Creating connected representations of cortical gray matter for functional MRI visualization. *IEEE Trans. Med. Imag.* 16, 852–863.
- Tyler, C.W., Baseler, H.A., Kontsevich, L.L., 2001. A global symmetry model predicts both fMRI time series and psychophysical responses to symmetry patterns. Program No. 11.11.2001. Abstr. Soc. Neurosci.
- Tyler, C.W., Kontsevich, L.L., Baseler, H.A., 2002. Rapid independent spatiotemporal components in the fMRI BOLD response. Program No. 759.7. 2002 Abstract Viewer Itinerary Planner. Society for Neuroscience, Online.
- Waldvogel, D., van Gelderen, P., Immisch, I., Pfelffer, C., Hallett, M., 2000. The variability of serial fMRI data: correlation between a visual and a motor task. *NeuroReport* 11, 3843–3847.
- Wandell, B.A., Chial, S., Backus, B.T., 2000. Visualization and measurement of the cortical surface. *J. Cogn. Neurosci.* 12, 739–752.
- Worsley, K.J., Poline, J.B., Friston, K.J., Evans, A.C., 1997. Characterizing the response of PET and fMRI data using multivariate linear models. *NeuroImage* 6, 305–319.