Comparing functional connectivity via thresholding correlations and singular value decomposition

Keith J. Worsley^{1,2,*}, Jen-I. Chen², Jason Lerch² and Alan C. Evans²

¹*Department of Mathematics and Statistics, and* ²*Montreal Neurological Institute, McGill University, Montreal, Canada*

We compare two common methods for detecting functional connectivity: thresholding correlations and singular value decomposition (SVD). We find that thresholding correlations are better at detecting focal regions of correlated voxels, whereas SVD is better at detecting extensive regions of correlated voxels. We apply these results to resting state networks in an fMRI dataset to look for connectivity in cortical thickness.

Keywords: connectivity; correlation; singular value decomposition; fMRI

1. INTRODUCTION

The idea behind functional connectivity is to establish connectivity between two regions of the brain on the basis of similar functional response. For example, if two regions of the brain show similar functional magnetic resonance imaging (fMRI) measurements over time, then we could say that they are functionally connected, even though there may be no direct neuronal connection between these two regions. We can extend this idea to any image measurements. For example, if two regions of the brain show similar anatomical features (over subjects), such as cortical thickness, then we could say that they are ‘functionally’ connected as well.

The purpose of this paper is to compare two common approaches for detecting functional connectivity of image data (Koch *et al.* 2002; Horwitz 2003). The first and most direct method is to calculate a correlation of image values between pairs of voxels, then threshold these correlations to reveal the statistically significant connectivity (Cao & Worsley 1999).

The second is to do a singular value decomposition (SVD). SVD is equivalent to principal components analysis (PCA; Baumgartner *et al.* 2000), and similar to independent components analysis (ICA; van de Ven *et al.* 2004). In non-mathematical terms, SVD seeks to express the correlation structure with a small number of ‘principal components’, multiplied by random weights that vary randomly over time or subject. Voxels with high principal component values clearly covary together and are therefore positively correlated; voxels with high opposite signed components covary in opposite senses and are therefore negatively correlated. In practice, we extract the first few principal components, then threshold these components at an arbitrary level (since there are as yet no *p*-value results for local maxima of principal or independent

components). These regions are then our estimate of the connected voxels.

A third method, clustering, attempts to form clusters of voxels whose values over time or over subject are similar (Cordes *et al.* 2002). This is closely related to the first method, thresholding correlations, since correlation is used by many clustering methods as a measure of similarity and thresholding correlations simply clusters together all voxels whose similarity exceeds a threshold value. A fourth method, structural equations models (Gonçalves & Hall 2003), attempts to model the connectivity, but this is only feasible for connectivity between a small number of pre-selected voxels or regions. A thorough treatment of these last two methods is beyond the scope of this paper.

Our first step is to establish a common notation. It will be convenient to generalize slightly and assume that we have two sets of *N* images on each experimental unit (e.g. time or subject), denoted by the matrices *X* and *Y*. The rows are the image values, and the *N* columns of *X* and *Y* are the units (the number of columns *N* is identical, but the number of rows may differ). For example, *X* may be fMRI images mapped onto the visual cortex, and *Y* may be fMRI images of the frontal lobe, or the two sets of images may be different modalities such as positron emission tomography (PET) and fMRI across subjects. The most common case is when there is only one set of images, in which case we shall set *X* = *Y*. Usually the columns of *X* and *Y* are centred by subtracting their mean value (over units), or in general by removing the effect of a common linear model. To investigate correlations, and not covariances, we shall assume that each column of *X* and *Y* is normalized by dividing by its root sum of squares, so that the diagonal elements of the correlation matrices *X*'*X* and *Y*'*Y* is 1. The correlation between *X* and *Y* is defined as:

$$C = XY' \quad (1.1)$$

* Author for correspondence (keith.worsley@mcgill.ca).

One contribution of 21 to a Theme Issue ‘Multimodal neuroimaging of brain connectivity’.

Table 1. Resels for a convex search region R in D dimensions ($FWHM$ is the effective full width at half maximum of a Gaussian kernel used to smooth the white noise errors in the image data X . The diameter of a convex three-dimensional set is the average distance between all parallel planes tangent to the set. For a ball this is the diameter; for a box it is half the sum of the sides.)

d	D			
	0	1	2	3
0	1	1	1	1
1		$\frac{\text{length}(R)}{FWHM}$	$\frac{(1/2) \text{ perimeter length}(R)}{FWHM}$	$\frac{2 \text{ diameter}(R)}{FWHM}$
2			$\frac{\text{area}(R)}{FWHM^2}$	$\frac{(1/2) \text{ surface area}(R)}{FWHM^2}$
3				$\frac{\text{volume}(R)}{FWHM^3}$

If the two sets of images are the same, that is $X=Y$, we refer to C as an auto-correlation matrix, otherwise it is a cross-correlation matrix.

2. CORRELATIONS

Their (null) degrees of freedom (d.f.) n is the residual degrees of freedom of the linear model, with $n=N-1$ in the case of centred data. The correlation can then be converted (element-wise) to a t statistic with $m=n-1$ d.f. in the usual way:

$$t = \frac{\sqrt{m}C}{\sqrt{1-C^2}}. \quad (2.1)$$

If the images are smooth, then p -values for local maxima of the maximum auto- or cross-correlation statistical parametric map (SPM) C can then be found using random field theory. If the search regions R and S of X and Y have D and E dimensions, respectively, then

$$P\left(\max_{R,S} C > t\right) \approx \sum_{d=0}^D \text{Resels}_d(R) \sum_{e=0}^E \text{Resels}_e(S) \text{EC}_{d,e}^C(t), \quad (2.2)$$

where Resels and EC are the resels of the search region and the Euler characteristic (EC) density of the correlation SPM. The EC densities $\text{EC}_{d,e}^C(t)$ can be found in [Cao & Worsley \(1999\)](#) and for convenience they are reproduced in Appendix A. The resels are given in [table 1](#).

In practice, it is often forbidding to calculate all the correlations, so a common practice is to take a 'seed' voxel and correlate all other voxels with the observations at the seed voxel. In this case the threshold can be determined as above but with R replaced by a point in $D=0$ dimensions. The resulting threshold is identical to that obtained by applying the usual random field theory to the t statistic SPM t from equation (2.1)—see [Worsley *et al.* \(1996\)](#).

[Hampson *et al.* \(2002\)](#) propose iterating this procedure: use the statistically significant global maximum of the t SPM as a new seed and repeat the analysis until no further seeds are found. Each seed is tested for connectivity at the same (corrected) p -value of say 0.05. [Hampson *et al.* \(2002\)](#) show that the false positive rate of the overall procedure is still controlled at roughly 0.05. The argument is that when we get to the last seed,

Table 2. Thresholds at $p=0.05$ for $n=100$ null d.f., and three-dimensional search regions of size 1000 cm^3 with 10 mm smoothing.

method	dimensions		SPM	
	D	E	C	t
one voxel–one voxel	0	0	0.165	1.66
one 'seed' voxel–volume	0	3	0.448	4.99
volume–volume	3	3	0.609	7.64
(auto-correlation)				
volume–volume	3	3	0.617	7.81
(cross-correlation)				

the chance of finding any further seeds is 0.05, and the chance of finding any others beyond that is roughly 0.05^2 which is very small.

Of course the main criticism of such an approach is that it will find only one network of connected voxels that happen to contain the initial seed voxel. Other disjoint networks will not be detected. To detect all networks, we must resort to thresholding the correlation between all pairs of voxels, setting the threshold as above.

The amount of storage can be forbidding, since C is a voxels $\times$ voxels matrix, but we can avoid storing it altogether by first calculating the threshold, then only keeping those correlations that exceed the threshold. This can be reduced still further by only retaining local maxima, that is, pairs of voxels whose correlation is larger than the correlation between one voxel and any neighbour of the other.

The price to pay for searching all voxels is not as great as might be expected. For $n=100$ null d.f., three-dimensional spherical search regions of size 1000 cm^3 and 10 mm smoothing, the $p=0.05$ thresholds are given in [table 2](#). The auto-correlation thresholds are obtained by halving the six-dimensional search region (since correlations between r and s are the same as those between s and r), or equivalently, doubling the p -value to 0.1. As we can see, the increase in thresholds from 3 to 6 dimensions is not nearly as great as from 0 to 3 dimensions.

3. SINGULAR VALUE DECOMPOSITION

The SVD of the cross-correlation matrix C is:

$$C = UWV', \quad (3.1)$$

where U and V are orthonormal matrices whose rows are voxels and whose columns are components for X and Y , respectively, and W is a diagonal matrix of component weights. We then approximate C by equation (3.1) with the smaller weights in W set to zero.

Note that if $X=Y$ then this is just a PCA. Note also that if X is not an image but a matrix of covariates, this is commonly called partial least squares ([McIntosh & Lobaugh in press](#)).

The size of C (voxels $\times$ voxels) is usually much larger than its rank, again making storage prohibitive. Fortunately there is an easier way of finding the SVD of C from the eigenvectors A of the following much smaller

Figure 1. Focal connectivity by (a) t statistic for correlation with a seed voxel and (b) PCA. The thresholded t statistic (equivalent to thresholding the correlation) clearly shows the correlated focal regions, but the PCA shows nothing.

units $\times$ units matrix:

$$\left. \begin{aligned} (X'XY'Y)A &= AL, \quad B = A(A'Y'YA)^{-1/2}, \\ W &= L^{1/2}, \quad U = X(Y'YBW^{-1}), \quad V = YB. \end{aligned} \right\} \quad (3.2)$$

Note that L and $A'Y'YA$ are positive diagonal matrices. We can also write X and Y in terms of U and V as follows

$$X = U(BW) ', \quad Y = V(Y'YB) ', \quad (3.3)$$

so that we can regard the columns of BW and $Y'YB$ as (orthogonal) unit components that weight the spatial components in U and V to produce the observed correlation structure of X and Y , respectively.

4. COMPARISON OF THRESHOLDED CORRELATIONS AND SINGULAR VALUE DECOMPOSITION

Thresholding correlations directly detects those pairs of voxels that are highly correlated. SVD indirectly detects regions whose voxels behave in a similar way, that is, all moving up or down together. However, the

two methods do not always detect the same regions. To illustrate this, we simulated $N=120$ smooth standard Gaussian white noise (full width at half maximum (FWHM)=8 mm) images, chosen to represent the noise component of typical fMRI data. We added connectivity by reversing the SVD, that is, we made up a spatial component to reflect the regions to be connected, modulated this with a zero-mean Gaussian temporal component whose variance was chosen to induce the desired correlation, and added it to the noise component. We chose two types of connected regions.

- (i) Focal—two Gaussian-shaped regions (FWHM=8 mm) placed in the right frontal and left occipital region.
- (ii) Extensive—a rough binary mask of the right frontal and occipital region, smoothed with a Gaussian-shaped filter (FWHM=8 mm).

Here, we are looking for auto-correlations so $X=Y$. We then applied SVD (here PCA, since $X=Y$) and thresholding correlations with a seed voxel, chosen to be the voxel of maximum added spatial component in

Figure 2. Extensive connectivity by (a) t statistic for correlation with a seed voxel and (b) PCA. The thresholded t statistic (equivalent to thresholding the correlation) shows no statistically significant evidence of connectivity, but the PCA clearly shows the connected regions (red).

the frontal region. The results are shown in figures 1 and 2.

Note that the data have mean zero, so there is no need for centring, and so $n=N=120$ and the d.f. of the t statistic is $m=119$. The search regions were approximated by spheres with a volume of 1184 cm^3 containing 30 786 voxels. For correlation with a seed voxel and searching over all target voxels, the $p=0.05$ threshold for the t statistic is 4.89; for searching over all pairs of voxels, the $p=0.05$ threshold for the t statistic is 6.95.

For the focal correlation, the maximum t statistic for correlation with the seed voxel is 7.78, which is significant at the $p=0.05$ level even if we allowed for searching over all pairs of voxels. However, the PCA analysis shows no evidence of this connectivity (figure 1).

For the extensive correlation, the t statistic for correlation of the seed voxel with the voxel at the centre of the anterior component is 1.56, which is not significant even without correcting for searching.

The maximum t statistic is still not significant at 4.23. However, the PCA analysis clearly shows the connected regions (figure 2).

The lesson we learn from this is that SVD (PCA) is good at detecting extensive regions of connected voxels, whereas thresholding correlations is good at detecting focal regions of connected voxels.

5. APPLICATION

We apply the above methods to functional data, fMRI resting state networks and anatomical data of cortical thickness.

(a) *fMRI resting state network*

A subject was given a 9 s painful heat stimulus, followed by 9 s rest, then 9 s warm (neutral) stimulus, followed by 9 s rest, repeated 10 times as fully described in Worsley *et al.* (2002). The sample size of 120 frames was acquired at $\text{TR}=3 \text{ s}$; the first three were discarded leaving $N=117$. A linear model was fitted to account for the hot and warm block stimuli

Figure 3. fMRI resting state network. Inside the mid-cortical surface (transparent), the first principal component of the whitened residuals is thresholded at ± 0.5 of its maximum value (yellow, blue-green blobs). The ends of the rods join voxels where the auto-correlation SPM C of fMRI residuals exceeded $t = \pm 0.7$ (higher than the $p = 0.05$ threshold of $t = \pm 0.563$). Only six-dimensional local maxima are shown. Red rods indicate positively correlated voxels; blue rods indicate negatively correlated voxels (there is only one). Note that the red rods tend to join similarly coloured principal component regions, mostly the blue-green blobs.

(convolved with an hemodynamic response function (HRF)), and the drift was modelled as a cubic in the acquisition time. The residuals from this linear model, whitened to remove temporal correlation (Worsley *et al.* 2002), were used for further analysis, leaving $n = 111$ null d.f.

The search region $R = S$ is the same region of the brain used in the simulation in §4. From equation (2.2) with $n = 111$ null d.f., the $p = 0.05$ two-sided threshold for C is $t = \pm 0.563$ (since we are interested in both positive and negative correlations). There were too many correlations above this threshold to display, so figure 3 shows only six-dimensional local maxima above $t = \pm 0.7$ ($p < 5 \times 10^{-9}$).

Also shown in figure 3 is the first principal component of the time $\times$ voxel matrix of the whitened residuals, thresholded at ± 0.5 of its maximum value. This also reveals regions that are similarly correlated. In fact, the positive auto-correlations tend to join regions with either high (yellow) or low (blue-green) principal

components, and negative auto-correlations tend to join regions with different principal component values.

The main feature present in these patterns of connectivity is a strong right–left connection between regions close to the auditory cortex, and between left and right occipital regions. An explanation is that the subject might be processing random auditory (such as scanner noise) and visual information that is activating both right and left regions simultaneously, thus inducing a positive correlation between these regions. Some of the short-range in-plane connectivity, most evident on the left anterior outer cortex, may be owing to uncorrected head motion, which would induce apparent correlations between neighbouring outer cortex voxels.

(b) Cortical thickness

We illustrate the method on the cortical thickness of $N = 321$ normal adult subjects aged 20–70 years, smoothed by 20 mm FWHM; part of a much larger

Figure 4. Connectivity of cortical thickness. (a) Cortical thickness of one subject, smoothed by 20 mm, plotted on the average of the $N=321$ mid-cortical surfaces. (b) First principal component of the subject $\times$ node matrix of residuals removing a gender effect. The ends of the rods join nodes where the auto-correlation SPM C of cortical thickness exceeded $t = \pm 0.338$ ($n=319$ d.f., $p=0.05$, corrected). Only four-dimensional local maxima inside the same hemisphere are shown. (c) Back view, and colour bar: yellow to red rods indicate positively correlated nodes; blue rods indicate negatively correlated nodes. (d) Top view. Note that red rods tend to join similarly coloured principal component regions, whereas blue rods tend to join differently coloured principal component regions. (e,f) same as (c,d) but removing an age effect and age-gender interaction ($n=317$), which also removes some of the effective connectivity.

dataset fully described in Goto *et al.* (2001) and analysed in Worsley *et al.* (in press). The data on one subject are shown in figure 4a. We removed a gender effect, and then calculated the auto-correlation SPM, C , for all pairs of the 40 962 triangular mesh nodes, ignoring pairs of nodes that were too close. The search region $R=S$ is the whole cortical surface, with $D=E=2$, $\text{Resels}_0(S)=2$, $\text{Resels}_1(S)=0$ (since a closed surface has no boundaries) and $\text{Resels}_2(S)=759$ (see Worsley *et al.* 1999). From equation (2.2) with $n=319$ null d.f., the $p=0.05$ two-sided threshold for C is $t = \pm 0.338$ (since we are interested in both positive and negative correlations). Figure 4b–d shows only four-dimensional local maxima above t inside the same hemisphere.

The cortical surface is colour coded by the first principal component of the subject $\times$ node matrix of residuals, removing a gender effect. This also reveals regions that are similarly correlated. In fact, the positive auto-correlations tend to join regions with either high (red) or low (blue) principal component, and negative auto-correlations tend to join regions with different principal component values. Figure 4e,f shows the same analysis but removing a linear age effect and an age-gender interaction, so that $n=317$ null d.f. It is noticeable that many of the correlations now disappear, which demonstrates that they were induced by age effects. For example, if two regions become thinner over time, then this will induce an apparent positive correlation between these two regions. This illustrates

the importance of removing all fixed explanatory variables before carrying out a functional connectivity analysis.

6. DISCUSSION

Our comparison of detecting connectivity by thresholding correlations and SVD is only qualitative. From a theoretical point of view, very little is known about the stochastic behaviour of SVD and, in particular, there is no known way of thresholding SVD components to control the specificity. The only thing known about thresholding correlations is its specificity, not its sensitivity. A quantitative treatment would have to resort to extensive simulations, which are beyond the scope of this paper.

Nevertheless, our limited examples show that thresholding correlations is good at picking up focal, highly localized networks of connectivity, which can be completely missed by SVD. On the other hand, extensive regions of voxels that covary together and are correlated either positively or negatively with other extensive regions are unlikely to be picked up by thresholding correlations, but they can be qualitatively detected by SVD. We expect ICA to perform in a similar way to SVD with respect to focal and extensive regions of connected voxels, whereas we expect clustering, particularly the single-linkage types, to behave in a similar way to thresholding correlations.

Why do techniques like SVD and PCA give better detection for extensive regions of connected voxels? The answer to this is not yet fully understood because no theoretical results are available. But, on the other hand, it is clear why thresholding correlations should be good at detecting focal regions of connected voxels. The answer is obvious—if, for example, two small regions are perfectly correlated ($C=1$) then they will *always* be detected by thresholding correlations, no matter what threshold is used. But if they are very focal, their contribution to the SVD will be drowned by the random contributions from all other unconnected pairs of regions, and it is unlikely that they will emerge in the first few components.

A natural question to ask is how many SVD components we should look at. Of course the more we take, the more complete a description we obtain, but then the interpretation becomes harder. In principle, we should need at least as many components as there are isolated sub-networks. In other words, if three regions R_1, R_2, R_3 are all connected, and three further regions R_4, R_5, R_6 are connected, but there are no connections between $\{R_1, R_2, R_3\}$ and $\{R_4, R_5, R_6\}$, then we would at need at least two SVD components to capture this connectivity structure, and more if the strengths of the connectivities differ within the two sets of regions. Since we never know the number of such sub-networks in advance, then we never know how many components to retain. A formal test is difficult because of the presence of spatial correlation. In practice, we tend to look for the turning point in the plot of per cent variance explained versus component, beyond which the per cent variance explained does not seem to decrease markedly.

It is worth noting what SVD produces for pure noise data, where the only correlations are local spatial

correlations. It can be shown that if the spatial correlations are stationary (the same everywhere) and we have enough data (time points or subjects), then the SVD components are nothing but the Fourier basis functions in order of the power spectrum of the spatial correlation. Typically what happens in well-Gaussian-smoothed brain imaging data is that the first component is an anterior–posterior trend, the second is a right–left trend, the third is a superior–inferior trend and the rest are high-frequency spatial trends. This is because the dominant frequencies of Gaussian-smoothed data are the lowest frequencies. The conclusion is that if you see such SVD components there is good reason to suspect that there is no connectivity in the data (other than spatial smoothness).

We might ask what neuroscientific questions are better addressed using SVD and which are better addressed using thresholded correlations. The answer is to use both, since they are sensitive to different things. The main advantage of thresholding correlations over all other methods is that we can rigorously set the threshold to control the specificity.

APPENDIX A

EC density of the cross-correlation SPM

The EC density of the cross-correlation SPM in $d=0$, $e=0$ dimensions is just the upper tail probability of the Beta distribution with parameters $1/2$, $(n-1)/2$, where n is the null d.f., evaluated at t^2 . For $d>0$, $e\geq 0$ and $n>d+e$, it is

$$\begin{aligned} EC_{d,e}^C(t) = & \left(\frac{\log 2}{\pi} \right)^{(d+e)/2} \frac{(d-1)!e!2^{n-2}}{\pi} \\ & \times \sum_{k=0}^{\lfloor (d+e-1)/2 \rfloor} (-1)^k t^{k(d+e-1-2k)} (1-t^2)^{(n-1-d-e)/2+k} \\ & \times \sum_{i=0}^k \sum_{j=0}^k \left(\Gamma\left(\frac{n-d}{2}+i\right) \Gamma\left(\frac{n-e}{2}+j\right) \right) \\ & \times (i!j!(k-i-j)!(n-1-d-e+i+j+k)!)^{-1} \\ & \times ((d-1-k-i+j)!(e-k-j+i)!)^{-1}, \end{aligned}$$

where $\lfloor \cdot \rfloor$ rounds down to the nearest integer, terms with negative factorials are ignored and $EC_{d,e}^C(t) = EC_{d,e}^C(t)$. Note that the EC density in Cao & Worsley (1999) differs from that given here by a factor of $(4 \log 2)^{(d+e)/2}$. This is because the p -value approximation in Cao & Worsley (1999) is expressed in terms of Minkowski functionals (intrinsic volumes) whereas here the p -value (equation (2.2)) is expressed in resels. The summations have also been rearranged for easier numerical evaluation.

REFERENCES

- Baumgartner, R., Ryner, L., Richter, W., Summers, R., Jarmasz, M. & Somorjai, R. 2000 Comparison of two exploratory data analysis methods for fMRI: fuzzy clustering vs. principal component analysis. *Magn. Reson. Imaging* **18**, 89–94.
- Cao, J. & Worsley, K. J. 1999 The geometry of correlation fields, with an application to functional connectivity of the brain. *Ann. Appl. Probab.* **9**, 1021–1057.

- Cordes, D., Haughton, V., Carew, J. D., Arfanakis, K. & Maravilla, K. 2002 Hierarchical clustering to measure connectivity in fMRI resting-state data. *Magn. Reson. Imaging* **20**, 305–317.
- Gonçalves, M. S. & Hall, D. A. 2003 Connectivity analysis with structural equation modelling: an example of the effects of voxel selection. *NeuroImage* **20**, 1455–1467.
- Goto, R. *et al.* 2001 Normal aging and sexual dimorphism of Japanese brain. *NeuroImage* **13**(Suppl. 1), 794.
- Hampson, M., Peterson, B. S., Skudlarski, P., Gatenby, J. C. & Gorel, J. C. 2002 Detection of functional connectivity using temporal correlations in MR images. *Hum. Brain Mapp.* **15**, 247–262.
- Horwitz, B. 2003 The elusive concept of brain connectivity. *NeuroImage* **19**, 466–470.
- Koch, M. A., Norris, D. G. & Hund-Georgiadis, M. 2002 An investigation of functional and anatomical connectivity using magnetic resonance imaging. *NeuroImage* **16**, 241–250.
- McIntosh, A. R., Lobaugh, N. J. 2004 Partial least squares analysis of neuroimaging data: applications and advances. *NeuroImage* **23**(Suppl. 1), S250–S263.
- van de Ven, V. G., Formisano, E., Prvulovic, D., Roeder, C. H. & Linden, D. E. 2004 Functional connectivity as revealed by spatial independent component analysis of fMRI measurements during rest. *Hum. Brain Mapp.* **22**, 165–178.
- Worsley, K. J., Andermann, M., Koulis, T., MacDonald, D. & Evans, A. C. 1999 Detecting changes in nonisotropic images. *Hum. Brain Mapp.* **8**, 98–101.
- Worsley, K. J., Liao, C., Aston, J. A. D., Petre, V., Duncan, G. H., Morales, F. & Evans, A. C. 2002 A general statistical analysis for fMRI data. *NeuroImage* **15**, 1–15.
- Worsley, K. J., Marrett, S., Neelin, P., Vandal, A. C., Friston, K. J. & Evans, A. C. 1996 A unified statistical approach for determining significant signals in images of cerebral activation. *Human Brain Mapping* **4**, 58–73.
- Worsley, K. J., Taylor, J. E., Tomaiuolo, F., Lerch, J. 2004 Unified univariate and multivariate random field theory. *NeuroImage* **23**(Suppl. 1), S189–S195.