

Cognitive Conjunction: A new approach to brain activation experiments

Price CJ and Friston KJ

The Wellcome Dept of Cognitive Neurology, Institute of Neurology,
Queen Square, London WC1N 3BG, UK

Address for correspondence:

Dr Cathy J Price
The Wellcome Dept of Cognitive Neurology,
Institute of Neurology,
12, Queen Square,
London WC1N 3BG, UK
Tel: 44+ (0) 171 833 7472
Fax: 44+ (0) 171 813 1420
email: cprice@fil.ion.ucl.ac.uk

Abstract

This paper introduces the concepts and procedures behind ‘cognitive conjunctions’, a new approach to designing cognitive activation experiments. Cognitive conjunction compliments categorical approaches such as cognitive subtraction and requires a specific form of statistical inference that involves the conjunction of several hypotheses. Whilst cognitive subtraction studies are designed such that a pair of tasks differ only by the processing component(s) of interest, cognitive conjunction studies are designed such that two or more distinct task pairs each share a common processing difference. The neural correlates of the process of interest are then associated with the common areas of activation for each task pair.

There are two main advantages of cognitive conjunction relative to cognitive subtraction. The first is that it provides a greater latitude for selecting baseline tasks because it is not necessary to control for all but the component of interest. The only constraint on selecting the baseline task is that the component of interest is the only process that differs in each task pair. The second advantage is that cognitive conjunction does not depend on ‘pure insertion’ - the assumption that the addition of an extra processing component in the activation task has no effect on the implementation of processes that are also engaged by the baseline task.

The differences between the design and statistical analysis of experiments based on cognitive subtraction, cognitive conjunction and factorial designs are illustrated with a study of phonological retrieval. Cognitive conjunction analysis indicates that irrespective of whether subjects name words, objects, letters or colours, there is activation of the left posterior basal temporal lobe, the left frontal operculum, the left thalamus and midline cerebellum.

Introduction

Cognitive subtraction has been the mainstay of experimental design since the inception of cognitive activation studies in the late eighties. Subtraction designs involve selecting an activation task that engages the cognitive component of interest and a baseline task that activates all but the component of interest. The experimental design can be elaborated further by the addition or deletion of separable cognitive components to the tasks. Brain regions associated with the added cognitive components are identified by serial subtraction of scans obtained during the different tasks, by assuming that (i) the extra activations are due to the added cognitive component and (ii) the brain's implementation of previous components remains unchanged.

One of the problems with designing cognitive subtraction studies is finding baseline tasks that activate all but the cognitive process of interest. For instance, even when a baseline task does not require the explicit involvement of a cognitive process, there can be implicit processing, beyond the demands of the task, that reduce the difference between the activation and baseline tasks (Price et al. 1996). Another problem with cognitive subtraction is the effect that added components have on previous components. In general, for cognitive subtraction to work, we have to assume that the interaction between new and existing components can be ignored. This assumption is known as *pure insertion* wherein a new cognitive component can be purely inserted without affecting the expression of previous ones (ie those shared by both activation and baseline tasks). More generally, however, the expression of shared components is affected when new components are added and the difference between two tasks will comprise the added task component and the interaction between the added and shared components (see Friston et al. in press for further discussion).

The approach presented in this paper addresses and attempts to resolve both of these issues by introducing the idea of cognitive conjunctions. Cognitive conjunctions are an extension to the cognitive subtraction paradigm. While cognitive subtraction looks for activation differences between a pair of tasks that share all but the component of interest, cognitive conjunction looks for the commonality in activation differences (ie subtractions) between two or more pairs of tasks that share only the component of interest. Figures 1a and 1b illustrate this difference.

Figure 1 about here

Figure 1a represents a cognitive subtraction hierarchy. The process of interest (PI) is revealed by subtracting activity during the baseline task (B) from that during the activation task (A). Figure 1b represents a cognitive conjunction design which has two task pairs (IA,IB and IIA,IIB) designed such that the cognitive differences between the tasks of each pair both contain the process of interest (PI). Regional activation associated with PI is revealed by finding areas activated in both independent subtractions (IA - IB and IIA - IIB). For IA - IB, the differences are {P2 P4} and for IIA - IIB the differences are {P3, P4}. P4 (= PI) is activated in both contrasts, P2 and P3 are distinct but arbitrary task components.

This paper is divided into two sections. The first section introduces and discusses the concept of cognitive conjunction. The second section describes the procedures of analysis using a PET study of stimulus naming to identify the brain regions that implement phonological retrieval. We will show that conjunction analysis (i) allows the identification of the functional anatomy of cognitive processes, without relying on pure insertion and (ii) provides a greater latitude for task selection by removing the constraints on the choice of baseline task imposed by cognitive subtraction.

Cognitive Conjunction

Cognitive conjunction uses a series of activation and baseline task pairs to define a set of differences. These differences can include many cognitive components but only the component of interest is common to all task pairs. The activation conjunction is identified by the conjoint testing of several hypotheses, each pertaining to individual subtractions. By identifying the areas of common activation we can associate these regional effects with the common processing component. An important point here is that cognitive components that are not common to all task pairs can include interaction terms (ie the interaction between an added component and the shared components). Because these effects are discounted in the conjunction analysis, one does not need to rely on pure insertion. Further, the only restriction on the baseline tasks is that differences between task pairs both contain the process of interest (PI). They do not have to control for all the non-interesting components of the activation task. This allows for a flexible and unconstrained choice of baseline tasks. For example; (i) the baseline tasks could be the same for all pairs (eg an independent rest scan for each pair) or specific to each pair; and (ii) the baseline tasks can share a greater or lesser number of processes with the activation tasks. In other words the differences between tasks do not have to be matched exactly. They could be very similar or very different depending on the experimental question. For example baseline tasks could differ substantially from activation tasks that involve some explicit processing, so as to avoid implicit processing in the former.

Cognitive conjunction differs fundamentally from cognitive subtraction at both the level of experimental design and analysis. One design difference is that the baseline and activation tasks are not constrained to differ by a single cognitive component (see above). Another is that cognitive conjunction is not serial (where each successive task serves as a baseline for the next); like factorial designs each activation task has its own baseline, such that each subtraction is independent of the other subtractions included in the conjunction. These design differences allow one to eschew assumptions like pure insertion and therefore lend the analysis a greater degree of neurobiological validity. The reason that cognitive conjunctions do not rely on pure insertion is that the conjunction discounts interaction terms whether they exist or not. Cognitive subtraction on the other hand assumes these interaction terms do *not* exist. This can be seen clearly in Figure 1b. Above we

assumed that P2 in the first task pair and P3 in the second task pair were distinct and novel cognitive processes. However, they could also be construed as interaction effects; ie P2 could represent the interaction between P1 and PI and in the second task pair P3 could represent the interaction between P2 and PI. By virtue of the fact that the conjunction involves PI, and only PI, the interaction effects are discounted even if they are substantial. Allowing for interaction effects means we do not have to assume pure insertion.

At the level of statistical analysis there are differences between the way one tests for subtractions and conjunctions. In cognitive subtraction, the area of activation is identified by subtracting one task from another, whereas in cognitive conjunction an area of common activation is identified once the subtractions have been effected. More exactly, cognitive subtraction relies on the rejection of one null hypothesis (ie no difference between two conditions). Cognitive conjunction, on the other hand, relies upon the conjoint rejection of multiple null hypotheses (e.g. a significant activation in the first and the second task pairs). We expand briefly on this distinction below and more formally in the appendix.

Functional neuroimaging studies are usually analyzed using some form of statistical parametric mapping (SPM). This involves the construction of an image of a statistic that tests a hypothesis about differences in the distribution of brain activity between tasks. In what follows, we effectively combine a series of statistical parametric maps (SPMs) to identify areas of common or conjoint activation. There are many ways that one could test for these conjunctions, we have chosen a particular approach that is implemented easily in the context of SPM. In this approach, we essentially create an SPM of the sum of all the activations and eliminate voxels where differences among these activations are significant. The rationale for this can be seen in relation to factorial designs. For example, consider two task pairs, whose activations can be thought of as reflecting the presence of a common cognitive component under two levels (ie the level of the first task pair and the level of the second task pair). These activations can be thought of as two simple main effects. A conjunction is defined as the presence of a main effect in the absence of an interaction. In other words, the activation or main effect is significant and the simple main effects are not significantly different. Using this definition of a conjunction, we distinguish between main effects with and without interactions. The latter do not constitute conjunctions and normally call for reporting the two simple main effects separately. By virtue of this, a conjunction analysis can be used to complement the usual analysis of factorial designs.

The mathematical details of the conjunction analysis are presented in the appendix for the interested reader.

A PET study of naming

This section describes the cognitive model, experimental design and functional analysis (i.e. how each task decomposes into separable components) of a PET study which uses a cognitive conjunction design to identify the brain regions involved in phonological retrieval. We ascribe the term phonological retrieval to the activation of the verbal label (ie the name) attached to a visual stimulus or concept. The stimuli chosen for the study were words, letters, objects and colours.

Cognitive model

Figure 2 is the cognitive model on which we have based our experimental design, it specifies the processing components involved in each naming task and emphasis the convergence of different routes to phonological retrieval.

Figure 2 here

For each of the perceptual categories, there will be a specific set of processes required to retrieve the name. For instance, during object naming phonological retrieval is dependent on the activation of both structural and semantic memories; during reading, phonological retrieval can proceed via sublexical pathways (for instance when we read pseudowords like "neeb" which have no semantic content); and during colour naming, phonological retrieval is not dependent on either structural or semantic memories. The aim of the cognitive conjunction is to identify activation that is common to naming each perceptual category irrespective of the processing route taken. From a cognitive perspective, the conjunction of these tasks includes visual analysis and the retrieval and execution of phonology.

Experimental design

Four pairs of activation (A) and baseline (B) tasks were used to identify the brain regions implicated in phonological retrieval. Illustrations of the stimuli are shown in Figure 3.

Figure 3 here

Below we list the constituent cognitive components of the activation and baseline tasks and the differences for each task pair. It will be seen that phonological retrieval is the only cognitive component common to all task pair differences.

Task pair I: Word naming

This task pair comprised (A) reading single familiar monosyllabic words (task 1) and (B) saying the same prespecified word to strings of false font (task 2). The cognitive components common to both these tasks were early visual analysis and articulation. The differences between these two tasks includes orthographic, semantic and sublexical processing, phonological retrieval and the interactions amongst these components.

Task pair II: Letter naming

This task pair comprised (A) naming single arabic letters (task 3) and (B) saying the same prespecified word to single false font characters (task 4). The cognitive components common to both these tasks were early visual analysis and articulation. The differences between these two

tasks includes orthographic letter processing, phonological retrieval and the interaction between these processes.

Task Pair III: Object naming

This task pair comprised (A) naming visually presented easily identifiable objects (task 5) and (B) saying “yes” to the same stimuli (task 6). The cognitive components common to both these tasks were early visual analysis, object processing (3D form processing and activation of structural and semantic memories) and articulation. The differences between these two tasks include explicit phonological retrieval and the interaction between object recognition and explicit phonological retrieval.

Task pair IV: Colour naming

This task pair comprised (A) naming the colour of 2D patterns (task 7) and (B) saying “yes” to the same stimuli to acknowledge the stimulus had been seen (task 8). The cognitive components common to both these tasks were early visual analysis of form and colour, and articulation. The differences between these two tasks includes phonological retrieval and the interaction between visual analysis and phonological retrieval.

Functional analysis of tasks

Figure 4 represents a graphic task analysis where we have represented each of the four task pairs in terms of their constituent processing components. The format corresponds to that adopted in Figure 1 but the constituent processing components used are those illustrated in the cognitive model of Figure 2. The grey filled regions indicate task components that distinguish between activation and baseline tasks, the black filled regions indicate where these differences were common for every pair. It is immediately apparent that the only cognitive component that distinguishes between each task pair is phonological retrieval.

Figure 4 here

Data Acquisition and analysis

Identification of the brain regions significantly activated across all four task pairs involved the following procedure. Each task pair was replicated three times in one of two groups of six subjects. One group of subjects named words and letters with the respective baselines, the other group of subjects named pictures and colours. The data were acquired using PET and a bolus $H_2^{15}O$ technique as previously described, realigned, spatially normalised (Friston et al. 1996) and analyzed with statistical parametric mapping using ANCOVA with global activity as a subject specific confounding covariate. Four orthogonal contrasts were specified, each testing for activations within each of the four task pairs. The resulting Z values were subjected to a conjunction analysis as described in the appendix.

Results

The SPM illustrating the maximum intensity projection of the Z statistic for the conjunction of the task pair differences is shown in Figure 5. It can be seen that three regions surviving a corrected threshold of significance are identified: the left posterior basal temporal lobe (Brodmann's Area 37), the left frontal operculum (from the left anterior insula to the lateral inferior prefrontal cortex) and midline cerebellum. There were also conjoint activations, that only reached an uncorrected threshold of significance, in the left thalamus and the left lateral inferior occipital lobe. Figure 6 plots the adjusted mean activity for all 8 conditions in these regions to illustrate the activation differences for all four independent task pairs.

Figures 5 and 6 here.

Discussion

In this paper we have introduced the idea of ascribing a function to a particular brain region on the basis of conjoint activations by pairs of tasks that each share a common processing difference. This can be thought of as an extension of categorical analysis in the sense that we combine many independent subtractions in order to identify a conjunction of activations. For conjunctions to be implemented, the activation task in each pair must (i) engage the process of interest and (ii) have its own baseline. These design prerequisites contrast to those of subtraction designs (where each successive task serves as a baseline for the next) but resemble factorial designs. Below we reiterate and expand upon the differences between cognitive subtraction, factorial designs and cognitive conjunctions.

Cognitive subtraction

The simplest form of subtractive design involves a task that activates the cognitive process of interest and a second task that controls for all but the process of interest. Subtraction designs can also be elaborated into serial or hierarchical subtraction when the baseline task involves one more cognitive component than a third task. A good example of a hierarchical subtractive design is the study of reading by Petersen et al. (1990). This study had one variable with five different levels relating to the type of visual stimuli viewed: words, pseudowords, consonant letter strings, false font and visual fixation. A cognitive subtraction analysis identifies differences in activation between different levels of the hierarchy and associates these with the differences in psychological processes between levels. As discussed above, there are two main limitations of cognitive subtraction studies. The first is that they rely on pure insertion which assumes that an extra cognitive component can be purely inserted without affecting the expression of pre-existing components. The second relates to the difficulty of finding baseline tasks that activate all but the process of interest. This is particularly relevant when there is implicit processing of a stimulus beyond the demands of a task (see Price et

al., 1996).

Factorial designs

In factorial designs, there are two or more variables (or factors) and the different levels of each variable are matched. For example, if the two variables were phonological retrieval and object recognition, the simplest factorial design would have four tasks which evaluated (i) phonological retrieval in the presence of object recognition, (ii) phonological retrieval in the absence of object recognition, (iii) object recognition in the absence of phonological retrieval and (iv) neither object recognition nor phonological retrieval. This design allows the effect that one variable has on the expression of the other variable to be measured explicitly. By convention, the analysis of factorial designs involves calculating the main effects of each variable and the interaction between them. The main effects in activation studies identify the brain areas where there is more activation in the sum of the activation tasks than in the sum of the baseline tasks (eg tasks with phonological retrieval - tasks without phonological retrieval and tasks with object recognition - tasks without object recognition). The interaction between variables identifies areas where the effect of one variable varies depending on the presence or absence of another variable.

Factorial designs have several important advantages over simple subtraction designs. Firstly, they allow greater generalizability of the results because the level effects can be specified for each factor (as in pure subtraction) or generalised for all factors. Secondly, and most importantly, when the effect of one factor level varies according to the level of another factor, factorial designs allow us to verify the significance of this difference with the interaction term (see Friston et al., in press).

Cognitive conjunctions

Although cognitive conjunction designs resemble those of factorial designs in so far as each activation task has its own baseline task, there are also distinct differences in the design and analysis of these approaches. When the experimental design is factorial, there are two or more variables and the different levels of each variable must be matched across task pairs (see above). In contrast, the different levels of each task do not need to be matched in cognitive conjunction designs and the activation and baseline tasks are not constrained to differ by a single cognitive component. The baselines could, theoretically, all involve the same task (e.g. a rest condition for each activation condition) providing that the difference between any two tasks, in cognitive terms, includes the component of interest. In terms of analysis, the main effect in a factorial design includes the sum of task pair differences irrespective of whether there is differential activation between task pairs. A conjunction analysis varies from this convention by identifying areas where there is a significant main effect in the absence of an interaction. In other words, each task pair shows activation and those areas in which task pair effects are not significantly different are identified.

By excluding areas that differentially activate, even though they may be significantly

activated in each task pair, the analysis of cognitive conjunctions is more conservative or restrictive than that of cognitive subtractions and factorial designs. The conservative approach of cognitive conjunction, however, enables us to distinguish areas that are functionally segregated from areas that are not. For instance, if the activation of an area by phonological retrieval varies in the presence or absence of colour (ie there is a significant interaction between colour naming and other naming tasks), this would imply that the area was specialised for both phonological retrieval and the integration of colour rather than being dedicated to phonological retrieval per se. On the other hand if an area responds to, and only to, phonological retrieval (ie it is unaffected by the type of naming task) then a true conjunction will ensue implying that the area is specialised for phonological retrieval irrespective of other processing requirements. In summary, the conjunction of activations that are not significantly different, defines functionally segregated areas whereas simply demonstrating common activations (in the main effect) implies specialisation but only in relation to some other processes or system. On the basis of this argument, it can be seen that a conjunction analysis identifies a very specific sort of functional attribution that can be powerfully complimented by an analysis of interaction effects where one can treat the task pairs as a factorial experiment (see Friston et al. in press).

Another important aspect of cognitive conjunctions is that, like factorial designs, they do not depend on ‘pure insertion’. In general, the difference between an activation and baseline task pair involves not only the extra cognitive component with the activation task but also the interaction of this added component with processes shared in the baseline task. Cognitive subtraction assumes these interactions do not exist, factorial designs measure the interaction term explicitly and cognitive conjunctions discount interaction terms whether they exist or not (because the interactions are unique to each task pair or can be made so by experimental design).

The experiment we have used to illustrate the design and analysis of cognitive conjunction studies investigates the brain regions activated when subjects retrieve the phonology of visually presented stimuli. The activation tasks involved naming words, letters, objects and colours and each of these tasks had a corresponding baseline that did not involve explicit naming. The processing differences between task pairs varied but the one consistent difference between pairs was phonological retrieval. The conjunction of activations for each task pair was identified by the conjoint testing of several hypotheses. Areas activated equivocally by each task pair were the left posterior basal temporal lobe (Brodmann’s Area 37), the left frontal operculum and midline cerebellum. Lesion studies have shown that damage to any of these areas impairs naming (Mesulam 1990). Using a cognitive conjunction study, we have demonstrated that these areas are functionally segregated for phonological retrieval from visually presented stimuli.

Previous studies that have used cognitive subtraction to identify areas involved in phonological retrieval have been hindered by the difficulty of finding task pairs that differ only in terms of phonological retrieval. For instance, any task that presents word like stimuli to control for

orthographic and semantic processing also activates phonological processes implicitly. Using a cognitive conjunction approach, orthographic and semantic processing do not have to be controlled in the word task pair because the areas associated with phonological retrieval are identified by finding the processing areas that are shared by reading and other naming tasks.

Conclusion

In this paper we have presented a simple and accessible variant on experimental design and data analysis of functional imaging data that facilitates the identification of brain systems implementing specific cognitive components. This approach is based upon, but less constrained than, cognitive subtraction and should provide for a greater latitude of experimental design and possible retrospective analysis of previously reported studies. We hope that this reasonably simple embellishment of existing experimental and analysis strategies will further refine our understanding of functional specialisation and segregation in the brain.

Acknowledgements

CJP and KJF were funded by the Wellcome Trust. We would like to thank all our colleagues for invaluable discussion particularly Richard Frackowiak. We would also like to thank Caroline Moore for her help acquiring the data and John Herries for his help with the figures.

Appendix

In this appendix we present the details of how we construct SPMs to test for the conjunction of two or more hypotheses. Although there are a number of ways in which one could test for the conjoint expression of two or more effects, we have the special problem of formulating such a test so that it can be used in the context of statistical parametric mapping (and implicitly the theory of Gaussian fields). In brief the solution we have adopted consists of creating a $\text{SPM}\{Z\}$ that reflects the sum of all the effects one is interested in and then eliminate regions where there are significant differences among these effects. This second step confers the essence of a conjunction: For example; consider two effects evidenced by high values of z_1 and z_2 . The sum of these numbers would be an appropriate statistic for the assessment of the first *or* second hypothesis because either a large z_1 or z_2 can give a high value of $z_1 + z_2$. However a conjunction requires the first *and* second hypotheses to be true. This is the case if z_1 and z_2 are high and are not significantly different. In terms of factorial designs this can be construed as identifying main effects in the absence of an interaction. It should be noted that conjunctions discount activations that are significantly different even if they are all significant in their own right.

By taking appropriate linear combinations of the SPMs pertaining to each effect (i.e. component $\text{SPM}\{z_i\}$) it is possible to ensure that the linear combination representing the sum is orthogonal to the differences. Because the test uses $\text{SPM}\{Z\}$ s we can still use the standard inference procedures developed for statistical parametric mapping. Furthermore because of the above orthogonality, or independence, the elimination of regions where significant differences are observed can be used to reduce the search volume, rendering the correction for multiple comparisons less severe, and the analysis more sensitive.

In what follows we present the mathematical details of the approach for the general problem of testing for the conjunction of N effects. The problem can be formulated as follows: In the context of the general linear model the N hypotheses can be specified in terms of a set of contrasts (e.g. a matrix \mathbf{C} with one contrast per row i.e. $\mathbf{C} = [\mathbf{c}_1; \mathbf{c}_2; \dots; \mathbf{c}_N]$) that specify [normalized] linear compounds $\mathbf{z} = [z_1, z_2, \dots, z_N]$ of parameter estimates (e.g. task condition means) where the compounds are weighted by the elements in the row vectors $\mathbf{c}_1, \mathbf{c}_2, \dots$. We now wish to construct a SPM that reflects the conjoint expression of the effects specified by $\mathbf{c}_1, \mathbf{c}_2, \dots$ or, in other words, when z_1, z_2, \dots have jointly high values.

The general linear model, for one voxel, can be written in matrix format as:

$$\mathbf{x} = \mathbf{G}\boldsymbol{\beta} + \mathbf{r}$$

where \mathbf{x} is the response variable (e.g. rCBF); a column vector with one element for each scan. \mathbf{G} is the design matrix modelling the effects, with one effect in each column and one row for each scan. The parameter column vector \mathbf{b} contains one element for each effect in \mathbf{G} . \mathbf{r} is a column vector of identically and independently Gaussian residuals with variance σ^2 . The parameter estimates $\hat{\mathbf{b}}$ of \mathbf{b} are obtained as usual using standard least squares.

The first step is to construct a set of Z fields [i.e. $\text{SPM}\{z_i\}$], one for each contrast. An approximation to these fields is obtained using the linear compound $\mathbf{c}_i \cdot \hat{\mathbf{b}}$ divided by its standard deviation σ_j under the null hypothesis:

$$z_i = \mathbf{c}_i \cdot \hat{\mathbf{b}} / \sigma_j$$

where $\sigma_j^2 = \sigma^2 \mathbf{c}_i \cdot (\mathbf{G}^T \mathbf{G})^{-1} \cdot \mathbf{c}_i^T$

and σ^2 is estimated with $R(\epsilon) / r$. $R(\epsilon)$ are the sum of squares due to error, at the voxel in question, and r corresponds to the degrees of freedom due to error [number of scans - rank(\mathbf{G})]. In what follows we assume that r is sufficiently large that the z_i can be described with a unit normal distribution $N(0,1)$ (otherwise these statistics have the Student's t distribution). z_i are computed for every voxel giving an SPM for each of the N contrasts. These component $\text{SPM}\{z_i\}$ has a smoothness characterized by $W = |\Lambda|^{-1/2D}$. Λ is the covariance matrix of the $\text{SPM}\{z_i\}$'s derivatives and D is the dimensionality.

The next step is to combine these $\text{SPM}\{z_i\}$ into a single SPM that reflects the conjunction of all the N hypotheses. This is achieved by creating a new SPM that is the [normalized] sum of all the component $\text{SPM}\{z_i\}$ and eliminating regions where there are significant differences among the z_i . This is effected in such a way that for any number of component $\text{SPM}\{z_i\}$ the elimination is independent of the SPM representing the sum. This allows us to subtract the regions eliminated from the search volume when making a correction to the final p values (see Friston *et al* 1995). For two hypotheses this is easy because the sum of two z_i is independent of the differences (even if they are correlated).

Recall $\mathbf{z} = [z_1, z_2, \dots]$. Under the null hypothesis \mathbf{z} has a multivariate normal distribution $N(0, \mathbf{L})$ with covariance:

$$\mathbf{L} = \text{diag}(\mathbf{J})^{-1/2} \cdot \mathbf{J} \cdot \text{diag}(\mathbf{J})^{-1/2}$$

where $\mathbf{J} = \mathbf{C} \cdot (\mathbf{G}^T \mathbf{G})^{-1} \cdot \mathbf{C}^T$

The sum of z_i can be thought of as the 'projection' of the point defined by \mathbf{z} on plotting the z_1, z_2, \dots against each other onto the $z_1 = z_2 = \dots = z_N$ line. This iso-effect line will be referred to as the *conjunction line* (the white line in Figure 7). The differences among z_1, z_2, \dots are reflected in the distance of the point \mathbf{z} from this line. In order to ensure that the sum (distance along the conjunction lines) and differences (distance from the conjunction line) are independent we first transform the component z scores with a matrix \mathbf{h} such that the new component z scores ($\mathbf{z}^* = \mathbf{z} \cdot \mathbf{h}$) are independently and identically distributed. I.e. $\mathbf{h}^T \mathbf{L} \mathbf{h} = \mathbf{1}$. \mathbf{h} is simply determined using the eigenvector solution of \mathbf{L} where $\mathbf{L} \cdot \mathbf{e} = \mathbf{e} \cdot \lambda$ and $\mathbf{h} = \mathbf{e} \cdot \lambda^{-1/2}$. \mathbf{e} is a unitary orthogonal matrix of eigenvectors and λ is a diagonal matrix of eigenvalues. Let $\mathbf{s} = (\mathbf{e} \cdot \lambda^{-1/2})^{-1} \cdot [1, 1, \dots, 1]^T$ and $\mathbf{s}^* = \mathbf{s} / \mathbf{s}^T \cdot \mathbf{s}$ then $\mathbf{z}^* \cdot \mathbf{s}^* = (z_1 + z_2 + \dots + z_N) / \mathbf{s}^T \cdot \mathbf{s} = Z$, the [normalized] sum. The squared distance from the conjunction line is $D = \mathbf{z}^* \cdot \mathbf{z}^{*T} - Z^2$. $\mathbf{z}^* \cdot \mathbf{z}^{*T}$ represents sum of squares of new z scores \mathbf{z}^* .

The weighted sum $Z = \mathbf{z}^* \cdot \mathbf{s}^*$ has the unit Gaussian distribution under the null hypothesis and a high value of Z provides evidence for one or more of the [alternative] hypotheses defined by the N contrasts. A SPM of this statistic $\text{SPM}\{Z\}$ has the same smoothness W as the component $\text{SPM}\{z_i\}$. The statistic D has a chi-squared distribution with degrees of freedom $N - 1$, under the null hypothesis. This means that we can eliminate voxels that show a significant deviation from the conjunction line according to $D > \chi^2(N - 1)_{0.05}$. After these voxels are removed from $\text{SPM}\{Z\}$, it is subject to the usual inference procedures to give corrected and uncorrected p values pertaining to voxels, clusters of voxels or sets of clusters (Friston *et al* 1994).

In this the final section we present some simulation results undertaken to ensure the validity of the above procedure. We have chosen to use the probability distribution of the size of clusters as a validating measure. This distribution depends on the expected number of maxima, which, in turn, is the basis of tests based on peak height. Inference based on spatial extent is predicated on the distribution examined below. The theoretical distribution is given by Eq (11) in Friston *et al* (1994). We simulated Gaussian SPMs by convolving $64 \times 64 \times 64$ random Gaussian fields with a Gaussian kernel corresponding to $W = 2, 2$. We then took pairs of simulated component SPMs and, by taking appropriate linear combinations, rendered then correlated (0.1). Each pair was subject to the conjunction analysis as described above. These simulated conjunction SPMs were thresholded with $u = 3.2$ and the size of the ensuing clusters recorded. We continued simulating SPMs in this fashion until we obtained at least 10^4 clusters. The empirical distribution of cluster sizes was compared with the theoretical results. The expected and observed distributions were in good agreement (Figure 8).

References

- Friston KJ Worsley KJ Frackowiak RSJ Mazziotta JC & Evans AC (1994) Assessing the significance of focal activations using their spatial extent *Human Brain Mapping* 1:214-220
- Friston, K.J., Holmes, A. Worsley, K.J., Poline, J.B., 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., Ashburner, J., Poline, J.B., Frith, C.D., Heather, J.D. & Frackowiak, R.S.J. (1996) Spatial realignment and normalization of images. *Hum Brain Mapp* .
- Friston KJ, Price CJ, Fletcher P, Moore C, Frackowiak RSJ, Dolan RJ (1996) The trouble with cognitive subtraction. *Neuroimage* (in press)
- Hasofer AM (1978) Upcrossings of random fields. *Suppl. Adv Appl. Prob.* 10:14-21
- Mesulam MM (1990) Large scale neurocognitive networks and distributed processing for attention, language and memory. *Annals of Neurology* 28: 597-613.
- Petersen SE, Fox PT, Snyder AZ, Raichle ME (1990) Activation of extrastriate and frontal cortical areas by visual words and wordlike stimuli. *Science*, 249: 1041-1044.
- Price, C.J., Wise, R.J.S. & Frackowiak, R.S.J. 1996a Demonstrating the implicit processing of visually presented words and pseudowords. *Cerebral Cortex* 6, 62-70.

Figure Legends

Figure 1: Figure 1a represents a cognitive subtraction hierarchy. PI is the process of interest, A is the activation task and B is the baseline task. Figure 1b represents a cognitive conjunction design which has two task pairs (I and II) each with an activation (A) and baseline (B) task. P1, P2, P3 and P4 are distinct but arbitrary task components.

Figure 2: A cognitive model of the processing components involved in reading, letter naming, object naming and colour naming.

Figure 3: Illustrations of the stimuli used in the experiment.

Figure 4: A graphic task analysis of the constituent processing components of the experimental conditions. The format corresponds to that adopted in Figure 1 but the constituent processing components used are those illustrated in the cognitive model of Figure 2. The grey filled regions indicate task components that distinguish between activation and baseline tasks, the black filled regions indicate where these differences were common for every pair.

Figure 5: The SPM illustrating the maximum intensity projection of the Z statistic for the conjunction of the task pair differences.

Figure 6: The adjusted mean activity for all 8 conditions in the regions identified by the conjunction analysis.

Figure 7: Acceptance areas for conjunctions for two hypotheses (upper panel) and three hypotheses (lower panel). Note that the lower panel represents a section through a three dimensional space and that the acceptance area is in reality a cylinder centred on the conjunction line.

Figure 8: Expected (solid line) and observed (broken line) distribution of cluster sizes of supra-threshold clusters in simulated conjunction SPMs. See main text for the details of the simulations.