ELSEVIER

NeuroImage

www.elsevier.com/locate/ynimg
NeuroImage 25 (2005) 661–667

Technical Note

Conjunction revisited

Karl J. Friston, William D. Penny, and Daniel E. Glaser*

The Wellcome Department of Imaging Neuroscience and Institute of Cognitive Neuroscience, University College London, 17 Queen Square, London, WC1N 3AR, UK

Received 23 August 2004; revised 23 December 2004; accepted 12 January 2005
Available online 3 March 2005

The aim of this note is to revisit the analysis of conjunctions in imaging data. We review some conceptual issues that have emerged from recent discussion (Nichols, T., Brett, M., Andersson, J., Wager, T., Poline, J.-B., 2004. Valid Conjunction Inference with the Minimum Statistic.) and reformulate the conjunction of null hypotheses as a conjunction of k or more effects. Analyses based on minimum statistics have typically used the null hypothesis that $k = 0$. This enables inferences about one or more effects ($k > 0$). However, this does not provide control over false-positive rates (FPR) for inferences about a conjunction of $k = n$ effects, over n tests. This is the key point made by Nichols et al., who suggest a procedure based on supremum P values that provides an upper bound on FPR for $k = n$. Although valid, this is a very conservative procedure, particularly in the context of multiple comparisons. We suggest that an inference on a conjunction of $k = n$ effects is generally unnecessary and distinguish between congruent contrasts that test for the same treatment and incongruent contrasts of the sort used in cognitive conjunctions. For congruent contrasts, the usual inference, $k > 0$, is sufficient. With incongruent contrasts it is sufficient to infer a conjunction of $k > u$ effects, where u is the number of contrasts that share some uninteresting effect. The issues highlighted by Nichols et al., have important implications for the design and analysis of cognitive conjunction studies and have motivated a change to the SPM software, that affords a test for the more general hypothesis $k > u$. This more general conjunction test is described.

Introduction

The central distinction, which has been highlighted by recent discussions, is whether conjunction refers to the activation (i.e., consistently large activation) or the underlying effects (i.e., consistently significant activation). Activation is an attribute of the data, usually defined through a statistic. An effect is an attribute of the real world, which we cannot observe. Declaring a voxel to be 'activated' allows one to infer the effect is present with some sensitivity and specificity. The distinction is formulated in Nichols et al. (2004) in terms of a *global null hypothesis* and a *conjunction null hypothesis* about effects. SPM tests the global null using the minimum T statistic. A test for the conjunction null, proposed by Nichols et al. (2004), has the same form but uses a much higher threshold. Put simply, the difference rests on whether the conjunction refers to the observed statistics or to the effects one is trying to infer. Both approaches are valid but have different uses. Conjunction has been clearly defined as "the joint refutation of multiple null hypotheses" (Friston et al., 1999). In other words, a conjunction of activations allows one to infer a conjunction of *one or more* effects. However, Nichols et al. have pointed to examples in the literature where the inference is misinterpreted as a conjunction of *all* effects. This was their motivation for highlighting the issue and proposing a new test.

In what follows, we make three points. First, conjunctions based on the global null remain valid and exact. Second, although valid, the alternative proposed in Nichols et al. (2004) is conservative. In fact, it is often sufficiently conservative to render it powerless in neuroimaging. This is especially pronounced when considering the multiple comparisons problem. The third point is that rejection of the conjunction null, although sufficient, is usually unnecessary. We discuss this separately for congruent and incongruent contrasts, testing the same and different treatments, respectively. When the contrasts test the same effect, one can assume a binomial prior on the number of effects. This was the starting point for the meta-analysis presented in Friston et al. (1999) and connects the current analysis with previous work. The considerations for cognitive conjunctions are more complicated and are usefully informed by Nichols et al. (2004). The key revision here is that tests for a number $k > u$ of effects may be called for, depending on the conjunction design and assumptions about the regional deployment of treatment effects. We conclude with a section on how conjunctions are specified in the next release of the SPM software that allows one to test for $k > u$ effects. This more general specification subsumes tests of the global null $k = 0$ and tests of the conjunction null $k > n - 1$.

* Corresponding author. Fax: +44 20 7916 8517.
E-mail address: daniel.glaser@ucl.ac.uk (D.E. Glaser).
Available online on ScienceDirect (www.sciencedirect.com).

1053-8119/$ - see front matter © 2005 Elsevier Inc. All rights reserved.
doi:10.1016/j.neuroimage.2005.01.013

Definitions

It is important to be clear on basic terminology used in this discussion. Consider a test which is supposed to give a certain false-positive rate α. If it gives exactly as many false-positives as expected it is called an exact test. If it gives more false-positives than expected, it is invalid. If it gives fewer false-positives than expected, it is still valid but conservative.

Activations and effects

In classical imaging statistics, one declares a voxel to be activated if its statistic exceeds some threshold. This statistic reflects the likelihood of the effect being truly present as opposed to being absent. Activation is therefore an attribute of the data. By declaring a voxel activated, we infer the effect is present in a probabilistic sense. The effect will be absent in some proportion of activated voxels. The relationship between activation and effect is best characterised in terms of conditional probabilities, namely the specificity $1 - \alpha$ and sensitivity β and their complements false-positive and negative rates. See the upper panel of Fig. 1. The threshold is chosen to ensure the false-positive rate $FPR = \alpha$ is small.

A conjunction of activations can be related to conjunctions of effects in exactly the same way. See lower panel of Fig. 1. Here, we have a variable k representing the number of effects that are truly present in a conjunction of n contrasts. The lower row of the probability table in Fig. 1 represents $p(n|k)$ the probability of a conjunction conditional on there being k effects. These can be used to specify the FPR for any null hypothesis.

Global and conjunction null hypotheses

Conventionally, conjunction analyses are based on the global null hypothesis that there are no activations $k = 0$. According to Fig. 1, for a single test $FPR = \alpha^n$. Therefore, a significant conjunction allows one to say with a specificity of $1 - \alpha^n$ that there is a conjunction of one or more effects (i.e., $k > 0$). However, this does not mean that all the effects are present. To make an inference that there is a conjunction of all effects (i.e., $k = n$), one has to include all the alternative outcomes under the conjunction null $k < n$. This is the basis of an alternative procedure, advocated in Nichols et al. (2004). The difference between the two procedures is based on which outcomes the null hypotheses encompass (see Fig. 1). Below, we deal briefly with implementation and issues surrounding both these null hypotheses.

Testing the global null

If the objective of conjunction analyses is to reject the global null, why not use an F test spanning the contrasts in the conjunction? The answer is that conjunctions allow one to focus on a specific departure from the global null; namely, outcomes that are consistent. This increases sensitivity markedly, when the effects are consistent, as depicted schematically in the upper panel of Fig. 2. Consider two contrasts. The bivariate distribution of two T values, under the global null, is shown in the upper panels (for eight degrees of freedom). The false-positive rate corresponds to the integral of this density over some region. If the two T values fall in this *rejection region*, we reject the global null. The deployment of this region defines the departure from the null hypothesis that is considered interesting. For example, an F test considers any departure interesting, as reflected by the circular region surrounding the null distribution on the upper left of Fig. 2. This is a large region and therefore the threshold (radius of the circle) has to be high to maintain a low FPR. The rejection region corresponding to the conjunction is the upper [hyper-]quadrant, in which all the T values are greater than some minimum T value. This is a more restricted

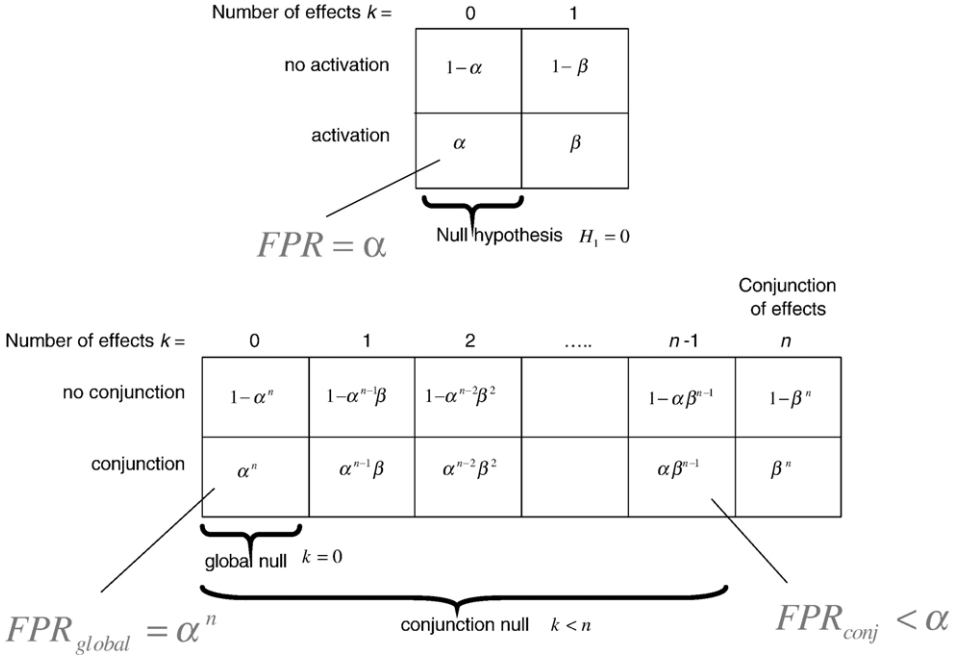

Fig. 1. Conditional probabilities for activation and their conjunction given the number of true effects. These probabilities are for a single voxel with known specificity and sensitivity. For simplicity, we have assumed that sensitivity is the same for all effects, and that they are independent.

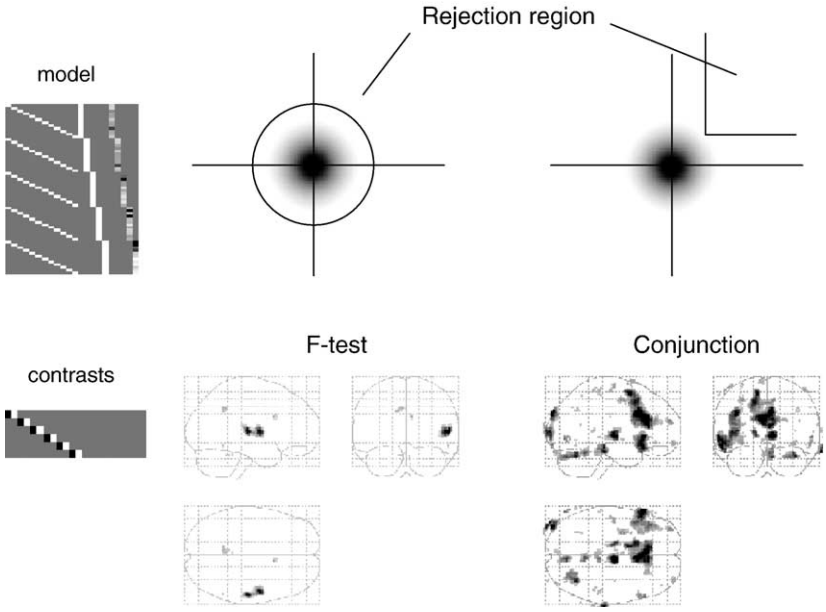

Fig. 2. Schematic comparing conjunction analyses with those based on conventional F tests. The upper row shows bivariate T distributions under the global null hypothesis and the rejection regions associated with each of the two tests. If the two T values fall in these regions, one can reject the null hypothesis. Lower left: the contrasts used to specify the tests. Lower row: the ensuing SPMs based on a verbal fluency PET data set. The SPMs have been thresholded at $P = 0.001$ (uncorrected).

region and the thresholds can be more relaxed to maintain the same specificity.

Heuristically, conjunctions generalise one-sided t tests to multiple dimensions. For a single contrast, a two-tailed t test is the same as an F test, and tests for both positive and negative effects. A one-tailed t test, with the same specificity, uses a lower threshold and has a higher sensitivity to positive [or negative] effects. A conjunction of t tests is the multidimensional equivalent of the one-sided t test and allows one to prescribe an increase in sensitivity, to consistent effects, at the expense of missing inconsistent responses (i.e., large positive and negative). In short, testing the global null is like performing a one-tailed F test.

Like F tests, conjunctions are simply inference devices that allow one to test a multidimensional hypothesis, specified by a contrast weight matrix. The only difference is that the conjunction is only sensitive to consistent effects. Conjunction analyses, using the minimum T statistic, are useful when one knows, a priori, the direction of the effect. For example, the search for bilateral effects in voxel-based-morphometry is an established use of conjunctions (Belton et al., 2003; Salmond et al., 2000). Because of their plastic potential, children seldom develop severe neuropsychological deficits unless homologous regions in both hemispheres are damaged. A conjunction of effects (loss of grey matter density) in homologous voxels has therefore been used to detect anatomical correlates in developmental neuropsychology. Here, the contrasts testing for changes in both hemispheres are orthogonal and test for the same signed effect. In these analyses, finding a significant region allows one to infer regional grey matter loss in one or both hemispheres.

Conjunctions are not as sensitive as a single contrast testing for the average effect over all contrasts (by the Neyman–Pearson Lemma). However, rejecting a single hypothesis about the average is not equivalent to rejecting a multidimensional hypothesis (i.e., multiple null hypotheses) with a conjunction analysis. This is because inferring the average effect is greater than zero is not equivalent to a conjunction. In the example above, declaring a significant reduction in grey matter density in one or both hemispheres is not the same as saying the loss, averaged over both hemispheres is significant. The latter could occur with a profound decrease in one hemisphere and an increase in the other. A conjunction analysis would not find this inconsistent effect.

The nature of inference, afforded by conjunctions, is central to the issues addressed here, particularly with the conventional use of conjunctions to test the global null. These inferences mean the evidence for consistent effects is significant, not the evidence for significant effects is consistent. This distinction may be semantic but speaks directly to people's misconception about conjunctions that Nichols et al. (2004) have highlighted. In the example above, a significant conjunction does not mean that the right hemisphere has lost grey matter and the left hemisphere has lost grey matter. To infer this would require tests of each hypothesis separately (c.f. reporting post hoc contrasts after finding a significant effect with an F test). This separate hypothesis testing is essentially what Nichols et al. are proposing. In short, a significant conjunction is not a conjunction of significances.

Minimum statistic tests

The fact that conjunctions can be formulated in terms of a minimum test statistic is useful because the null distribution of the minimum T statistic can be computed analytically. Furthermore, there are analytic expressions for the maximum of this minimum T statistic over a spatial search volume based on random field theory (Worsley and Friston, 2000). These expressions can be used to provide a FPR that is adjusted for the search volume in the usual way.

The lower panels of Fig. 2 highlight the relative power of conjunction analyses, using the minimum T statistic. The SPMs

were based on the verbal fluency data used to illustrate SPM procedures over the years. They were acquired from 5 subjects each responding to a heard letter by repeating the letter or producing a word that started with that letter. Each of the two conditions were repeated six times. The SPM$\{F\}$ on the left shows the F statistic testing for a condition-specific effect in any (one or more) of the six replications. The SPM$\{\min(T)\}$ shows the minimum T values over the six replications. Both SPMs are testing the global null hypothesis and all are thresholded to give the same false-positive rate ($P = 0.001$ uncorrected). The key thing to note is that the SPM$\{\min(T)\}$ discloses more significant voxels than the conventional SPM$\{F\}$. In fact, the SPM$\{F\}$ found no activations at all (all the voxels correspond to deactivations). The reason that the conjunction SPM$\{\min(T)\}$ is more powerful is that the treatments, in each of the six contrasts, were congruent. This was assured by experimental design in which the same treatment was delivered six times.

Congruent vs. incongruent effects

At this point, it is worth introducing a distinction between contrasts that test the same thing and those that test different treatments. We will refer to these as congruent and incongruent contrasts. The verbal fluency example above used congruent contrasts in the sense that the treatment was the same over each replication. Another common example of congruent contrasts would be contrasts testing for the same effect over subjects. In the context of congruent contrasts, it is sufficient to test the global null because its rejection allows us to infer that this treatment effect was detected on one or more occasions. It does not matter if the effect was not detected in some contrasts; a significant effect has been demonstrated. The analysis of congruent contrasts was the subject of Friston et al. (1999) and will be reprised below.

Incongruent contrasts (e.g., a contrast for object naming and a contrast for word naming) are more problematic. In this instance, it may be relevant that the effect was absent in one of the contrasts. Incongruent contrasts were the focus of cognitive conjunctions (Price and Friston, 1997), the original motivation for conjunction analyses. The idea here was to demonstrate region-specific correlates of a cognitive component that was common to a set of incongruent contrasts. Nichols et al. (2004) note that rejection of the conjunction null, rather than the global null, is indicated in this context.

Testing the conjunction null

If the objective is to infer a conjunction of effects, then it should be sufficient to test each contrast separately and establish they are all significant. This is precisely what Nichols et al. (2004) conclude, although their derivation is a little more involved: in rejecting the conjunction null one has to control the false-positive rate over all outcomes that constitute that null. It may be useful to consider a distribution over the number of effects.; c.f. a Bayesian perspective, where we consider k as a random variable with prior distribution $p(k)$. From Fig. 1, this is:

$$FPR_{conj} = \sum_{k=0}^{n-1} p(n|k)p(k)$$

$$= \sum_{k=0}^{n-1} \alpha^{n-i}\beta^{i}p(k) \tag{1}$$

The problem here is that we do not know $p(k)$. However, we can establish an upper bound by noting

$$FPR_{conj} < \sup_{i} \alpha^{n-i}\beta^{i} = \alpha \tag{2}$$

This means that if we set the specificity of each test to some suitably small value, we can be assured that the FPR is controlled for inferences about a conjunction of effects. This is exactly the same as showing each contrast is significant in its own right.

Practically, the P value, or false-positive rate for the global null is the probability of obtaining the minimum T value by chance $p(\min(T_i)|k = 0)$. The corresponding FPR for the conjunction null is the supremum P value over contrasts $\max(p(T_i)|k = 0))$. From now on, we will refer to the two approaches as the minimum statistic and supremum P value approaches.

Fig. 3 tries to illustrate the heuristic behind the supremum P value procedure. Again, consider two incongruent contrasts. Because we want to infer both effects are present, the null includes the situation where only one is present. The ensuing distribution of T statistics is shown on the right. Clearly, to control FPR, in this worst case scenario, the threshold adopted must be more conservative than for the global null (left panel).

An over-conservative test

Although valid the alternative procedure is very conservative. This is because Eq. (2) provides only an upper bound on the FPR.

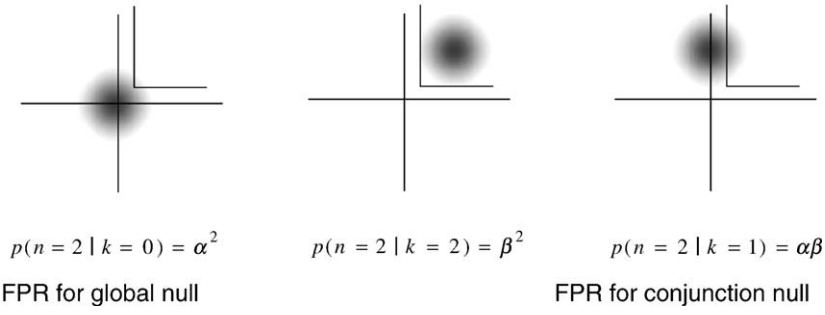

Fig. 3. Schematic illustrating departure from the global null. Note that the density of the T statistics (shaded region) encroaches on the rejection region (bounded by the dotted lines) when one of the effects is present (right panel). This is a violation of the global null but not of the conjunction null.

The actual FPR may be much smaller, depending on k. This can be seen easily with the following example: assume we have two incongruent contrasts (one for object-naming and one for word-naming). We construct an SPM{min(T)} and identify the most significant voxel as being in the left occipito-temporal region. The minimum T value at this voxel was 2.4. The supremum P value adjusted for the search volume was $\max(p(T_i|k = 0)) = p\ (T \geq 2.4|k = 0) = 0.99$ after correction for multiple comparisons. We have therefore failed to show a conjunction of effects. We now give our data to a colleague who has never heard of the supremum procedure. He examines the contrast for object-naming at a corrected level and finds a significant activation with a T value of 4.8 with an adjusted $p(T \geq 4.8|k = 0) = 0.03$ in the same region. To see if word-naming activates this region, he examines the second contrast, searching over a sphere of 8 mm radius, centred on the maximum of 4.8. He then identifies our original voxel that now has a P value of $p(T \geq 2.4|k = 0) = 0.02$ adjusted for the small search volume. He concludes, properly, that both object- and word-naming cause effects in this region. Where did we go wrong?

The problem is that the supremum P value assumes a worst case scenario to provide the upper bound FPR. In the context of neuroimaging, this is conservative because the FPR is based on the null hypothesis that every voxel in the entire search volume expresses an effect in all but one contrast. Not only is this supremum very conservative, but it includes a null hypothesis that cannot occur. It cannot occur because region-specific effects cannot, by definition, exist everywhere. In other words, if all the brain activated, there would be no region-specific response. This is why whole-brain activation is treated as a confounding effect in global normalisation procedures.

What do we want to test?

Congruent contrasts

The difficulty with the supremum P value approach is that we do not know the probability distribution of the number of effects $p(k)$, therefore, we assume the worst case

$$p(k) = \begin{cases} = 1 & k = n - 1 \\ = 0 & k \neq n - 1 \end{cases} \quad (3)$$

The problem is that this distribution could not be realised by any plausible generative model. For example, suppose the treatment tested by some congruent contrasts produced an effect with a frequency or probability γ. The ensuing distribution of effects would have a binomial distribution

$$p(k) = \binom{n}{k}(1-\gamma)^k\gamma^{n-k} \quad (4)$$

There is no value of γ that gives a distribution conforming to the null in Eq. (3). So what is the null distribution for k? In the case of congruent contrasts, we want to know whether our treatment produces an effect with non-zero probability. Therefore, the null distribution obtains when $\gamma = 0$. In this case

$$p(k) = \begin{cases} = 1 & k = 0 \\ = 0 & k \neq 0 \end{cases} \quad (5)$$

This is simply the global null as used conventionally. In short, it is entirely sufficient to use the minimum statistic test to reject $k = 0$ to infer $\gamma > 0$. In some cases, it may be interesting to supplement the inference quantitatively, with a confidence interval on γ. This was described in some detail in Friston et al. (1999). However, in practice, these confidence intervals have not been much used. Note that, for congruent contrasts, we do not need to reject the conjunction null that $k < n$.

Incongruent contrasts and cognitive conjunctions

In the case of incongruent contrasts, the situation is more complicated. Here, the treatments differ and γ will be treatment specific. However, there are constraints which allow us to define the null hypothesis. For example, the aim of conjunction analysis is to identify responses to a common treatment (e.g., a cognitive component of interest or CCI). This common component speaks to the fact that the treatments are compound, with unique components and common components. The logic of cognitive conjunctions is that functionally specialised brain regions respond selectively to one component. The aim is to find the brain region that responds to the common component. Therefore, any voxel will respond to the common component, to a unique component or to no components. The null outcomes here are responses to a unique component $k = 1$ or no component $k = 0$. It is therefore sufficient to infer $k > 1$. Because we do not know the outcome probabilities $p(k)$, we can resort to the supremum approach of Nichols et al. (2004). This entails assuming the worst case scenario that every region responds to a unique component with probability one, that is, $k = 1$. For a single voxel this ensures $FPR < \alpha^{n-1}$. In the context of statistical parametric mapping, this corresponds to assuming a null distribution for the minimum statistic based on $n-1$ contrasts. Heuristically, we still use the minimum statistic over all n contrasts but assume one of them was highly significant everywhere, a priori.

This approach can be generalised to conjunction designs in which the cognitive components are unique to a subset of u contrasts. The previous paragraph assumed $u = 1$. The example presented in Nichols et al. (2004) used four working memory tasks and a visual task. In this instance, unique components are found in four of the five contrasts. This calls for an inference that $k > u = 4$. This is, of course, a test of the conjunction null that Nichols et al. were promoting. However, this design is very inefficient for finding the responses to common components. A good conjunction design should ensure that u is small.

In short, cognitive conjunctions allow one to test for effects due to common treatment components. In carefully designed conjunction studies, where each treatment comprises a common component and unique components in u or fewer contrasts, it is sufficient to infer that $k > u$ by using the global null distribution of the minimum statistic for n–u contrasts. Again, note that one does not have to reject the conjunction null (unless $u = n - 1$). It is worth noting that defining the components of a treatment, especially in cognitive science, is not always straightforward. Debates about conjunction often rest on the task analysis and the deeper issues usually pertain to interpretation, as opposed to statistics.

Practical implications for analysis

In this section, we describe the changes to SPM that allow one to test the null hypothesis that $k > u$. We then review the analysis of

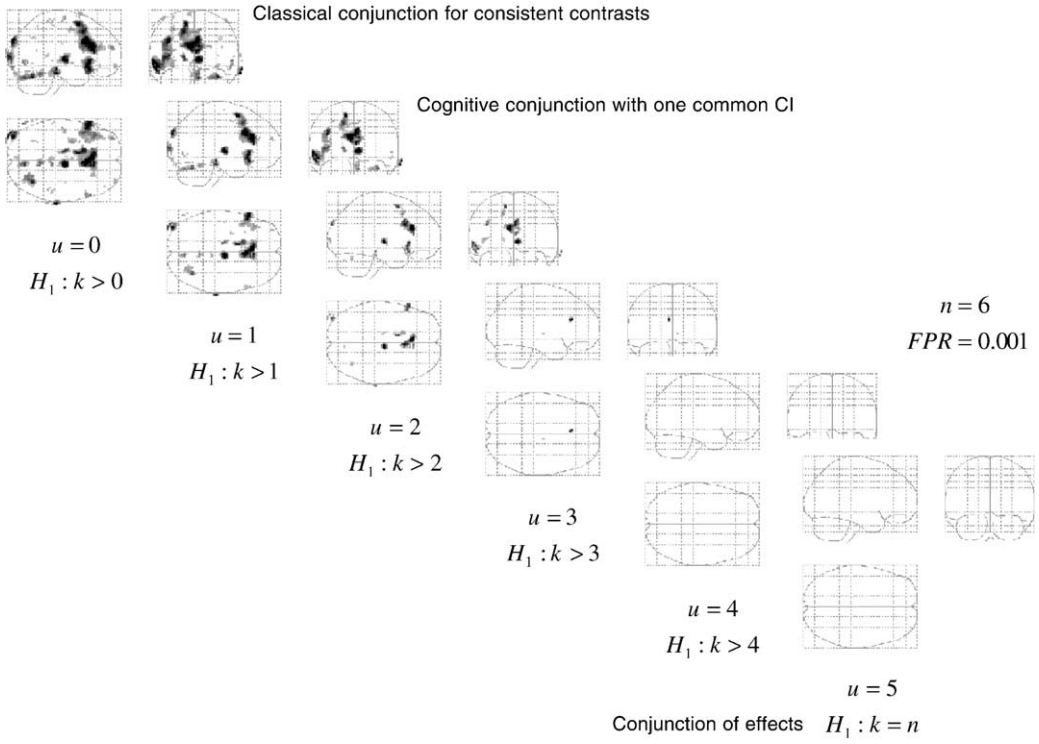

Fig. 4. Conjunction SPMs based on the PET data of Fig. 2. Here the threshold ($P = 0.001$, uncorrected) is based on null distributions for 6, 5, 3, etc contrasts, using the minimum T statistic over all six. These represent tests of the conjunction of $k > 0$, $k > 1$ and $k > 2$ effects, respectively, specified by $u = 0$, $u = 1$ and $u = 2$ (see main text).

congruent and incongruent contrasts. We conclude with comments on how to qualify conjunction analysis, when reporting results.

Changes to SPM

The key contribution of Brett et al. (2004) and Nichols et al. (2004) is to enforce a re-evaluation of inferences required in the context of cognitive conjunctions. This has clarified the constraints on conjunction designs, namely that the components that are not common should be unique to a small number of contrasts, ideally $u = 1$. Furthermore, valid inference with the minimum statistic requires one to infer that $k > u$. This can be effected simply using the null distribution of the minimum statistic under the global null for $n–u$ contrasts. A special case of incongruent contrasts, used in cognitive conjunctions, is when all the treatments are the same and the contrasts are congruent. In this instance, $u = 0$.

To enable tests of conjunctions of $k > u$ effects, SPM will now ask you to specify u after selecting the n contrasts. The prompt is 'number of effects under null'. For a test of the global null, enter '0'. For a test of the conjunction null, enter $n-1$. Notice that the global and conjunction nulls are now both extreme cases of the more general null hypothesis that $k > u$. We will therefore refer to this as the conjunction null from now on. The ensuing SPM is indexed by the number of contrasts assumed for the null distribution of the minimum statistic. For example, SPM$\{T^4_{32}\}$ refers to an SPM whose P values are based on the minimum T statistic with 32 degrees of freedom over 4 contrasts. This would be obtained with 5 contrasts and $u = 1$. The actual changes to the code are trivial and involve the addition of one line that effectively reduces the number of contrasts SPM thinks are in the conjunction $n \leftarrow n-u$. The nice thing about the revision to SPM2 is that the both the conventional and alternative procedures can be implemented within the same framework. A conventional test of $k > 0$ is specified with $u = 0$. The analysis proposed by Nichols et al. (2004) corresponds to making $u = n - 1$.

When $u = 0$ inference is valid and exact. When $u > 0$, one is implicitly invoking a supremum P value test, which provides a conservative upper bound on the false-positive rate. As one might expect, sensitivity falls quickly as the number of effects u grows. This is illustrated in Fig. 4 using the verbal fluency data of Fig. 2. Here we have shown the SPMs for $u = 0, 1, ...$ It can be seen that all voxels disappear by $u = 4$. If we treated this design as congruent, then the test $k > 0$ is sufficient and allows us to declare the verbal fluency treatment significant in voxels that show a conjunction. However, we could pretend that the six contrasts were incongruent with an unspecified contrast-specific component to each treatment. In this instance, we would need to show $k > 1$. The resulting voxels are seen in the second SPM (Fig. 4) and show that this analysis is still more sensitive than a conventional F test (c.f. Fig. 2).

Congruent contrasts

The procedure and inference for congruent contrasts remain unchanged and is based on inferring that $k > 0$. This is an implicit inference that $\gamma > 0$, where γ is the probability that the treatment causes an effect. Operationally, one selects n contrasts and specifies $u = 0$.

In the context of hierarchical observations (e.g., multi-subject studies), one should not "substitute conjunction analyses for random-effect analyses. Where the latter are indicated there is no

alternative" (Friston et al., 1999). The argument in this paper, and in Friston et al. (1999), concerning congruent contrasts, reduces to the following assertion; if all subjects show an effect with probability γ and one or more subjects show an effect (i.e., $k > 0$), then γ must be bigger than zero. However, positing the existence of random effects, by invoking γ, falls short of actually estimating the between-subject variation in activation and making a population-based inference.

Incongruent contrasts

Here, there are two ways to proceed. First, using small volume adjusted P values centred on the maximum of the first contrast as outlined in Testing the conjunction null. Second, one can use a minimum statistic procedure with $u > 0$, where u is the number of contrasts containing unique treatment components. The former approach is exact but requires you to specify an order in which the contrasts enter the conjunction. The second approach is based on a supremum P value and is therefore conservative. It does however retain some sensitivity for small u and is commutative.

Qualifying and reporting

In reporting subsequent conjunction analyses, it might be good practice to describe the inference with something like the following:

We performed a conjunction analysis using SPMs of the minimum T-statistic over n orthogonal contrasts. Inference was based on P-values adjusted for the search volume using random field theory. The null distribution for the minimum statistic was based on $n - u$ statistics. This enabled us to infer a conjunction of $k > u$ effects at significant voxels.

For those people who have used the global null for inferences about cognitive conjunctions, and simply want to qualify their inference. An appropriate passage might be:

It should be noted that our significant conjunction does not mean all the contrasts were individually significant (i.e., a conjunction of significance). It simply means that the contrasts were consistently high and jointly significant. This is equivalent to inferring one or more effects were present.

Conclusion

We hope that this note provides a clear framework for the range of uses of conjunction analyses, and serves as a guide to their interpretation.

Acknowledgments

We would like to acknowledge the patience and constructive input of the SPM co-authors, especially Keith Worsley. This work was funded by the Wellcome Trust and the ICN MRC Co-op Imaging grant.

References

Belton, E., Salmond, C.H., Watkins, K.E., Vargha-Khadem, F., Gadian, D.G., 2003. Bilateral brain abnormalities associated with dominantly inherited verbal and orofacial dyspraxia. Hum. Brain Mapp. 18, 194–200.

Brett, M., Wager, T., Poline, J.-B., Nichols, T., 2004. When is a conjunction not a conjunction. NeuroImage (Abstract 2419.1055001).

Friston, K.J., Holmes, A.P., Price, C.J., Bűchel, C., Worsley, K.J., 1999. Multi-subject fMRI studies of conjunction analyses. NeuroImage 10, 385–396.

Nichols, T., Brett, M., Andersson, J., Wager, T., Poline, J.-B., 2004. Valid Conjunction Inference with the Minimum Statistic. NeuroImage 25, 653–660.

Price, C.J., Friston, K.J., 1997. Cognitive conjunction: a new approach to brain activation experiments. NeuroImage 5, 261–270.

Salmond, C.H., Ashburner, J., Vargha-Khadem, F., Gadian, D.G., Friston, K.J., 2000. Detecting bilateral abnormalities with voxel-based morphometry. Hum. Brain Mapp. 11, 223–232.

Worsley, K.J., Friston, K.J., 2000. A test for a conjunction. Stat. Prob. Lett. 47, 135–140.