ELSEVIER

NeuroImage

www.elsevier.com/locate/ynimg
NeuroImage 38 (2007) 422–438

Population-level inferences for distributed MEG source localization under multiple constraints: Application to face-evoked fields

R.N. Henson,[a,*] J. Mattout,[b] K.D. Singh,[c] G.R. Barnes,[d] A. Hillebrand,[d] and K. Friston[e]

[a]*MRC Cognition and Brain Sciences Unit, 15 Chaucer Road, Cambridge, CB2 2EF, UK*
[b]*INSERM U821-Brain Dynamics and Cognition, Lyon, France*
[c]*CUBRIC and School of Psychology, Cardiff University, Wales, UK*
[d]*Wellcome Trust Laboratory for MEG Studies, Aston University, England, UK*
[e]*Functional Imaging Laboratory, University College London, England, UK*

Received 14 May 2007; revised 10 July 2007; accepted 17 July 2007
Available online 7 August 2007

We address some key issues entailed by population inference about responses evoked in distributed brain systems using magnetoencephalography (MEG). In particular, we look at model selection issues at the within-subject level and feature selection issues at the between-subject level, using responses evoked by intact and scrambled faces around 170 ms (M170). We compared the face validity of subject-specific forward models and their summary statistics in terms of how estimated responses reproduced over subjects. At the within-subject level, we focused on the use of multiple constraints, or priors, for inverting distributed source models. We used restricted maximum likelihood (ReML) estimates of prior covariance components (in both sensor and source space) and show that their relative importance is conserved over subjects. At the between-subject level, we used standard anatomical normalization methods to create posterior probability maps that furnish inference about regionally specific population responses. We used these to compare different summary statistics, namely; (i) whether to test for differences between condition-specific source estimates, or whether to test the source estimate of differences between conditions, and (ii) whether to accommodate differences in source orientation by using signed or unsigned (absolute) estimates of source activity.

Keywords: Magnetoencephalography; Inverse problem; Restricted maximum likelihood

Introduction

This paper is about the analysis of multi-subject MEG data using distributed source estimates. We use a summary statistic approach (Holmes and Friston, 1998) to perform mixed-effects analyses for population inference. This approach passes first level (within-subject) estimates of evoked magnetic responses to a second (between-subject) level for inference. Our focus is on (i) the specification of model components at the first level, in terms of constraints on the distributed source solution, and (ii) the appropriate summary statistic to pass to the second level, given subject differences in source location and orientation.

First level: source estimation

Locating the generators of electro- or magnetoencephalographic data – the "MEG/EEG inverse problem" – is an under-constrained problem (Nunez, 1981). The best one can achieve is the most probable location given a number of constraints. For equivalent current dipole (ECD) solutions, it is assumed that activity can be modeled as a small number of dipoles with no spatial extent; iterative algorithms are then used to move the dipoles around source space (i.e., the brain) until the difference between the predicted and actual data is minimized (e.g., Scherg and Berg, 1991). For ECD solutions, the constraint is the number of dipoles used. For distributed source solutions, electrical activity is estimated at a number of fixed locations within the brain, usually with more locations than there are sensors. For distributed solutions, constraints can be "hard", such as fixing the source locations and orientations based on anatomical information from MRI (e.g., Dale and Sereno, 1993), or "soft", such as favoring solutions with minimal overall current – so-called "minimum-norm" solutions (Hämäläinen and Ilmoniemi, 1984) – or solutions with the greatest smoothness (e.g., Pascual-Marqui et al., 1994).

We previously proposed a minimum-norm method in which multiple soft constraints can be employed, the contributions of which are estimated using restricted maximum likelihood (ReML) (Phillips et al., 2002b). This method employs a hierarchical linear model with Gaussian errors that can be formulated in a "Parametric Empirical Bayes" (PEB) framework (Friston et al., 2002). With only one constraint, the PEB method reduces to the classical "weighted minimum (L2) norm" (WMN) solution (Hauk, 2004). However, the important advantage of ReML is that multiple constraints can be used. The relative weightings of each constraint,

* Corresponding author. Fax: +44 1223 359 062.
E-mail address: rik.henson@mrc-cbu.cam.ac.uk (R.N. Henson).
Available online on ScienceDirect (www.sciencedirect.com).

1053-8119/$ - see front matter.
doi:10.1016/j.neuroimage.2007.07.026

or "hyperparameters", are analogous to the degree of regularization in WMN. These hyperparameters are optimal from a maximum likelihood perspective; indeed, they have been shown to be superior (and less computationally expensive) to the equivalent degree of regularization estimated using conventional "L-curve" techniques (Mattout et al., 2006). Another advantage of the PEB approach is that ReML maximizes the "model log evidence" (also known as the negative free energy). This means that the hyperparameters optimize the model itself; under suitable (non-negative) constraints on the hyperparameters, this optimization can result in some constraints being switched off, when the corresponding hyperparameter estimate is very small. This behavior enables ReML to perform automatic model selection (Friston et al., 2006b) and is particularly useful when one knows neither the ground truth (i.e., true source location), nor which prior assumptions (i.e., constraints) are relevant.

In terms of first level analyses (source reconstruction), this paper extends our previous work by addressing the reproducibility of empirical priors on source reconstruction over subjects. This is an important issue because empirical priors are estimated from subject-specific data in order to maximize the marginal likelihood or evidence, given those data. This means the only external validation of the priors rests on showing that the same priors are selected over independent data, i.e., subjects. Although the conservation of important priors over subjects was our primary focus, we also investigate another key technical advance in source reconstruction; namely the use of hyperpriors on prior spatial covariance components. Hyperpriors are used commonly to finesse ill-posed problems, including the MEG inverse problem (see Sato et al., 2004 for a nice example using variational Bayes). Here we show how hyperpriors can be important for hierarchical models with multiple constraints, allowing us to explore more simultaneous constraints than is typical for previous inverse schemes.

Second level: population inference

Solutions to the MEG/EEG inverse problem, particularly those constrained by MRI, are normally estimated on a subject-specific basis. The second focus of this paper is how one combines these solutions to make inferences about the population. This is relatively simple for ECD models because anatomical differences among subjects (in dipole location and orientation) can be discounted by pooling "homologous" ECDs over subjects. In distributed source reconstructions, the problem is less simple because inferences about regionally specific effects are based on pooling estimated activity at each voxel over subjects. Population inference on distributed sources is more ambitious because it tests not only for responses that are expressed systematically over time but also in the same place. This means a significant finding implies that the functional anatomy is conserved over subjects. Here, we estimate solutions in each subject's native MRI space and use normalization procedures developed for MRI (Ashburner and Friston, 1999) to map them into a common stereotactic space (based on Talairach and Tournoux, 1988). This allows application of well-established mass-univariate approaches (statistical parametric mapping — SPM) to locate regionally specific responses over subjects.

We also consider two issues relating to group-level inferences. The first concerns whether the difference between two or more conditions should be estimated by localizing each condition separately, and then contrasting the source solutions, or by contrasting the mean evoked responses, and localizing this differential effect. We will refer to the former as the "difference of localizations" (DoL) and the latter as the "localization of the difference" (LoD). Although the present hierarchical forward model is linear, the ReML estimates of the hyperparameters are not linear functions of the data; therefore, the two approaches can give different results. Generally, researchers have tended to adopt the DoL approach. One rationale often given is that the signal-to-noise ratio (SNR) for an evoked component versus prestimulus baseline is normally greater than that for the difference in that component across conditions. In fact, this need not be the case: differences in evoked responses do not necessarily coincide with a peak in those responses versus prestimulus baseline, and it is possible to have a higher SNR for the difference between conditions than for each condition alone (vs. prestimulus baseline). In the Results section, we report both approaches and consider the pros and cons of each in the Discussion section.

The second issue pertains to the features of the individual source localizations on which a group inference is based. Specifically, we consider whether to make inferences on signed or unsigned source estimates. The sign of the estimated source amplitude depends on the orientation of each dipole with each subject's cortical mesh (see below), and this depends on local gyral anatomy, which may differ across subjects. One might therefore prefer to discount such individual differences in local orientation by taking the absolute value of each source estimate. Alternatively, one might wish to enforce consistent orientations (as well as strengths) at corresponding locations across subjects, choosing to reject source estimates of opposite signs in homologous regions as noise. Again, we compare both approaches in the Results and Discussion) sections.

Preview of paper

In this paper, we apply ReML and SPM to MEG data from 9 subjects in order to make population-level inferences about the functional anatomy of the "M170", here defined as the difference in fields evoked by intact and scrambled faces between approximately 140 and 200 ms post-stimulus. The general approach is shown in Fig. 1. In brief, a T1-weighted MRI from each subject was segmented to provide a tessellated cortical graymatter mesh of approximately 7000 vertices. At each vertex a dipole was placed, oriented normal to the cortical surface. This mesh was coregistered with the MEG sensors via fiducial markers, and a single-shell spherical forward model created. The amplitude of each dipole (source) within this mesh was then estimated using a number of constraints on the data covariance within the time window of interest in both measurement (sensor) space and source space: this involves computing the optimum mixture of covariance components with ReML and using the ensuing covariances to evaluate the conditional or a posteriori estimates of the sources (this is formally equivalent to the M and E-steps in Expectation Maximization). The source estimates were then interpolated from the mesh to a 3D image, and this image normalized to a template space using parameters derived from the subject's T1-weighted MRI. After spatial smoothing, the resulting images were used to create posterior probability maps (PPMs) (Friston et al., 2002), in which the mean and variance of a Gaussian posterior distribution over subjects were estimated at each voxel, and the resulting map thresholded for a certain probability of "activation".

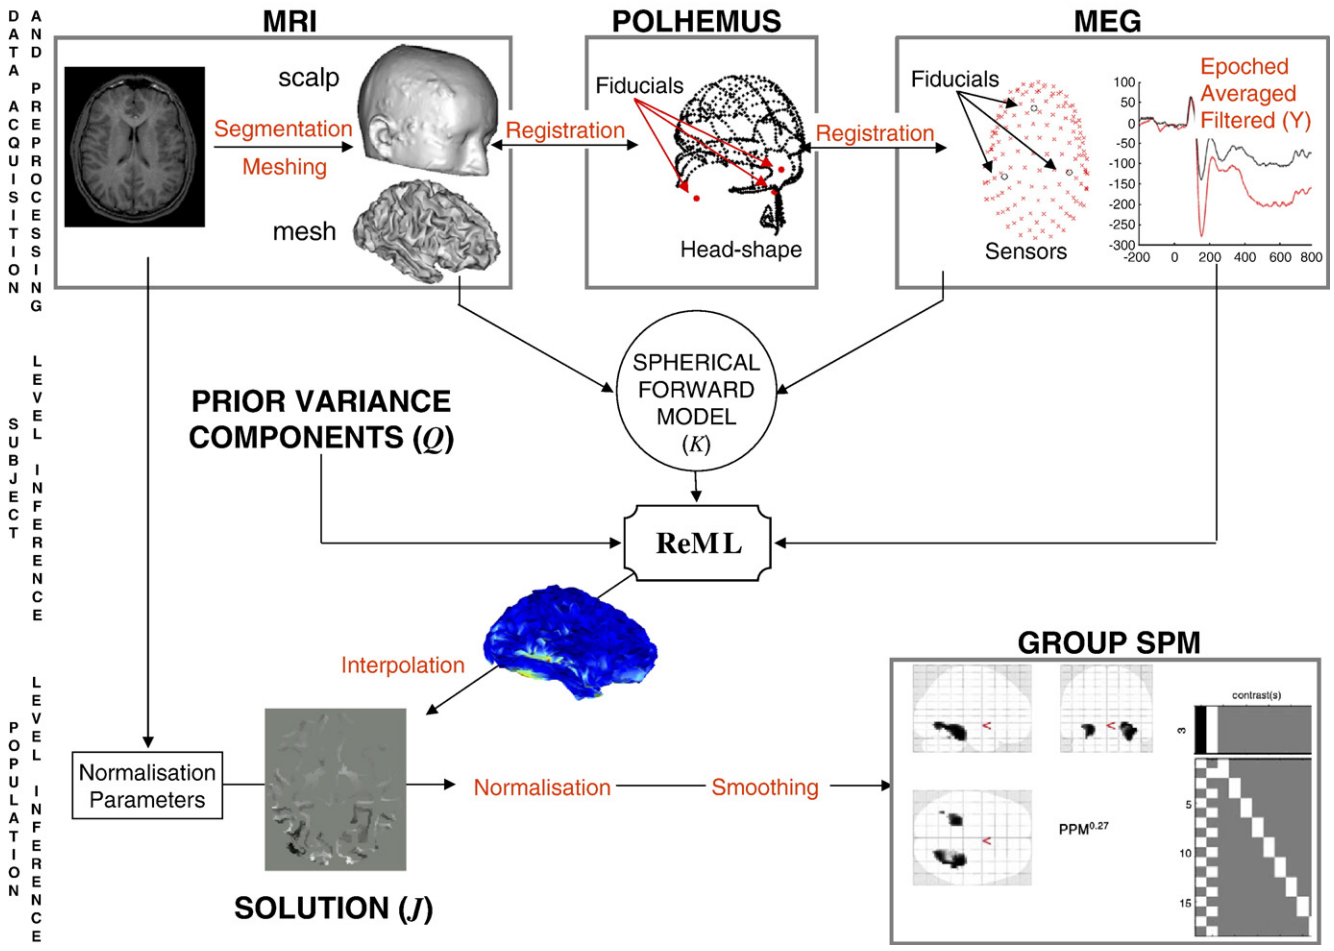

Fig. 1. Schematic of general processing pathway.

In the following sections, we expand on this approach, beginning with a brief reprisal of the PEB model and ReML approach to the inverse problem, in order to introduce the various covariance constraints and then consider the statistical analysis of between-subject effects.

A Parametric Empirical Bayes model of the inverse problem

The PEB formulation of the MEG/EEG inverse problem is based on a two-level hierarchical, linear model. At the first level, the inverse problem is expressed by the linear equation:

$$Y = KJ + E_1 \qquad E_1 \sim N(0, C_1) \tag{1}$$

where Y is an n (sensors) $\times t$ (time points) matrix of sensor data; K is an $n \times p$ (sources) matrix representing the "forward model" that encodes how a current at each source would be expressed as a magnetic field over the sensors; J is the $p \times t$ matrix of unknown dipole currents, i.e., the within-subject model parameters that we wish to estimate, and E_1 is an $n \times t$ matrix of multivariate Gaussian noise (hence the "parametric" part of PEB), with zero mean and covariance C_1. Note that the data are mean-corrected over sensors by treating the mean as a confound (the "restricted" part of ReML). We will refer to C_1 as the "sensor-level covariance".

The WMN solution to this equation obtains by assuming that the expected value of J is zero (see Phillips et al., 2005). This is expressed through the second level of the model:

$$J = 0 + E_2 \qquad E_2 \sim N(0, C_2) \tag{2}$$

where E_2 is a $p \times t$ matrix of zero-mean Gaussian noise with covariance C_2. We will refer to C_2 as the "source-level covariance". The error covariance E_2 can be also be regarded as a Bayesian prior on the parameters J in the first level (hence the "empirical Bayes" part of PEB); their mean of zero renders them so-called "shrinkage priors".[1]

Covariance components

The constraints on the inverse solution are embodied in the covariance matrices C_1 and C_2 (see also Hillebrand et al.,

[1] Note that shrinkage priors shrink the conditional estimates towards their prior expectation of zero. A prior expectation of zero does not mean we think all sources have zero value; it means that, in the absence of data, this is the value the sources would be assigned. Deviations from zero, both positive and negative, are penalized in proportion to their prior precision. Although we do not know in advance whether this deviation should be positive or negative we can influence the amplitude of this deviation though the prior covariance. Optimizing this covariance is the essence of empirical Bayes and the inverse solution we employ.

2005).[2] These are related to the expected data covariance simply by (Phillips et al., 2005):

$$E[YY'] = C_1 + KC_2K'$$

where the superscript $'$ indicates the matrix transpose operator. C_1 and C_2 can be (hyper)parameterized as a linear combination of covariance matrices, P:

$$C_1 = \sum_{i=1}^{n} \lambda_i P_{1,i}$$
$$C_2 = \sum_{i=1}^{m} \lambda_{n+i} P_{2,i} \quad (3)$$

Informally, a high variance for a given source or sensor implies that nonzero values for that source/sensor are less unlikely. Similarly, a higher covariance between two sources or two sensors makes their values less likely to be independent. After projecting the source covariance matrices ($P_{2,i}$) to the sensor space, this results in a linear combination of "covariance components", Q:

$$E[YY'] = \sum_{i=1}^{n+m} \lambda_i Q_i = \lambda_1 P_{1,1} + \ldots \lambda_n P_{1,n} + \lambda_{n+1} K P_{2,1} K' + \ldots \lambda_{n+m} K P_{2,m} K' \quad (4)$$

where λ_i are the unknown hyperparameters that control the relative contribution of the different components. Here, the Q_i matrices were scaled to have an L2 norm of unity; this allows quantitative comparison of the associated hyperparameters. The data were also scaled by a fixed factor for all subjects to avoid numerical overflow problems.

Seven different covariance constraints were considered here: three at the sensor level ($P_{1,i}$) and four at the source level ($P_{2,i}$). In what follows, we describe and motivate each constraint. Examples of the corresponding covariance components are provided in Fig. 2.

Sensor constraint 1: "IID"

The first constraint at the sensor level was simply the assumption of independent and identically distributed (IID) sensor noise, i.e.: $P_{1,1}=I_n$, where I_n is an $n \times n$ identity matrix.

Sensor constraint 2: "baseline"

The second constraint was an estimate of the sensor noise covariance obtained from the prestimulus baseline period. This period was −100 ms to 0 ms (63 sample points): given an $n \times t$ matrix, B, of the baseline MEG data at each sensor, averaged across all trials (intact and scrambled faces), then $P_{1,2}=\mathrm{cov}(B)$.

Sensor constraint 3: "anti-averaging"

The final constraint at the sensor level was an estimate of the noise covariance obtained during the critical time window of the M170, based on "anti-averaging". In anti-averaging, one half (selected randomly) of trials are subtracted from the other half. On average, this removes the evoked (signal) component, but leaves an estimate of the random (noise) component (Schimmel, 1967). Given an $n \times t$ matrix A of the MEG data anti-averaged across all trials (intact and scrambled), then $P_{1,3}=\mathrm{cov}(A)$.

Source constraint 1: "IID"

The first constraint at the source level was simply the assumption that the p sources were independent and identically distributed (IID), i.e., $P_{2,1}=I_p$. Note that, when passed through the forward model K, the resulting $n \times n$ covariance component in sensor space is no longer an identity matrix since it contains covariance (off-diagonal) terms induced by correlated forward fields.

Source constraint 2: "Laplacian" smoothness

This constraint reflects the assumption that currents in nearby dipoles in source space have similar values. This source covariance was modeled by a Gaussian kernel: $P_{2,2}(i,j)=\exp(-d_{ij}^2/2s^2)$; where d_{ij} is the geodesic distance between dipoles i and j within the cortical mesh, and s is a spatial smoothness parameter, which was fixed at 16 mm here. This constraint encourages spatially smooth solutions (cf. the LORETA method; Pascual-Marqui et al., 1994).

Source constraint 3: "MSP"

This constraint is based on Multivariate Source Prelocalization (MSP) (Mattout et al., 2005). Briefly, MSP entails a multivariate correlation between the normalized data $\bar{Y}$ and the normalized forward model $\bar{K}$. This returns a coefficient for each dipole β_i that captures the potential contribution of that dipole, via its forward field, to the data, independent of its scaling. This can be expressed in the diagonal matrix: $P_{2,3}(i,i)=\beta_i$. As with other constraints derived from the data, this constraint was estimated by averaging across intact and scrambled faces, so as not to bias estimation of the difference.

Source constraint 4: "depth weighting"

This constraint gives greater "weight" to deeper sources. Despite the fact that we use a forward model based on the physical properties of magnetic fields (i.e., their dependence on depth), this weighting is believed important for minimum L2 norm solutions, which tend to favor superficial, low intensity solutions to deep, high intensity solutions (Gorodnitsky et al., 1995). This constraint was coded by the diagonal covariance matrix: $P_{2,4}=\mathrm{diag}(K'K)^{-1/2}$ (Fuchs et al., 1999).

ReML and conditional estimates

The conditional source estimate $\hat{J}$ is the Maximum A Posteriori (MAP) estimate of J and is given by:

$$\hat{J} = C_2K'(C_1 + KC_2K')^{-1}Y \quad (5)$$

(cf. the weighted minimum norm, or Tikhonov solution; Phillips et al., 2005). The covariances required for this estimate come from Eq. (3), using ReML estimates of the hyperparameters. ReML estimates of λ_i differ from the standard ML estimates by accounting for the loss in degrees of freedom due to the conditional uncertainty about the parameters (source estimates).

[2] For the beamformer case, for example, the a priori source covariance is given directly by $C_{2,ii}=(k_i' C_d^{-1} k_i)^{-1}$ with non-diagonal elements set to zero, C_d the measured data covariance and k_i a column of K.

As evident from Eq. (4), the data covariance is hyperparameterized as a linear mixture of proper covariance components. This means that the hyperparameters are (non-negative) scale parameters. We ensured non-negativity by using a log-transform and estimating $\alpha_i = \ln(\lambda_i) \Leftrightarrow \lambda_i = \exp(\alpha_i)$. By imposing a Gaussian hyperprior on α:

$$p(\alpha) = N(\eta, \Omega) \tag{6}$$

this transformation is equivalent to placing a log-normal hyperprior on the hyperparameters. The resulting ReML objective function, or "free energy", F, is given by:

$$\begin{aligned} F = &-\frac{1}{2}(Y - K\hat{J})C_1^{-1}(Y - K\hat{J}) - \frac{1}{2}\hat{J}C_2^{-1}\hat{J} - \frac{t}{2}\ln|C_1^{-1}| \\ &- \frac{t}{2}\ln|C_2^{-1}| + \frac{t}{2}\ln|K'C_1^{-1}K + C_2^{-1}| \\ &- \frac{1}{2}(\alpha - \eta)'\Omega^{-1}(\alpha - \eta) - \frac{1}{2}\ln|\Omega^{-1}| + \frac{1}{2}\ln\left|\sum{}^{-1}\right| + \text{const} \end{aligned} \tag{7}$$

where Σ is the conditional covariance of the hyperparameters (see Friston et al., 2006b, for details). F is related to the log evidence for the model, i.e., log-probability of the data given the hyperparameters $\ln(p(Y|\lambda))$.

In the present paper, we used "sparse" hyperpriors, i.e., with a relatively small value for their expectation $\eta = -8$ (i.e., prior expectation of $\lambda_i = \exp(-8)$) and a relatively large value for their variance $\Omega = vI, v = 32$. The value of the prior expectation η was chosen to maximize the free energy F (Eq. (7)); see Results. The use of sparse hyperpriors encourages the hyperparameters of redundant covariance components to be switched off since $\alpha_i \to -\infty \Leftrightarrow \lambda_i \to 0$ (Friston et al., 2006b). Because the prior variance of α_i is constant, the variance of λ_i shrinks to zero. This means that the hyperparameter estimate is effectively zero with no uncertainty; i.e., it is not relevant. Therefore, ReML can be used to implement automatic model selection by noting that the optimal model will comprise only those covariance components with non-negligible hyperparameters. The use of these sparse hyperpriors thus means that formal model selection is unnecessary, provided all potential model components are included.

Paradigm and MEG data

The present paradigm is identical to that used previously with EEG and fMRI (Henson et al., 2003). Here, we analyze data from a single, eleven-minute session in which subjects saw intact or scrambled faces, subtending vertical and horizontal visual angles of approximately 4°. One half of the intact faces were famous; one half were non-famous. Given that famous and non-famous faces did not differ over the critical M170 time window in these data, or in previous EEG data (Henson et al., 2003), famous and non-famous faces were pooled. Scrambled versions of each face were created by phase-shuffling in Fourier space and masking by the outline of the original image (to match size). The scrambled faces were therefore approximately matched for spatial frequency power density (see Fig. 3 for an example). Subjects made left–right symmetry judgments about each stimulus by pressing one of two keys with either their left or right index finger (range of reaction times was 1031 ms–1798 ms). There were 86 intact and 86 scrambled face trials.

Nine subjects were tested, four female, ranging from young to middle-aged adults. Their involvement complied with the Code of Ethics of the World Medical Association (Declaration of Helsinki) and the standards established by the local review board. The MEG data were sampled at 625 Hz on a 151-channel axial gradiometer CTF Omega system at the Wellcome Trust Laboratory for MEG Studies, Aston University, England. The total amount of movement across the session ranged from 0.2 to 5.6 mm (median = 1.1 mm).

The MEG data were pre-processed using the SPM5 software package.[3] The data were epoched from −100 to +600 ms and averaged over trials to produce the event-related field (ERF) for each trial-type. The ERFs were baseline corrected from −100 to 0 ms and low-pass filtered to 40 Hz in both forward and reverse directions using a 5th-order Butterworth digital filter.

ReML estimates of the hyperparameters used the covariance of the data over 26 sample points, defined from −20 ms to +20 ms of the peak of the difference between the ERFs to intact and scrambled faces for each subject. The MAP estimates of the sources were evaluated at the mid-point of this window.

MRI and forward model

A T1-weighted MPRAGE-MRI scan was acquired for each subject with voxel size 1 × 1 × 1 mm. These structural scans were segmented, and the graymatter segment was used to construct a continuous triangular mesh representing the neocortex using Anatomist.[4] This mesh contained about 7200 vertices (ranging from 7204 and 7211 across subjects). The mean inter-vertex spacing ranged from 4.3 mm to 5.3 mm across subjects. The normal to the surface at each vertex was calculated from an estimate of the local curvature of the surrounding triangles (Dale and Sereno, 1993).

Three sensor coils were placed on the nasion and left and right peri-auricular point of each subject. The position of these coils was detected by the MEG machine. The coil positions were also digitized with 3D Polhemus Isotrak digitizer. These points were then used to co-register the MEG and MRI spaces. Brainstorm[5] (Baillet et al., 2001) was then used to construct a single-shell, spherical forward model. This model was used with Sarvas' formula (Sarvas, 1987) to map the contribution of each dipole on the mesh to each MEG sensor (i.e., to create the matrix K).

Normalization and smoothing

Each subject's MRI was normalized to the Montreal Neurological Institute (MNI) T1 template within the coordinate system of Talairach and Tournoux (1988) using SPM5. The linear and nonlinear normalization parameters were estimated as part of the segmentation process, which uses a unified generative model for intersubject variations in anatomy (Ashburner and Friston, 2005).

The estimated dipole strengths in each subject's mesh were interpolated into a 3D space of 1 × 1 × 1 mm voxels using trilinear interpolation (Baillet et al., 2001). The anatomical normalization parameters estimated above were then used to normalize these images with a final voxel size of 3 × 3 × 3 mm. The normalized images were smoothed with a 16 mm FWHM isotropic Gaussian

[3] http://www.fil.ion.ucl.ac.uk/spm.
[4] http://brainvisa.info/doc/html/brainvisa/en/processes/aboutAnatomist.html.
[5] http://neuroimage.usc.edu/brainstorm.

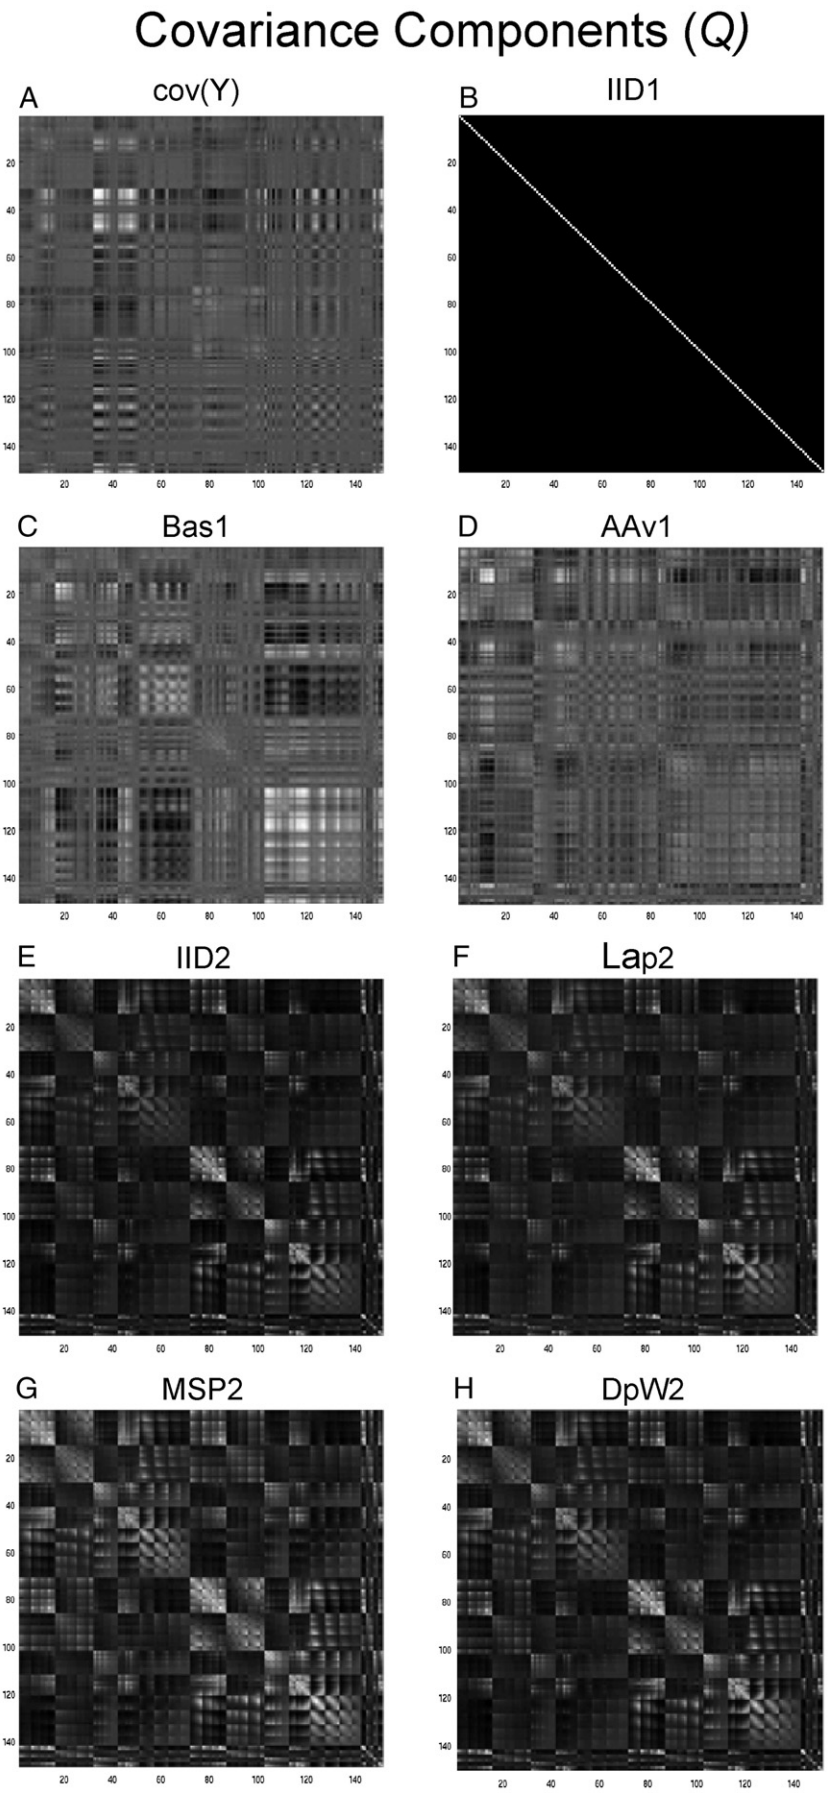
Covariance Components (Q)
A
cov(Y)
B
IID1
C
Bas1
D
AAv1
E
IID2
F
Lap2
G
MSP2
H
DpW2
20
40
60
80
100
120
140

kernel to accommodate residual differences in functional anatomy. This (relatively large) Gaussian smoothing also renders inference more robust under the Gaussian assumptions behind PPMs (see below). Finally, these images were masked to exclude voxels outside the solution space, namely those voxels that did not appear in at least one subject's normalized mesh in 3D space.

Note that we have chosen to evaluate responses at the between-subject level in three-dimensional space, as opposed to averaging the responses in a two-dimensional cortical manifold (Fischl et al., 1999). There are two principled reasons for this: first, variations in the spatial deployment of activity from subject to subject can be accommodated using the matched filter theorem by smoothing the subject-specific estimates in three dimensions. This is an important principle, used in intersubject analyses of fMRI data. The matched filter theorem requires the data to be smoothed with a kernel that has a transfer function which matches the 'signal'. In multi-subject studies, the 'signal' is dispersed spatially by variations in functional and gyral anatomy. It can be seen easily that random variations in gyral anatomy are not confined to a 2D manifold but are expressed in a 3D space; this calls for a 3D smoothing kernel. Second, we prefer to display and report our results in the standard anatomical space that had been established by the imaging neuroscience community (most imaging modalities are not confined to a cortical surface).

Group analyses

Our aim was to localize the M170, i.e., the difference between the magnetic field evoked by intact versus scrambled faces that is maximal between approximately 140 and 200 ms post-stimulus. As described in the Introduction section, we compared the DoL approach – of localizing the evoked response to intact and scrambled faces separately (versus prestimulus baseline) and contrasting the source solutions – with the LoD approach — of localizing the difference in the mean evoked responses for intact versus scrambled faces. For each such approach, we performed two analyses (i.e., 3D interpolation, normalization and smoothing) using either the raw (signed) source estimates or their absolute value. We refer to the latter as the "unsigned" analyses.

For the DoL approach, both signed and unsigned analyses have the expectation that, on average, voxels will show no difference between intact and scrambled faces. For the unsigned LoD approach, the inference is defined less clearly because the expected average absolute value is greater than zero (given nonzero noise). We therefore decided to ask where the absolute value of the source estimates was greater than average for the mesh; this corresponds to subtracting from each source estimate, for a given subject, the mean absolute source estimate across all dipoles in the mesh. This unsigned LoD analysis is therefore qualitatively different to the other three analyses.

Posterior probability mapping

The factorial combination of LoD/DoL and signed/unsigned approaches resulted in four PPMs. These PPMs also derive from a two-level PEB framework, in which the first level corresponds to a conventional linear model at each voxel (e.g., a paired *t*-test) and the second level represents the distribution of activity over all voxels (Friston et al., 2002). This distribution is modeled by a Gaussian with zero mean. This acts as an empirical "shrinkage prior" on the voxel-specific estimates at the first level, in exactly the same way that the empirical priors in source space provide the posterior estimates of sources, as described earlier. In both cases the requisite covariances are estimated with ReML. The resulting posterior probability distributions for each voxel (PPMs) can then be thresholded at a given probability that the difference in conditional source estimates for intact and scrambled faces exceeded one standard deviation of the empirical prior distribution over voxels. Here we used a threshold of 95%, taking into account both tails of the posterior distributions in the case of the signed tests and the unsigned DoL tests (the unsigned LoD test had only one tail of interest, i.e., greater than average activity).

Results

MEG data

The ERFs for intact and scrambled faces are shown for each subject in Fig. 3A from a left temporal sensor (MLT24). The largest divergence between intact and scrambled faces coincided with the peak of an evoked component for both stimulus types versus prestimulus baseline, most likely corresponding to the M170 component that is believed to be specific to faces (Liu et al., 2000). (Intact and scrambled faces also differed at later times, but this more sustained difference is not considered here.) The time window of interest for each subject was centered on the peak of the global field power (GFP) of the difference between their evoked response to intact and scrambled faces in the range 140–200 ms post-stimulus. The mean time of the peak GFP was 169 ms, with a range across subjects of 149–195 ms. A time window was selected from 20 ms before to 20 ms after each subject's peak (26 samples). The 2D topographies of the differential evoked field at the time of the maximal differential GFP for each subject are shown in Fig. 3B. (The covariance of the channel data for one subject is shown in Fig. 2A.)

Hyperparameter and parameter estimates (1st level)

To optimize the hyperprior expectation, the value of the ReML objective function F (Eq. (7)) was plotted as a function of η from −32 to +32. As can be seen in Fig. 4A, when averaging across subjects, the optimal value of η was around −8. For lower values of η, there was evidence of discrete jumps in F for each subject (Fig. 4B), suggesting a "phase transition" in the ReML objective function (e.g., emergence of local minima). Thus our subsequent results were computed with $\eta=-8$, for which the results were confirmed as robust to different initializations of λ.

The ReML estimates of the hyperparameters are shown for each subject in Fig. 5, for each of the three effects of interest (intact faces, scrambled faces and their difference). Each graph shows the hyperparameter estimate for each of the nine subjects (1–9) and

Fig. 2. The covariance of the MEG data across the 151 MEG sensors from the time window of interest, cov(Y), and the seven normalized covariance components, Q_i from one subject (Subject 9). The first three components ("IID1", "Bas1", "AAv1") live at the sensor level, the next four ("IID2", "Lap2", "MSP2", "DpW2") are projected from the source space via the forward matrix K (see text for more details). Note that the latter four components are very similar owing to this shared projection (itself a reflection of the underdetermination of the inverse problem). IID = independent and identically distributed; Bas = baseline, AAv = anti-averaging, Lap = Laplacian (smoothness), MSP = Multivariate Source Prelocalization, DpW = depth weighting.

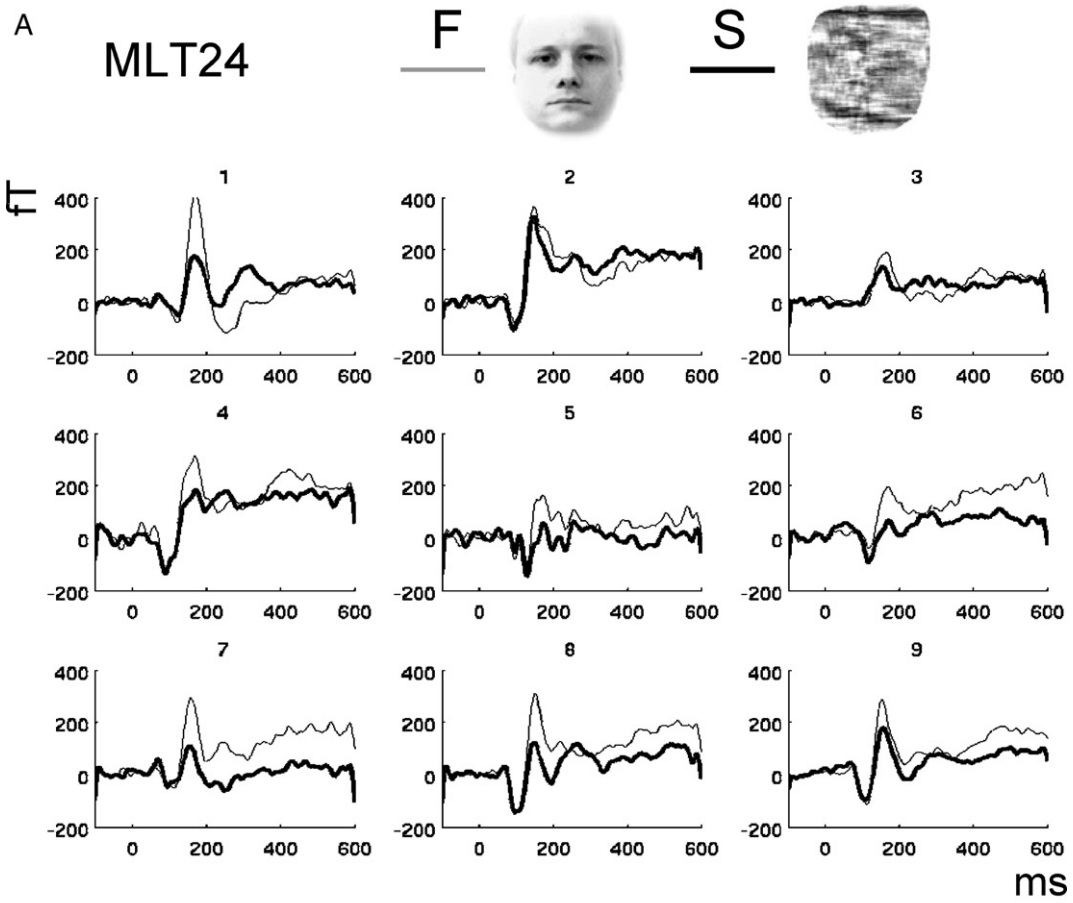
A
MLT24
F
S
fT
1
2
3
4
5
6
7
8
9
400
200
0
-200
0
200
400
600
ms

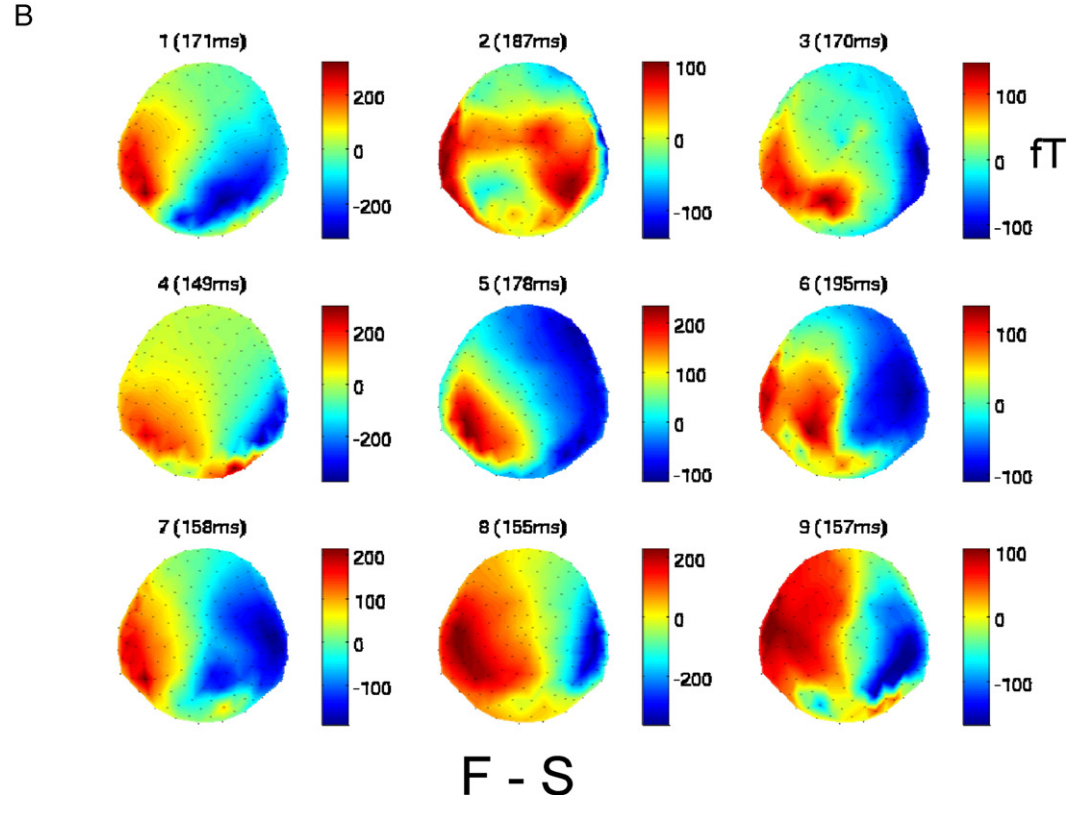
B
1 (171ms)
2 (187ms)
3 (170ms)
4 (149ms)
5 (178ms)
6 (195ms)
7 (158ms)
8 (155ms)
9 (157ms)
200
100
0
-100
-200
fT
F - S

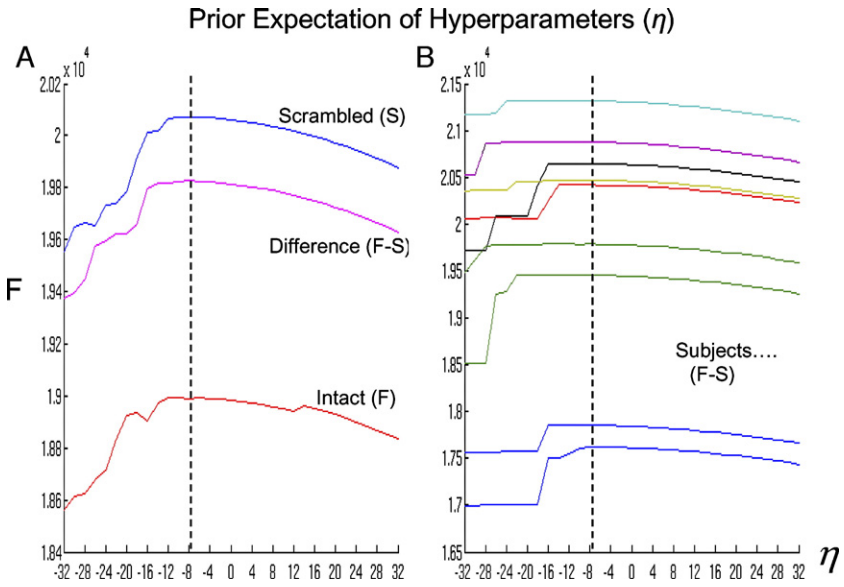

Fig. 4. Plots of ReML log evidence, or free energy (F), against expectation of log-normal hyperprior (η) for (A) each effect of interest, averaged across subjects, and (B) the differential effect between intact vs. scrambled faces, separately for each subject. The broken line shows the value used in this paper.

each of the seven covariance components. The pattern is similar for all three effects of interest: for a given subject, only two constraints tended to have non-negligible hyperparameters, one at the sensor level and one at the source level. At the sensor level, this corresponded most often to the baseline estimate of sensor noise (and less occasionally, the anti-averaging constraint). At the source level, the most relevant hyperparameter corresponded to either the MSP constraint (for four subjects) or the depth weighting constraint (for the other five).

The estimates of the parameters (i.e., source solutions) for the differential response between intact and scrambled faces (LoD) are shown for ventral and right lateral views of each subject's native mesh in Fig. 6. Most subjects showed evidence of bilateral ventral occipitotemporal activity, particularly in a right posterior region.

Population inference (2nd level)

Normalized and smoothed images of the above source solutions were entered into four different group analyses, corresponding to the DoL and LoD approaches, crossed with signed versus unsigned source strengths (see Methods). The DoL and LoD signed analyses gave very similar results: both showed evidence of differential activity in bilateral fusiform/parahippocampal and medial orbitofrontal cortex (Figs. 7A and B; Table 1). Interestingly, the differences were of opposite sign (orientation) in the ventral temporal and frontal regions.

Inspection of the inverse solutions showed a tendency for source amplitudes close in source space to have opposite signs when on opposite sides of a sulcus (reflecting the same dipole orientation, given that dipoles on opposite sides of sulci have near-opposite orientations). We therefore expected that spatial smoothing of the resulting 3D images would attenuate estimated source activity. Surprisingly however, the unsigned DoL analysis showed similar results to the signed analysis, at least in ventral temporal regions, though the orbitofrontal region no longer exceeded the 95% PPM threshold, and an anterior medial temporal region appeared instead (Fig. 7C; Table 1). These voxels all showed greater absolute source strength for intact relative to scrambled faces (no voxels showed evidence of a greater absolute source strength for scrambled faces at this threshold).

Finally, the unsigned LoD analysis showed a different pattern (Fig. 7D), with evidence for "greater-than-average" activity restricted to right fusiform and occipital regions (Table 1). The difference between this analysis and the other analyses can be attributed to the qualitatively different type of inference (see Discussion).

For comparison, Fig. 7E shows a thresholded PPM for the difference between intact and scrambled faces from the identical paradigm (Henson et al., 2003), but using fMRI with 18 (different) subjects. While there are more activated regions in the fMRI results, there is reasonable concordance between the fMRI and localized MEG data (see Discussion).

Discussion

The MEG/EEG inverse problem is underdetermined and so requires constraints. The more constraints one applies, assuming they are valid, the more accurate the solution. We described a method for automatically weighting multiple constraints as a function of their relevance in maximizing the log evidence of a single model, in which all constraints are encoded as covariance matrices.

Constraints can be considered as "hard" or "soft". For the present analyses, hard constraints were placed on the location of

Fig. 3. (A) The MEG data for each subject from the channel MLT24 that showed the largest differential evoked field between intact faces, F (light lines), and scrambled faces, S (dark lines), over subjects. (B) Topographies of the differential ERF over sensors for each subject from the time point with the maximal difference (shown above) using piecewise bilinear interpolation between sensors. Viewed from the top, nose pointing upwards.

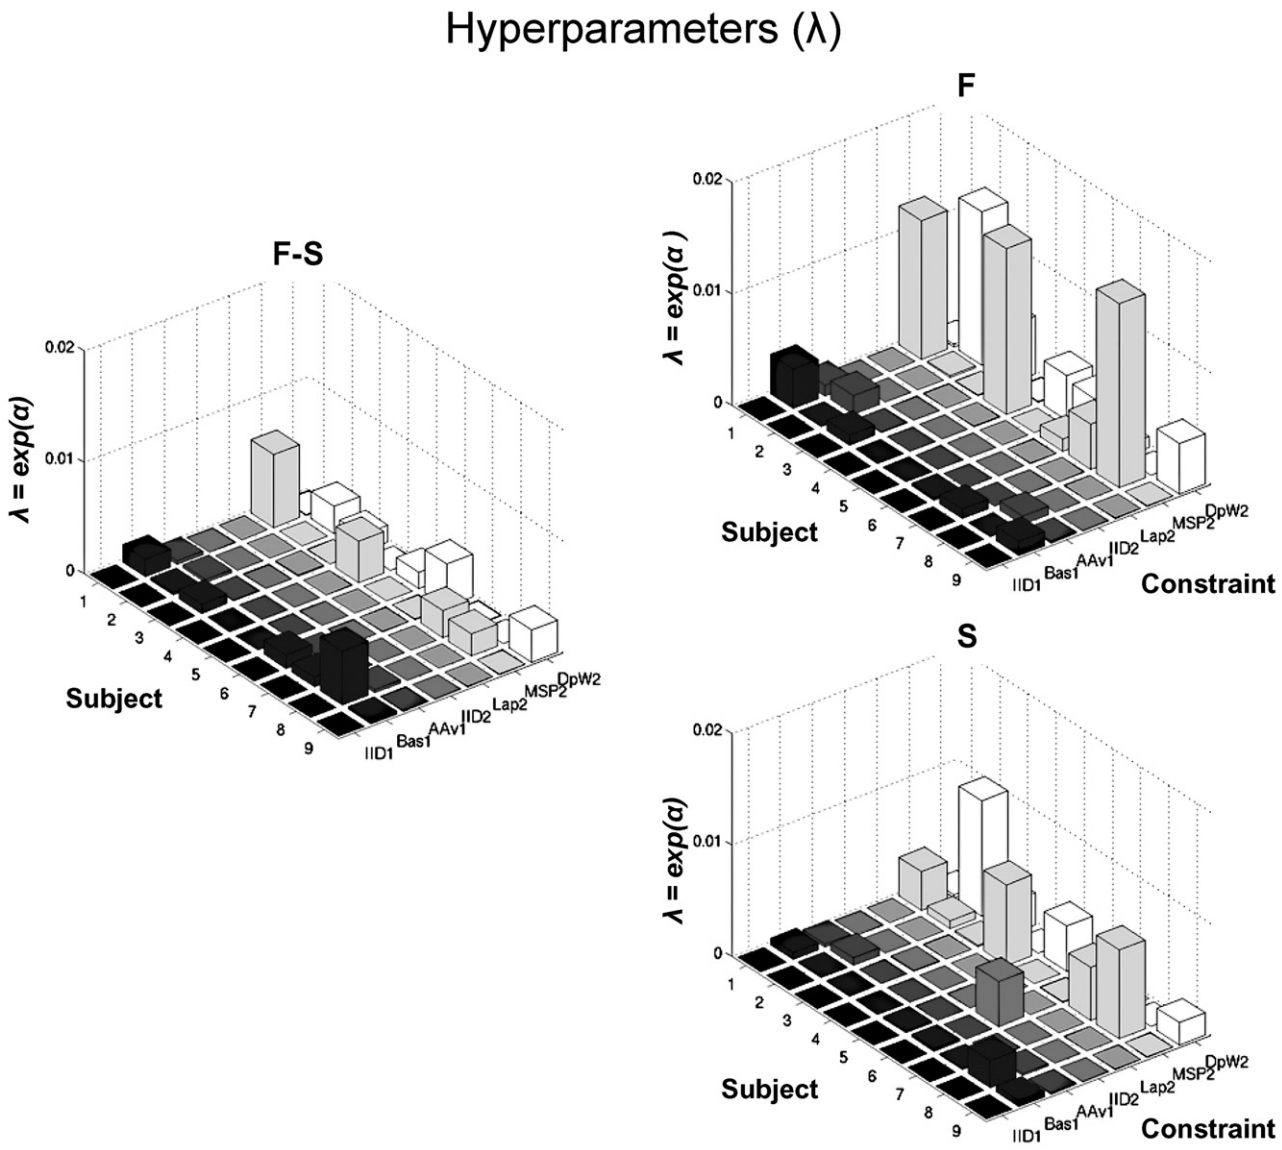

Fig. 5. ReML estimates of the hyperparameters λ for each effect of interest: the difference between intact and scrambled faces (F − S, left), intact faces alone (F, right) and scrambled faces alone (S, right) faces versus prestimulus baseline. See Fig. 2 legend for more details.

electrical sources in the form of a reasonably high-density mesh of about 7200 vertices, constrained to the cortical graymatter segmented from individual MRIs. We also employed a hard constraint on the orientation of those dipoles, enforcing them to be normal to mesh surface. It is possible that localization errors were introduced by assuming the dipole orientations were fixed (Lin et al., 2006). Indeed, for solutions with strong priors (like beamformers) slight errors in these fixed locations or orientations result in strongly degraded performance (Hillebrand and Barnes, 2003). One solution is to estimate additional current components in the two directions orthogonal to the surface normal (Phillips et al., 2005); another is to use the average and deviation of dipole orientations over patches of cortex (the "loose orientation constraint", Lin et al., 2006). More generally, there will be little difference between a forward model with a high density of local dipoles with different orientations and a smaller number of sparse dipoles with free orientations. This is because the same effective regional dipole on the high-density mesh can assume any orientation that is permitted by a linear mixture of [three] differently orientated local dipoles. Of course, this assumes that the precision of the localization is evaluated at a scale that does not encompass a single dipole; this precision is often precluded by the coherent or smooth source estimates afforded by constraints on the inverse solution. In this case, our experience is that topographically similar solutions are obtained with meshes that have between 3000 and 10,000 vertices (Mattout et al., in press).

Our main interest was in the soft constraints, which are naturally implemented as priors within a Bayesian framework, or more precisely, as covariance components within a Parametric Empirical Bayes (PEB) framework. The form of these covariance components was either determined by theoretical assumptions or derived from orthogonal partitions of the data. The relative importance of each component (as indicated by its associated

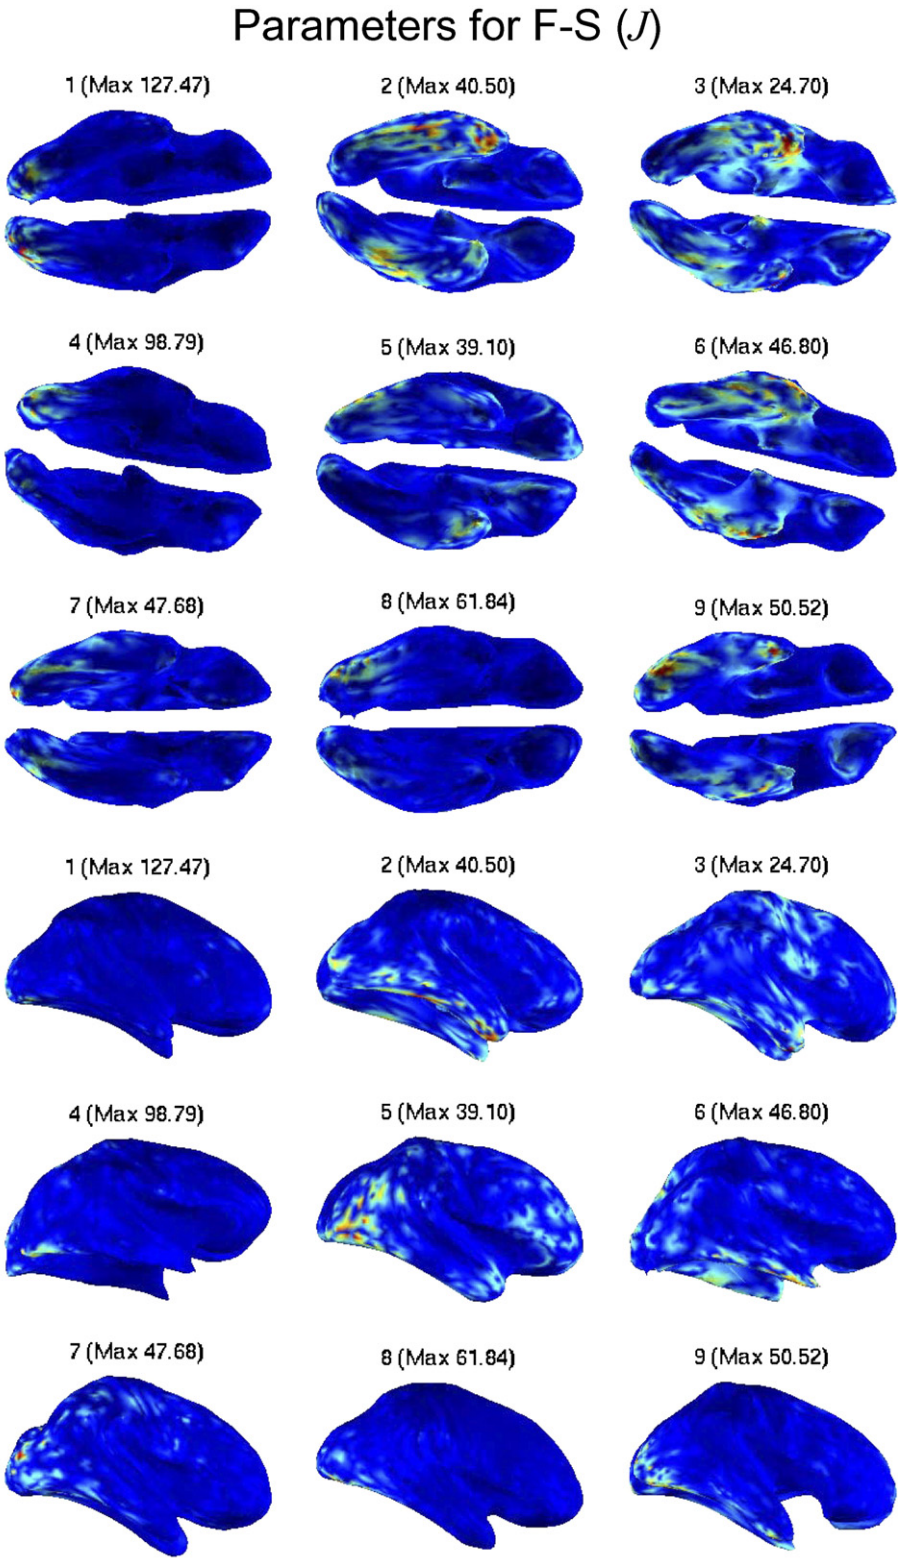

Fig. 6. Within-subject inverse solutions (i.e., parameter estimates) for the differential effect between intact (F) and scrambled (S) faces (LoD) for each subject, shown on ventral (top panel) and right lateral (bottom panel) views of each subject's cortical mesh. The meshes have been "inflated" slightly to aid visualization of sulcal activity. Color indicates the absolute value of dipole current at each vertex of the mesh, scaled to maximum value in the mesh (shown above each mesh). The units are arbitrary.

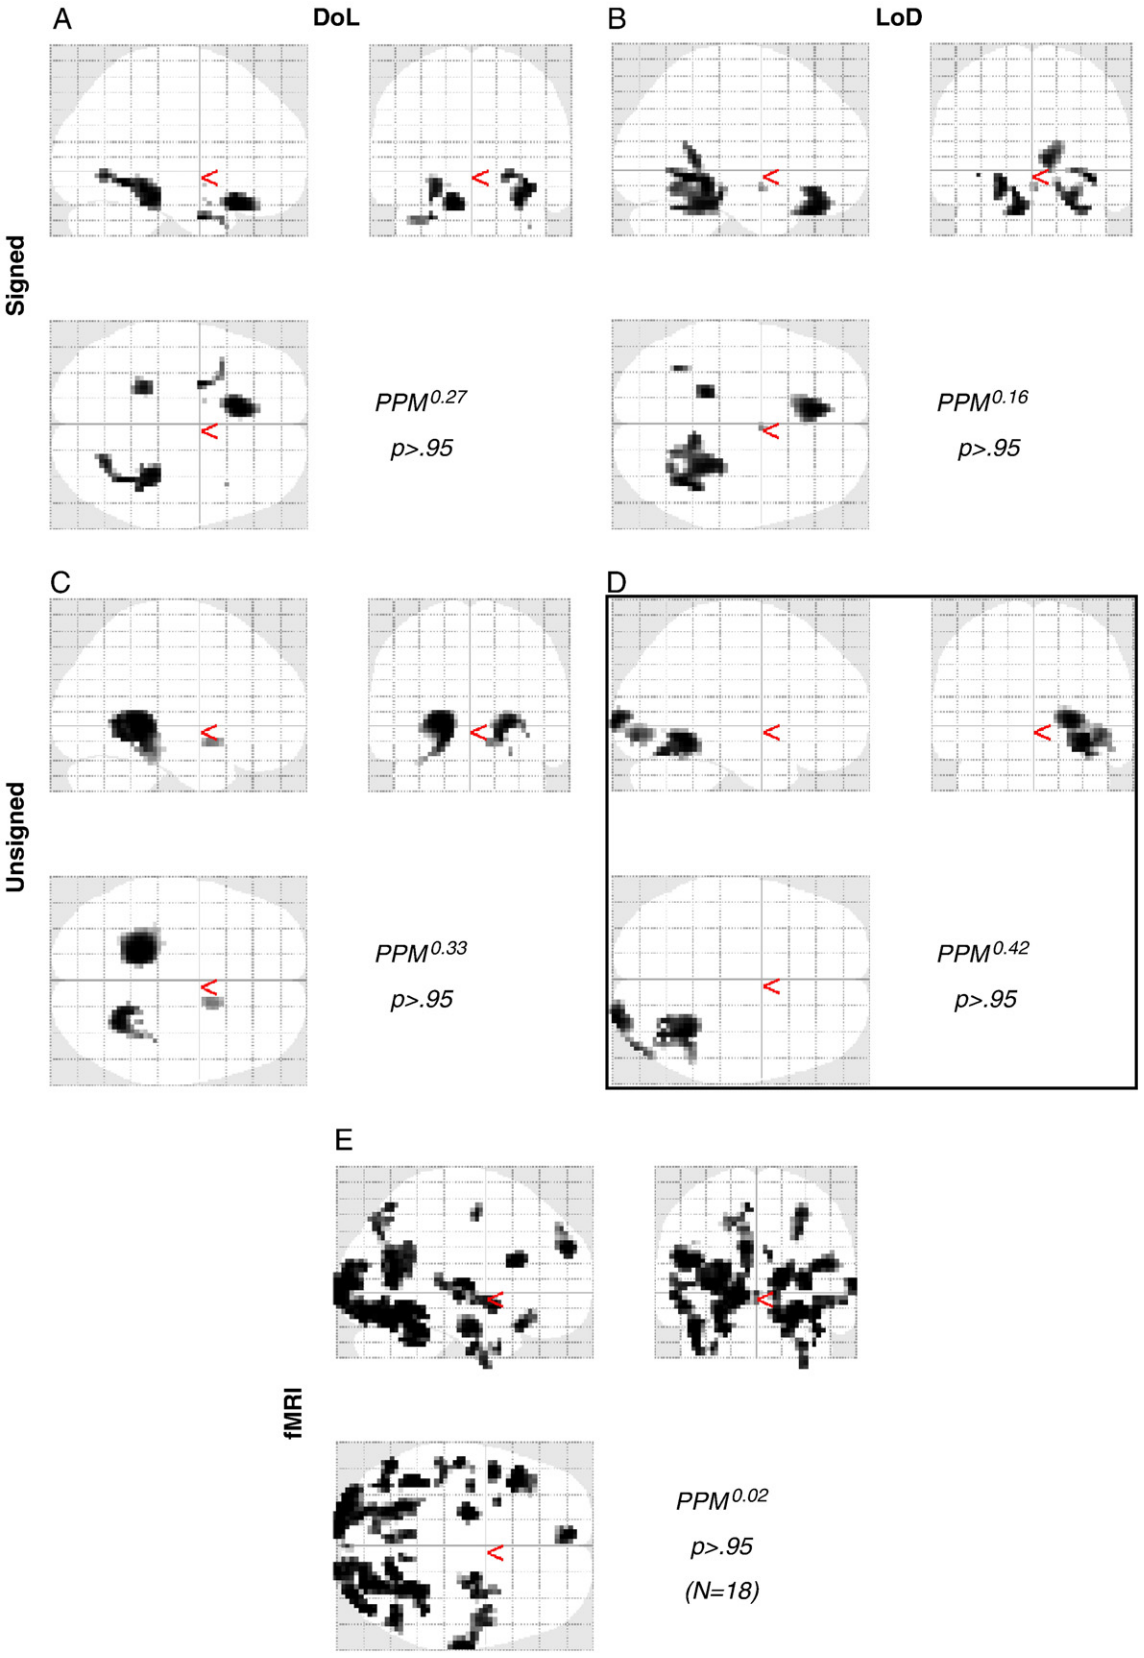

Fig. 7. Between-subject PPMs thresholded at 95% probability of differential activity between intact and scrambled faces shown as maximal intensity projections (MIPs) in MNI space, for panel A signed DoL analysis, (B) signed LoD analysis, (C) unsigned DoL analysis and (D) unsigned LoD analysis. Border round panel D indicates a qualitatively different inference from the other panels (see text). Panel E shows corresponding thresholded PPM for fMRI data from 18 (different) subjects within exactly the same paradigm, regardless of sign of difference (replotted from Henson et al., 2003).

Table 1
Regions of 10 or more voxels with a 95% probability or greater of differential activity between intact and scrambled faces (i.e., considering both tails of the posterior distributions in the case of the signed analyses and unsigned DoL tests; the unsigned LoD test is a qualitatively different inference relative to the mean over voxels, which has only one tail of interest)

Region	# Vox	*x*	*y*	*z*
Signed DoL L(F)−L(S)				
Right fusiform/parahippocampal	124	+39	−39	−12
		+36	−54	−6
Left fusiform/parahippocampal	51	−24	−39	−9
Left medial orbitofrontal	93	−12	+24	−21
Left temporal pole	22	−27	0	−30
Signed LoD L(F−S)				
Right fusiform	42	+39	−48	−6
		+30	−60	0
Right fusiform/parahippocampal	134	+27	−33	−21
Left fusiform/parahippocampal	57	−21	−36	−9
Left medial orbitofrontal	111	−9	+36	−24
Posterior cingulate	53	+12	−45	9
Unsigned DoL abs(L(F))−abs(L(S))				
Right fusiform/parahippocampal	301	+27	−54	−3
		+36	−39	0
Left fusiform/parahippocampal	128	−18	−39	0
		−27	−36	−18
Right anterior medial temporal	26	+16	6	−12
Unsigned LoD abs(L(F−S))−mean				
Right fusiform	190	+36	−51	−9
		+36	−66	−18
Right inferior occipital	120	+27	−96	−3

LoD=localization of difference, DoL=difference of localizations, F=intact faces, S=scrambled faces. Coordinates refer to MNI space. Anatomical labels are only approximate.

hyperparameter) was estimated using restricted maximum likelihood (ReML). These constraints are soft in the sense that they can be down-weighted (their hyperparameter reduced) if the data indicate that they are not relevant.

We also demonstrated the use of spatial normalization of the resulting source solutions to make population inferences about the location of MEG responses from a sample of subjects. Such population inference is still uncommon in this field (though see Dale et al., 2000; Park et al., 2002; Singh et al., 2003; for exceptions) and will minimize the reporting of spurious sources that can arise in single subject inversions. The distributed source solutions were estimated using individually defined meshes and forward models, derived from each subject's anatomical MRI. These solutions were then spatially normalized to a 3D template space, to compare across individuals, as is common practice in fMRI data analysis (though see also work by Fischl et al., 1999, in which intersubject registration is based on sulcal patterns in a two-dimensional manifold). After smoothing to accommodate residual functional-anatomical differences, the resulting images were used to construct posterior probability maps (PPMs), which show the probability that the source estimate at each voxel (the difference between intact and scrambled faces in the present case) exceeds a certain size (one standard deviation of the posterior distribution over voxels in the present case). This gives an idea of the confidence with which differential activity can be claimed at each location in the cortex. Our results illustrate that MEG can be an effective 3D imaging technique whose results (on only 9 subjects) show reasonable concordance with fMRI results on twice that number, including significant activation in reasonably "deep" cortical structures.

Analyzing group results in 3-space, as opposed to the 2-dimensional cortical manifold, may seem counterintuitive, given that sources are deployed on the cortical surface and this is where intersubject variability is expressed. However, irrespective of the form of this variability, there is an imperative for smoothing the source reconstructions in 3-space, prior to mass-univariate statistical analysis. This is because the probability of spatial overlap among activated sources will always be greater in 3 dimensions than in 2 dimensions; it is this overlap that represents a statistically reliable regional effect. Smoothing on the cortical manifold restricts overlap to a subspace of all voxels (volume elements) and will miss a proportion of regional effects. It should be noted that this argument applies only to regions where the cortical surface shows high curvature. In regions where the curvature of the cortical sheet is large in relation to smoothing, the geodesic and Euclidean distances will be the same (approximately) and smoothing in two and three dimensions will give the same results (approximately).

Covariance components at the within-subject level

The use of log-normal hyperpriors on the hyperparameters, a form of automatic model selection (Friston et al., 2006b), shrinks the hyperparameters for redundant covariance components towards zero. At the sensor level, the empirical estimate of noise covariance from the prestimulus baseline period was helpful, i.e., most consistently nonzero across subjects and effects of interest.[6] Noise characteristics during the critical M170 time window that were estimated using anti-averaging helped only occasionally, while the simple assumption of independent and identically distributed (IID) noise appeared redundant with these two empirical estimates. Thus automatic model selection (Friston et al., 2006b) appeared to operate well over the sensor-level constraints in emphasizing only those that were useful in fitting the data.

For source-level constraints, only the Multivariate Source Prelocalization (MSP) constraint or the depth weighting constraint had non-negligible hyperparameters on more than one occasion, and only one of these dominated in any one case (i.e., the source-level IID and Laplacian smoothness constraints were generally redundant in their presence). The dominance of only one source-level constraint is not surprising since the associated covariance components (in sensor space) are very similar. This is because, although the covariance matrices associated with each constraint differ in the source space, such differences are diminished on projection into the lower dimensional sensor space (via the forward matrix K). In a sense, this is a reflection of the underdetermination of the inverse problem.

[6] When the baseline measure of sensor noise covariance was estimated from averaging the baseline covariances across trials, rather than from the covariance of the average baseline, the results were similar. The theoretical difference between these two approaches is in the relative contribution of evoked and induced responses, and given that an evoked response was not expected for the (pre-stimulus) baseline period, similar results were expected.

This high correlation (in sensor space) between some of the covariance components makes it difficult to find the best set of hyperparameters. This process was eased by making the hyperpriors sparse (though evidence of local minima emerged for very small prior expectations, e.g., $\eta<-12$). In one sense, the difficulty in distinguishing the relative importance of the present source-level constraints does not matter because similar source solutions (parameters) will obtain if any single source-level constraint is nonzero (given their similarity). However it also means that one cannot place much emphasis on the relative size of hyperparameters as an indication of the general "importance" of such source-level constraints. Thus while either MSP or depth weighting dominated across different subjects, this should not be taken as evidence against the more general usefulness of, for example, a source-level smoothness constraint (e.g., Pascual-Marqui et al., 1994).[7]

Note that some of the constraints were derived from the data themselves: MSP at the source level and anti-averaging and baseline at the sensor level. One might regard such empirical constraints as inappropriate to use as covariance components on the same data. However, the baseline constraint derives from a different time window – the prestimulus baseline – which is likely to be independent of the time window localized. Anti-averaging derives from the same time window but should reflect a partition of the data (i.e., noise) that is orthogonal, assuming a sufficiently large number of trials, to the partition of interest (i.e., the signal). MSP however derives from the same time window and is closely related to the data partition of interest. Nonetheless, all of these empirically derived components were estimated from the average of the intact and scrambled face conditions and so should not bias the contrast of interest, i.e., the difference between intact and scrambled faces.

Feature selection at the between-subject level

Population inference on distributed EEG/MEG source solutions using spatial normalization has been performed previously (Dale et al., 2000; Park et al., 2002; Singh et al., 2003). Here we used "posterior probability maps" (PPMs), which allow one to threshold posterior densities for the probability of an effect of a given size (rather than the probability of falsely rejecting the null hypothesis of no effect, as in classical inference). PPMs thus eschew the multiple comparison problem (over voxels) associated with classical inference, as described in Friston and Penny (2003). They do assume Gaussian distributions, here of the summary statistics over subjects, which is expected under the Central Limit Theorem and with the Gaussian spatial smoothing employed. Nonetheless, we repeated the group analyses using non-parametric permutation tests, as implemented in SnPM (Nichols and Holmes, 2002), with a threshold corrected for multiple comparisons using family-wise error rate of 0.05. The results were very similar to those reported here (available on request). However, the emphasis here was on two "feature selection" issues: whether to contrast localizations or to localize contrasts and how to deal with the sign (orientation) of sources.[8]

Localization of difference vs. difference of localization

While one might expect the differences in condition-specific localized responses (DoL) and the localization of response differences (LoD) to give the same answer, the use of spatial constraints means that the source estimates are different. More precisely, the DoL and LoD approaches will not give identical inverse solutions because the hyperparameters under the two approaches are not linear functions of the data. For the signed analyses, the solutions were nonetheless very similar (Fig. 7). For the unsigned analyses, the solutions differed markedly because the inferences for the DoL and LoD approaches were different: in the former case, the inference relates to the difference in absolute values of source estimates for intact vs. scrambled faces; in the latter case, it relates to the absolute value of the source estimate vs. the mean absolute value across all dipoles in the mesh.

One way of looking at the localization of the difference (LoD) is to regard it as assuming a priori that the difference is localized. This may be tenable when the two conditions only evoke differences in the temporal form of the response; for example, latency differences or frequency-specific differences. However, two conditions may also engage processing systems that are distinct spatially. For example, imagine that face and non-face stimuli elicit responses in different cortical areas, such that a comparison of face and non-face reconstructions would, ideally, show two foci. If there are prior constraints that enforce sparse spatial reconstructions (e.g., a single ECD), localization of the difference may only reveal one differential source. Moreover, there are other reasons for preferring the DoL approach. These include the opportunity to localize "induced" responses (Friston et al., 2006a), which are not present in the mean response over trials, and the ability to use conventional parametric statistics when contrasting absolute source strengths (i.e., for unsigned analyses).

The distinction between DoL and LoD approaches can also be explored through different formulations of beamformer algorithms. Standard beamformer approaches follow a DoL approach. The beamformer weights/image characteristics are determined from a covariance matrix computed based on both active and passive data windows (Barnes and Hillebrand, 2003). An additional metric (e.g., pseudoT, Robinson and Vrba, 1999) is then calculated to contrast power between the two states. If one considers a source that is active in the active state and silent in the passive state, the estimate of its variance (using this method) will be the mean, not the maximum, of the two periods. That is, the prior for this source will be underestimated. A recent study (Sekihara et al., 2006) has attempted to effectively project out the data within a passive state before the calculation of the covariance matrix, giving rise to an LoD implementation of the beamformer. This means that the priors are effectively based only on sources that have changed in

[7] Having said this, the smoothness constraint in the current context applies to dipoles whose orientation is constrained to be normal to the mesh. This means that dipoles that are situated either side of a gyrus or sulcus are often estimated with opposite sign. Even though our measure of proximity was based on distance within the mesh (rather than 3D Euclidean distance), such dipoles may still be close enough that constraining them to have similar signs works against the data.

[8] Another selection issue in such group-based analyses is whether to scale individual solutions by, for example, some function of their mean value (Park et al., 2002). We do not think this is appropriate here because our forward models should accommodate any artifactual scaling differences owing to, for example, different distances of subjects' heads from the sensors, and we would expect that any remaining proportional differences across subjects reflect true intersubject variation and therefore should not be discounted.

amplitude from the active to the passive state. However, the strong assumption of the LoD approach that the passive state (or control condition) contains no data of interest suggests that, while the LoD approach may be appropriate in well-controlled experiments that elicit different temporal responses, it does not have the generality of the DoL approach in looking for spatiotemporal differences in localized responses.

Signed vs. unsigned

We had expected the signed analyses to attenuate the expected response over subjects because the sign of the dipole activity encodes the subject-specific orientation of the cortical surface (at each dipole), and this could vary randomly from subject to subject. Moreover, the 3D spatial smoothing might be expected to diminish the amplitude of nearby sources that have opposite signs (e.g., owing to opposing orientations when on opposite sides of a gyrus or sulcus). However, we were wrong; the signed summary statistics appeared robust to the problems we had anticipated. For example, while there were slightly more suprathreshold voxels in the unsigned DoL results than the signed DoL results, when using a PPM threshold of 95% (Fig. 7), the only additional region in the unsigned analysis was a right anterior medial temporal region; and any differences in left orbitofrontal cortex or left temporal pole, as found in the signed DoL analysis, did not survive the threshold in the unsigned DoL analysis. These empirical observations suggest that, at least in the context of the analysis framework used here, the heuristic 'smoothing distributed dipoles in 3D space will attenuate and misplace source estimates' need not always hold.

Localization of M170

The main purpose of this paper was to establish the face validity of our forward models and the summary statistics they furnish rather than to make strong neuroscientific claims about the generators of the M170 in the population of healthy adults (particularly given the small number of subjects here). Indeed, we focused on the hyperparameters because they determine the covariances at the source or sensor level, and any differences among subjects are expressed explicitly in terms of differences in the hyperparameters. Nonetheless, it is worth considering the group solutions in relation to previous findings.[9] The main loci of interest were bilateral fusiform cortex (apparently more extensive on the right), more anterior ventral temporal cortex and medial orbitofrontal cortex (Table 1 and Fig. 7). The former region is consistent with fMRI findings that have consistently identified face-responsive regions in fusiform cortex, which are generally stronger and more extensive in the right hemisphere. The more anterior temporal regions are also consistent with some fMRI findings, including those from the same paradigm (Henson et al., 2003). The orbitofrontal difference is less often reported in fMRI studies, though such an early differential orbitofrontal response has been attributed to initial processing of visual objects (Bar et al., 2006), and neurons that respond to faces have been reported in Macaque orbitofrontal cortex (Scalaidhe et al., 1999). We were surprised not to observe strong lateral posterior temporal differences in our group analyses since the posterior superior temporal sulcus was implicated by our previous fMRI and EEG data from the same paradigm (Henson et al., 2003). One possibility is that this source has a large radial component, which is more easily visible with EEG than with MEG (Watanabe et al., 2005; though see Hillebrand and Barnes, 2002).

Finally, despite some similarities between the group statistics for the MEG and fMRI data (Fig. 7), there are differences not only in the number, but also the location, of activated regions. One obvious reason for the greater number of "activated" regions in the fMRI data is the fact that the BOLD signal integrates neural activity across several seconds, meaning that the fMRI results include differences between intact and scrambled faces that arose subsequent to the M170. One could increase the time window of the MEG data localized to explore this. Another reason for the differences (aside from the different number and identity of subjects and possible threshold effects) may relate to the asymmetry between electromagnetic and metabolic responses. This asymmetry rests on the notion that some neuronal responses may be seen by fMRI but not MEG, due to a radial orientation or incoherent dynamics, while some electromagnetic sources may not be expressed metabolically (e.g., phase synchronization with no changes in the spectral density of ongoing dynamics). This asymmetry might explain differences in locations of activations (e.g., the medial orbitofrontal activity in the signed MEG results but not fMRI results).

Future directions

Here we have used the reproducibility over subjects to validate subject-specific forward models (i.e., individually defined meshes and forward models and individually defined time windows). In other work, we introduced the possibility of using a canonical mesh (in MNI space) that is deformed via inverse spatial normalization to match individual brains (Mattout et al., in press). Within this framework, one can apply, for example, further constraints that live only in template space, such as the results of previous group-based SPMs of fMRI data. Furthermore, one can apply constraints (such as MSP, for example) that could be derived from the average evoked response across subjects. Indeed, one could even use the source solution estimated from the average data (when also using an average mesh and forward model) and re-enter the conditional covariance of these sources as an additional covariance component for the individual subject models. In principle, this will stabilize the results in terms of intersubject variability and improve the between-subject inference. Note that inferences about between-condition effects are unbiased by this procedure because every condition and subject-specific source reconstruction would have the same covariance components. This would be a first step towards a proper hierarchical model, in which a third (subject) level is added to the two (sensor and source) levels considered in this paper. These models can be estimated using PEB in exactly the same way as described above to provide direct Bayesian inference about grand mean responses over subjects. In other words, the procedure described above is a Bayesian inversion at the within-subject level, with separate inference at the between-subject level; a full hierarchical model would embody intersubject

[9] It is important to note that we defined the M170 as the earliest difference between ERFs to intact versus scrambled faces. Some researchers have defined the M170 (or N170 in EEG) by the peak response around 170 ms to faces alone and localized this peak relative to pre-stimulus baseline. In this case, the localization is likely to include early visual regions that respond to any nonspecific change in the visual field. Other researchers have defined the M170 as the difference between faces and other non-face objects. This has the advantage of controlling for the presence of a perceived object. However, non-face objects (such as houses) usually have the disadvantage that they entail differences in low-level visual properties such as spatial frequency spectra.

variations about a grand mean and require a single inversion. A further important development will be to relax some of the hard constraints. For example, the orientation of dipoles in each mesh could be constrained by a prior distribution of orientations, centered on the normal to the mesh surface (Phillips et al., 2005; Lin et al., 2006). The meshes themselves could also be refined, perhaps involving a prior reduction using spatial basis functions (Phillips et al., 2002a) or MSP (Mattout et al., 2005) or some form of automatic relevance detection (Friston et al., 2006b).

Conclusion

In summary, we have addressed some key issues in population inference about responses evoked in distributed brain systems using MEG. At the within-subject level, we focused on the constraints or priors required for Bayesian inversion of distributed source models. We used restricted maximum likelihood (ReML) estimates of these components to show, qualitatively, that their relative importance reproduces over subjects. Specifically, baseline estimates of sensor noise proved important in maximizing the model log evidence. Constraints at the source level were very similar when projected to the sensor level, making it difficult to adjudicate between, for example, MSP vs. depth weighting constraints. This suggests that source level constraints may need to be more distinct (e.g., using sparse spatial constraints). At the between-subject level, we focused on whether to contrast localizations or to localize contrasts and whether to test the directed (signed) estimates of the sources or their absolute (unsigned) values. There are reasonable arguments for contrasting unsigned estimates of condition-specific localizations. However, the difference between contrasting localizations and localizing contrasts for signed values was small. One possible reason is that, at around 170 ms, faces evoke activity in the same areas as scrambled faces, but to a greater level. Surprisingly, the use of unsigned estimates was not clearly superior to the use of signed estimates, suggesting that differences in source orientation across subjects are not major factor within the context of the normalization, smoothing and group averaging employed here.

Acknowledgments

This work is funded by the UK Medical Research Council and Wellcome Trust. We thank Jean-Francois Mangin, Stefan Kiebel and Stef Hassel.

References

Ashburner, J., Friston, K.J., 1999. Nonlinear spatial normalization using basis functions. Hum. Brain Mapp. 7 (4), 254–266.

Ashburner, J., Friston, K.J., 2005. Unified segmentation. Neuroimage 26 (3), 839–851.

Baillet, S., Mosher, J.C., Leahy, R.M., 2001. Electromagnetic brain mapping. IEEE Signal Process. Mag. 18, 14–30.

Bar, M., Kassam, K.S., Ghuman, A.S., Boshyan, J., Schmid, A.M., Dale, A.M., Hamalainen, M.S., Marinkovic, K., Schacter, D.L., Rosen, B.R., et al., 2006. Top–down facilitation of visual recognition. Proc. Natl. Acad. Sci. U. S. A. 103 (2), 449–454.

Barnes, G.R., Hillebrand, A., 2003. Statistical flattening of MEG beamformer images. Hum. Brain Mapp. 18, 1–12.

Dale, A.M., Sereno, M., 1993. Improved localization of cortical activity by combining EEG and MEG with MRI surface reconstruction: a linear approach. J. Cogn. Neurosci. 5, 162–176.

Dale, A.M., Liu, A.K., Fischl, B.R., Buckner, R.L., Belliveau, J.W., Lewine, J.D., Halgren, E., 2000. Dynamic statistical parametric mapping: combining fMRI and MEG for high-resolution imaging of cortical activity. Neuron 26, 55–67.

Fischl, B., Sereno, M.I., Tootell, R.B., Dale, A.M., 1999. High-resolution intersubject averaging and a coordinate system for the cortical surface. Hum. Brain Mapp. 8, 272–284.

Friston, K.J., Penny, W.D., 2003. Posterior probability maps and SPMs. NeuroImage 19 (3), 1240–1249.

Friston, K.J., Penny, W., Phillips, C., Kiebel, S., Hinton, G., Ashburner, J., 2002. Classical and Bayesian inference in neuroimaging: theory. NeuroImage 16 (2), 465–483.

Friston, K., Henson, R., Phillips, C., Mattout, J., 2006a. Bayesian estimation of evoked and induced responses. Hum. Brain Mapp. 27 (9), 722–735.

Friston, K., Mattout, J., Trujillo-Barreto, N., Ashburner, J., Penny, W., 2006b. Variational free energy and the Laplace approximation. NeuroImage 34, 220–234.

Fuchs, M., Wagner, M., Kohler, T., Wischmann, H.A., 1999. Linear and nonlinear current density reconstructions. J. Clin. Neurophysiol. 16, 267–295.

Gorodnitsky, I.F., George, J.S., Rao, B.D., 1995. Neuromagnetic source imaging with FOCUSS: a recursive weighted minimum norm algorithm. Electroencephalogr. Clin. Neurophysiol. 95 (4), 231–251.

Hämäläinen, M., Ilmoniemi, R. 1984. Interpreting measured magnetic fields of the brain: estimates of current distributions: Helsinki University of Technology.

Hauk, O., 2004. Keep it simple: a case for using classical minimum norm estimation in the analysis of EEG and MEG data. NeuroImage 21 (4), 1612–1621.

Henson, R.N., Goshen-Gottstein, Y., Ganel, T., Otten, L.J., Quayle, A., Rugg, M.D., 2003. Electrophysiological and haemodynamic correlates of face perception, recognition and priming. Cereb. Cortex 13, 793–805.

Hillebrand, A., Barnes, G.R., 2002. A quantitative assessment of the sensitivity of whole-head MEG to activity in the adult human cortex. NeuroImage 16, 638–650.

Hillebrand, A., Barnes, G.R., 2003. The use of anatomical constraints with MEG beamformers. NeuroImage 20, 2302–2313.

Hillebrand, A., Singh, K.D., Holliday, I.E., Furlong, P.L., Barnes, G.R., 2005. A new approach to neuroimaging with magnetoencephalography. Hum. Brain Mapp. 25 (2), 199–211.

Holmes, A.P., Friston, K.J., 1998. Generalisability, random effects and population inference. NeuroImage 7 (4), 754.

Lin, F.-H., Belliveau, J.W., Dale, A.M., Hämäläinen, M.S., 2006. Distributed current estimates using cortical orientation constraints. Hum. Brain Mapp. 27, 1–13.

Liu, J., Higuchi, M., Marantz, A., Kanwisher, N., 2000. The selectivity of the occipitotemporal M170 for faces. NeuroReport 11 (2), 337–341.

Mattout, J., Pelegrini-Issac, M., Garnero, L., Benali, H., 2005. Multivariate Source Prelocalization (MSP): use of functionally informed basis functions for better conditioning the MEG inverse problem. NeuroImage 26, 356–373.

Mattout, J., Phillips, C., Penny, W.D., Rugg, M.D., Friston, K.J., 2006. MEG source localization under multiple constraints: an extended Bayesian framework. NeuroImage 30 (3), 753–767.

Mattout, J., Henson, R., Friston, K., in press. Canonical source reconstruction for MEG. Computational Intelligence and Neuroscience.

Nichols, T.E., Holmes, A.P., 2002. Nonparametric permutation tests for functional neuroimaging: a primer with examples. Hum. Brain Mapp. 15, 1–25.

Nunez, P., 1981. Electric Fields of The Brain: The Neurophysics of EEG. Oxford Univ. Press, New York.

Park, H.J., Kwon, J.S., Youn, T., Pae, J.S., Kim, J.J., Kim, M.S., Ha, K.S., 2002. Statistical parametric mapping of LORETA using high density EEG and individual MRI: application to mismatch negativities in schizophrenia. Hum. Brain Mapp. 17 (3), 168–178.

Pascual-Marqui, R.D., Michel, C.M., Lehmann, D., 1994. Low resolution electromagnetic tomography: a new method for localizing electrical activity in the brain. Int. J. Psychophysiol. 18 (1), 49–65.

Phillips, C., Rugg, M.D., Friston, K.J., 2002a. Anatomically informed basis

functions for EEG source localization: combining functional and anatomical constraints. NeuroImage 16 (3 Pt 1), 678–695.
Phillips, C., Rugg, M.D., Friston, K.J., 2002b. Systematic regularization of linear inverse solutions of the EEG source localization problem. NeuroImage 17 (1), 287–301.
Phillips, C., Mattout, J., Rugg, M.D., Maquet, P., Friston, K.J., 2005. An empirical Bayesian solution to the source reconstruction problem in EEG. NeuroImage 24 (4), 997–1011.
Robinson, S.E., Vrba, J., 1999. Functional neuroimaging by synthetic aperture magnetometry (SAM). In: Yoshimoto, T., Kotani, M., Kuriki, S., Karibe, H., Nakasato, N. (Eds.), Recent Advances in Biomagnetism. Tohoku Univ. Press, Sendai, pp. 302–305.
Sarvas, J., 1987. Basic mathematical and electromagnetic concepts of the biomagnetic inverse problem. Phys. Med. Biol. 32, 11–22.
Sato, M.A., Yoshioka, T., Kajihara, S., Toyama, K., Goda, N., Doya, K., Kawato, M., 2004. Hierarchical Bayesian estimation for MEG inverse problem. NeuroImage 23 (3), 806–826.
Scalaidhe, S.P., Wilson, F.A., Goldman-Rakic, P.S., 1999. Face-selective neurons during passive viewing and working memory performance of rhesus monkeys: evidence for intrinsic specialization of neuronal coding. Cereb. Cortex 9 (5), 459–475.
Scherg, M., Berg, P., 1991. Use of prior knowledge in brain electromagnetic source analysis. Brain Topogr. 4, 143–150.
Schimmel, H., 1967. The (±) Reference: accuracy of estimated mean components in average response studies. Science 157, 92–94.
Sekihara, K., Hild, K.E., Nagarajan, S.S., 2006. A novel adaptive beamformer for MEG source reconstruction effective when large background activities exist. IEEE Trans. Biomed. 53, 1755–1764.
Singh, K.D., Barnes, G.R., Hillebrand, A., 2003. Group imaging of task-related changes in cortical synchronisation using nonparametric permutation testing. NeuroImage 19 (4), 1589–1601.
Talairach, J., Tournoux, P., 1988. Co-Planar Stereotaxic Atlas of the Human Brain. George Thieme Verlag, Stuttgart, pp. 1–122.
Watanabe, S., Miki, K., Kakigi, R., 2005. Mechanisms of face perception in humans: a magneto- and electro-encephalographic study. Neuropathology 25 (1), 8–20.