Information-based functional brain mapping

Nikolaus Kriegeskorte*†, Rainer Goebel‡, and Peter Bandettini*

*Section on Functional Imaging Methods, Laboratory of Brain and Cognition, National Institute of Mental Health, Building 10, Room 1D80B, 10 Center Drive MSC 1148, Bethesda, MD 20892-1148; and ‡Department of Cognitive Neuroscience, Faculty of Psychology, Universiteit Maastricht, Universiteitssingel 40, 6229 ER, Maastricht, The Netherlands

The development of high-resolution neuroimaging and multielectrode electrophysiological recording provides neuroscientists with huge amounts of multivariate data. The complexity of the data creates a need for statistical summary, but the local averaging standardly applied to this end may obscure the effects of greatest neuroscientific interest. In neuroimaging, for example, brain mapping analysis has focused on the discovery of activation, i.e., of extended brain regions whose average activity changes across experimental conditions. Here we propose to ask a more general question of the data: Where in the brain does the activity pattern contain information about the experimental condition? To address this question, we propose scanning the imaged volume with a ''searchlight,'' whose contents are analyzed multivariately at each location in the brain.

neuroimaging | functional magnetic resonance imaging | statistical analysis

Functional brain mapping has evolved from the idea that the brain consists of functionally specialized macroscopic regions. In early neuroimaging experiments using positron emission tomography, brain activity was measured at a spatial resolution in the centimeter range. At this resolution, the volume elements (voxels) were similar in size to the putative functional regions, so only the spatial-average activity of a region could be studied. In the classical approach to functional brain mapping, therefore, the experiment is designed to activate a functional region as a whole. The region is then localized by computing an activation statistic for each location of the imaging volume and thresholding the resulting statistical map. We refer to this approach as activation-based.

With the advent of functional magnetic resonance imaging (fMRI), spatial resolution increased. Standard functional measurements were performed with voxel widths of ≈4 mm in each dimension. Although a typical functional region at this resolution is covered by multiple voxels, standard fMRI analysis to this day has remained true to the activation-based approach, in which a region is assumed to become active as a whole. This approach manifests itself in the widespread investigation of the spatially averaged activity for regions of interest. Event-related average time courses and bar graphs depicting the activity across conditions, for example, reflect a region's spatially averaged activity.

The assumption that functional regions extended across multiple voxels will become activated as a whole also plays a key role in statistical inference at the level of whole maps in several established methods, including the widespread statistical parametric mapping (refs. 1–3; see also ref. 4). The extended-activations assumption motivates the spatial smoothing of the data, which is standardly performed. Spatial smoothing accentuates extended activations by removing the ''salt-and-pepper'' fine structure of the activity patterns, which is treated as noise. As a positive side effect, the resulting reduction of the data's spatial complexity alleviates the multiple-comparisons problem: With smoothing, controlling the false-positives rate (when performing a statistical test at each location of a functional volume) requires weaker correction, which entails greater statistical power for detection of extended activations. Furthermore, spatial smoothing helps combine data from different subjects

coregistered in Talairach space (see ref. 5), where corresponding functional regions can be off by many millimeters between subjects. In a typical group analysis, data are spatially smoothed by convolution with a Gaussian kernel of 8-mm full width at half maximum.

Although smoothing greatly reduces the information content of the data, the local combination of signals it provides is necessary. If smoothing is omitted in a standard voxelwise univariate fMRI analysis, statistical sensitivity suffers, and less of the activated volume is detected (Fig. 2). Upon lowering the threshold, the activation maps show salt-and-pepper patterns, which are hard to distinguish from noise, inconsistent across subjects, and impossible to report verbally. Nevertheless, these fine-scale patterns of weak effects may contain neuroscientifically relevant information.

The amount of information removed by smoothing fMRI data to the scale of functional regions increases with growing spatial resolution of the measurement. Today, a voxel width of 2 mm in each dimension is robustly achievable with standard clinical scanners at 3-T field strength. Using high field strengths, fMRI is invading the submillimeter range. However, so long as theory, experiment, and analysis are based on the idea of functional regions activated as a whole, the high-resolution information fMRI now provides will not be used. The neuroscientific exploitation of this information poses new conceptual, experimental, and statistical challenges.

Unsmoothed fMRI activity patterns are often patchy and inconsistent across subjects. They are not well accommodated by considering the patches as activated functional regions at a smaller scale. Should fine-scale activity patterns then be thought of as distributed representations? How can informative patterns be sensitively detected, distinguishing them from noise of similar salt-and-pepper appearance? How can data be related between subjects when the available common spaces (e.g., Talairach space; ref. 5) lack precision and, more fundamentally, fine-scale activity patterns, like fingerprints, may be unique to each individual? We propose to approach these challenges by abstracting from the actual patterns of activity in a local neighborhood and considering the information they convey about the experimental condition.

Information contained in distributed fMRI activity patterns has been analyzed for extended predefined regions (see refs. 6–8), heuristically chosen discontiguous sets of voxels (see refs. 9, 10), and global patterns (e.g., refs. 11 and 12). Here we show how information can be continuously mapped throughout a functional volume. Instead of searching the functional volume for regions whose spatially averaged activity changes across conditions (activation-based approach), we ask, more generally,

---

NEUROSCIENCE

A Detection performance of mapping techniques
(receiver operating characteristics for unsmoothed data)

size of region containing the effect

90, 270 voxels

10, 30 voxels

sensitivity

1-specificity

2 SE

low (0.1, 0.2)   high (0.3, 0.4)
functional-contrast-to-noise ratio

B The effect of data smoothing on detection
(area under ROC as a function of the degree of smoothing)

size of region containing the effect

90, 270 voxels

10, 30 voxels

area under ROC

2 SE

smoothing-Gaussian FWHM [mm]

low (0.1, 0.2)   high (0.3, 0.4)
functional-contrast-to-noise ratio

C Summary

area under ROC

region size

functional-contrast-to-noise ratio

high (0.3, 0.4)

low (0.1, 0.2)

D Color coding of mapping techniques

searchlight statistic:

average absolute t value   Mahalanobis distance

size of spherical searchlight

radius [mm]	number of voxels
6	123
5	81
4	33
3	19
2	7
1	1

locally multivariate (information-based)

univariate t contrast (activation-based)

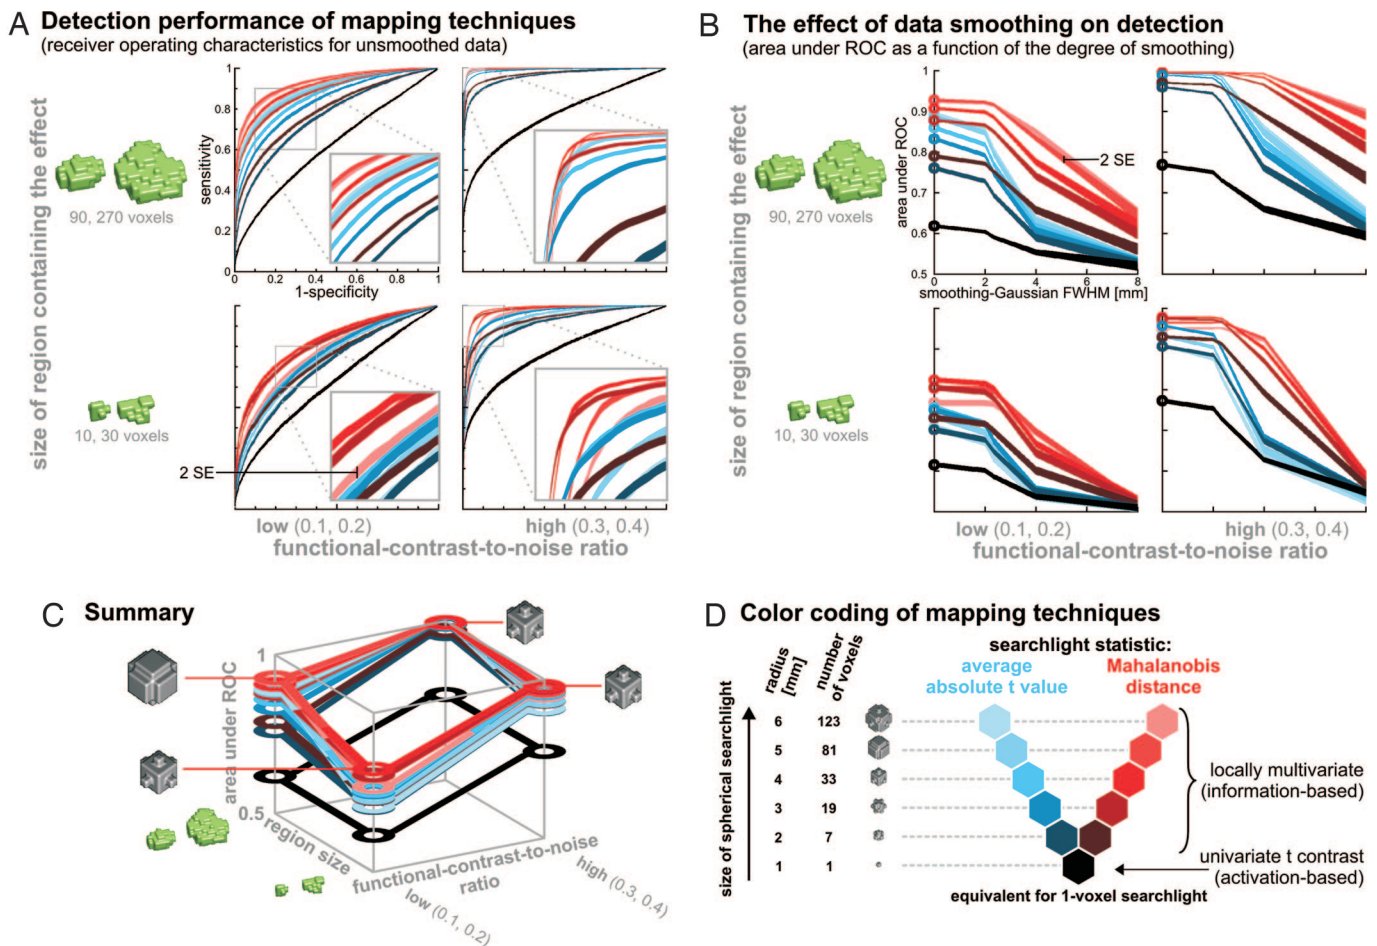equivalent for 1-voxel searchlight

Fig. 1. Simulated fMRI data. (*A–C*) The performance of activation- (black lines) and information-based mapping (colored lines) at detecting focally distributed effects (for color coding, see *D*). *A* and *B* each show an arrangement of four plots, where the top row displays average results obtained for large effect regions, and the bottom row displays average results obtained for small effect regions. The left and right columns display average results obtained for the two lower and the two higher functional contrast-to-noise ratios, respectively. *A* shows ROCs for the case of unsmoothed data. *B* shows the effect of spatial smoothing. The vertical axis here represents the area under the ROC. For the included case of no smoothing (i.e., full width at half maximum = 0), the areas under the ROC (marked as circles) correspond to the ROCs shown in *A*. Note that smoothing degrades performance for all techniques. For the crucial case of unsmoothed data, *C* summarizes the essential results by visually relating the detection performances afforded by the different techniques for small and large regions and low and high functional contrast-to-noise ratio. The searchlights yielding optimal performance in each case are shown in gray (4- or 5-mm radius). The circles in *C* replicate the circles in *B* reflecting the areas under the ROCs shown in *A*. In *A* and *B*, the line thickness measured vertically is 2 SEMs obtained by repeating the simulations and analyses 40 times with fresh noise. The shapes of the regions shown in green in *A–C* are exactly those used in the simulations. *D* illustrates the color coding in *A–C* and shows the searchlights used.

where in the brain the local activity pattern changes. Our approach is information-based, in that it localizes regions containing information about the experimental condition. This information can lie in a change of the spatial-average activity level or in a change of the activity pattern.

An Implementation of Information-Based Functional Mapping

The combination of signals from all voxels within a region can afford greater statistical power. When the target of the analysis is to find activated regions, local signals can be combined by local averaging, i.e., by smoothing of the data. Here the target of the analysis is to find informative regions, and smoothing would remove information in the spatial fine structure of the data. We therefore combine local signals using a multivariate statistic that compares the activity patterns among conditions.

Spherical Multivariate Searchlight. To obtain a continuous map in which informative regions are marked, we move a spherical multivariate "searchlight" through the measured volume. The searchlight is centered on each voxel in turn. A searchlight of

4-mm radius, for example, will comprise 33 voxels of 2-mm width in each dimension (see Fig. 1*D*). To combine the signals from all voxels falling into the searchlight, a multivariate effect statistic is computed at each location. For each voxel in the volume, the resulting map shows how well the multivariate signal in the local spherical neighborhood differentiates the experimental conditions.

Multivariate Effect Statistics. In the present implementation, we focus on information carried by the spatial pattern of activity within each region. Furthermore, we assume that the activity pattern, although covered in noise, is replicable and can, thus, be estimated as the average across trials within each condition. Instead of averaging single-trial response patterns, we estimate the response amplitude at a given voxel as a least-squares fit of a predictor time course for each condition. This more general approach can handle all kinds of experimental design, including rapid event-related ones. The predictors are based on a linear model of the temporal structure of the hemodynamic response (see ref. 13). Assuming replicable spatial response patterns and

Kriegeskorte *et al*.

a temporal response model as we do here is standard in fMRI analysis.

Consider the simple case of two experimental conditions. What statistic should be used to represent the difference between the two activity patterns within the searchlight?

Average absolute t value. One approach would be to first perform a conventional linear-regression contrast analysis to obtain a t value for each voxel. The t value represents the activity difference between the two conditions in that voxel. One may then average the absolute t values within the searchlight to obtain a measure of the difference between patterns. Taking the absolute values is crucial here, because opposite effects would cancel out if the raw t values were averaged. For instance, the average t value might be zero if activity in half the voxels is greater during condition A and activity in the other half greater during condition B, a case in which the activity pattern contains a lot of information about the condition. The average absolute t value within the searchlight combines local signals with sensitivity to the spatial fine structure. However, it does not take the covariance structure of the noise into account.

Mahalanobis distance. A natural choice that accounts for noise covariance is the Mahalanobis distance, a multivariate generalization of the t value. To obtain a t value, the estimated difference between two activity levels is divided by its estimated variability. To obtain the Mahalanobis distance, the estimated difference vector between two activity patterns is normalized by the estimated multivariate noise covariance (Eq. **3**).

Randomization Test and False-Discovery-Rate (FDR) Thresholding.

The spherical-searchlight mapping of the Mahalanobis distance provides a descriptive information-based map. We perform statistical inference by a randomization technique to obtain a map of P values. This approach is attractive, because its validity does not depend on the multivariate distribution of the data, and because it correctly handles colored noise. To obtain an inferential map accounting for multiple comparisons, the P map from randomization is thresholded to ensure the average FDR (i.e., the expected proportion of falsely marked voxels among all marked voxels) will not exceed $q = 0.05$ (see refs. 14 and 15).

Validation of the Proposed Method

Simulated fMRI Data. *Rationale.* The classical activation-based approach is based on the assumption that effects are spatially smooth (i.e., in the low spatial-frequency band). Here we explore how fMRI activity-pattern information in any spatial-frequency band can be used for the localization of functional regions. We therefore simulate effects with no spatial-frequency bias: The effects have equal power in all spatial-frequency bands, containing information in their locally averaged component as well as in their spatial fine structure.

We simulate a time course of 3D volumes as might be obtained in a slow event-related fMRI experiment with two conditions. Each condition is associated with a spatial pattern of response, which fades in and out according to a linear model of the temporal structure of the hemodynamic response. The two spatial patterns are composed of values drawn independently from a Gaussian distribution and confined to realistically shaped regions of four different sizes (Fig. 1 *A* and *B*), which are embedded in a functional volume. The signal data are added to spatiotemporal noise with a realistic degree of correlation between adjacent voxels.

Results of simulation. The main results reported below are apparent by visual inspection of the statistical maps (Fig. 3, which is published as supporting information on the PNAS web site). To quantitatively assess how well different local descriptive statistics distinguish between the effect-region voxels and the surrounding pure-noise voxels, we use receiver-operating characteristics (ROCs; Fig. 1). Because the effect regions are known for the simulations, the proportion of correctly detected voxels among all effect-region voxels (the sensitivity) and the proportion of correctly rejected voxels among all pure-noise voxels (the specificity) can be determined for all possible thresholds. An ROC shows the resulting tradeoff relation between sensitivity and specificity (Fig. 1*A*). As a single-number measure of how well a given statistic distinguishes between effect regions and pure noise, we use the area under the ROC (Fig. 1 *B* and *C*; for more details on this standard measure, see *Supporting Text*, which is published as supporting information on the PNAS web site, *ROC analysis of statistical maps from simulated fMRI data*). For the broad parameter space tested in our simulation, we observe the following.

Information-based mapping using the Mahalanobis distance dominates the other techniques. With an appropriate searchlight radius, Mahalanobis mapping (red curves in Fig. 1 *A–C*) is most sensitive at all specificities for all region sizes, functional-contrast-to-noise ratios, and degrees of data smoothing. Best detection performance is achieved using the intermediate-size searchlight (4-mm radius), except when the regions are large (90 or 270 voxels), and the functional-contrast-to-noise ratio is low. In that case, a slightly larger searchlight (5-mm radius) yields slightly better diagnosticity (Fig. 1*C*).

Activation-based mapping performs worse than information-based approaches. Univariate t mapping (black curves in Fig. 1 *A–C*) performs worse than all information-based statistics (colored curves) for all combinations of region size (small or large) and functional contrast (high or low) for unsmoothed as well as moderately smoothed data. This is because the univariate approach fails to benefit from local spatial combination of signals.

Spatial smoothing of the data monotonically degrades diagnosticity of all statistics. For all combinations of region size and functional contrast-to-noise ratio, smoothing degrades performance of all statistics tested, including the univariate t value (Fig. 1*B*). This reflects the fact that effects in our simulation were equally strong in all spatial-frequency bands up to the Nyquist limit imposed by voxel size. Data smoothing filters out the high spatial-frequency component of the effect patterns and, thus, reduces the overall effect energy.

The benefit of spatial signal combination within the searchlight is greater for larger regions. For univariate mapping, detection is equally poor independent of region size (Fig. 1*C*). The information-based techniques using a spherical searchlight of 2-mm radius or larger (colored curves) perform markedly better than the activation-based univariate mapping (black curves) for small regions, and the benefits of spatial combination of signals within the searchlight are even greater for larger regions. This holds for low and high functional contrast and for all searchlight sizes (except the degenerate one-voxel searchlight, for which both information-based approaches become diagnostically equivalent to the activation-based univariate t mapping).

A searchlight of 4-mm radius yields optimal or near-optimal detection performance. When the size of the spherical searchlight is very inappropriate to the size of the region (small region scanned with large searchlight or vice versa), detection performance suffers. The optimal searchlight radius depends on the shape and size of the regions to be detected. But for all region shapes and sizes simulated, the 4-mm searchlight performed optimally or near optimally (Fig. 1*C*). Detection, thus, does not strongly depend on an exact match between region and searchlight.

Even a small searchlight affords a large increase of detection performance. Spatial signal combination with even a 2-mm-radius searchlight comprising only seven voxels drastically improves performance in comparison to univariate t mapping (except for highly smoothed data with performance near chance level).

A multivariate noise model can improve detection performance.
The Mahalanobis distance outperforms the average absolute *t* value for any given searchlight size. This is because the Mahalanobis distance is based on a multinormal noise model, which picks up on the realistic correlation present in the simulated noise. The noise model helps separate signal from noise.

Real fMRI Data. *Rationale.* The simulated activity-pattern effects were better detected by information-based mapping. But do such effects actually occur in fMRI data? To test this, we analyzed real fMRI data using activation- and information-based techniques. Subjects watched photos of faces and houses while being scanned. If the extended activations expected from the literature (face and place regions; refs. 16–18) are the only effects present, both approaches will find the major blobs. However, activation-based mapping will benefit from its smoothness assumption and mark additional weaker extended activations. So the activation-based analysis should mark more voxels, essentially a superset. If, on the other hand, fine-grained pattern effects are widespread alongside the expected extended activations, information-based mapping, with its sensitivity to effects in all spatial frequency bands, should find many voxels not highlighted in the activation-based map. Furthermore, we again expect activation-based mapping, with its implicit smoothness assumption, to detect some weak spatially smooth effects undetected by information-based mapping.

For the information-based analysis, we chose a searchlight radius of 4 mm, because this radius yielded the best detection performance in the simulation. For the activation-based analysis, the data were spatially smoothed with a spherical kernel of 4-mm radius, to exactly match the range of spatial combination of signals between the two approaches. (The resulting degree of smoothing is comparable to that obtained with the conventional choice of a Gaussian of 8-mm full width at half maximum.)

The activation-based analysis consisted in a conventional *t* contrast mapping. The information-based analysis used the Mahalanobis distance between the response patterns within the searchlight as the effect statistic. For activation- and information-based analyses, statistical inference was performed by the same randomization technique (see *Methodological Details*) requiring no distributional assumptions and correctly handling colored noise in either case. (For details on preprocessing of the data and design-matrix construction, see *Supporting Text*.)

Results of fMRI analysis. The activation-based map (Fig. 2*B*) highlights regions more strongly active during face than house perception and vice versa. The information-based map (Fig. 2*C*) highlights regions whose activity pattern distinguishes the two categories. Qualitatively, we observe that the information-based map highlights extended swaths of cortex, including the category-selective regions (see refs. 16–18) found by the activation-based mapping. This result is consistent with the findings of Haxby *et al.* (ref. 6; see also ref. 19). A quantitative comparison of the maps (Fig. 2 *D* and *E*; see also Fig. 4, which is published as supporting information on the PNAS web site) shows three sets of voxels: The set marked only by information-based mapping (red), the set marked by both techniques (yellow), and the set marked only by activation-based mapping (green).

Information-based mapping marks many voxels not marked by activation-based mapping. Of all voxels marked in either map, 35% are marked only in the information-based map (red in Fig. 2 *D* and *E*) on average for the group of 11 subjects analyzed. This indicates that there are many regions containing category information in the fine-grained structure of the activity patterns. The activation-based mapping does not detect these regions, because the category information is lost when the data are smoothed. Note, however, that performing conventional univariate analysis without smoothing (Fig. 2*A*) also does not lead to the detection of these regions, because the effects are weak in single voxels and therefore detected only with local spatial combination of signals.


Fig. 2. Real fMRI data. (*A–C*) Activation- and information-based mapping for a single subject. *A* and *B* show univariate *t* maps contrasting activity during perception of face and house images. The color scale linearly reflects the *t* value for the contrast between faces and houses (see color bar) for voxels above the FDR threshold. *A* shows the *t* map for unsmoothed data and *B* for data smoothed with a 4-mm-radius spherical kernel. Note that smoothing increases the number of voxels marked. *C* shows the information-based map of *P* values. The effect statistic is the Mahalanobis distance between face and house response patterns computed at each voxel for the contents of a 4-mm-radius spherical searchlight. The color scale linearly reflects the *P* value (see color bar). *D* shows which voxels are marked only by activation-based mapping (green), which only by information-based mapping (red), and which by both (yellow) for this subject. The sizes of these sets of voxels are related to each other in *E Left*. The analysis has been performed with similar results for all 11 subjects (*E Center* and *Right*). The slices shown in *A–D* are slices 4, 6, 8, 10, and 12 (in anatomically ascending order) of 15 axial slices acquired. The right side of each slice represents the right hemisphere.

Large overlap of the maps. Of all voxels marked in either map, 44% are marked in both maps (yellow in Fig. 2 *D* and *E*) on average for the group. This overlap is expected, because both techniques are sensitive to the extended activations known to be associated with the perception of faces and buildings.

Kriegeskorte *et al.*

Activation-based mapping marks voxels not marked by information-based mapping. Of all voxels marked in either map, 21% are marked only in the activation-based map (green in Fig. 2 *D* and *E*) on average for the group. This result is expected, because activation-based mapping utilizes a smoothness assumption, which makes it more sensitive to weak extended activations. That the information-based technique fails to detect these activations represents the statistical cost of broadening the focus of sensitivity to a more general class of effects.

Discussion

We propose information-based functional brain mapping, an alternative to the classical activation-based approach. The core idea is to localize functional regions carrying a particular type of information by scanning the entire volume with a multivariate searchlight. By simulation, we show that information-based mapping is more sensitive to focally distributed effects with equal power in all spatial-frequency bands than activation-based mapping. By experiment, we show that focally distributed effects better detected by information-based mapping actually occur in fMRI data. Information-based mapping is attractive whenever one is interested in finding where in the brain the regional spatial activity pattern differs across experimental conditions.

What Experiments Are Amenable to Information-Based Mapping?
Information-based mapping can be applied in the context of any fMRI experiment. For block designs and event-related designs with a nonrandom condition sequence, the randomization scheme will need to be adapted to the circumstances. Despite this general applicability, the added value of information-based analysis depends on the neuroscientific questions to be addressed and the experimental design chosen. If the experiment is designed to activate a particular functional region as a whole, an activation-based mapping may be a conceptually more natural and statistically more sensitive means of localizing the region. By contrast, experiments that target distributed representations are naturally suited for information-based mapping. The technique can be applied in all domains, including sensation, perception, motor control, and higher cognitive function.

Information Theory and the Generality of Multivariate Effect Measures.
It is natural to use information-theoretic measures to quantify activity-pattern information. For example, an estimate of the mutual information between the experimental conditions and the concomitant local activity patterns can serve as an effect measure. However, we refer to any analysis that detects activity-pattern information as information-based, whether effects are expressed in bits or by a different multivariate measure.

Inspired by the generality of the concept of information, the method may be extended to detect differences in the spatiotemporal pattern of activity and differences regarding aspects of the pattern distribution other than its centroid. However, broadening the focus of an analysis to a more general class of effects comes at a cost in sensitivity. We feel that the modest generalization of the classical activation-based approach proposed here is realistic for fMRI data and well worth its cost.

Spherical or Cortex-Patch Searchlight.
Information-based mapping statistically combines local signals within a multivariate searchlight. It is, thus, most sensitive to information in contiguous macroscopic regions well sampled by the searchlight. For optimal joint sampling of informative voxels, the radius of the spherical searchlight should reflect the size and compactness of the regions. However, our simulation shows that a radius of 4 mm yields near-optimal performance for small as well as large regions of realistic shape. When cortical information is targeted, using an explicit representation of each subject's cortical sheet (see refs. 20–22) and replacing the spherical searchlight by a cortex patch promises to improve the joint sampling of informative signals and enhance sensitivity.

Group Analysis.
A group-analytical extension of the information-based approach to functional mapping has the potential to address an important challenge functional brain mapping is currently faced with: Although human fMRI already operates at a resolution of 1–2 mm, the spatial reference frames relating locations in different individual brains have a much lower spatial precision. Information-based group analysis should therefore follow a two-scale approach. At the fine spatial scale of millimeters, activity patterns are assumed to be unique to each individual and therefore analyzed separately for each subject, as described here. At the coarse spatial scale of centimeters, single-subject information-based maps are combined in standard space to obtain a group summary and to increase statistical power.

Methodological Details

Information-Based Mapping. *Estimation of spatial activity patterns.*
To estimate the spatial activity pattern elicited during each condition, we fit a linear model to each voxel's time course by ordinary least squares. The model comprises a hemodynamic response predictor (ref. 13) for each condition and, optionally, additional predictors to model artifactual components such as trends, head-motion effects, and baseline shifts between measurement runs. For the volume as a whole, the linear model follows the usual trend:

$$Y = XB + E, \qquad [1]$$

where Y is the time-by-voxel data matrix, X is the time-by-predictor design matrix, B is the predictor-by-voxel matrix of beta weights, and E is the time-by-voxel matrix of errors. Least-squares minimization of $sum(diag(E^T E))$ leads to

$$\hat{B} = (X^TX)^{-1}X^TY, \qquad [2]$$

where $\hat{B}$ is the predictor-by-voxel matrix of β estimates. Each row of $\hat{B}$ that represents an experimental condition (as opposed to trend or head motion) contains an estimated spatial activity pattern.

Mahalanobis distance. Consider two conditions and their associated spatial activity patterns within the searchlight. The rows of $\hat{B}$ corresponding to the two conditions provide the activity-pattern estimates a_1 and a_2 (row vectors with one element per searchlight voxel). The Mahalanobis distance is defined as

$$\Delta^2 = (a_2 - a_1)\hat{\Sigma}^{-1}(a_2 - a_1)^T, \qquad [3]$$

where $\hat{\Sigma}$ is an estimate of the error covariance matrix Σ of the voxels within the searchlight. At each searchlight position, Σ is estimated with optimal shrinkage toward the diagonalized sample covariance using results of Ledoit and Wolf (see refs. 23 and 24). (This estimator is more stable than the sample covariance. However, results obtained using the latter instead are very similar for the data we analyzed.)

Randomization test. To obtain a map of P values, we simulate a null-hypothesis distribution of Mahalanobis distances for each subject and contrast of interest separately. First we compute a set of 1,001 temporal condition sequences that meet design-optimization constraints. A random element of this set is actually used for the rapid event-related experiment. The mapping is then performed 1,001 times by using the actual and the 1,000 alternative condition sequences. We collect the Mahalanobis distances of all brain voxels in all maps ($\approx 3 \times 10^7$ values for each contrast and subject), which form the randomization distribution, a model of the null distribution. For each voxel, the rank of

NEUROSCIENCE

its actual Mahalanobis distance in the randomization distribution determines its *P* value, as follows: *P* = percent rank/100. For the statistical justification of this approach and similar techniques applicable to nonrandom and block designs (with a few reasonable additional assumptions), see ref. 25.

Control of false positives. The *P* map (uncorrected for multiple comparisons) is thresholded to control the average FDR (see refs. 14, 15). The technique here was designed to assume independent or positively dependent tests. We checked the validity of our inferential mapping on resting-state human fMRI data, where all positives are false positives. The uncorrected false-positives rate and the average FDR were found to match their nominal levels.

Simulated fMRI Data. A simulated data set was generated by Eq. 1, where *X* embodies the temporal, *B* the spatial structure of the signal, and *E* the spatiotemporal noise.

Temporal structure of the signal. We simulate a two-condition slow event-related experiment. Events last 500 ms, and their onsets are separated by 16 s. There are 20 events per condition. The entire simulated experiment thus lasts 10 min 40 s. The condition sequence is random. The temporal resolution is one volume per 2 s. The time course of the fMRI signal associated with each condition is simulated by a linear hemodynamic-response model (see ref. 13), assuming instantaneous rectangular neural responses to each simulated event.

Spatial structure of the signal. The simulated functional volume consists of nine slices of a size of 128×128 voxels. The voxels are assumed to be 2 mm wide in each dimension. Along one dimension of the slice, we vary the size of the effect regions in four levels (10, 30, 90, and 270 voxels). Along the other dimension, we vary the functional contrast-to-noise ratio in four levels (0.1, 0.2, 0.3, and 0.4). Each slice is accordingly subdivided into 4×4 subblocks of dimensions $32 \times 32 \times 9$ voxels. Within each subblock representing the smallest and second-smallest region size, the region shapes are repeated four times, to provide sufficient data for the case of small regions. Within the effect regions, each condition is associated with a Gaussian white-noise effect pattern (rows of *B*). Outside the regions, there are no effects. The functional contrast-to-noise ratio is defined as the spatial average within the effect region of the absolute activity level at the maximum of the hemodynamic response divided by the temporal standard deviation of the noise. The shapes of the regions (exactly as shown in Figs. 1 and 3) were created using a region growing process prioritized by a Gaussian random field.

To obtain realistically compact randomly shaped regions whose thickness does not grossly exceed what one might expect of functional regions in the human cortex, the region growing was spatially biased by means of a disk-shaped pedestal embedded in the Gaussian random field.

Spatiotemporal structure of the noise. To match the correlation found between the residual time courses of neighboring voxels in real fMRI data, the noise is generated by slight spatial smoothing of spatiotemporal Gaussian white noise with a Gaussian kernel of 2.35-mm full width at half maximum. This leads to an adjacent-voxel linear correlation across time of 0.25, approximately matching that found in the residuals of our real fMRI data, which were acquired at the same resolution as simulated here.

Real fMRI Data. *Stimuli, design, and task.* Subjects continually fixated on a central cross while viewing images of faces and houses. We used a rapid event-related design with a basic trial duration of 3 s. For improved contrast estimation, the random stimulus-sequence generator ensured some temporal clustering of trials within each condition. Each image was presented for 400 ms. To enforce and monitor attentive viewing, we asked subjects to perform an anomaly detection task.

Subjects. The 11 subjects (five male, six female, ages 18–30) all had normal or corrected-to-normal vision. All received information about fMRI and gave their informed consent. The experiments and consent form were approved by the ethical committee of the Academisch Ziekenhuis (university hospital) associated with the Katholieke Universiteit Nijmegen (The Netherlands).

Measurements. We acquired 15 transversal functional slices with a Siemens Magnetom Trio scanner (3 T) using a single-shot gradient-echo echo-planar-imaging sequence for blood–oxygen-level-dependent fMRI. The pulse-sequence parameters were as follows: in-plane resolution, 2×2 mm^2; slice thickness, 2 mm (no gap); time between functional volumes, 1,500 ms; slice acquisition order, interleaved; field of view, 256×256 mm^2; acquisition matrix, 128×128; time to echo, 32 ms; flip angle, 75°. The analysis for each subject is based on one measurement run lasting 14.8 min. Imaging was performed at the Donders Centre for Cognitive Neuroimaging (Nijmegen, The Netherlands).

1. Worsley, K. J., Evans, A. C., Marrett, S. & Neelin, P. (1992) *J. Cereb. Blood Flow Metab.* **12,** 900–918.
2. Friston, K. J., Worsley, K. J., Frackowiak, R. S. J., Mazziotta, J. C. & Evans, A. C. (1994) *Hum. Brain Mapp.* **1,** 214–220.
3. Friston, K. J., Holmes, A. P., Worsley, K. J., Poline, J. P., Frith, C. D. & Frackowiak, R. S. J. (1995) *Hum. Brain Mapp.* **2,** 189–210.
4. Poline, J. B., Worsley, K. J., Evans, A. C. & Friston, K. J. (1997) *NeuroImage* **5,** 83–96.
5. Talairach, J. & Tournoux, P. (1988) *Co-Planar Stereotaxic Atlas of the Human Brain* (Thieme, New York).
6. Haxby, J. V., Gobbini, M. I., Furey, M. L., Ishai, A., Schouten, J. L. & Pietrini, P. (2001) *Science* **293,** 2425–2430.
7. Kamitani, Y. & Tong, F. (2005) *Nat. Neurosci.* **8,** 679–685.
8. Haynes, J. D. & Rees, G. (2005) *Nat. Neurosci.* **8,** 686–691.
9. Edelman, S., Grill-Spector, K., Kushnir, T. & Malach, R. (1998) *Psychobiology* **26,** 309–321.
10. Cox, D. D. & Savoy, R. L. (2003) *NeuroImage* **19,** 261–270.
11. Carlson, T. A., Schrater, P. & He, S. (2003) *J. Cognit. Neurosci.* **15,** 704–717.
12. Formisano, E., Esposito, F., Di Salle, F. & Goebel, R. (2004) *Magn. Reson. Imaging* **22,** 1493–1504.
13. Boynton, G. M., Engel, S. A., Glover, G. H. & Heeger, D. J. (1996) *J. Neurosci.* **16,** 4207–4221.
14. Genovese, C. R., Lazar, N. A. & Nichols, T. (2002) *NeuroImage* **15,** 870–878.
15. Benjamini, Y. & Hochberg, Y. (1995) *J. R. Stat. Soc. Ser. B* **57,** 289–300.
16. Puce, A., Allison, T., Gore, J. C. & McCarthy, G. (1995) *J. Neurophysiol.* **74,** 1192–1199.
17. Kanwisher, N., McDermott, J. & Chun, M. M. (1997) *J. Neurosci.* **17,** 4302–4311.
18. Epstein, R. & Kanwisher, N. (1998) *Nature* **392,** 598–601.
19. Ishai, A., Ungerleider, L. G., Martin, A., Schouten, J. L. & Haxby, J. V. (1999) *Proc. Natl. Acad. Sci. USA* **96,** 9379–9384.
20. Van Essen, D. C., Drury, H. A., Joshi, S. & Miller, M. I. (1998) *Proc. Natl. Acad. Sci. USA* **95,** 788–795.
21. Dale, A. M., Fischl, B. & Sereno, M. I. (1999) *NeuroImage* **9,** 179–194.
22. Kriegeskorte, N. & Goebel, R. (2001) *NeuroImage* **14,** 329–346.
23. Ledoit, O. & Wolf, M. (2003) *J. Empir. Finance* **10,** 603–621.
24. Schaefer, J. & Strimmer, K. (2005) *Stat. App. Genet. Mol. Biol.* **4,** No. 1, Article 32.
25. Nichols, T. E. & Holmes, A. P. (2001) *Hum. Brain Mapp.* **15,** 1–25.
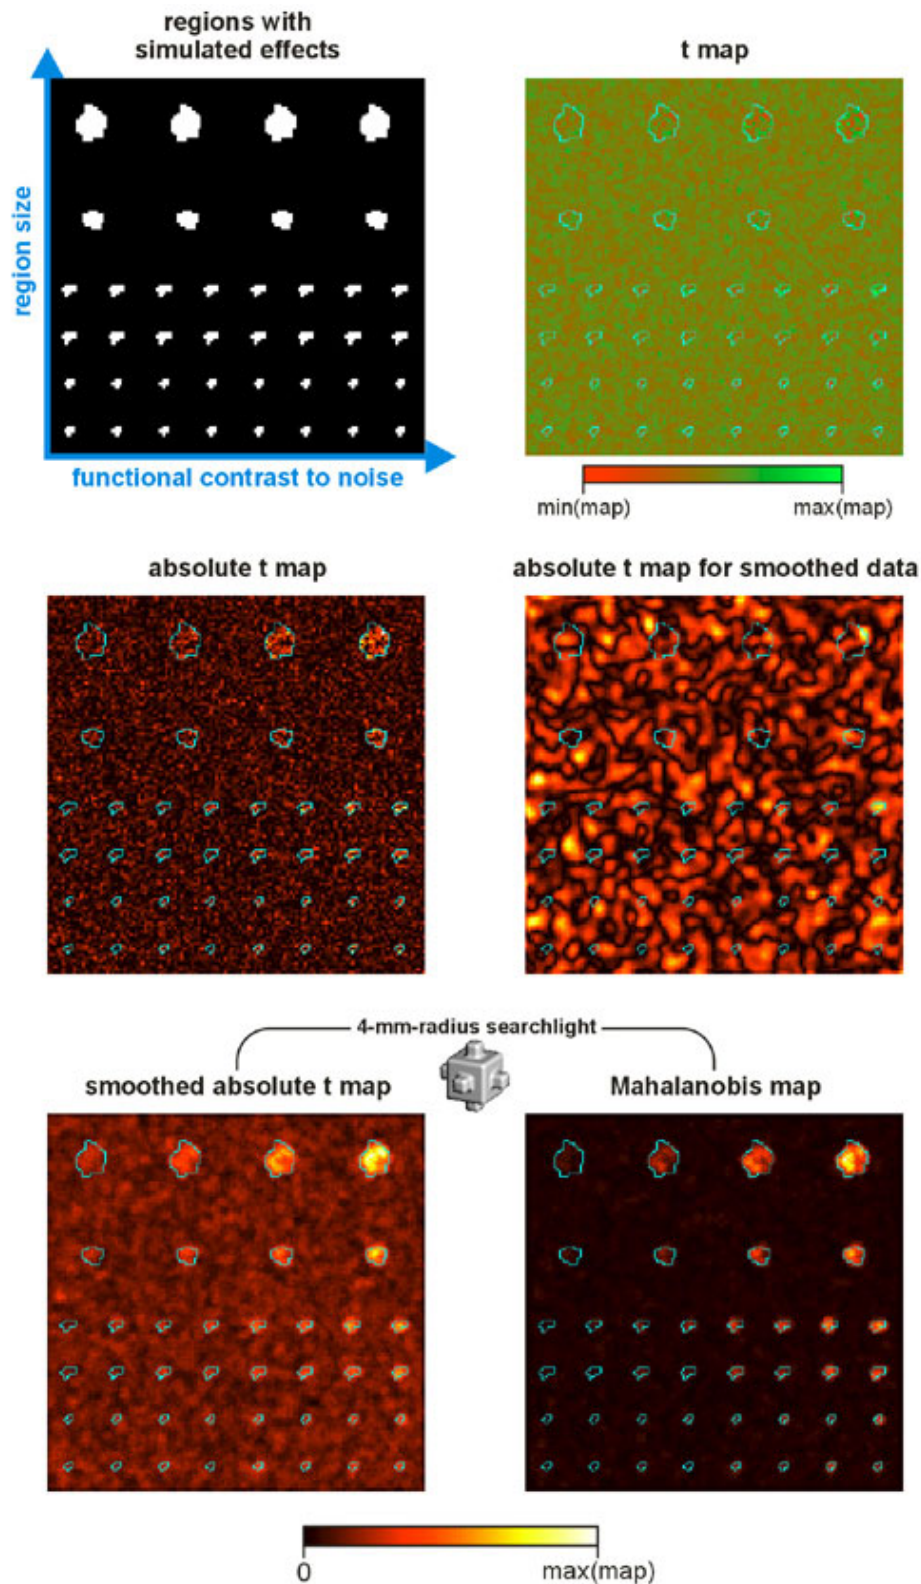
Supporting Information

Fig. 3. Simulated fMRI data, statistical maps. Unthresholded mapping results obtained with different statistics for the simulated fMRI data. This figure promotes an intuitive understanding of the quantitative ROC results presented in Fig. 1. All maps show the central slice of the simulated 3D volume. (*Top Left*) The regions containing the simulated effect patterns. Functional contrast to noise varies horizontally, region size vertically within the slice (blue axes; see *Methodological Details*). (*Top Right*) The *t* map obtained by conventional voxelwise-univariate contrast analysis (condition 1 versus condition 2) performed on

1

unsmoothed data. Note that this approach does not highlight the regions very clearly, even for high contrast-to-noise ratios. (*Center Left*) The absolute t map. Large effect regions of higher contrast to noise are visually apparent here. When the same map (absolute t) is computed for smoothed data (*Center Right*), some correct hot spots remain, but many incorrect ones appear in random locations outside the effect regions. This reflects the fact that smoothing removes the information in the high-spatial-frequency band of the effects, i.e., adjacent opposite effects cancel out. Smoothing the absolute t map (instead of the data) provides a much better map (*Bottom Left*), which correctly highlights even smaller regions of lower contrast to noise. This reflects the fact that adjacent subtle effects now accrue (instead of canceling when they have opposite signs in the t map). The absolute t map has been smoothed by convolution with a sphere of 4-mm radius. Equivalently, this map can be thought of as the result of spherical-searchlight mapping, where the value at each voxel is the average absolute t value within the spherical neighborhood. An even better map is obtained by spherical-searchlight mapping using the Mahalanobis distance between the activity patterns within the spherical neighborhood as the statistic at each voxel (*Bottom Right*). The Mahalanobis map as well was obtained using a spherical searchlight of 4-mm radius. These visually apparent results are quantitatively substantiated by the ROC analyses in Fig. 1. The red-green color bar applies to the t map (*Top Right*). The color scale linearly spans the range from the minimum (negative) to the maximum (positive) value of the t map. The black-orange-yellow-white colorbar at the bottom applies to the lower four maps. The color scale linearly spans the range from 0 to the maximum value of each map. In all statistical maps the effect-region contours are shown in cyan. Two fine points should be noted: (*i*) Maps very similar to the Mahalanobis map can be obtained by either squaring each value of the smoothed absolute t map (this does not change the ROC) or by smoothing the map of squared t values. However, visual inspection with thresholds and ROC analysis reveal that the Mahalanobis map is superior to both of these alternatives. This indicates that the lower bias of the noise model implicit to the Mahalanobis map (i.e., a multinormal model) is worth its cost in terms of variance for the present combination of searchlight dimensionality and degrees of freedom. (*ii*) (*Middle Right*) Smoothing has been performed by convolution of each data volume with an 8-mm full-width-at-half-maximum Gaussian. Similar results are obtained using a 4-mm-radius sphere, i.e., the same spatial integration window used to obtain the maps of the bottom row.


Fig. 4. Real fMRI data, control analysis. Comparison of voxels marked in activation-(green) and information-based (red) mapping analyses. Voxels marked in both maps are shown in yellow. Fig. 4 replicates the comparison shown in Fig. 2 *D* and *E* for the same subject with one modification: Instead of using the false discovery rate (FDR) threshold obtained separately for each map, the FDR threshold obtained for the activation-based map (*P* = 0.0032) has been used for both maps. As a result, fewer voxels than before are marked in the information-based map and the expected FDR is lower than 5% for that map (conservative). Note that this does not entail a substantial change to the information-based map or the comparison (refer to Fig. 2 *D* and *E* for details.)

Supporting Text

Conventional Aspects of the Analyses.

Preprocessing of real functional MRI (fMRI) data. The fMRI data sets were subjected to a series of preprocessing operations: (*i*) slice-scan-time adjustment, (*ii*) head-motion correction, (*iii*) removal of temporal drifts of frequencies below 0.009 Hz, and (*iv*) selection of brain voxels (those shown as anatomical background in Fig. 2) for mapping analyses and false discovery rate (FDR) thresholding. Steps *i-iii* were performed by using the brainvoyager 2000 software package (Brain Innovation, Maastricht, The Netherlands; Version 4.8). Step *iv* was performed in matlab (MathWorks, Natick, MA) by thresholding the spatially smoothed temporal average of the fMRI data.

Design-matrix construction and linear modeling of real fMRI data. In both the activation- and the information-based analyses, we assumed an instantaneous rectangular neural response to the stimuli and obtained hemodynamic response predictors from a linear model (1). The design matrix included these cognitive predictors along with six head-motion-parameter time courses, the global spatial-mean time course, and a confound-mean predictor. The activation-based analysis consisted in standard univariate linear regression and contrast *t* analysis. All activation- and information-based analyses were performed using custom software developed in matlab.

Receiving-operator characteristic (ROC) analysis of statistical maps from simulated fMRI data. We use ROC analysis to quantify how well different local effect statistics distinguish between the effect regions and the surrounding pure noise. An ROC depicts the tradeoff relation between sensitivity and specificity as the threshold is varied over the complete range of map values. In Fig. 1 *B* and *C*, the area under the ROC is used as a summary measure indicating to what extent high sensitivity and specificity can simultaneously be achieved, i.e., the quality of the effect statistic. More precisely, the area under the ROC represents three intuitively meaningful quantities: (*i*) the average sensitivity across all specificities, (*ii*) the average specificity across all sensitivities and (*iii*) the probability that upon drawing a random effect-region voxel and a random pure-noise voxel, the effect-region voxel will have the higher effect statistic. If the area under the ROC is 1, the effect statistic allows perfect classification of effect-region and pure-noise voxels by choice of an appropriate threshold. If the area under the ROC is 0.5, then the statistic does not order effect-region and pure-noise voxels better than chance.

1. Boynton, G. M., Engel, S. A., Glover, G. H. & Heeger, D. J. (1996) *J. Neurosci.* **16,** 4207-4221.