

An Independent Component Analysis Based Tool for Exploring Functional Connections in the Brain

Sara M. Rolfe

A thesis submitted in partial fulfillment of
the requirements for the degree of

Master of Science in Electrical Engineering

University of Washington

2007

Program Authorized to Offer Degree: Electrical Engineering

University of Washington
Graduate School

This is to certify that I have examined this copy of a master's thesis by

Sara M. Rolfe

and have found that it is complete and satisfactory in all respects,
and that any and all revisions required by the final
examining committee have been made.

Committee Members:

Linda Shapiro

James Brinkley

Date: _____

TABLE OF CONTENTS

	Page
List of Figures	ii
Chapter 1: Introduction	1
Chapter 2: Related Techniques	4
2.1 Methods for Evaluating Voxel Similarity	4
2.2 Finding and Characterizing Clusters of Activations	17
Chapter 3: Methodology for Answering Queries	24
3.1 Query 1.	24
3.2 Query 2.	33
3.3 Query 3.	39
3.4 Query 4.	44
Chapter 4: Evaluation and Discussion of Results	48
4.1 Relationship to fMRI Raw Data	48
4.2 Relationship to SPM Results	50
4.3 Contributions	55
4.4 Future Work	59
Bibliography	60

LIST OF FIGURES

Figure Number	Page	
2.1	The colors in this image represent the correlation coefficients for one two-dimensional slice of the three-dimensional fMRI scan. Each voxel time-course from the two-dimensional subset is correlated with the time-course of the voxel shown at position (15, 15) in the image. The higher coefficient values, which indicate more similarity, are shown in dark red.	6
2.2	The three plots of the SPM user interface show the X, Y, and Z planes of the brain scan. The dark areas mapped on the brain show activations of high statistical significance to the model.	7
2.3	Examples of the uncorrelated spatial patterns in the fMRI data set identified by PCA	9
2.4	PCA finds two orthogonal basis for the data, but ICA is able to correctly identify a statistically independent basis, revealing the true structure of the distribution [7].	10
2.5	(a) Original four signals, three frequencies and noise, (b) Four randomly weighted combinations of the original four signals, (c) Estimate of the underlying signals using the FastICA algorithm.	12
2.6	(a) Original ICA volume with activation threshold applied, (b) ICA volume weighted by number of activated neighbors	18
2.7	(a) Random placement of fifty bins in the ICA volume (b) Bin positions after clustering is performed (c) Final assignment of activated voxels to significant bins, where bins are differentiated by color	21
2.8	(a) Cluster query volume showing bin locations (b) Similarity ranking based on feature vector distances to 108 activation patterns from a second patient (c) Closest match to query volume (d) Second closest match to query volume (e) Third closest match to query volume. . . .	23
3.1	A conceptual example of spatial independent components. [9]	25
3.2	Graphical user interface used to query the database.	28

3.3	Sample results from query 1 showing a selected region of interest and the returned IC's with significant activation in that region.	31
3.4	Detailed view of the first match, IC 113, showing the mapping onto the patient's structural MRI image and the associated time-course. . .	32
3.5	Sample results from query 1 showing a selected region of interest and the returned IC's with significant activation in that region.	35
3.6	Detailed view of the first match, IC 87, showing the mapping onto the patient's structural MRI image and the associated time-course. . . .	36
3.7	Sample results from query 2 showing IC's in the database that are similar to selected volume 1 from the query in Figure 3.5.	37
3.8	Detailed view of the first match from query in Figure 3.5, IC 94 from patient <i>P189</i> , showing the mapping onto the patient's structural MRI image and the associated time-course.	38
3.9	Sample results from query 3 showing the loaded SPM results from the query patient, and the returned similar SPM results from other patients in the database.	42
3.10	Detailed view of the first match, patient <i>P186 spmF0001</i> , showing the mapping onto the patient's structural MRI image.	43
3.11	Sample results from query 4 showing a selected region of interest and the returned SPM similarity volumes from other patients in the database with significant activations in that region.	46
3.12	Detailed view of the first match, patient <i>P186, spmT0009</i> , showing the mapping onto the patient's structural MRI image.	47
3.13	Detailed view of the fifth match, patient <i>P190, spmT0008</i> , showing the mapping onto the patient's structural MRI image.	47
4.1	IC map relationship to fMRI raw data time-courses	49
4.2	SPM results map shown (a) mapped onto the patients structural MRI and (b) as a three-dimensional volume.	51
4.3	Query showing ROI-based search for IC maps with activation in the same region as the selected SPM results map.	52
4.4	Detail browser for patient <i>P185</i> IC map 55 showing best match for SPM results map in Figure 4.2.	53
4.5	Search results for cross-patient SPM results that are similar to the SPM results for patient <i>P185</i> , which are displayed on the structural image. The first page of results is shown in (a) and the second page of results is shown in (b).	54

4.6	Results for ROI-based search for patient <i>P185</i> . The IC map similar to the SPM results displayed on the structural image is shown in (a). Details for the returned IC map 110 are shown in (b).	56
4.7	Results for ROI-based search for patient <i>P190</i> . IC maps similar to the SPM results displayed on the structural image are shown in (a). Details for the first match, IC map 74 are shown in (b).	57
4.8	Results for ROI-based search for patient <i>P187</i> . IC maps similar to the SPM results displayed on the structural image are shown in (a). Details for the first match IC map 90 are shown in (b).	58

Chapter 1

INTRODUCTION

Functional Magnetic Resonance Imaging (fMRI) is a non-invasive technique used to detect neural activations, which has been widely applied to mapping functions of the human brain [11]. When neurons in the brain are activated, the increase in electrical activity causes an increase in the local metabolic rate. The increased consumption of oxygen results in fluctuations in the levels of paramagnetic deoxyhemoglobin in the blood which is sensed using a magnetic field. The changing level of deoxyhemoglobin measured in the brain is referred to as the blood oxygenation level dependent (BOLD) signal.

During an fMRI scan session, a series of three-dimensional images are captured consecutively, usually two to three seconds apart. The value of the image at each small volume unit, called a voxel, is the BOLD signal intensity. The patient is usually presented with a stimuli to induce neural activations. While the types of stimuli vary, the standard experimental design is a simple ‘on-off’ or ‘boxcar’ design, where the patient is repeatedly presented with a task followed by a pause.

Functional connectivity is defined as correlations between spatially remote neural events. If activity in two brain regions is observed to covary, it suggests that the neurons generating that activity may be interacting [4]. This covariance can be measured by observing the BOLD signal from two locations in an fMRI scan. The time-series from different voxel locations can be compared in several different ways, some of which will be discussed in Chapter 2.

The Structural Informatics Group at the University of Washington is developing

software tools for processing, integrating, and visualizing multimodality data for surgical planning and the study of language mapping in the brain [12]. Data for this project is collected from patients who are about to undergo surgery for intractable epilepsy in order to remove the foci of the seizures. Three main data types are collected for each patient. Before surgery, fMRI volumes showing brain activations for a language task and structural MRI volumes showing the detailed anatomy of the brain are collected. During surgery, data is collected using Cortical Stimulation Mapping (CSM) to pinpoint language areas to be avoided during surgery. For this procedure, the patient is woken up and given a language task to perform while electrical stimulation is applied to areas of his or her brain.

One goal of the project is to integrate these data types to allow for analysis of the relationships between the data sources. Identification of regions in an fMRI image volume associated with the language areas identified by CSM could provide a non-invasive alternative to CSM for epileptic patients and would advance the study of language mapping in the normal brain. Since it is expected that language regions behave as a functionally connected network [3], it is useful to be able to search for functional connectivity between voxels in fMRI image volumes and to compare networks of connectivity across patients.

This work introduces a MATLAB-based tool developed for exploring functional connectivity in the brain. Four queries have been designed to allow the user to find and compare these connections. These queries are executed from an interface that allows the user to choose the query type and interact with the data. The queries supported by the user interface designed for this work are:

1. Starting with SUR, CSM, or fMRI data, select an (x, y, z) coordinate in the brain. Get the raw data fMRI time series of the voxel at that location. Use signal similarity measures to find correlated voxels within the patient's brain.
2. Starting with one patient's voxel correlations, search for patients who have a

similar correlation pattern for a voxel in the same region.

3. For a given patient fMRI and an activation threshold, find patients with similar activation patterns, by searching the fMRI images showing statistically significant activations.
4. For a given patient fMRI and a given location, find other patients who have greater than or equal activation values at that location, by searching the fMRI images showing statistically significant activations.

This thesis describes the use of independent component analysis (ICA) as a measure of voxel similarity, which allows the user to find and view statistically independent maps of correlated voxel activity. The tool developed in this work uses a specialized clustering technique, designed to find and characterize clusters of activated voxels, to compare the independent component spatial maps across patients. This same method is also used to compare SPM results across patients.

Chapter 2 discusses in detail the techniques used to perform measurements of voxel and spatial map similarity. Chapter 3 describes the methods used to answer each of the four queries. Instructions for executing each query and the algorithms used to find the query results are presented. Chapter 4 provides an evaluation of the methodology and discusses the results.

Chapter 2

RELATED TECHNIQUES

There are two main tasks associated with the four query types. The first task is the measurement of voxel similarity. This is required by the first query, which allows the user to select a region of interest (ROI) in a patient's brain and returns a map of other regions in the patient's brain with similar activation patterns. Although many techniques exist for finding similar voxel behavior, this chapter will address the use of direct correlation, Statistical Parametric Mapping (SPM), principal component analysis (PCA), and independent component analysis (ICA).

The second task needed to perform the queries is the measurement of similarity between spatial activation maps. This measurement is needed for the second and third queries, which allow the user to select an ICA or SPM similarity map, and return similar maps of the selected data type from the database. For this application, a specialized clustering algorithm, which will be discussed in detail in this chapter, has been developed.

2.1 Methods for Evaluating Voxel Similarity

2.1.1 Direct Correlation

The simplest method of finding similarity between voxels is to use standard correlation to compare their time series. The correlation coefficient measures the amount of similarity between two random processes. It is defined for two random processes X and Y as

$$c_{X,Y} = \frac{E[(X - \mu_X)(Y - \mu_Y)]}{\sigma_X \sigma_Y}, \quad (2.1)$$

where μ is the sample mean of a process and σ is its standard deviation. The higher the correlation coefficient, $c_{X,Y}$, the more similarity there is between X and Y .

When a region of interest in the brain is selected, its voxel time-course can be extracted from the raw data of an fMRI scan and correlated with the time-courses of all other voxels in that scan. A high correlation coefficient for two voxel time-courses indicates shared activation patterns, which may be a result of a functional connectivity. In Figure 2.1, one slice of a patient's brain scan is shown, color coded by correlation coefficients. The time-courses from the voxels at each position shown were correlated with the time-course from a voxel at position (15, 15). The correlation values were then normalized for the image. It can be seen from Figure 2.1 that the area around the voxel at position (15, 15) is highly similar. This is expected since the targeted functional activations occur in clusters of 2×2 voxels (about 3 mm square) or larger [11]. A narrow area of activation along the front of the brain can also be observed on the left side of Figure 2.1. It is difficult to assess the accuracy of this region, since there is no ground truth that can be used to evaluate the results.

While this technique is straightforward and intuitive, it has an inherent drawback that impacts its usefulness. The fMRI signal from one voxel is actually the sum of the signals from the tissue, blood, spinal fluid, and other components in an approximately 1.5 mm cube [11]. The signals can be related to task and non-task related BOLD response, spontaneous metabolic changes, or physiological influences such as cardiac and respiratory fluctuations. The signals that are generally most interesting to researchers are the BOLD responses which indicate neural activity. It has been demonstrated that even in activated areas, the task related component of a voxel time-course makes up less than ten percent of the total magnitude [9]. This suggests that other fluctuations are responsible for the majority of the signal. Direct correlation of these raw time-courses will tend to mask interesting signal components and provide similarity information about less interesting physiological fluctuations.




Figure 2.1: The colors in this image represent the correlation coefficients for one two-dimensional slice of the three-dimensional fMRI scan. Each voxel time-course from the two-dimensional subset is correlated with the time-course of the voxel shown at position (15,15) in the image. The higher coefficient values, which indicate more similarity, are shown in dark red.

2.1.2 Statistical Parametric Analysis

Statistical Parametric Analysis (SPM) is a software package developed by the Wellcome Department of Imaging Neuroscience (University College London) that uses statistics to analyze brain activity from fMRI and other functional imaging experiments. It provides a library of utilities, such as head movement correction, cross patient spatial normalization, and smoothing, which significantly improve the quality of raw imaging data. These correction routines were used on all fMRI data presented in this paper.

SPM is also used to map regions in the brain that work together on a task; this is another method that can be used to identify functional connections. The approach used by SPM circumvents the mixed source problem by using a combination of statistical processes to find voxel activation patterns related to the exam stimulus. A test vector modeled after the hypothesized response to the exam stimulus is compared to

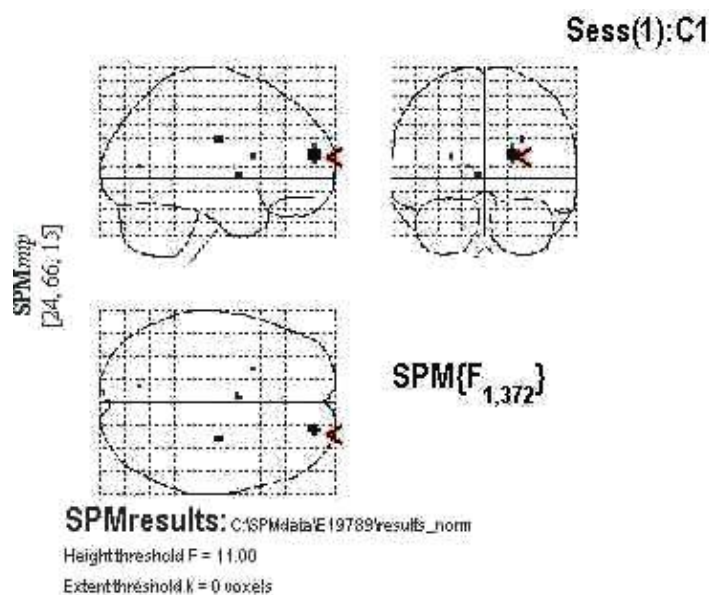


Figure 2.2: The three plots of the SPM user interface show the X, Y, and Z planes of the brain scan. The dark areas mapped on the brain show activations of high statistical significance to the model.

each voxel using univariate measures (f-tests and t-tests) to look for statistical significance. Each voxel is assigned a value based on the significance of its relation to the test vector. These values are then displayed as a three-dimensional volume, as shown in Figure 2.2.

While SPM solves the problem of separating the small-magnitude, task-related signal from the raw data, it has limited effectiveness for identifying functional connections. The first weakness is its assumption of a model for the response to experimental stimuli. This implies that the response is identical across all voxels from one patient scan and across all patients. In reality these may differ, especially in the case of patients with a pathological brain condition [9]. The use of a model also ignores transient task related changes that may be caused by learning or fatigue.

Secondly, SPM is limited to evaluating only task-related similarity between voxel

time-courses. While these components are more interesting than known non-task related components such as cardiac and vascular fluctuations, this approach ignores unknown networks of voxel activations that may not be directly dependent on the task vector. For example, there has been recent interest in low frequency (less than .1 Hz) activation patterns showing complex spatial structure that have been observed both at rest and during experimental stimulation[1]. These activation maps may be the result of baseline activity indicating the presence of functional networks. To better explore these regions of shared activation, a data driven technique is needed that does not require the prediction of an activation vector.

2.1.3 Principal Component Analysis

Principal component analysis (PCA) is a linear transformation that decorrelates the dependent variables in a data set. It projects the data onto a basis set where the elements are orthogonal to each other. To perform PCA, the first step is to calculate the the covariance matrix of the data set, which represents the tendency of all possible voxel pairs to covary. The covariance matrix of an $m \times n$ data set X containing m observations of the n voxels in each of the column vectors x_1, x_2, \dots, x_m , is defined as

$$Cov(X) = \begin{pmatrix} E[(x_1 - \mu_1)(x_1 - \mu_1)] & E[(x_1 - \mu_1)(x_2 - \mu_2)] & \cdots & E[(x_1 - \mu_1)(x_n - \mu_n)] \\ E[(x_2 - \mu_2)(x_1 - \mu_1)] & E[(x_2 - \mu_2)(x_2 - \mu_2)] & \cdots & E[(x_2 - \mu_2)(x_n - \mu_n)] \\ \vdots & \vdots & \ddots & \vdots \\ E[(x_n - \mu_n)(x_1 - \mu_1)] & E[(x_n - \mu_n)(x_2 - \mu_2)] & \cdots & E[(x_n - \mu_n)(x_n - \mu_n)] \end{pmatrix},$$

where μ_i is the mean of the column vector x_i . Once the covariance matrix is found, its eigenvectors are computed. The eigenvectors represent the new orthogonal basis of the data set. The eigenvectors y are defined as the non-trivial solution

$$C * y = \lambda * y, \tag{2.2}$$




Figure 2.3: Examples of the uncorrelated spatial patterns in the fMRI data set identified by PCA

where C is the square covariance matrix, y is a column vector containing the eigenvectors, and λ contains the scalar eigenvalues. The eigenvalue associated with an eigenvector quantifies its contribution to the variance of the data, and therefore its significance. The PCA algorithm ranks the eigenvectors in order of decreasing eigenvalues. By definition, the eigenvectors are all orthogonal to each other, so the ranked eigenvectors provide the new uncorrelated bases for the data. In this application, the eigenvectors contain uncorrelated spatial patterns of activated voxels. Examples of the three-dimensional spatial patterns extracted from one scan session from one patient are shown in Figure 2.3.

PCA is commonly used to reduce the dimensionality of the data since bases with insignificant contributions to the overall variance can be removed with little information loss. In this application some of the most interesting components of the data contribute very little to the overall variance, so the original dimensions should be retained. The goal of PCA in this application is the projection of the raw data from the scans onto an uncorrelated basis set. Since the time-course observation from each voxel is known to contain a mixture of unknown underlying signals, if these signals have separate causes, they should be uncorrelated and separable by PCA. This is true when the distribution statistics of the data are second order. In this case independence




Figure 2.4: PCA finds two orthogonal basis for the data, but ICA is able to correctly identify a statistically independent basis, revealing the true structure of the distribution [7].

and decorrelation are equivalent. However for higher order statistics, decorrelation is not a strong enough condition.

2.1.4 Independent Component Analysis

Independent component analysis (ICA) is a technique used to separate statistically independent source signals from a set of observations. It is superficially related to PCA. Where PCA attempts to identify underlying components in the data set that are orthogonal and therefore uncorrelated, ICA identifies underlying components that are statistically independent. Statistical independence is a stronger condition than decorrelation, since it can use any order of statistics rather than being limited to the two orders available to PCA (from the covariance matrix). In addition, ICA does not restrict the components to orthogonal dimensions, so the underlying bases are allowed to be very similar, as long as the independence requirement is met. An example of the effect of removing these restrictions is shown in Figure 2.4.

In the ICA model, T observations of a random variable x , $[x_1, x_2, \dots, x_T]$, are

assumed to be a linear mixture of T statistically independent components, c_1, c_2, \dots, c_T . Each component x_j in the observation set can be represented as

$$x_j = m_{j1}c_1 + m_{j2}c_2 \dots + m_{jT}c_T, \quad (2.3)$$

where the coefficient m_{ji} is a random variable describing the contribution of c_i to x_j . In matrix form, this becomes

$$X = MC \quad (2.4)$$

where X is a $T \times V$ matrix containing T observations of V variables, M is a $T \times T$ matrix of the mixing coefficients m_{ji} , and C is a $T \times V$ matrix containing T observations of the V underlying components. M is generally assumed to be a square matrix, since the ICA algorithm requires the calculation of its inverse, $W = M^{-1}$.

In this model, both the mixing matrix, M , and the independent components, C , are unknown. ICA estimates the values of both, using the assumption that all c_i are independent. The algorithm finds an initial estimate of \tilde{M} and its inverse, $\tilde{W} = \tilde{M}^{-1}$, is calculated. The independent components are then estimated to be


$$\tilde{C} = X\tilde{W}. \quad (2.5)$$

The algorithm measures the independence of the components \tilde{c}_i and uses this measurement to estimate a new mixing matrix \tilde{M} . This process is repeated iteratively until the estimates for C converge.


There are several different algorithms available for performing ICA. Two algorithms that have been widely used and are explored in this paper are FastICA [6] and Infomax [10].

FastICA


FastICA, developed by Hyvarinen and Oja in 1998, is a computationally efficient fixed-point algorithm for maximizing the statistical independence of the estimated components by maximizing their non-Gaussianity [6]. It is reasonable to assume that



(a)



(b)



(c)

Figure 2.5: (a) Original four signals, three frequencies and noise, (b) Four randomly weighted combinations of the original four signals, (c) Estimate of the underlying signals using the FastICA algorithm.

the underlying distributions in the data set, modeled as C , do not have Gaussian distributions. When any non-Gaussian random variables are added together, their joint distribution becomes increasingly Gaussian in shape. Therefore, any weighted sum of the independent components will be more Gaussian than the actual non-Gaussian components. When the estimate of one component \tilde{c}_j is dependent on more than one underlying component, its distribution will be more Gaussian than the actual c_j . The FastICA algorithm iteratively calculates the rows of \tilde{W} to maximize the non-Gaussianity, and therefore independence, of the rows of \tilde{C} .

There are several ways to quantify the non-Gaussianity of a random variable. FastICA maximizes non-Gaussianity by maximizing the negentropy of the estimated components [5]. Negentropy, named for an abbreviation of *negative entropy*, is a statistical measurement based on entropy. For a random variable, entropy describes the amount of information that can be gained from observation. Observation of variables with low entropy will allow the identification of a narrow range of values the variable is likely to take. If a variable has high entropy, it will have a higher degree of randomness, and observation over time will provide less predictive power. Entropy of a random variable y is defined as

$$H(y) = \sum_i p(y = a_i) \log(p(y = a_i)), \quad (2.6)$$

where a_i are the possible values that y can take. A Gaussian random variable has the largest possible entropy of all variables with an equal variance, so high degree of entropy can be associated with a high degree of gaussianity. Negentropy is a measurement of the entropy of a random variable that is designed to be always non-negative and equal to zero when the distribution is Gaussian. Negentropy is defined in terms of entropy as

$$J(y) = H(y_{gauss}) - H(y), \quad (2.7)$$

where $H(y)$ is defined as the entropy of the random variable y and $H(y_{gauss})$ is the entropy of a Gaussian random variable with the same mean and variance as y . In practice, as developed by Hyvriinen [5], this is approximated by the calculation of

$$J(y) \approx \sum_{i=1}^k (E[G_i(y)] - E[G_i(y_{gauss})])^2, \quad (2.8)$$

where y is the random variable, y_{gauss} is the Gaussian random variable with equal mean and variance, k is an arbitrary number, generally chosen to be 2, and G_i is the i 'th non-quadratic function. An estimate of the negentropy can be achieved with almost any non-quadratic G_i , but the most accurate results will be gained by choosing a function that grows slowly. The iterative formula for estimating the weights w by maximizing the approximate negentropy of the estimated components $\tilde{c} = w^T x$, using Equation 2.8 is

$$w_{new} = E[xG'(w_{old}^T x)] - E[G''(w_{old}^T x)]w_{old}. \quad (2.9)$$

In the FastICA algorithm, the values in the weighting matrix are calculated one at a time to find the optimal weighting for each independent component, $w^T x$. The iterative steps are:

1. Select random weight, w_{old} .
2. $w_{new} = E[xG'(w_{old}^T x)] - E[G''(w_{old}^T x)]w_{old}$
3. $w_{old} = w_{new}/|w_{new}|$
4. If not converged, return to step 2.

Since they are found independently, it is important to decorrelate the estimated components after every iteration of the algorithm. This is necessary to prevent the

convergence to two of the same underlying components. The weights are iteratively estimated and after each estimation n , the previously estimated components c_1, \dots, c_{n-1} are projected onto and subtracted from the current component, c_n . The current component is then renormalized.

Infomax

The Infomax algorithm is a stochastic gradient ascent algorithm developed by Nadal and Parga [10] and Bell and Sejnowski [2]. This method uses mutual information rather than non-Gaussianity to maximize the statistical independence of the components. Mutual information for two random variables is defined as

$$I(X, Y) = p(X, Y) \frac{\log(p(X, Y))}{p(X)p(Y)}, \quad (2.10)$$

where $p(X)$ and $p(Y)$ are the probability density functions for X and Y , and $p(X, Y)$ is their joint probability density function. Mutual information measures the extent to which observation of one variable reduces the uncertainty of the second. This is minimized when X and Y are independent and one variable provides no knowledge about the other. In this case,

$$p(X, Y) = p(X)p(Y), \quad (2.11)$$

and $I(X, Y) = 0$. Mutual information can also be expressed in terms of the entropies of X and Y , where

$$I(X, Y) = H(X) + H(Y) - H(X, Y) \quad (2.12)$$

$$= H(X) - H(X|Y), \quad (2.13)$$

where $H(X)$ and $H(Y)$ are the marginal entropies of X and Y , $H(X, Y)$ is their joint entropy, and $H(X|Y)$ is their conditional entropy. The Infomax algorithm maximizes

the joint entropy of the outputs X and Y , where the joint entropy is defined in terms of mutual information as

$$H(X, Y) = H(X) + H(Y) - I(X, Y). \quad (2.14)$$

Maximizing this quantity maximizes the individual entropy, while simultaneously minimizing the mutual information of the two variables to reduce redundancy. This results in maximizing the independence of X and Y . This system is implemented with a non-linear neural network which uses stochastic gradient learning rules to maximize the joint entropy of the outputs to estimate the independent components of the ICA model from Equation 2.4.

Comparison of techniques

Comparisons of these and other ICA algorithms have shown that while some data sets may achieve slightly better results with one algorithm, their performance is nearly equivalent [7]. For the fMRI data sets used in this work, an evaluation of the algorithms is difficult because there is no ground truth to which the results can be compared. The advantages of the differences between results are very difficult to assess. For these data sets, the Infomax and FastICA algorithms were applied, and an analysis of the results showed no clear advantages. However, FastICA did have better processing speed. Infomax makes use of an adaptive learning algorithm, which can add time to the processing, and is highly dependent on the learning rate and other network parameters. FastICA avoids this learning problem entirely since it uses a fixed-point, rather than adaptive learning. This can result in a speed-up factor of between 10 and 100, when compared to adaptive algorithms [8]. For this application, the increase in speed was significant enough to make the FastICA algorithm a better choice.

2.2 Finding and Characterizing Clusters of Activations

The second task required by the queries is the measurement of similarity between spatial maps. This measurement is used in the second query to find similar independent component maps and in the third query to find similar SPM results maps. For this application, similarity between two spatial maps is defined as similarity between the location and characteristics of their activated voxel clusters. The first step in measuring spatial map similarity is to identify the voxel clusters in a spatial map and to extract their relevant cluster features.

Before the clustering algorithm is applied to a volume image, the image is preprocessed to improve identification of cluster centers. A threshold is applied to the image to locate the activated voxels. Voxels determined to be activated are assigned a value of one. This binary labeling is used to eliminate the influence of activation level when identifying the clusters. The activations which are most interesting often make up only about five to ten percent of the raw data, so weighting by activation value is not useful for finding cluster centers. After the threshold has been applied, the image is filtered to emphasize the significance of voxels located in groups. A brain activation is expected to involve a cluster of voxels, so an activated voxel is more likely to be significant when its neighbors are also activated. Since this algorithm is aimed at identifying activation centers, voxels which are surrounded by activations are more likely to be near the center. To utilize this information, each activated voxel x_i is assigned the value

$$x_i = e^n, \tag{2.15}$$

where n is the number of neighboring voxels activated. The non-linear function further exaggerates the significance of voxels with a large number of neighbors. This preprocessing improves the speed and accuracy of the clustering step. The effect of the weighting can be seen in Figure 2.6. The weighted image clearly shows that the

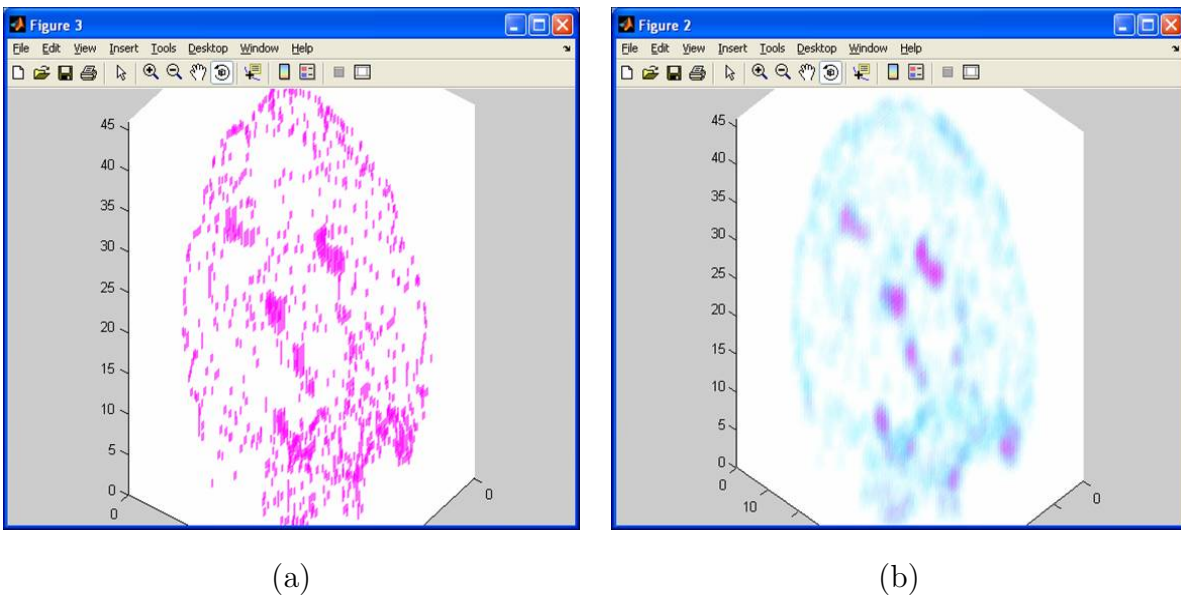


Figure 2.6: (a) Original ICA volume with activation threshold applied, (b) ICA volume weighted by number of activated neighbors

significant voxel clusters in pink have been differentiated from the noisy activations.

After the preprocessing, modified k-means clustering is used to find the activation centers. Fifty random bin locations are generated in the image. For each activated point in the image, the Euclidean distance to the bins is calculated, and the point is assigned to the nearest bin. When each point has been assigned to a bin, the bin center locations are replaced with the mean of the locations of the activated points in each bin. Unlike standard k-means clustering, the contribution of each point to the mean is dependent on the value at that location. This modification to k-means clustering makes use of a priori information about activation structure to locate the centers. Each activation location i in a bin is weighted by

$$w_i = \frac{x(i)}{\sum_{j=1}^n x(j)}, \quad (2.16)$$

where n is the number of activation locations in the bin, and $x(j)$ is the value at voxel location j in the weighted, thresholded image. The weights in each bin sum to one,

so the sum of the weighted locations is the weighted mean. The iterative algorithm for this procedure is as follows:

1. For all bins b_j and each activated point p_i , calculate the distance between the three-dimensional locations of bin $b_j(x, y, z)$ and activation point $p_i(x, y, z)$,

$$d_{i,j} = \text{abs}[b_j(x) - p_i(x)] + \text{abs}[b_j(y) - p_i(y)] + \text{abs}[b_j(z) - p_i(z)]$$

and assign activated point p_i to the bin that minimizes d_i .

2. For each bin b_j , find sum_j , the sum of the weights at the n locations assigned to b_j

$$sum_j = \sum_{i=1}^n w_i,$$

where the index i represents each activation point p_i assigned to bin j .

3. Normalize the weight at each activation location p_i in b_j ,

$$norm_weight_{i,j} = \frac{weight(p_i)}{sum_j}.$$

4. Weight each activation location in bin j where

$$weighted_location(i, j) = norm_weight(i, j) * p_i(x, y, z).$$

5. Calculate the new location of each bin j , where

$$new_bin_j = \sum_i weighted_location(i, j),$$

and index i represents each point p_i in bin j .

This process is repeated until the bin centers converge within a threshold, or until twenty iterations have been completed. Twenty was experimentally determined to be sufficient.

Once the bin locations have converged, the excess bin locations are removed. First, bins with less than ten points are removed. After this step, each bin location is stored sequentially. Before a bin location is stored, the distance to each stored bin is found. Chebyshev, rather than Euclidean distance is used, defined as the maximum of the differences along every dimension,

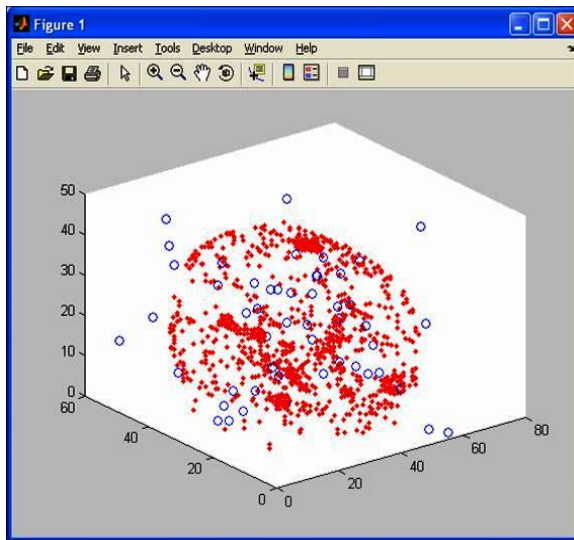
$$d_{x,y} = \max(\text{abs}(x_i - y_i)). \quad (2.17)$$

This penalizes differences along one dimension. If the distance is less than a threshold value, the stored bin location is replaced with the mean of the two bin locations, and the activation locations in the first bin are added to the values in the stored bin. Example results from this clustering algorithm are shown in Figure 2.7.

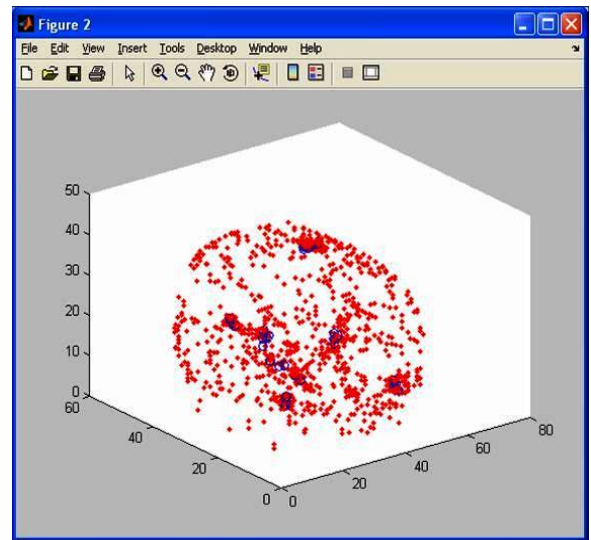
2.2.1 *Extracting Features from the Clusters*

Once the activation clusters have been located, five features are extracted from each cluster and are stored in a feature vector.


1. **Location of bin center:** This feature is the (x, y, z) coordinates of the cluster center found by the clustering algorithm.
2. **Bin size** The size of the bin is equal to the number of activated voxels assigned to that bin.
3. **Average distance from bin center:** The average distance of the activated voxels in the bin to the bin center captures differences that may arise in cluster location for activated regions with the same bin centers. Chebyshev distance is used to penalize differences along one dimension.
4. **Average bin weight: (centrality)** In the clustering algorithm, the IC volume is filtered by weighting each activated voxel by the number of its activated



(a)



(b)



(c)

Figure 2.7: (a) Random placement of fifty bins in the ICA volume (b) Bin positions after clustering is performed (c) Final assignment of activated voxels to significant bins, where bins are differentiated by color

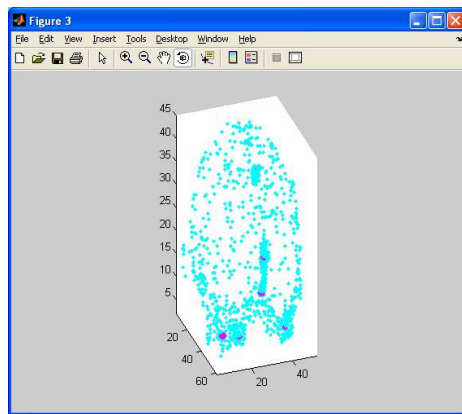
neighbors. This emphasizes the significance of voxels in clusters, since the targeted functional activations are expected to occur in clusters of 2×2 voxels or more. The average weight of the values in the bin provides a measure of how closely the activated locations are clustered, or the bin “centrality”. It is found by summing the weights assigned to the bin and normalizing it by the total number of points.

5. **Weighted variance of distances:** This feature captures the variance of the Chebyshev distances between the bin center and the activation locations assigned to the bin. The variance is weighted so that locations with larger weights contribute more to the variance calculation and is given by

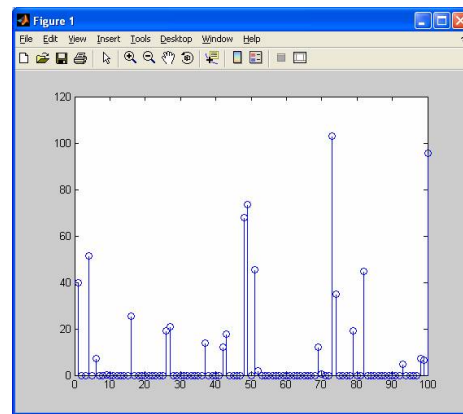
$$\sum_i \frac{dist_i * weight_i - E[dist_i * weight_i]}{\sum_j dist_j * weight_j}, \quad (2.18)$$

where $dist_i$ is the Chebyshev distance between the bin center and activation location i and $weight_i$ is the value of the weight at location i .

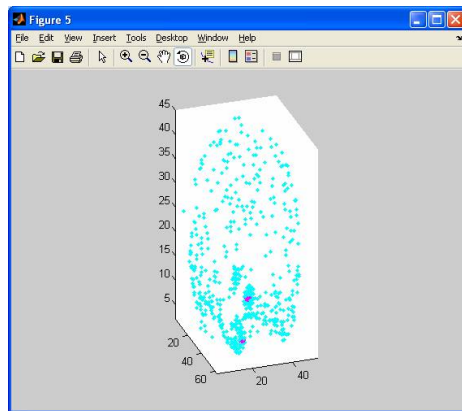
These features are stored in a feature vector. An activation pattern will have one feature vector for each cluster identified by the clustering algorithm. These vectors are used to compare activation patterns across patients. An example of this similarity measurement is shown in Figure 2.8, where the spatial map in (a), shown with the bin locations highlighted in pink, is compared to 108 different spatial maps. It can be seen from this figure that the spatial maps with the highest similarity ranking, shown in (c) through (e), have bins in locations similar to the query image.



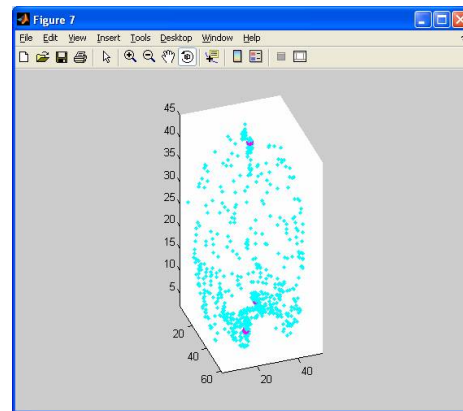
(a)



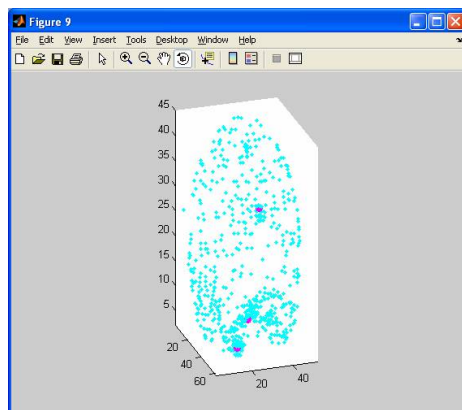
(b)



(c)



(d)



(e)

Figure 2.8: (a) Cluster query volume showing bin locations (b) Similarity ranking based on feature vector distances to 108 activation patterns from a second patient. (c) Closest match to query volume. (d) Second closest match to query volume. (e) Third closest match to query volume.

Chapter 3

METHODOLOGY FOR ANSWERING QUERIES

3.1 Query 1.

Starting with SUR, CSM, or fMRI data, select a three-dimensional coordinate in the brain and retrieve the raw data fMRI time series of the voxel at that location. Use signal similarity measures to find correlated voxels within the patient's brain.

This query is used to locate regions in the brain that may be functionally connected to a region of interest. Functional connectivity, the correlations between spatially remote neural events, is observed when activity in two brain regions covaries. This covariance is measured by comparing the voxel time-courses, the raw data value at a voxel position at each sample point in one scan, aligned in time.

3.1.1 Application of ICA to fMRI

When applying ICA to fMRI data, temporal or spatial independence can be used to model the independent components. *Temporal ICA* models a sample of M voxel time-series from one brain location as a mixture of M temporally independent source signals. The observed data X_{ji} is an $M \times N$ matrix containing M voxel time-series of length N , and can be modeled as

$$X_{ji} = \sum_{k=1}^M M_{jk} C_{ki}, \quad (3.1)$$

where M_{jk} is an $M \times M$ random mixing matrix and C_{ki} is an $M \times N$ matrix containing the M independent time series.

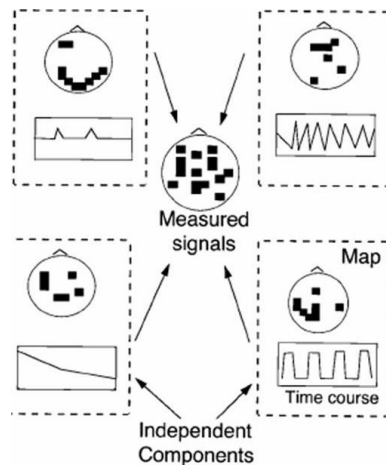


Figure 3.1: A conceptual example of spatial independent components. [9]

C and M can be obtained using ICA. The independent source signals may be the result of vascular pulsation, breathing induced motion, or the task related hemodynamic response which may have been summed together to create the observed data from the M voxel locations.

Spatial ICA models the entire observed data set of N volume images as a mixture of N independent voxel maps. Each voxel map has one time-course that it is associated with. An example of this concept is shown in Figure 3.1.

The observed data from one scan session is in the form of T three-dimensional volumes, where T is the number of time-points in the scan. For the calculations, each of these three-dimensional volumes is unwrapped into a vector of length V where V is the number of voxel positions in the three-dimensional volume. Each position in the one dimensional vector corresponds to one three-dimensional voxel location. The data from the entire scan, X , of size $T \times V$, contains one unwrapped volume vector in each row. In this form, the observed spatial patterns from each scan-time point have only one dimension, and so the dataset's independent components, which are statistically independent activation maps, will also be one dimensional. When the independent components are displayed for the user, the one-dimensional volume

vector can be re-wrapped into a three dimensional volume, since each position in the vector corresponds to a specific three-dimensional location.

The observed data X can be modeled as the independent spatial component maps C multiplied by a random mixing matrix M ,

$$X_{ji} = \sum_{k=1}^V M_{jk} C_{ki}. \quad (3.2)$$

where $C_{i,j}$ is a matrix of size $T \times V$ containing the T independent component maps and M is the mixing matrix. In matrix form, this is

$$\begin{pmatrix} x_{11} & x_{12} & x_{13} & \cdots & x_{1V} \\ x_{21} & x_{22} & x_{23} & \cdots & x_{2V} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ x_{T1} & x_{T2} & x_{T3} & \cdots & x_{TV} \end{pmatrix} = \begin{pmatrix} m_{11} & m_{12} & \cdots & m_{1T} \\ m_{21} & m_{22} & \cdots & m_{2T} \\ \vdots & \vdots & \ddots & \vdots \\ m_{T1} & m_{T2} & \cdots & m_{TT} \end{pmatrix} * \begin{pmatrix} c_{11} & c_{12} & c_{13} & \cdots & c_{1V} \\ c_{21} & c_{22} & c_{23} & \cdots & c_{2V} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ c_{T1} & c_{T2} & c_{T3} & \cdots & c_{TV} \end{pmatrix}.$$

When the matrix multiplication is carried out to find X in terms of M and C , this can be expressed as

$$X = \begin{pmatrix} m_{11}c_{11} + m_{12}c_{21}\dots + m_{1T}c_{T1} & m_{11}c_{12} + m_{12}c_{22}\dots + m_{1T}c_{T2} & \cdots & m_{11}c_{1V} + m_{12}c_{2V}\dots + m_{1T}c_{TV} \\ m_{21}c_{11} + m_{22}c_{21}\dots + m_{2T}c_{T1} & m_{21}c_{12} + m_{22}c_{22}\dots + m_{2T}c_{T2} & \cdots & m_{21}c_{1V} + m_{22}c_{2V}\dots + m_{2T}c_{TV} \\ \vdots & \vdots & \ddots & \vdots \\ m_{T1}c_{11} + m_{T2}c_{21}\dots + m_{TT}c_{T1} & m_{T1}c_{12} + m_{T2}c_{22}\dots + m_{TT}c_{T2} & \cdots & m_{T1}c_{1V} + m_{T2}c_{2V}\dots + m_{TT}c_{TV} \end{pmatrix}.$$

From this model, it can be seen that the columns of M contain the time-courses associated with each component map. The observed scan volume at time point t is row t of X . It can be seen that row t of M specifies the weight of contribution of each component at time t . Therefore, each column of M contains the weights for one of the component maps over the entire time interval.

3.1.2 Significance of the Independent Components

When using spatial ICA, the independent component map retrieved by the algorithm is a $T \times V$ matrix, where each of the T rows contains a spatial volume of V voxels that can be mapped back to its original three dimensions. Each volume represents a statistically independent pattern of activity. The voxels activated in this pattern

share the activation pattern contained in the associated time-course over the duration of the time-course. For example, vascular pulsations should be localized to large veins, breathing motion should appear in maps with strong tissue edges, and task related maps are expected to have time series corresponding to the stimulus onset times from the exam.

3.1.3 Query Results

This query returns independent spatial maps which show regions in the brain related to the volume of interest selected by the user. The independent components for each scan are preprocessed and stored in a database, along with the activation cluster features of each component volume.

To execute a query of this type using the the interface shown in Figure 3.2, the user will:

1. Select the patient number from the ‘PATIENT’ menu to load the patient’s structural volume into the main window.
2. Optionally select an SPM results file from the ‘SPM RESULTS’ menu to map voxels with time-courses statistically related to the experimental design onto the structural volume. This may be used to identify a region of interest to the user.
3. Select the query type, ‘Inter-patient IC Search’, using the radio buttons.
4. Browse the selected patient’s structural volume until a region of interest is identified.
5. Click the ‘ROI SELECT’ push button.

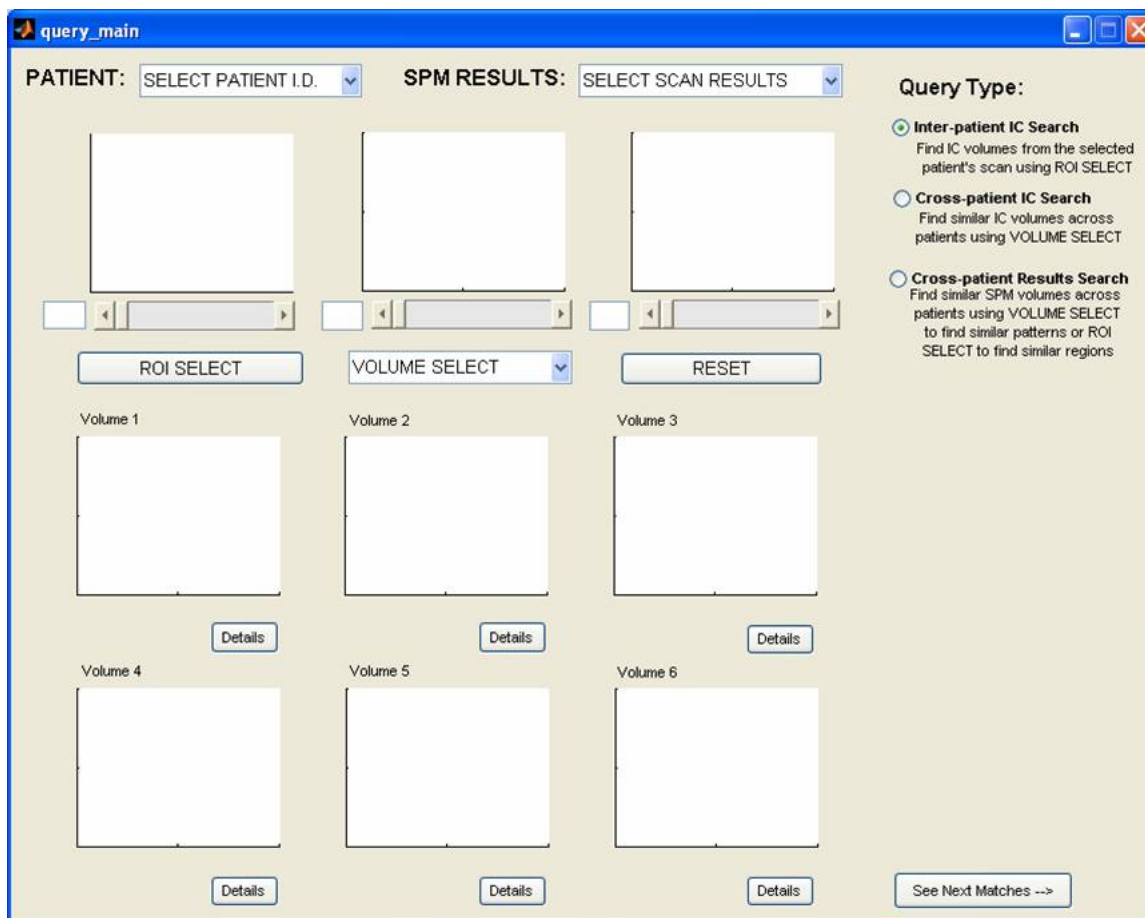


Figure 3.2: Graphical user interface used to query the database.

6. Use the mouse to draw a box around the region of interest on any of the three volume views.

Although the user can only select a two-dimensional box, the region of interest is a cube. The depth of the cube is pre-defined to be five samples, centered on the slice where the user has drawn the box. Once the user has selected a region of interest, all ICA volumes from the selected scan with an average activation in the region above a threshold are returned to the user in the re-wrapped, three dimensional format. These patterns describe independent activation maps that are significantly contributing to

the observed activations in that region. The algorithm used to find activated IC's and rank their significance is:

1. For each IC map i from the selected patient's scan, find the average activation value of IC map i in the selected ROI.
2. If the average activation is greater than the threshold of one third of the maximum IC map value, continue with the next steps to calculate the ranking. Else,

$$rank(i) = 0.$$

3. Find n , the number of activation clusters in IC map i .
4. Find d_{min} , the minimum of the distances between the ROI center and each activation cluster center in IC map i .
5. Assign IC map i the rank

$$rank(i) = \frac{1}{n^2 d_{min}}.$$

Steps three through five determine the order in which the IC maps with activations in the ROI are returned to the user. Activation maps with many clusters are more likely to be caused by noise, and are generally less interesting to the user than those with a few, large clusters. This information is included in the ranking with n , defined in step 3. Additionally, activations with a center close to the center of the ROI will be more relevant to the volume selected by the user. This information is included in the ranking with d_{min} , defined in step 4.

Once every IC map in the selected patient's scan has been ranked, by the assignment of the value $rank$, the maps with the highest non-zero rankings are returned to the user. The IC maps are displayed in their re-wrapped, three-dimensional format. Under each plot is a 'DETAILS' push button which the user can use to get more

information about the displayed IC map. Selecting this button launches a browser which displays the IC mapped back onto the patient's structural MRI volume and a plot of the IC's associated time-course. If there are more than six IC maps with an activation in the ROI, the 'NEXT' button can be used to browse through the remaining maps.

Sample results from query 1 are shown in Figure 3.3. The blue box outlining the user selected ROI can be seen on the first structural MRI plot. Four IC maps with activations in the ROI are displayed. The highest ranked match appears to have an interesting pattern, while the remaining three matches contain so many activation clusters that they are most likely the result of physiological noise. Figure 3.4 shows the detail browser for the first match, IC 113. Comparing the plots from the detail browser to the original query image, it can be seen that the returned result has a significant activation in the ROI.

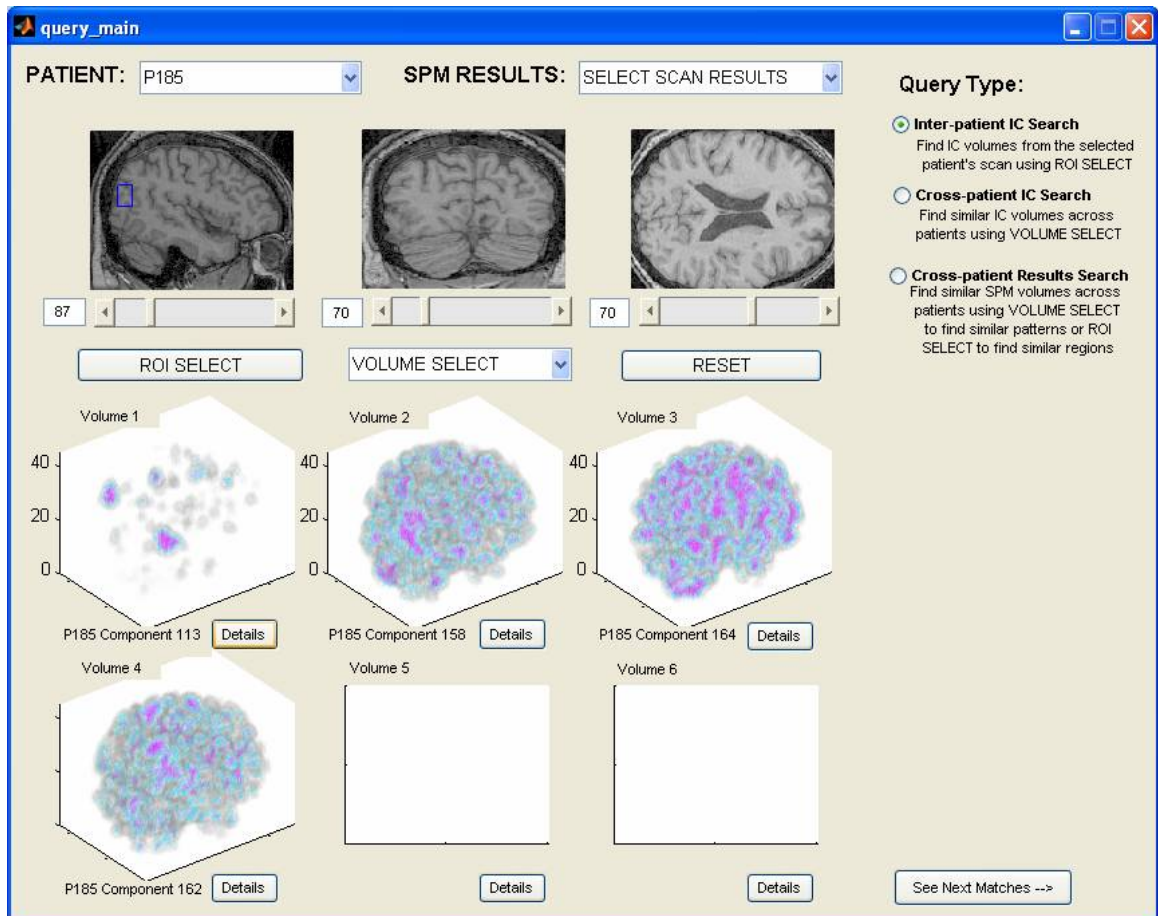


Figure 3.3: Sample results from query 1 showing a selected region of interest and the returned IC's with significant activation in that region.

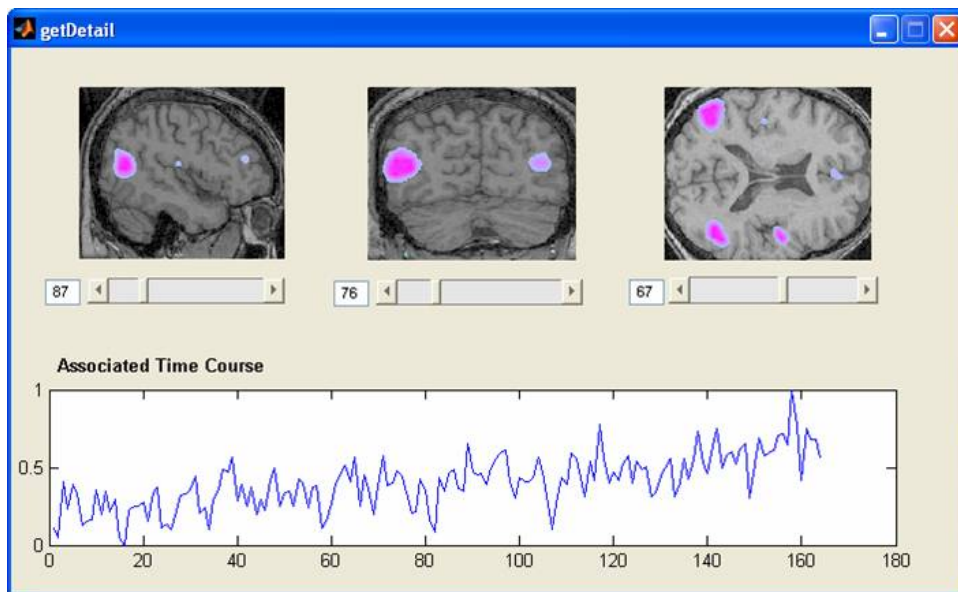


Figure 3.4: Detailed view of the first match, IC 113, showing the mapping onto the patient's structural MRI image and the associated time-course.

3.2 Query 2.

Starting with one patient’s voxel correlations, search for patients who have a similar correlation pattern for a voxel in the same region.

This query starts with results retrieved by the previous query. Once the query 1 algorithm has been used to find independent activation maps relevant to the user-selected region of interest, the user may select an interesting map and search for similar maps across patients. This can reveal whether this pattern is specific to the patient, or if there is evidence that other people may have had a similar activation during the scan. This information allows the user to explore the frequency with which this pattern occurs in the pool of patients in the database, and may look for other commonalities between the returned patients.

The steps to follow to execute this query, using the interface shown in Figure 3.2, are:

1. Complete steps 1 through 6 from query 1 to display relevant IC maps from the selected patient’s scan.
2. Change the query type to ‘Cross-patient IC Search’, using the radio buttons.
3. Use the ‘VOLUME SELECT’ menu to choose one of the six displayed IC maps.

The menu should be set to the title of the plot of the selected IC map.

The similarity of the ICA activation patterns is measured by an algorithm that uses the specialized clustering method described in Chapter 2 to find activated regions in the brain. For each patient in the database, there is one fMRI scan, with 168 corresponding IC maps. Each IC map has been pre-processed with the clustering algorithm, and n feature vectors, describing the characteristics of the n activation clusters identified in that map, are stored for each IC map. The number of activation clusters identified may be different for each IC. The algorithm for the comparison is:

1. Load the feature vector for each activation cluster k in the selected IC map from the query patient.
2. For each patient i in the database, and each IC map j from patient i , load the feature vector for each activation cluster f in IC map j , for patient i .
3. Find $mindist_{i,j,k}$, the minimum over f , of the distances between feature vector f from patient i , IC map j and feature vector k from the selected IC map of the query patient.

$$mindist_{i,j,k} = \min_f \left[\sum abs(feavec_{i,j,f} - feavec_k) \right]$$

4. If $mindist_{i,j,k} > 50$

$$clustermatch_{i,j,k} = mindist_{i,j,k} * binsize_{i,j} * averagebinweight_{i,j}$$

where $binsize$ and $averagebinweight$ are two elements of the feature vector. Else,

$$clustermatch_{i,j,k} = 0.$$

5. $ICmatch_{i,j} = \sum_k clustermatch_{i,j,k}$

These steps assign a ranking, $ICmatch_{i,j}$ to each IC map from each patient in the database. The six IC maps with the highest non-zero rankings, over all patients and all IC maps, are returned to the user, displayed in their re-wrapped, three-dimensional format. Like in the first query, the ‘DETAILS’ push button under each returned result can be used to get more information. Selecting this button launches a browser which displays the IC mapped back onto the patient’s structural MRI volume and a plot of the IC’s associated time-course. If there are more than six IC maps with a non-zero ranking, the ‘NEXT’ button can be used to browse through the remaining maps.

Sample results from query 2 are shown in Figures 3.5 through 3.8. Since query 2 requires the use of results from query 1, Figure 3.5 shows a sample query 1, ROI

based IC search for patient *P185*. The IC of interest which will be selected for query 2 is the first match returned, IC 87. The details from this IC map can be seen in Figure 3.6. The results of query 2 are shown in Figure 3.7, and it can be seen that ‘VOLUME 1’, the plot title of the first match from query 1, has been selected on the ‘VOLUME SELECT’ menu. The best match for patient *P185*, IC 87 is patient *P189*, IC 94. The details from patient *P189*, IC 94 can be seen in Figure 3.8. Comparing the match IC details to the query IC details in Figure 3.6, the similarities can be easily seen.

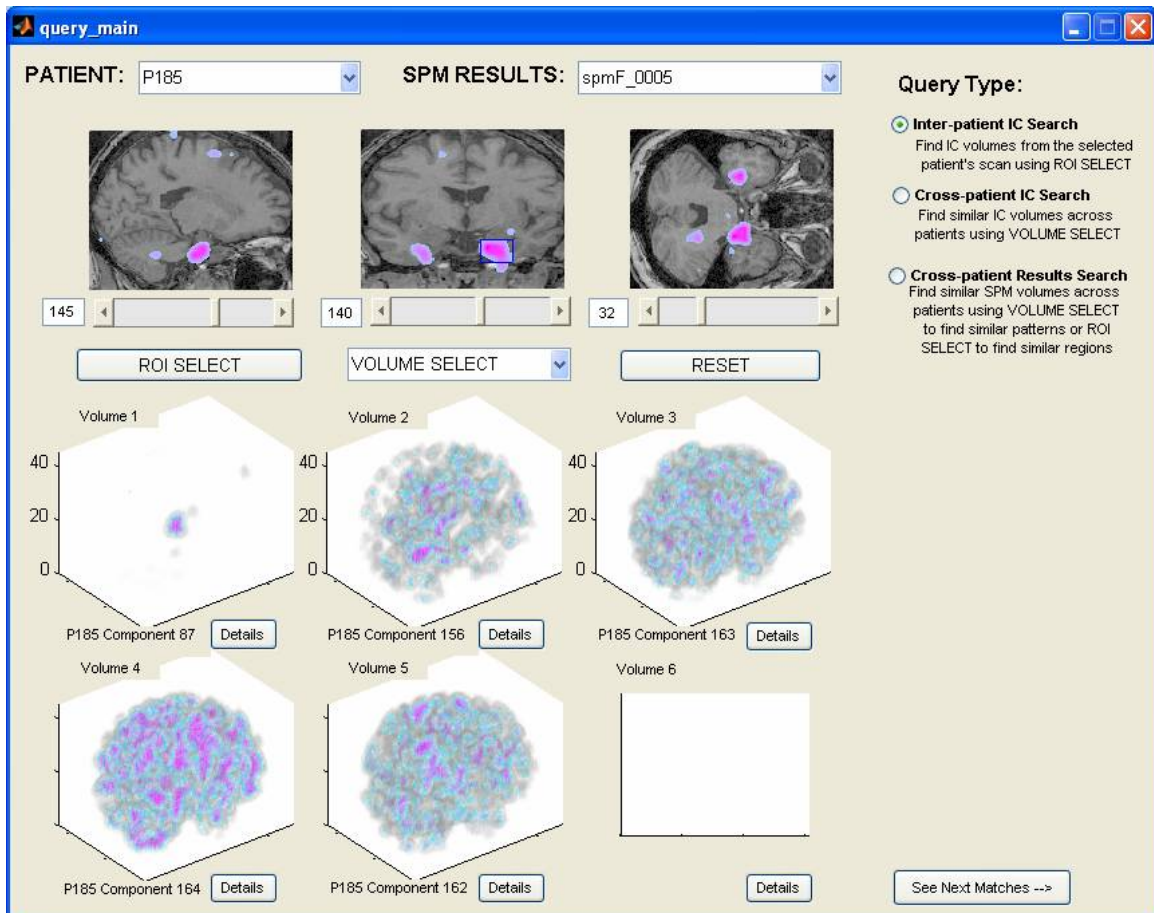


Figure 3.5: Sample results from query 1 showing a selected region of interest and the returned IC’s with significant activation in that region.

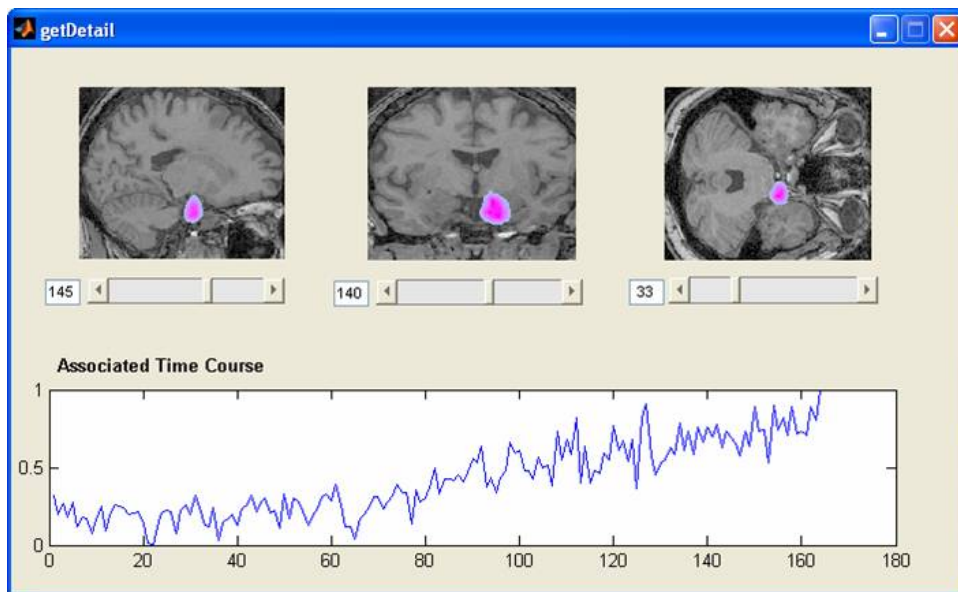


Figure 3.6: Detailed view of the first match, IC 87, showing the mapping onto the patient's structural MRI image and the associated time-course.

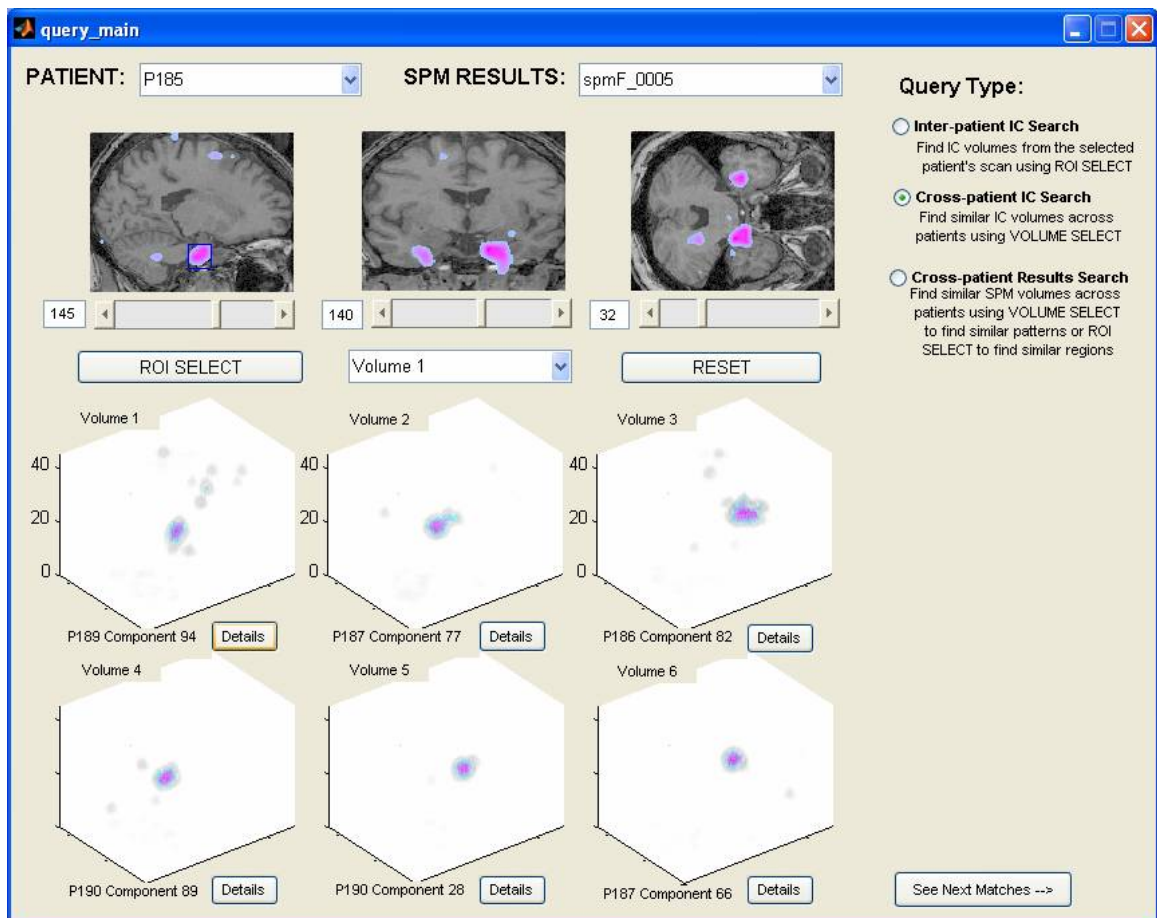


Figure 3.7: Sample results from query 2 showing IC's in the database that are similar to selected volume 1 from the query in Figure 3.5.

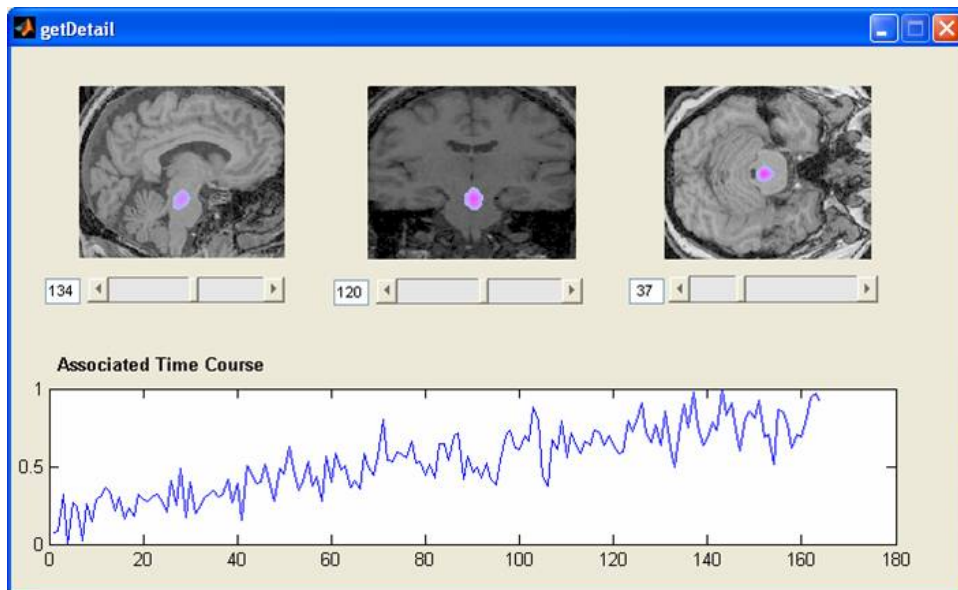


Figure 3.8: Detailed view of the first match from query in Figure 3.5, IC 94 from patient *P189*, showing the mapping onto the patient's structural MRI image and the associated time-course.

3.3 Query 3.

For a given patient fMRI and an activation threshold, find patients with similar activation patterns. This will query the fMRI images showing statistically significant activations.

This query searches the SPM results maps from all patients in the database to find those similar to a results map selected from the query patient. The SPM results map is a standard visual representation of statistically significant activations from an fMRI scan created by the software package SPM. Each voxel time series is correlated with the experimental design, and the value displayed at each voxel in the results map is the f-score or t-score indicating the statistical significance of its activation. The steps to execute this query, using the user interface shown in Figure 3.2, are:

1. Select the patient number from the ‘PATIENT’ menu to load the patient’s structural volume into the main window.
2. Select an SPM results file from the ‘SPM RESULTS’ menu to map voxels with time-courses statistically related to the experimental design onto the structural volume.
3. Change the query type to ‘Cross-patient Results Search’, using the radio buttons.
4. Use the ‘VOLUME SELECT’ menu to choose option ‘Current SPM Results’.

When the currently loaded SPM results map has been selected from the ‘VOLUME SELECT’ menu, the algorithm compares the results patterns from all patients, using the specialized clustering algorithm introduced in Chapter 2, to identify activations cluster locations and measure their similarity. For each patient entered in the database, there is one fMRI scan and sixteen corresponding SPM maps. Each

SPM map has been preprocessed with the clustering algorithm. For each SPM map in the database, there are k feature vectors describing the characteristics of k activation clusters in the map. Each map may have a different number of clusters. The algorithm for the SPM map comparisons is:

1. Load the feature vector for each activation cluster k in the query SPM results map.
2. For each patient i in the database and each SPM results map j for patient i , load the feature vector for each activation cluster f in SPM results map j , from patient i .
3. Find $mindist_{i,j,k}$, the minimum over f , of the distances between feature vector f from patient i , SPM results map j and feature vector k from the query SPM results map

$$mindist_{i,j,k} = \min_f \left[\sum abs(feavec_{i,j,f} - feavec_k) \right].$$

4. If $mindist_{i,j,k} > 50$

$$clustermatch_{i,j,k} = mindist_{i,j,k} * binsize_{i,j} * averagebinweight_{i,j},$$

where $binsize$ and $averagebinweight$ are two elements of the feature vector. Else,

$$clustermatch_{i,j,k} = 0.$$

5. $SPMmatch_{i,j} = \sum_k clustermatch_{i,j,k}$

These steps assign a ranking, $SPMmatch_{i,j}$ to each SPM map from each patient in the database. The six SPM maps with the highest non-zero rankings, over all patients and all SPM maps, are returned to the user, displayed in three-dimensional

format. To view these results mapped onto the structural fMRI of the appropriate patient, the ‘DETAILS’ push button should be selected. Like the other queries, if there are more than six IC maps with a non-zero ranking, the ‘NEXT’ button can be used to browse through the remaining maps.

The results from a sample query 3 are shown in Figure 3.9, where the SPM results map *spmF005* from patient *P185* is loaded onto that patient’s structural MRI images, and ‘Current SPM Results’ selected from the ‘VOLUME SELECT’ menu. The similar results are returned in volume plots one through six. The details from the first match, the SPM results map *spmF0001* from patient *P186*, are shown in the details browser in Figure 3.10. Comparing the structural mapping of the best match in Figure 3.10 to the structural mapping of the original query in Figure 3.9, it can be seen that they are very similar.

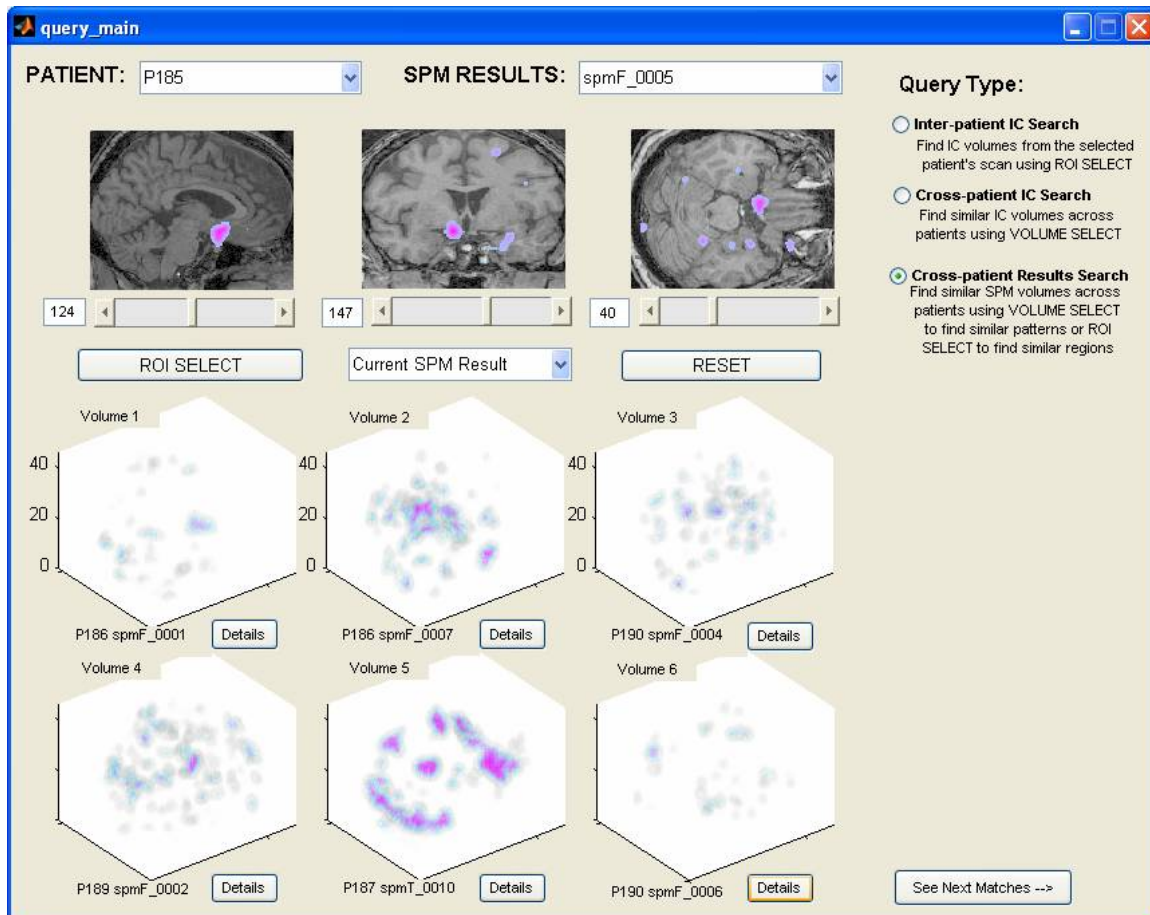


Figure 3.9: Sample results from query 3 showing the loaded SPM results from the query patient, and the returned similar SPM results from other patients in the database.

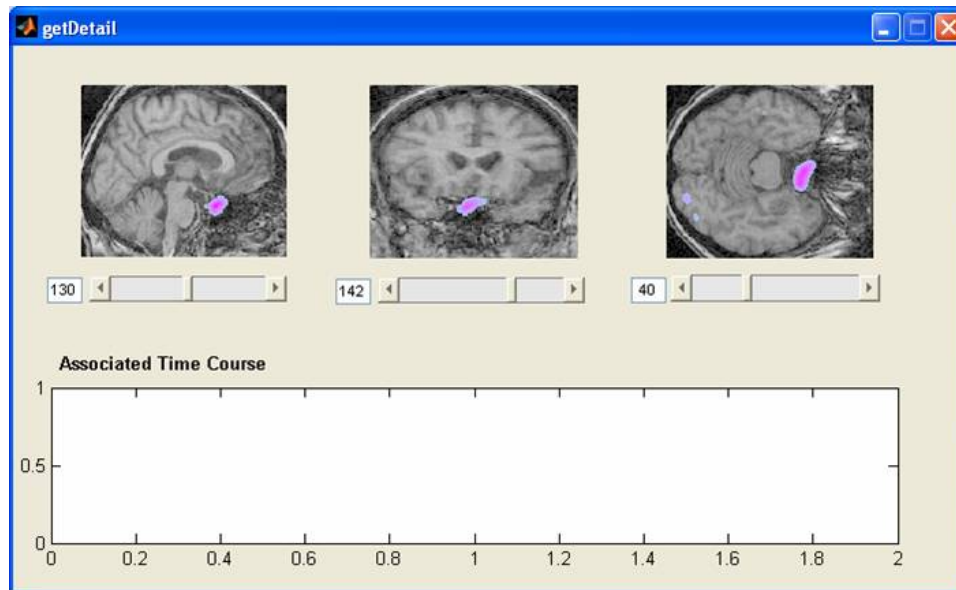


Figure 3.10: Detailed view of the first match, patient *P186 spmF0001*, showing the mapping onto the patient's structural MRI image.

3.4 Query 4.

For a given patient fMRI and a given location, find other patients who have greater than or equal activation values at that location. This will query the fMRI images showing statistically significant activations.

The last query is also based on the results image generated by SPM. To find other patients with a similar activation in the results image, the user selects the region of interest from the displayed results image, and the average statistical score of the area is calculated. The same area is extracted from each patient in the database, and the average scores are calculated. Patients with the same or higher average score are returned to the user, and their results images are displayed. This query differs from the previous query by using a region of interest, rather than the entire pattern to compare SPM results maps across patients. To execute a query of this type using the interface shown in Figure 3.2 the user will:

1. Select the patient number from the ‘PATIENT’ menu to load the patient’s structural volume into the main window.
2. Optionally select an SPM results file from the ‘SPM RESULTS’ menu to map voxels with time-courses statistically related to the experimental design onto the structural volume. This may be used to identify a region of interest to the user.
3. Select the query type, ‘Cross-patient Results Search’, using the radio buttons.
4. Browse the selected patient’s structural volume until a region of interest is identified.
5. Click the ‘ROI SELECT’ push button.

6. Use the mouse to draw a box around the region of interest on any of the three volume views.

As discussed in query 1, the region of interest is a cube, with length and width corresponding to the ROI box drawn by the user and the depth predefined as five samples centered on the slice the box has been drawn on. Once the ROI has been selected, the selected patient's SPM results maps with an average activation in the region above a threshold are sorted by relevance and returned to the user in re-wrapped, three dimensional format. The algorithm for finding and sorting the activated volumes is:

1. For each patient i in the database and each SPM results map j from patient i , find the average activation value of the SPM results map in the selected ROI.
2. If the average activation is greater than the threshold of one third of the maximum results map value, continue with the next steps to calculate the ranking. Else,

$$rank_{i,j} = 0,$$

and return to the first step.

3. Find n , the number of activation clusters in SPM results map j .
4. Find d_{min} , the minimum of the distances between the ROI center each activation cluster center in SPM results map j .
5. Assign SPM results map i , from patient j , the rank

$$rank_{i,j} = \frac{1}{n^2 d_{min}}.$$

Every SPM map in the selected patient's scan is been ranked by the assignment of the value $rank$, and the maps with the highest non-zero rankings are returned to

the user. The SPM maps are displayed in three-dimensional format. The ‘DETAILS’ push button under each returned result allows the user to display the selected SPM results volume mapped onto the appropriate patient’s structural MRI volume.

The results from a sample query 4 is shown in Figure 3.11. The region of interest can be seen outlined in blue on the first structural image. The detail browser for the first and fourth match can be seen in Figures 3.12 and 3.13. Comparing the plots from the detail browsers to the original query image, it can be seen that both results have a significant activation in the ROI.

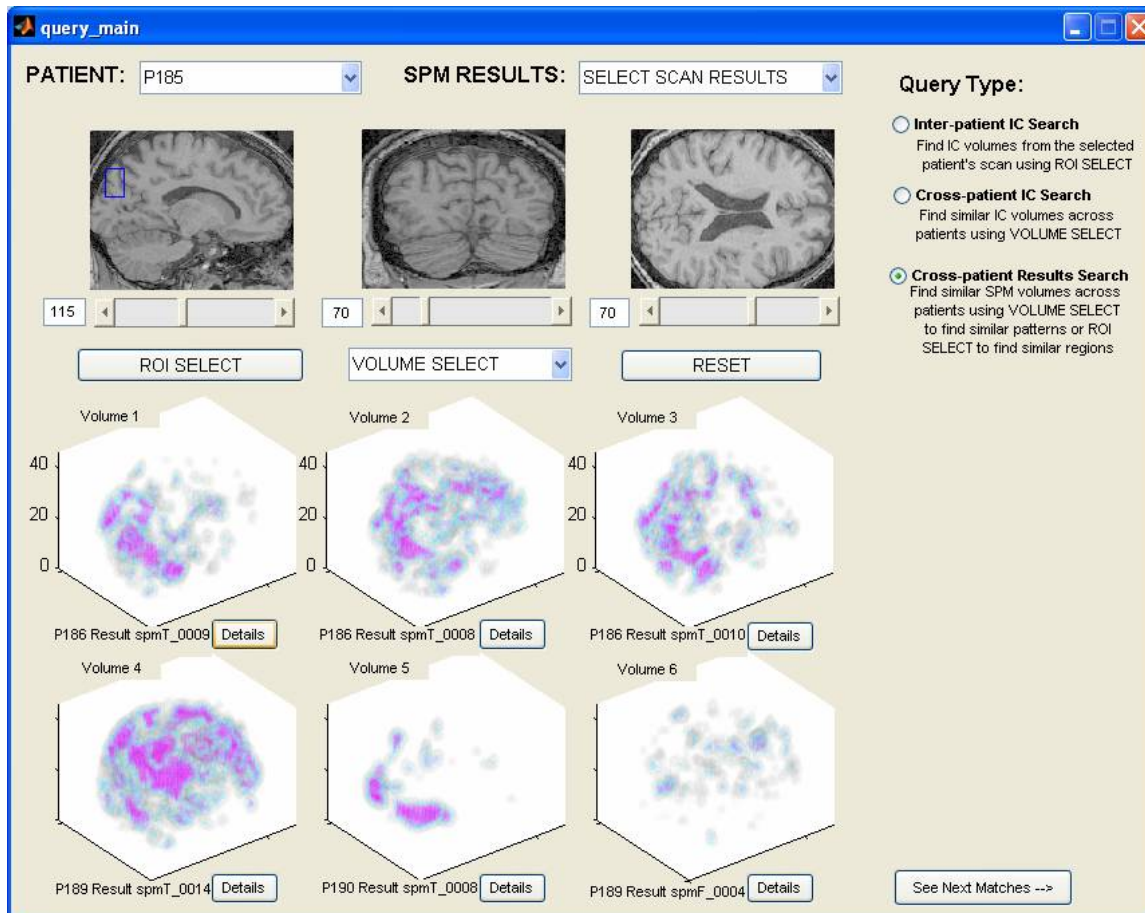


Figure 3.11: Sample results from query 4 showing a selected region of interest and the returned SPM similarity volumes from other patients in the database with significant activations in that region.

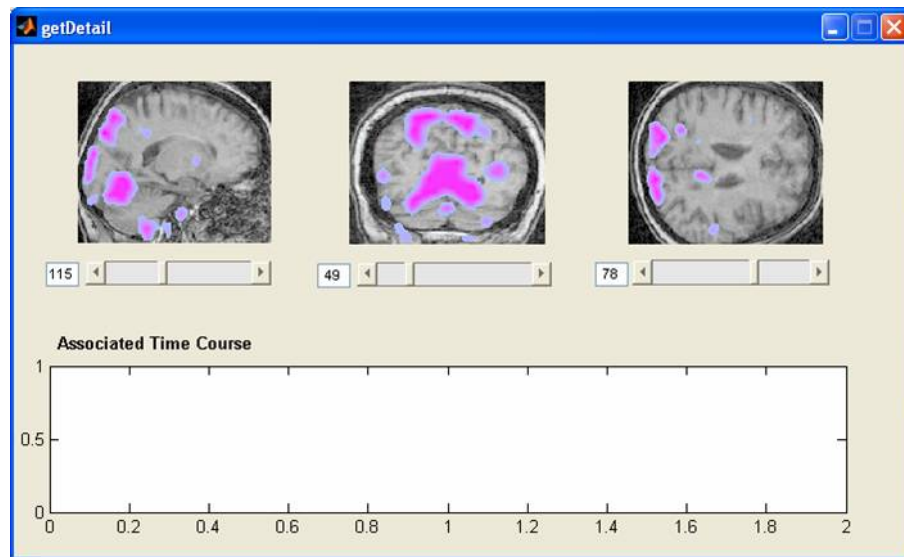


Figure 3.12: Detailed view of the first match, patient *P186*, *spmT0009*, showing the mapping onto the patient's structural MRI image.

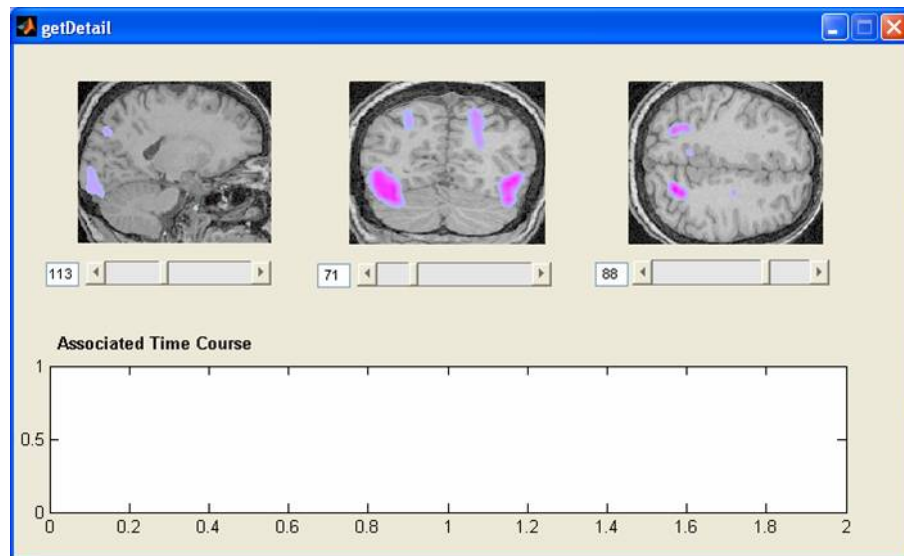


Figure 3.13: Detailed view of the fifth match, patient *P190*, *spmT0008*, showing the mapping onto the patient's structural MRI image.

Chapter 4

EVALUATION AND DISCUSSION OF RESULTS

In the previous chapter, several different sample queries were shown that used independent component analysis to evaluate similarity between voxel regions. A thorough evaluation of this technique would require knowing the actual relationships between all voxels. Since this ground truth is not available, we can evaluate the independent component maps based on their relationships to both the raw scan data and SPM results maps.

4.1 Relationship to fMRI Raw Data

One way to evaluate the IC maps is to look at how well they relate back to the original raw data from the fMRI scan. An IC map shows regions in the brain which covary according to the time-course associated with that IC map. The averaged raw time-courses from the voxels activated in one IC map would then be expected to relate back to that IC map's associated time-course. However, this is not a direct relationship. The maps are overlapping, so one IC's associated time-course will specify only one, statistically independent, component in the activated voxel time-courses.

One way to verify that the averaged raw data time-courses from an IC map's activated voxels contains a component relating to the associated time-course is to look at the frequency content of both signals. The raw data contains many statistically independent components, so it is expected that there will be several peaks in its frequency spectrum that are not present in the frequency spectrum of the IC associated time-course. However, if the associated time-course contains strong peaks at a few frequencies, corresponding peaks should be observable at the same frequencies in the

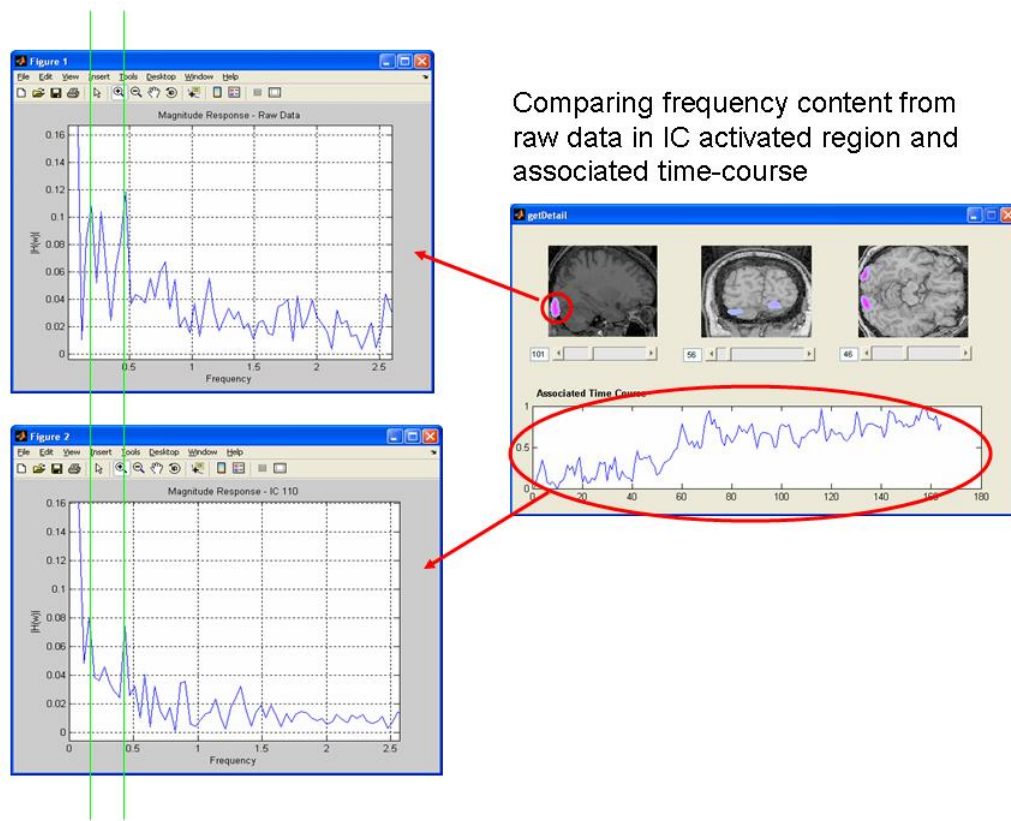


Figure 4.1: IC map relationship to fMRI raw data time-courses

raw data.

To perform this comparison, an IC map was selected, and the activated voxel locations were recorded. The raw time-courses for the voxels at these locations were then extracted and averaged. The discrete-time Fourier transform (DTFT) of this signal was plotted, and compared to the DTFT of the time-course associated with the IC map. This was done for several IC maps, and it was observed that the peaks in the associated time-course frequency spectrum corresponded to peaks in the averaged raw data frequency spectrum. An example of this comparison is shown in Figure 4.1.

4.2 Relationship to SPM Results

Statistical parametric mapping (SPM) is a standard way to identify correlated voxels in an fMRI scan. However, unlike the IC-based method developed in this work, SPM does not look at the relationships between voxels. Instead, it models the hypothesized response to the exam stimuli of the fMRI scan session, and each voxel is independently evaluated for a statistical relationship. This produces a map of the voxels that have a direct relation to the modeled stimuli response.

The IC-based method of finding voxel correlations is expected to be much more powerful than the SPM method, since it is not limited by the use of a hypothesized response, and can look for functionally connected regions which are indirectly dependent, or independent of the exam stimuli. However, while the IC-based method should return more information, it should still be able to replicate the results found by SPM, since the SPM map indicates voxel regions which vary according to the exam stimuli, a statistically independent source. It is expected that for each SPM result map, there would be a corresponding IC map with a similar activation pattern, and an associated time-course similar to the stimuli response model used by SPM. A ‘boxcar’ stimuli design was used for all scan sessions in the data base, so associated time-course is expected to show an ‘on-off’ pattern.

Using the user interface to browse through IC and SPM activation maps, it was concluded that the SPM results maps were a subset of the IC maps. An example of this finding is shown the following figures. In Figure 4.2, the selected SPM results map from patient *P189*, which will be compared to that patient’s IC maps, is shown. To find IC maps similar to the selected SPM map, an ROI based query is executed. The selected SPM results are mapped on to patient *P185*’s structural image in the main query window, as shown in Figure 4.3. The ‘Inter-patient IC Search’ query type is selected, and an ROI is drawn over the SPM map activations. The first IC map match returned by this query looks very similar to the SPM map shown in Figure

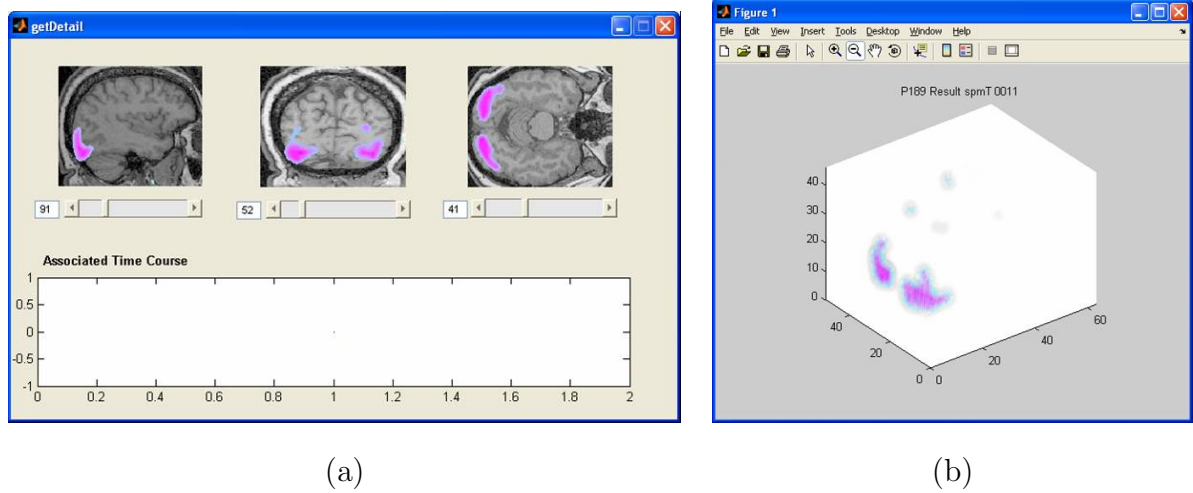


Figure 4.2: SPM results map shown (a) mapped onto the patient's structural MRI and (b) as a three-dimensional volume.

4.2 (b). The detail browser for this IC map is shown in Figure 4.4, and its structural similarity to the SPM map shown in Figure 4.2 (a) can be easily seen. In addition, the time-course associated with the IC map clearly has the expected 'on' and 'off' pattern associated with the stimuli response model.

To investigate whether this pattern could be found in other patients, the database was queried to find other patients with SPM results similar to those shown in Figure 4.2 for patient *P189*. Due to the fact that the patients in this database have intractable epilepsy, all patients are not expected to have normal responses to the experiment stimulus, and it cannot be assumed that the response to the task will be the similar over all patients.

Searching for patients with similar SPM results allows the identification of patients within the database who have a similar response to the experimental task. The results of this search are shown in Figure 4.5. It can be seen that the majority of the matches are from patients *P185* and *P190*, with a weaker match from *P187*. Next, a ROI-based search was performed to find IC maps corresponding to the SPM results for

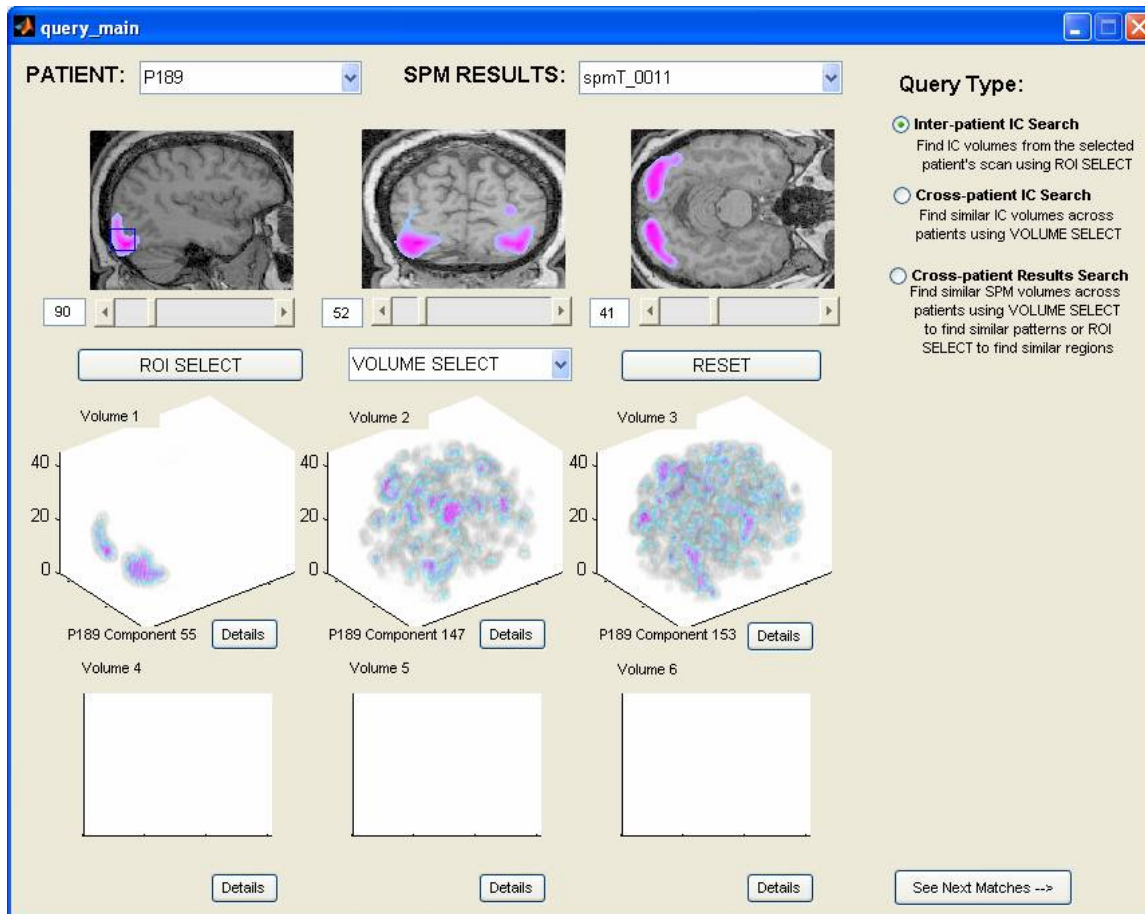


Figure 4.3: Query showing ROI-based search for IC maps with activation in the same region as the selected SPM results map.

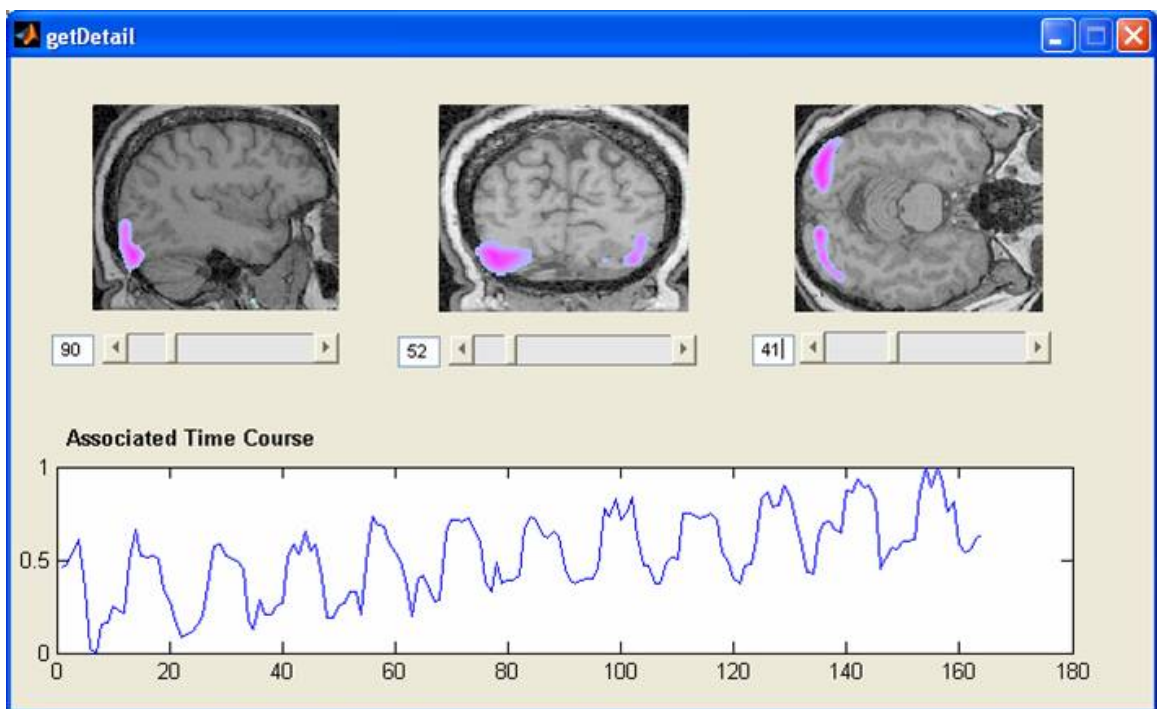
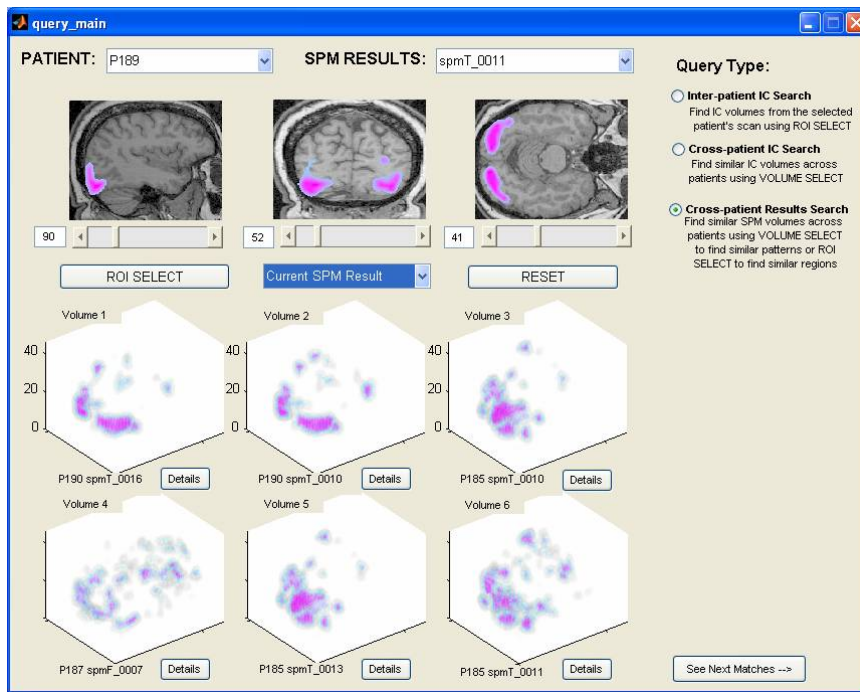
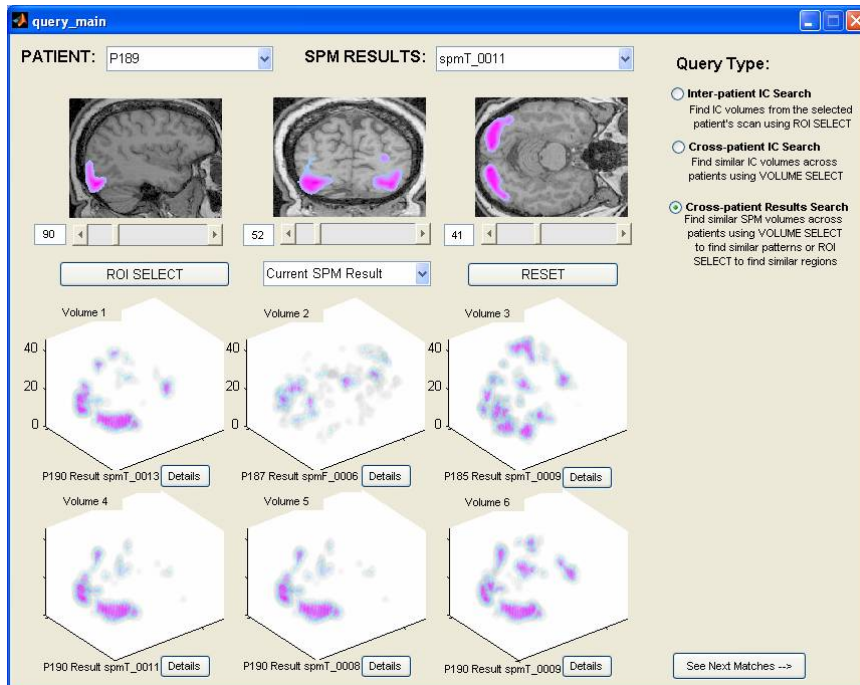


Figure 4.4: Detail browser for patient *P185* IC map 55 showing best match for SPM results map in Figure 4.2.



(a)



(b)

Figure 4.5: Search results for cross-patient SPM results that are similar to the SPM results for patient *P185*, which are displayed on the structural image. The first page of results is shown in (a) and the second page of results is shown in (b).

each of these patients.

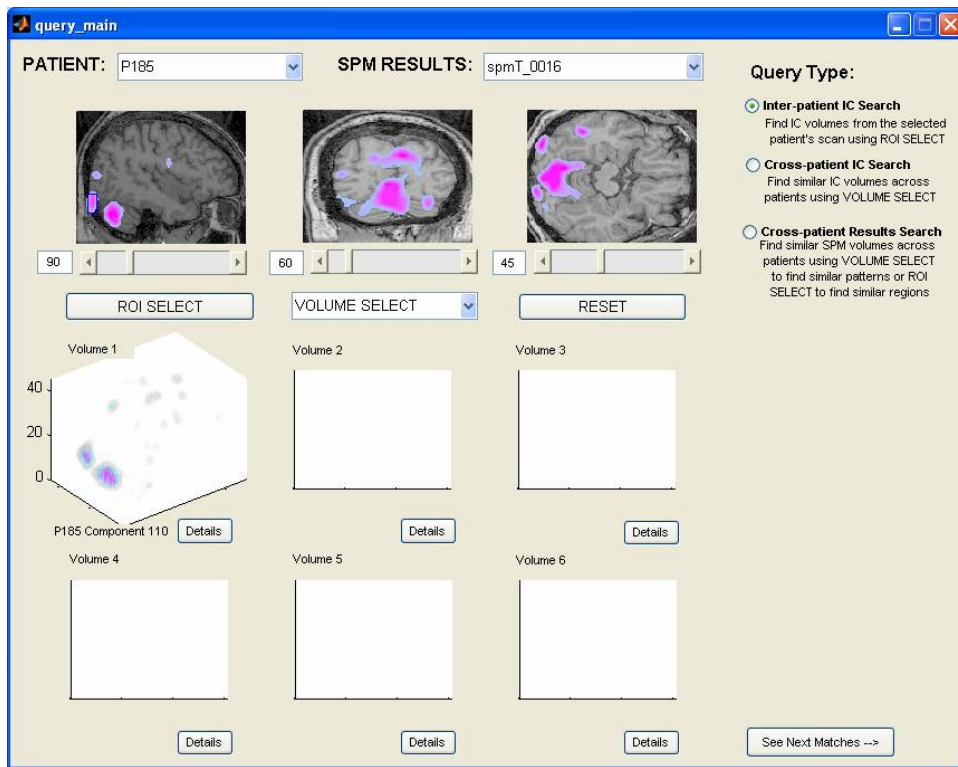
The results for these queries are shown in Figures 4.6 through 4.8. It can be clearly seen that the associated time-courses for patients *P185* and *P190* have a strong correlation to the ‘boxcar’ experimental stimulus. Patient *P187* has a slightly weaker correlation, but is still similar.

4.3 Contributions

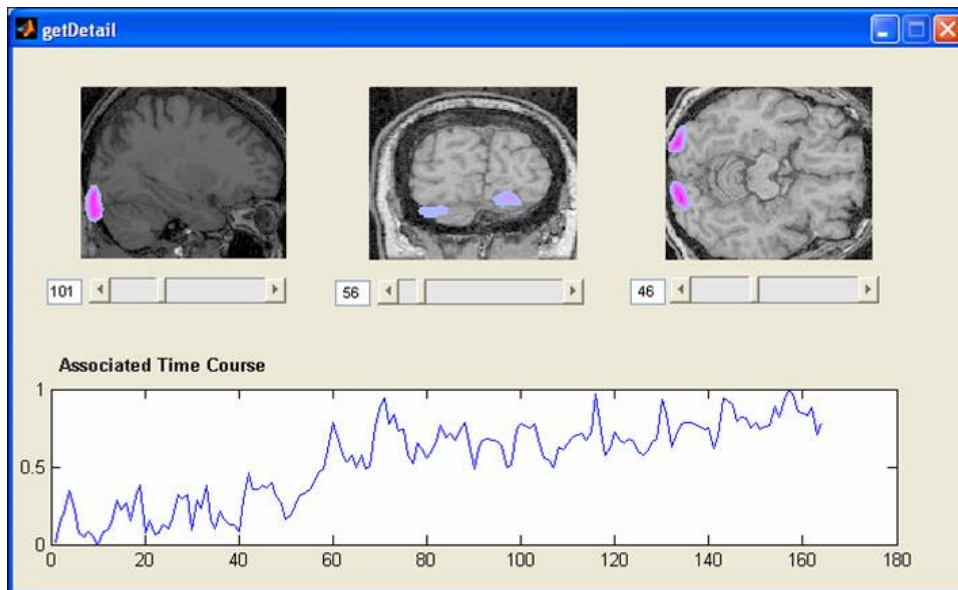
This thesis introduces a tool that uses independent component analysis as a measure of voxel similarity to allow the user to explore statistically independent maps of correlated voxel activity. A specialized clustering technique, designed to find and characterize clusters of activated voxels, allows comparison of the independent component spatial maps across patients. Similar search capabilities are provided for spatial maps produced by the popular fMRI analysis tool SPM, so that the relationship between the IC maps and the results produced by SPM may be investigated.

The main contributions of this work are:

1. a method for applying Independent Component Analysis as a voxel similarity measurement
2. a new algorithm for detecting and extracting characteristics from significant clusters of activations in three-dimensional activation maps.
3. a three-dimensional spatial similarity measurement based on cluster feature extraction
4. the development of algorithms for answering the four query types described in this work
5. construction of a user interface for executing these queries.

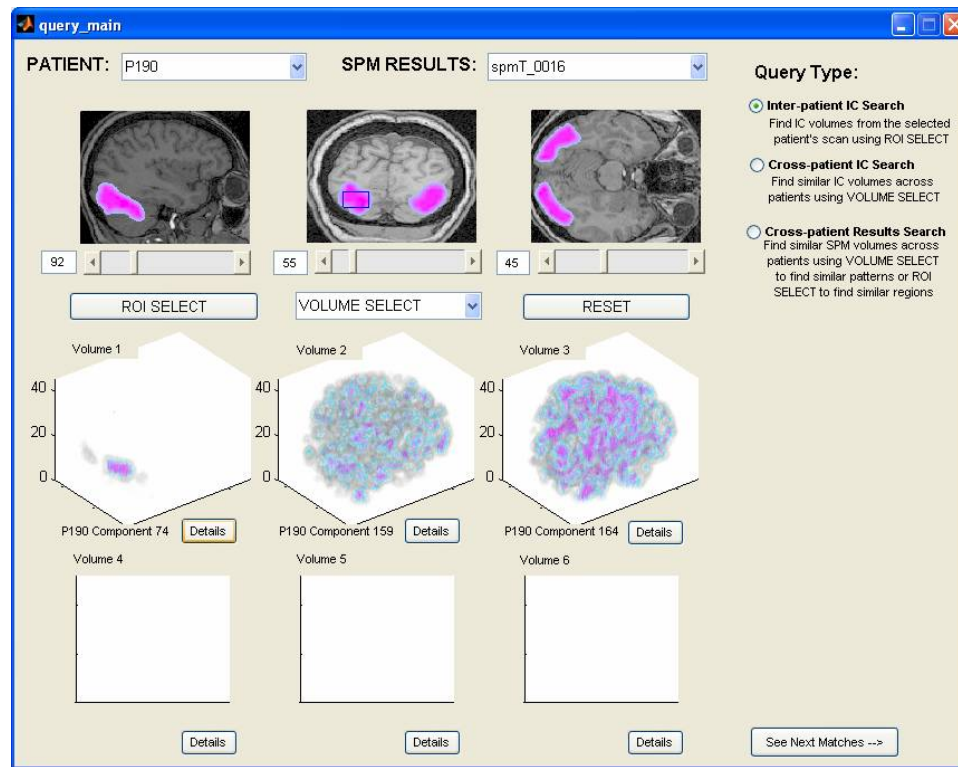


(a)

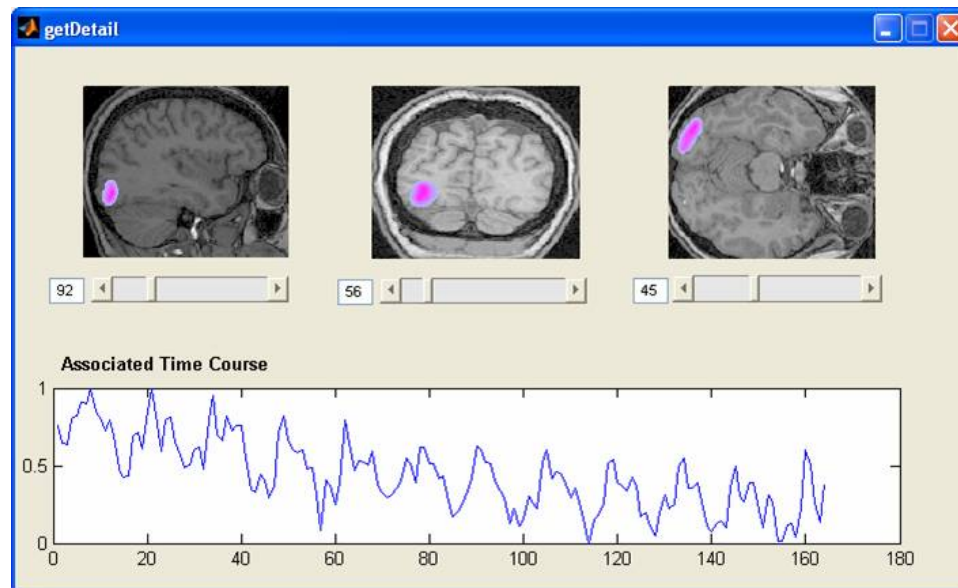


(b)

Figure 4.6: Results for ROI-based search for patient *P185*. The IC map similar to the SPM results displayed on the structural image is shown in (a). Details for the returned IC map 110 are shown in (b).

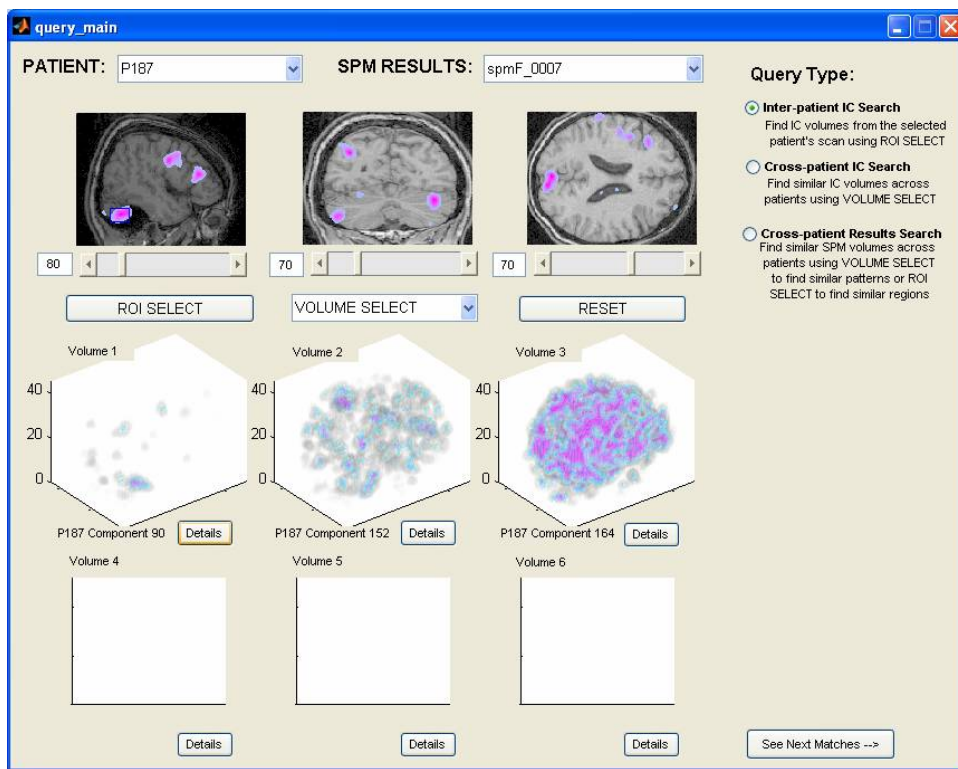


(a)

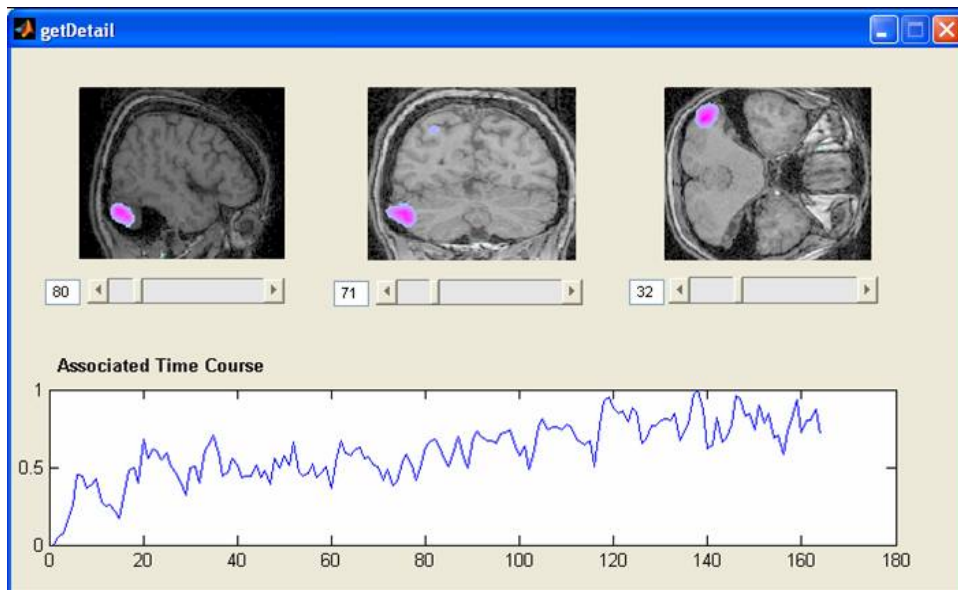


(b)

Figure 4.7: Results for ROI-based search for patient *P190*. IC maps similar to the SPM results displayed on the structural image are shown in (a). Details for the first match, IC map 74 are shown in (b).



(a)



(b)

Figure 4.8: Results for ROI-based search for patient *P187*. IC maps similar to the SPM results displayed on the structural image are shown in (a). Details for the first match IC map 90 are shown in (b).

4.4 Future Work

This work has introduced a tool for exploring functional connections in the brain using fMRI data. Most of the analysis that has been done on the independent components has been concentrated on the spatial maps. As seen in the example in Figure 4.4, the time-courses associated with the component maps can provide a significant amount of information about the source and importance of the map. Further work could be done to incorporate this information into the existing query algorithms, or to use it in the design of new queries.

BIBLIOGRAPHY

- [1] C.F. Beckmann, M. DeLuca, J.T. Devlin, and S.M. Smith. Investigations into resting-state connectivity using independent component analysis. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 360(1457):1001–1013, 2005.
- [2] A.J. Bell and T.J. Sejnowski. An information-maximization approach to blind separation and blind deconvolution. *Neural Computation*, 7(6):1129–1159, 1995.
- [3] D. Cordes, V.M. Haughton, K. Arfanakis, G.J. Wendt, P.A. Turski, C.H. Moritz, M.A. Quigley, and M.E. Meyerand. Mapping functionally related regions of brain with functional connectivity mr imaging. *American Journal of Neuroradiology*, 21(9):1636–1644, 2000.
- [4] B. Horwitz, B. Warner, J. Fitzer, M.A. Tagamets, F.T. Husain, and T.W. Long. Investigating the neural basis for functional and effective connectivity. application to fmri. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 360(1457):1093–1108, 2005.
- [5] A. Hyvärinen. *New approximations of differential entropy for independent component analysis and projection pursuit*. MIT Press Cambridge, MA, USA, 1998.
- [6] A. Hyvärinen and E. Oja. A fast fixed-point algorithm for independent component analysis. *Neural Computation*, 9(7):1483–1492, 1997.
- [7] T.P. Jung, S. Makeig, M.J. McKeown, A.J. Bell, T.W. Lee, and T.J. Sejnowski. Imaging brain dynamics using independent component analysis. *Proceedings of the IEEE*, 89(7):1107–1122, 2001.
- [8] T.W. Lee, M. Girolami, A.J. Bell, and T.J. Sejnowski. A unifying information-theoretic framework for independent component analysis. *Computers and Mathematics with Applications*, 39(11):1–21, 2000.
- [9] M.J. McKeown, S. Makeig, G.G. Brown, T.P. Jung, S.S. Kindermann, A.J. Bell, and T.J. Sejnowski. Analysis of fmri data by blind separation into independent spatial components. *Human Brain Mapping*, 6(3):160–188, 1998.

- [10] J.P. Nadal and N. Parga. Nonlinear neurons in the low-noise limit: a factorial code maximizes information transfer. *Network: Computation in Neural Systems*, 5(4):565–581, 1994.
- [11] S. Ogawa, RS Menon, SG Kim, and K. Ugurbil. On the characteristics of functional magnetic resonance imaging of the brain. *Annual Review of Biophysics and Biomolecular Structure*, 27(1):447–474, 1998.
- [12] A.V. Poliakov, K.P. Hinshaw, C. Rosse, and J.F. Brinkley. Integration and visualization of multimodal brain data for language mapping. *Proceedings of American Medical Informatics Association Fall Symposium*, pages 349–353, 1999.