

Machine learning classifiers and fMRI: A tutorial overview

Francisco Pereira^{a,*}, Tom Mitchell^b, Matthew Botvinick^a

^a Princeton Neuroscience Institute/Psychology Department, Princeton University, Princeton, NJ 08540, USA

^b Machine Learning Department, Carnegie Mellon University, Pittsburgh, PA 15213, USA

ARTICLE INFO

Article history:

Received 15 September 2008

Accepted 18 November 2008

Available online 21 November 2008

ABSTRACT

Interpreting brain image experiments requires analysis of complex, multivariate data. In recent years, one analysis approach that has grown in popularity is the use of machine learning algorithms to train classifiers to decode stimuli, mental states, behaviours and other variables of interest from fMRI data and thereby show the data contain information about them. In this tutorial overview we review some of the key choices faced in using this approach as well as how to derive statistically significant results, illustrating each point from a case study. Furthermore, we show how, in addition to answering the question of 'is there information about a variable of interest' (pattern discrimination), classifiers can be used to tackle other classes of question, namely 'where is the information' (pattern localization) and 'how is that information encoded' (pattern characterization).

© 2008 Elsevier Inc. All rights reserved.

Introduction

In the last few years there has been growing interest in the use of machine learning classifiers for analyzing fMRI data. A growing number of studies has shown that machine learning classifiers can be used to extract exciting new information from neuroimaging data (see Norman et al., 2006 and Haynes and Rees, 2006 for selective reviews). Along with the growth in interest and breadth of application, the methods underlying the use of classifiers with fMRI have continuously evolved and ramified (see O'Toole et al., 2007 for a historical overview). Given the novelty of the approach, there have been few attempts to organize and interrelate available methods in a single place. The present article strives to rectify that situation by providing a tutorial introduction to classifier methods in fMRI. Given that space constraints prevent us from being exhaustive, our focus is on introducing the methodology and the most common choices in its various aspects and, through that, perspective and references for further study.

Our presentation will be organized around the idea that classifier-based analyses, like traditional fMRI analyses, can be characterized in terms of a series of specific choices over a series of decision points in the analysis process, starting from selection of the scientific question to be asked and ending with choices among tests for hypotheses. We begin by laying out an illustrative example of a classifier-based analysis, and then dissect the analysis process, examining the set of choices it implicitly involves, and the alternatives available at each stage. There have been other proposals for a staged procedure for

using classifiers to analyze fMRI data (e.g. Strother et al., 2002), though with an emphasis on a single type of classifier and requiring the use of dimensionality reduction. We broaden this idea to include several kinds of classifier and also the use of feature selection. We conclude with a discussion of the sorts of scientific questions that may be fruitfully addressed using classifiers, and the ways in which the choice of question impacts subsequent analysis decisions.

Where possible, we discuss strengths and weaknesses of competing options. However, it is important to acknowledge from the outset that firm grounds for such evaluations are in many cases not yet available. Inasmuch as we can, we offer recommendations based on the results of formal principles or controlled empirical tests (from Pereira, 2007). Where these are lacking, we shall sometimes inject impressions drawn from our own personal experience with classifier-based analysis of fMRI data. Before entering into the discussion of the analysis proper, we begin in the next section by briefly introducing machine learning classifiers in their own right.

What is a classifier?

Classification is the analogue of regression when the variable being predicted is discrete, rather than continuous. In the context of neuroimaging, regression is most commonly used in the shape of a General Linear Model, predicting the time series of each voxel from many columns in the design matrix (Friston et al., 2007). Classifiers are used in the reverse direction, predicting parts of the design matrix from many voxels.

At a more detailed level, a classifier is a function that takes the values of various *features* (independent variables or predictors, in regression) in an *example* (the set of independent variable values) and predicts the *class* that that example belongs to (the dependent

* Corresponding author.

E-mail addresses: fpereira@princeton.edu (F. Pereira), tom.mitchell@cs.cmu.edu (T. Mitchell), matthewb@princeton.edu (M. Botvinick).

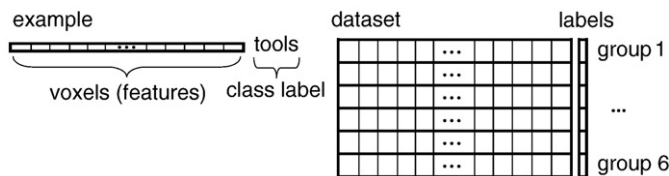


Fig. 1. An example where features are voxels arrayed as a row vector (left) and a dataset is matrix of such row vectors (right).

variable). In a neuroimaging setting, the features could be voxels and the class could be the type of stimulus the subject was looking at when the voxel values were recorded (see Fig. 1). We will denote an example by the row vector $\mathbf{x} = [x_1 \dots x_v]$ and its class label as \mathbf{y} . A classifier has a number of parameters that have to be learned from training data—a set of examples reserved for this purpose—similarly to how regression parameters are estimated using least squares. The learned classifier is essentially a model of the relationship between the features and the class label in the training set. More formally, given an example \mathbf{x} , the classifier is a function f that predicts the label $\mathbf{y} = f(\mathbf{x})$.

Once trained, the classifier can be used to determine whether the features used contain information about the class of the example. This relationship is tested by using the learned classifier on a different set of examples, the test data. Intuitively, the idea is that, if the classifier truly captured the relationship between features and classes, it ought to be able to predict the classes of examples it hasn't seen before. The typical assumption for classifier learning algorithms is that the training (and testing) examples are independently drawn from an 'example distribution'; when judging a classifier on a test set we are obtaining an estimate of its performance on any test set from the same distribution. This is depicted in Fig. 2. We will denote the training and test sets by $\mathbf{X}_{\text{train}}$ and \mathbf{X}_{test} , matrices with respectively n_{train} and n_{test} examples as their rows, and the example labels by the column vectors $\mathbf{y}_{\text{train}}$ and \mathbf{y}_{test} . The most commonly used measure of how well a classifier does on the test set is its accuracy. This is simply the fraction of examples in the test set for which the correct label was predicted, i.e. $\frac{\sum_{i=1}^{n_{\text{test}}} I(f(x_i), y_i)}{n_{\text{test}}}$, where $I(f(x_i), y_i) = 1$ if $f(x_i) = y_i$ (the label of the i th example was predicted correctly) or 0 otherwise.

As we shall discuss in Choosing a classifier section, there are several types of classifier. However, for reasons that will be introduced later, our emphasis will be on linear classifiers; in this type, the classification function relies on a linear combination of the features, i.e. $f(\mathbf{x}) = g(w_1 x_1 + \dots + w_v x_v)$, and the weights w_i in that combination are the parameters to be learned.

Classifier analysis: an illustrative example

In this section we will introduce an illustrative example of a classifier-based study, which shows the basic stages of bare-bones classifier analysis in order to provide a basis for the discussion in subsequent sections.

The experiment that originated our dataset¹ was designed to test whether a classifier could distinguish the activation as a result of seeing words that were either kinds of tool or kinds of building. The subject was shown one word per trial and performed the following task: the subject should think about the item and its properties while the word was displayed (3 s) and try to clear her mind afterwards (8 s of blank screen). There were 7 different items belonging to each of the two categories and 6 experimental epochs. An epoch in this setting was an experiment division where all 14 items were presented, without repetition; all 6 epochs had the same items.

Images were acquired with a TR of 1 s, with voxel dimensions $3 \times 3 \times 5$ mm. The dataset underwent a typical fMRI preprocessing stream from DICOM files using SPM (SPM, 2000). The steps followed were volume registration and slice acquisition timing correction, followed by a voxelwise linear detrending. A brain mask was extracted automatically and used to restrict voxels considered in the subsequent steps to roughly 16,000. Although the original dataset contained several subjects, we will focus on a single one.

Each trial was converted into an example by taking the average image during a 4 s span while the subject was thinking about the stimulus shown a few seconds earlier; these 4 s contained the expected peak of the signal during the trial, based on the standard BOLD haemodynamic response function. Each example was thus a vector containing the 16,000 feature values, corresponding to the 4 s average signal at each of the voxels. Using each trial to obtain an example yielded a total of $42 = 7 \text{ items} \times 6 \text{ repetitions}$ examples for each of the tools and buildings categories. Each example vector was normalized to have mean 0 and standard deviation 1. Examples were further grouped by which of the 6 epochs they belonged, as depicted in Fig. 1.

With examples in hand, the dataset was divided into a training and a test set, with a goal of learning a classifier whose task was to predict the category of an example. As mentioned previously, the parameters of a classifier are learned from the examples in the training set, and the classifier is then evaluated on the examples in a separate test set (see Fig. 2).

In the example study we divided the dataset for one subject into examples coming from either the even or the odd epochs, assigning those in epochs 1, 3 and 5 to be the training set (42 examples) and those in epochs 2, 4 and 6 to be the test set (42 examples). Given this, we trained a Gaussian Naive Bayes classifier (see Choosing a classifier section) and applied it to the test set. For 65% of the examples the label was predicted correctly.

Given our result, could we conclude that the classifier's accuracy was better than chance and, therefore, that the fMRI features contained relevant information about the stimulus category?

We considered the null hypothesis that the classifier was performing at chance level (e.g. the accuracy of prediction on average would be 50%) and tested it using the procedure we describe in Evaluating results section. For an accuracy of 65% with 42 test examples the test yielded a p -value of 0.0218 and the null hypothesis was rejected.

Classifier analysis: stages and options

The foregoing example, though presented as an unsegmented narrative, can be broken down into a set of stages, beginning with the conversion of raw data into a set of examples and proceeding through choice of classifier, training and test sets and interpretation of results. At each stage, the researcher faces a choice among multiple alternatives. In what follows we examine each stage and the choices

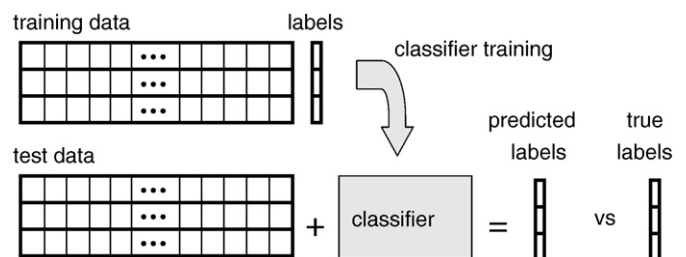


Fig. 2. A classifier is learned from the training set, examples whose labels it can see, and used to predict labels for a test set, examples whose labels it cannot see. The predicted labels are then compared to the true labels and the accuracy of the classifier—the fraction of examples where the prediction was correct—can be computed.

¹ Data was kindly provided by Robert Mason and Marcel Just, Center for Cognitive Brain Imaging, CMU, who ran the experiment.

it presents, as well as how those choices interact and are conditioned by practical factors.

Preparing data for classifier training

Creating examples

Prior to training (or applying) classifiers we need to transform the fMRI data into examples. Creating examples in general requires deciding what to use as features, how to extract their values from data and what we would like to predict. In our illustrative example we chose to average several TRs worth of images in a single trial to create an example; as a single word stimulus maps to a single trial this is a natural choice. However, many other options are available. We could instead have used each individual TR within a trial as an example by itself, or averaged many trials with words belonging to the same category into a single example. One is not limited to using voxels as features. We could use the average of several voxels in one ROI as a single feature or consider each voxel at each time point in a trial as a different feature. The latter is used in Mitchell et al., (2004) in a situation where what distinguishes the activation during processing of ambiguous and unambiguous sentences is the time course of activation in various voxels active in both conditions. Another possibility is to forsake creating examples directly from the activation signal and instead do it from linear model activity estimate images obtained by using a GLM to predict each voxel. An example of a given condition will then consist of the pattern of beta-values across voxels for the predictor corresponding to that condition. Finally, the decision of what to use as an example depends not just on what we would like to claim contains enough information to allow a classifier to predict but also on certain constraints described in the next section.

The question of what to use as class labels depends on the purposes of the researcher, but we note that one is not limited to using stimuli classes. One could also use subject responses or decisions in the face of a stimulus (e.g. 'did the subject get this right or wrong?'), or any measurement that can be tied to a particular example. A particularly creative instance of this is the study (Haynes and Rees, 2005), where the prediction is of which of two possible images a subject reports perceiving in a binocular rivalry situation.

In creating examples, there is an important tradeoff associated with the number of examples produced. This is the tradeoff between having many noisy examples (e.g. one per trial) or fewer, cleaner ones (e.g. one of each class per run), as a result of averaging images in the same class. Although there is no hard and fast number of examples necessary to train a classifier, having more is generally better, to the degree that the classifier can see through the noise present (whereas averaging eliminates both noise and natural variability of informative activation in the examples belonging to a class). Having more examples helps in the training set side in that some of the classifiers we consider require a certain number of examples to obtain good estimates of their parameters. In our experience, at least a few tens of examples in each class are needed. If this many are not available, there are types of classifiers that are particularly suitable for a situation with very few examples, and we discuss this at greater length in Choosing a classifier section. From the test set standpoint, having more test examples increases the power of the test for significance of the accuracy obtained. In Training and testing in cross-validation section we will introduce cross-validation, a procedure that makes very efficient use of examples for both training and testing.

Potential pitfalls

Most of the material in subsequent sections relies on a set of assumptions about examples that simplify the analysis of classifier performance. These are: 1) there is a source distribution of examples, from which they can be drawn, 2) these draws are independent (the example just drawn does not determine which will be drawn next)

and 3) training and test sets are independently drawn from this distribution. The creation of example needs to be done in a manner that ensures that these assumptions hold, and it is something that reviewers should be particularly attentive do. If examples are created from blocks or clearly separated trials in a slow event-related design, it's reasonable to assume that they do.

Conversely, using consecutive images or overlapping trials to produce individual training examples should be avoided on the grounds of both 2) and 3), as the strong temporal autocorrelation in the signal will virtually ensure those examples are not independent. A violation of 3) leads to very optimistic accuracy estimates, as a test example will not be a good instance of "any new example that could be drawn" but rather very similar to examples in the training set. A violation of 2) when 3) holds has more subtle effects. On the training set side, it can mean that having several very similar examples will not bring in much new information. On the test set, non-independent test examples cannot be used for the significance tests described in Evaluating results section (the binomial distribution outcome is the sum of the results of *independent* Bernoulli trials). Note that this does not bar us from using consecutive images in the test set, if the purpose is not just to obtain an accuracy estimate (see Applications of pattern classification section for Polyn et al., 2005, a case where training examples come from block data and test examples from fast event-related data).

Another issue to be cognizant of is the desirability of having the same number of examples in each class. When this is not the case a classifier learning algorithm may tend to focus on the most numerous class, to the detriment of the others. This, in turn, can affect the interpretation of an accuracy result (e.g. 80% accuracy may not be very good in a situation where 9 of 10 examples belong to one class and 1 of 10 to the other, if it means the classifier simply predicts the most numerous class by default). When there are different numbers of examples in each class balance can be achieved by using only a subset of the examples in the most numerous class, repeating the procedure for various subsets and averaging the results. Alternatively, one can work with the example imbalance and carefully monitor the prediction accuracy for the various classes.

Reducing the number of features

Given that there are generally many more features than examples, it is often advantageous to reduce the number of features considered to focus on a subset of particular interest; this is called *feature selection*. For instance, if using voxels as features, we may just want the voxels in a particular region of interest (ROI) or their values at a particular point in a trial. The crucial issue to keep in mind is that the choice of features at this stage must not depend on the labels that will be used to train a classifier. To see why, imagine a situation where we are interested in predicting the orientation of a stimulus grating. It is acceptable to restrict the voxels considered to those in visual cortex, demarcated anatomically. It is also acceptable to restrict them further to those showing some activity during task trials relative to a baseline. What is not acceptable is to select voxels that appear to distinguish one orientation from another in the *entire dataset*. The reason for this is that, in effect, it permits information from the test set to affect the learning of the classifier in the training set. As a result, it leads to optimistic accuracy estimates. Looking at the labels for the entire dataset is sometimes called 'peeking'. Note that the injunction against peeking does not mean that the class labels cannot be used at all in feature selection. They can be used once the data have been divided into training and test sets, considering solely the training set. We discuss the practical aspects of this approach further in Feature selection and cross-validation section.

An alternative path to reducing the number of features a classifier has to consider is *dimensionality reduction*. Here, a method such as Singular Value Decomposition/Principal Component Analysis (Hansen, 1999) or Independent Components Analysis (Calhoun et al., 2003)

is applied to the entire dataset matrix (rows are examples, columns are features). The commonality between these methods is that they transform the original feature space (e.g. voxels) into a new, low-dimensional feature space. This yields a new dataset matrix with the same number of rows but a reduced number of columns. This is often worth trying—on account of both ease and possible impact—and is a necessary step for particular classifiers, as we will see in Choosing a classifier section. However, it is not at all guaranteed to improve results, partially because most dimensionality reduction techniques ignore class labels in their criteria (though there are exceptions, e.g. SVDM Pereira and Gordon, 2006 and PLS McIntosh and Lobaugh, 2004 among others).

Example preprocessing

A final issue to consider in the construction of examples is that of preprocessing. By this we do not mean the usual preprocessing of neuroimaging data, e.g. motion correction or detrending, a topic covered in great depth in Strother, (2006). We mean, rather, that done on the examples, considered as a matrix where each row is an example (or each column is a feature). In the example study, we normalized each example (row) to have mean 0 and standard deviation 1. The idea in this case is to reduce the effect of large, image-wide signal changes. Another possibility would be to normalize each feature (column) to have mean 0 and standard deviation 1, either across the entire experiment or within examples coming from the same run. This is worth considering if there is a chance that some voxels will have much wider variation in signal amplitude than others. Although a linear classifier can in principle compensate for this to some degree by scaling the coefficient for each voxel, there are situations where it will not and thus this normalization will help. Our own informal experience suggests either row or column normalization is generally beneficial and should be tried in turn.

Choosing a classifier

Classifier types

Earlier we introduced the idea of a classifier as a function f that takes an example \mathbf{x} and generates a class label prediction $\mathbf{y} = f(\mathbf{x})$. The specific kind of function being learned—and the assumptions built into it—is what distinguishes among the various types of classifier. In this section we review and contrast the classifiers most commonly used with fMRI data, as well as the factors at stake in choosing which kind to use. We supply additional detail in Appendix A.

The simplest classification procedure is called ‘nearest-neighbour’ and it doesn’t even involve explicitly learning a classification function. Classification of a test example is done by finding the training set example that is most similar to it by some measure (e.g. lowest euclidean distance, considering the entire feature vector) and assigning the label of this nearest neighbour to the test example. Variants of this idea include averaging all the training set examples in each class into a single class ‘prototype’ example or assigning the majority label in the k -nearest neighbours (k odd). Both variants minimize the effect of noise (if the examples of different classes are all relatively close to each other) but can be susceptible to destroying some information if there is actual non-noise variability in the examples in each class (e.g. examples in one class come in two clusters that are relatively dissimilar to each other). Whereas nearest-neighbour classification can work very well if there is a small number of features, it tends to fare worse in situations where there are many features and only a few are informative. Hence, it is generally used in conjunction with some form of feature selection, as in Haxby et al., (2001) or Mitchell et al., (2004).

Classifiers that do learn an actual function divide into what are called *discriminative* and *generative* models (Hastie et al., 2001). In the former, the goal is to directly learn to predict from the training data; typically this entails learning a prediction function with a given

parametric form by setting its parameters. In the latter, what is learned is essentially a statistical model that could *generate* an example belonging to a particular class. This model can then be employed to produce a prediction. More formally, what is modelled is the distribution of feature values conditional on example class, $p(\mathbf{x}|\mathbf{y}=A)$ and $p(\mathbf{x}|\mathbf{y}=B)$, which is then inverted via Bayes Rule to classify (i.e. yields $p(\mathbf{y}=A|\mathbf{x})$ and $p(\mathbf{y}=B|\mathbf{x})$, and the prediction is the class for which that probability is largest.

In the typical fMRI study, we generally have many more features than examples; if using voxels in the whole brain as features, it’s not unusual to have a few tens of examples and tens of thousands of features (or at least a few hundreds, if using feature selection as described in Feature selection revisited section). The effect of this is that it will generally be possible to find a function that can classify the examples in the training set well, without this necessarily meaning that it will do well in the test set (for a more detailed explanation of the reason why see Mitchell, 1997): this phenomenon is called *overfitting*. One way of alleviating the danger of overfitting is to choose a simple function from among those that do well in the training set.

The most natural choice for a simple function is to have the prediction depend on a linear combination of the features that can dampen or increase the influence of each one of them, just as linear regression. To make the idea of a linear classifier more concrete, the classifier is parameterized by a set of weights \mathbf{w} and, for an example \mathbf{x} with V features

$$\mathbf{x}\mathbf{w} = x_1 w_1 + \dots + x_V w_V$$

the classifier predicts class A if $\mathbf{x}\mathbf{w} > 0$ or class B if $\mathbf{x}\mathbf{w} < 0$ (with ties broken randomly). One nice property of linear classifiers, beyond their simplicity, is the fact that each feature affects the prediction solely via its weight and without interaction with other features, giving us a measure of its influence on that prediction. We discuss ways to use this, as well as some caveats, in Pattern localization section.

A more geometric intuition for this is given in Fig. 3, in a rather simplified two voxel brain, with three examples of each class (each example is characterized by the value of the two voxels). Learning a linear classifier is equivalent to learning a line that separates points in the two classes as well as possible, the *decision boundary*. Instances of classifiers that learn this are Logistic Regression (LR, known as Multinomial Regression if there are more than 2 classes) and linear

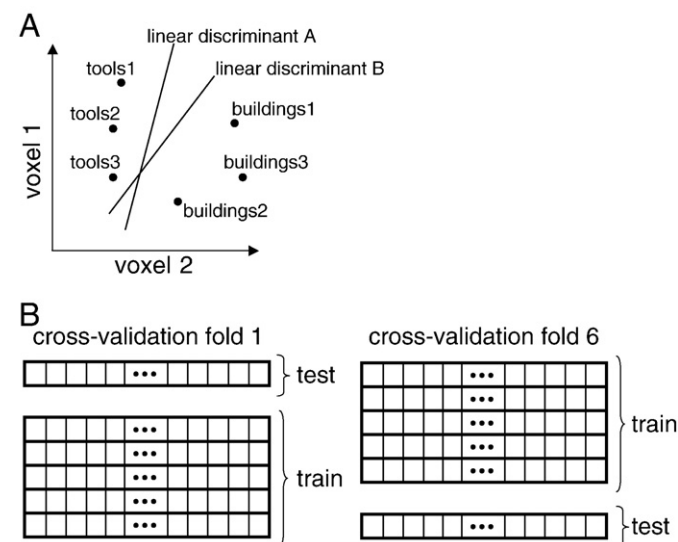


Fig. 3. (A) learning a linear classifier is equivalent to learning a line that separates examples in the two classes (vectors in a 2-voxel brain) as well as possible. (B) during cross-validation each of 6 groups of examples takes a turn as the test set while the rest serve as the training set.

Support Vector Machines (SVM). Whereas this might seem remote from the generative model idea introduced earlier, the classification decisions of Gaussian Naive Bayes (GNB) and Fisher's Linear Discriminant Analysis (LDA) can be expressed in this manner as well, and thus properly included in the linear classifier category.

Note that there are many possible linear discriminants that perfectly classify the six training examples shown in Fig. 3. The various classifiers have different rationales for choosing among these, which correspond to different definitions of 'separating points as well as possible'. Consider also that, even for a specific type of classifier, there may still be many possible settings of its parameters that lead to equally good predictions in the training set. The process of guiding the procedure that sets the parameters to a solution with desirable properties—small weights overall, for instance—is called *regularization*. In the classifiers we consider this takes the shape of an extra parameter that must be tuned also using the training set. Both the rationales and the way regularization is done for the various kinds of classifiers are described in Appendix A.

Classifier choice in practice

All this brings us to the question of which classifier to use, which we will attempt to answer based on our experience and a multi-study comparison of classifiers and feature selection methods (Pereira, 2007). In cases where there are a large number of features (e.g. all the voxels in the cortex), GNB is generally inferior to LR and linear SVM, as the regularization in the latter clearly help weigh down the effect of noisy features that are highly correlated with each other. The disadvantage mostly disappears if using feature selection, hence GNB is a useful classifier for procedures that need to be repeated many times, such as permutation tests, due to being much faster to train than the others. LDA needs to be used either in conjunction with extreme feature selection or with dimensionality reduction (Strother et al., 2002), as otherwise there are typically not enough examples to estimate its covariance matrix reliably.

LR and linear SVM tend to be roughly equivalent in terms of performance, other than for the fact that the first lends itself more naturally to cases where there are more than two classes. So far, and for the sake of simplicity, we have discussed classifiers in a situation where examples belong to one of two classes but, in general, we might have more (several different stimulus classes, subject responses, etc). Most of the classifiers described above work without modification in a multiple class scenario. The one exception is SVM; although multiclass versions do exist, most packages in current use require one to create several two class problems (e.g. each class versus all the others, or all pairs of classes) and combine them to produce a multiway classification decision. We describe some approaches to doing this in Appendix A. That said, often the object of interest in a multiple class situation is not so much the raw accuracy (do we believe that hundreds of classes corresponding to hundreds of object stimuli would be distinguishable based on fMRI activation, for instance?) as which classes can be distinguished from which others, and how. Hence, it might make more sense to consider all pairs of classes, train and test a classifier for each pairwise distinction and produce what is called a *confusion matrix*, a $\#classes \times \#classes$ matrix where entry (i,j) contains the accuracy of the distinction between classes i and j . Often this reveals groups of classes that are hard to distinguish from each other but that are distinguishable from another group (e.g. several small object stimuli versus animal and human faces).

There are more complex classifiers, such as nonlinear SVMs or artificial neural networks, which can let interactions between features and nonlinear functions thereof drive the prediction. In terms of Fig. 3 this corresponds to being able to have a nonlinear decision surface. However, we will not discuss such nonlinear classifiers at length here for two main reasons. The first is that it is not clear that they always provide a significant advantage in practical performance relative to

linear classifiers (in extant studies or the multi-study comparison in Pereira, 2007). In our opinion this is more a reflection of the fact that the number of examples available is so small than of the absence of complicated relationships between features. The second reason is that the relationship between features and the prediction can become harder to interpret when these classifiers are used (though see Hanson et al., 2004, which derives a class taxonomy from the hidden layer node activity in a neural network, for instance). This doesn't mean that they cannot be used fruitfully when the number of features is smaller, in particular when doing information-based mapping as discussed in Pattern characterization section. We describe these classifiers and their relationship to linear classifiers in more detail in Appendix A.

Training and testing in cross-validation

Cross-validation

In our illustrative example we divided the dataset into halves, one used for training and the other for testing. As we later discussed, we would like to train a classifier with as much data as possible. Sadly, we cannot train and test on the same data if the goal is obtain a useful estimate of the classifier's true accuracy, i.e. the probability that it would label any future test example correctly.

Furthermore, using just a few examples for testing will not lead to a good estimate: it will be inordinately variable, as well as constrained to a few possible values. Splitting the dataset in half to have more test examples, as we did in the example study, means we might be learning a much worse classifier than we potentially could with more training examples.

Fortunately, there is a procedure that allows us to have our cake and eat (almost) all of it too. The procedure is called *cross-validation*. In its most extreme variant, cross-validation follows these steps:

- leave one example out, train on the remaining ones, make a prediction for this example
- repeat for each example in turn
- compute the accuracy of the predictions made for all the examples.

This variant is called 'leave-one-out' cross-validation. Although each classifier trained is technically different, one can expect them to be similar and predict similarly, since they share so much training data. Hence we treat the accuracy estimate obtained in this manner as the expected accuracy of a classifier trained on a dataset with all but one example, following Kohavi, (1995). Note that this also provides a slightly conservative estimate of the accuracy we expect if we were to train a classifier using all the available examples.

In practice, leaving each example out can be computationally expensive because instead of a single classifier we are training as many classifiers as there are examples. One compromise is a procedure called a *k-fold* cross-validation, where k is the number of parts into which the dataset is divided; common choices are $k=10$ or $k=5$, corresponding to leaving out 10% or 20% of the examples on each fold. Fortunately, fMRI datasets are often naturally divided in a way that addresses all these issues, be it in runs or epochs within a run where all the possible classes or stimuli appear. Our illustrative dataset divides naturally into six groups, corresponding to the 6 epochs of stimuli in the experiment. Each group is used in turn as the test set in cross-validation, with the remaining groups used as the training set (on the right of Fig. 3). Hence we advise the reader to take advantage of those natural divisions at least on a first attempt, possibly moving to a finer grain division that still has the desired properties if more training examples are needed.

Aside from the number of examples left out there are other important considerations. The first is that the training data in each fold must contain examples of all classes, as otherwise the classifier for that fold will not be able to predict the absent ones. As mentioned earlier, classes should be balanced, i.e. have roughly the same number

of examples. Furthermore, examples that are correlated—by being close in time, for instance—should end up in the same fold. Otherwise, the classifier may be able to predict accurately for test examples with a correlated counterpart in the training set.

Feature selection revisited

In Preparing data for classifier training section, we introduced the idea of feature selection, in situations where we have some *a priori* reason for picking features, e.g. wishing to test whether voxels in a particular ROI have information. We return to the topic now in order to discuss together with the stage of analysis where it most likely will take place.

Feature selection is often deemed necessary in order to train classifiers in domains where datasets have very large numbers of features. Why might it help? Intuitively, the goal is to reduce the ratio of features to examples, decreasing the chance of overfitting, as well as to get rid of uninformative features to let the classifier focus on informative ones. Almost all the fMRI decoding results reported in the literature used either feature selection or some other form of dimensionality reduction (Norman et al., 2006; Haynes and Rees, 2006; Carlson et al., 2003; Strother et al., 2002). Feature selection in general is too vast a topic to broach in great detail here, but a useful review of both terminology and the thrust of various approaches is given in Guyon and Elisseeff, (2003).

One distinction often made is between *scoring/filtering* and *wrapper* methods. The former involves ranking the features by a given criterion—each feature is scored by itself, a bit like a univariate test of a voxel—and selecting the best in the ranking. The latter consists broadly of picking new features by how much impact they have on the classifier *given the features already selected* (or, in reverse, considering all the features to begin with and removing features while performance increases). This can be done by repeatedly training and applying the classifier in cross-validation within the training set. Given the often prohibitive computational expense of wrapper methods, due to the large number of features, this is in general unfeasible; there are, however, heuristic techniques that have been successfully used with fMRI data, such as Recursive Feature Elimination (Hanson and Halchenko, 2008). Given that feature selection is simpler to deploy initially and the approach used by the majority of extant papers, we focus primarily on that.

There are several generic methods for selecting informative features. In parallel, features in the fMRI domain are often voxels and there is a desire to be able to look for certain types of voxel behaviour, e.g. is a voxel very selective for one particular condition or discriminating between two groups of conditions. Hence most extant papers resort to methods that draw from both these sources, which we summarize below. Note that method rationales are described in terms of voxels and conditions (classes), for clarity, but could be adapted for other kinds of features:

- **Activity**—this method selects voxels that are active in at least one condition relative to a control-task baseline, scoring a voxel as measured by a *t*-test on the difference in mean activity level between condition and baseline.
- **Accuracy**—this method scores a voxel by how accurately a Gaussian Bayesian classifier can predict the condition of each example in the training set, based only on the voxel. It does a cross-validation within the training set for the voxel and the accuracy is the voxel's score.
- **Searchlight accuracy**—the same as Accuracy, but instead of using the data from a single voxel (GNB classifier with one feature) we use the data from the voxel and its immediately adjacent neighbours in three dimensions (a classifier with as many features as neighbours, e.g. 27 in the radius 1 neighbourhood, can be GNB, LDA or SVM).

- **ANOVA**—this method looks for voxels where there are reliable differences in mean value across conditions (e.g. A is different from B C D, or AB from CD, etc), as measured by an ANOVA.
- **Stability**—this method picks voxels that react to the various conditions *consistently* across cross-validation groups in the training set (e.g. if the mean value in A is less than B and greater than C, is this repeated in all groups).

Precise details of how these methods can be implemented and other alternatives are provided in Appendix A. Note that the methods above are, with the exception of searchlight accuracy, univariate. There are promising results that show that, in some circumstances, multivariate feature selection can work better, which are discussed in Applications of pattern classification section. Finally, the methods listed above provide rankings of voxels; one might still have to decide on how many of the top-ranked voxels to use from the ranking for each method. For methods that are based on a statistical criterion (e.g. *t*-tests of activation), it is common to use a multiple comparison criterion such as Bonferroni or False Discovery Rate (Genovese et al., 2002) and keep just the features deemed significant at a given level. This is reasonable but problematic if few or no voxels are significant (possible given the amount of noise and very small sample sizes) or the score produced by the method does not map to a *p*-value (though nonparametric methods can still be used to obtain one, see Appendix A).

Feature selection and cross-validation

Up till now we have discussed feature selection as a topic by itself. Often, however, feature selection is used in conjunction with cross-validation and this entails particular precautions. The first is that selection must happen *on the training set*. As mentioned earlier, for every fold in cross-validation we have a different training set and thus features have to be selected *separately* for each fold. A second reason is that, as mentioned earlier, the methods used in practice *rank* features and we still have to decide how many to use, unless the method provides a rationale (e.g. the number of voxels where a *t*-test is significant at a certain level). A typical approach is to not commit to a specific number, but to instead produce a feature ranking by each method for each fold. This can then be used to obtain, for each method, an accuracy score using each of the number of voxels desired (e.g. top 20, top 50, top 100, etc). The end product is a `number-of-methods × numbers-of-voxels` table of accuracy results, such as Table 1 (left) for our sample dataset.

If feature selection could be done prior to cross-validation (e.g. by selecting voxels from one ROI, or by some other method that avoids examining the labels of the test examples), the whole procedure can be run as if the dataset contained only the selected features. In this and the case where the method determines the number of features, the end product is a `number-of-methods × 1` vector of accuracy results.

An alternative to considering all numbers of voxels is to let the training set in each cross-validation suggest which number to pick. To do this, we can run a *nested cross-validation* (NCV) inside the training set. Consider our illustrative example, which had 6 groups of examples based on epochs. During the first cross-validation fold, we tested on group 1 and trained on groups 2–5. Using a nested cross-validation means performing the cross-validation procedure described above

Table 1

Classification results using each of several selection methods with either various fixed numbers of voxels (left) or with the number determined using nested cross-validation (NCV) inside the training set of each fold (right)

Method	Number of voxels						Method	NCV
	100	200	400	800	1000	All		
Accuracy	0.81	0.81	0.75	0.73	0.74	0.65	Accuracy	0.81
Searchlight	0.81	0.82	0.82	0.77	0.79	0.65	Searchlight	0.82
Activity	0.79	0.80	0.77	0.73	0.74	0.65	Activity	0.77
ANOVA	0.77	0.75	0.75	0.73	0.71	0.65	ANOVA	0.77

inside groups 2–5 by testing on 2 and training on 3–5, then testing on 3 while training on 2, 4, 5, 6, etc. The result for this training set is, again, `number-of-methods × number-of-voxels` table of results. For each method, this table will yield the number of voxels at which the accuracy is highest; this is the number that will be selected from the ranking produced from groups 2–5. The end product is a number of `methods × 1` vector of accuracy results, as shown in Table 1 (right) for our sample dataset.

Evaluating results

The person training a classifier on fMRI data is concerned with establishing that a variable of interest can be decoded from it, i.e. a classifier can be trained whose *true accuracy* is better than that of a classifier deciding at random. Formally, the “true accuracy” is the probability that the classifier will correctly label a new example drawn at random from the same distribution that the training examples came from; it can also be viewed as the accuracy one would get if one had an infinite number of examples in the test set. The accuracy on the test set is thus an *estimate* of the true accuracy of the classifier used. It is an unbiased estimate because the test examples have not been seen by the classifier during training. How precise an estimate it is depends on the size of the test set; the fewer examples used the greater the variability of the estimator, as we shall see below. This section is concerned with two major issues: determining how well a classifier performs (and whether it is better than random) and drawing conclusions from multiple classifiers trained and tested on the same data.

Result significance

A *statistically significant* classification result is one where we can reject the null hypothesis that there is no information about the variable of interest (the variable being predicted by the classifier) in the fMRI data from which it is being predicted. Establishing statistical significance is generally done by determining how improbable the observed classification accuracy would be were the null hypothesis to be true (this probability is called a *p-value*).

To illustrate, suppose we have a problem in which the classifier must output one of two classes that are equally frequent in the data, and that the classifier succeeds in correctly classifying k of the n test examples it is given. To determine whether this experimentally observed accuracy is statistically significant we must determine how likely it is that we would observe at least k successes out of n test trials if the classifier were in fact operating at chance (i.e., if the null hypothesis were satisfied). Clearly, if the null hypothesis were satisfied and the classifier could only guess labels at random, and if the test examples were drawn independently, then we would expect an accuracy of $k/n = 0.50$ on average; the actual observed accuracy might vary above or below 0.5 due to variations in the sample of test data we happen to be using. The point of determining the statistical significance is precisely to determine the probability that this experimentally observed accuracy might vary to a value at least as high as k/n even when the null hypothesis is satisfied. Note that the observed accuracy might also reach values substantially below 0.5 under the null hypothesis, again due to variation in the test sample, but we are not interested in the chance of coincidental deviations below the expected 0.5 because we would not try to claim our classifier was succeeding in that case.

In this situation the probability of obtaining a given classification result under this null hypothesis is easy to calculate. It can be viewed as the outcome of tossing a coin for each of the test examples to be classified, with the coin's bias reflecting the probability of labelling an example correctly by chance (50% in a two class problem, 25% in a four class problem, etc). More formally, each coin toss can be modelled as a Bernoulli trial with probability p of success. The probability of achieving k successes out of n independent trials is given by the binomial distribution. If we

define k to be the number of correctly labeled test set examples out of n , the p -value under the null hypothesis is simply $P(X \geq k)$, where X is a random variable with a binomial distribution with n trials and probability of success 0.5 (two class), 0.25 (four class), etc. If the p -value is below a certain threshold (typically 0.05 or 0.01) the result will be declared significant.

Note also that, thus far, we have been talking about a single classifier being applied to a test set. The accuracy estimate obtained through cross-validation is a combination of results produced by classifiers trained on mostly overlapping training sets, hence we will treat that estimate as if it came from a single classifier, following common practice (Kohavi, 1995), which then allows us to test it in exactly the same way.

The binomial test is exact (if the assumptions about example independence hold), and hence it is preferable to alternatives such as a χ^2 test (which would have the additional issues of being unreliable in the presence of only tens of examples, typical for our datasets). If the assumptions do not hold or there is some suspicion of optimistic bias, a viable alternative to testing result significance is a permutation test. Assuming there is no class information in the data, the labels can be permuted without altering the expected accuracy (chance level). In practice, this entails shuffling the labels within each training set, training a classifier and applying it to the test set. The procedure is then repeated many times, with a different shuffling each time. Over many repetitions, this yields a sample of accuracy results under the null hypothesis.

The p -value in this case is the fraction of the sample that is greater than or equal to the accuracy actually observed when using the correct labels. A good review of permutation tests in classification (with one functional neuroimaging application) is Golland and Fischl, (2003), and we provide additional details in Appendix A. An issue when doing this in combination with cross-validation is the need to keep track of fold structure when permuting example labels, as it is possible to have training sets where a given class is not represented and others where it is overrepresented. A practical solution is to do a stratified test, and permute within the partitions of the dataset used to build up the fold training sets.

Result confidence intervals

Instead of testing for significance one may want to obtain a confidence interval for the true accuracy of the classifier. Whereas this is an interval for the parameter of a binomial distribution, an approach often followed is to simply use a normal distribution confidence interval (also known as a Wald confidence interval). This is generally justified asymptotically and could be inappropriate simply on the ground of having too small a dataset, as judged by various textbook criteria. Brown et al., (1999) considers this issue and explains that, even in the situations where textbooks say the Wald interval can be used, it often fails to have the coverage probability desired. Hence there is little justification to use it and one should resort either to the exact interval (Langford, 2005 provides a good explanation of how to compute it numerically) or one of several closed-form approximations proposed in Brown et al., (1999).

Testing multiple results

Extending the significance test or confidence interval to multiple, independent results is straightforward, in that one can simply apply a multiple comparison correction criterion (e.g. Bonferroni) and adjust the significance threshold accordingly. This is the situation we have when we consider the nested cross-validation (NCV) results, as each one is the product of training a classifier using voxels selected with a different method. If, instead, we are just producing a `number-of-methods × number-of-voxels` table of results—such as Table 1—we could use the same approach to simultaneously test all the results in the table. This would be conservative, however, in that the results for a given method are correlated to some extent; this follows from the fact

that the half the voxels used by the classifier based on the top 100 voxels are also used by the one based on the top 50, say. An easy mistake to make in dealing with this issue is to select the maximum result among all those obtained using a particular selection method (the maximum of a table row); it is often accompanied by the mistake of testing that maximum result for significance in the way described in the previous section.

To see why this is a problem, consider a situation where the classifiers are using different numbers of voxels and, in each case, have a true accuracy (over the underlying example distribution) of 0.5. It is possible even in this situation that, at least for one entry in the table, the experimentally observed accuracy will be fairly high by chance; this becomes more likely the smaller the dataset is. If we are just deeming the 'best' number of voxels to be the one at which the maximum accuracy is attained, we would pick this result and possibly deem it significant. One possibility for addressing this issue is to report the entire table and the number of examples used and correct for all of them, at the risk of being conservative. The other possibility is to report the maximum accuracy result anyway but test it differently using a permutation test. The idea is exactly the same as the one used for testing a single result above, but used instead to obtain a distribution of the maximum accuracy result over label permutations. For a given permutation the entire result table could be generated and the maximum obtained. Over many permutations this yields a sample of accuracy results for each cell in the table under the null hypothesis but also, more importantly, for any derived measures such as the maximum of each table row. The p -value under the null hypothesis for each entry of the original result table is the fraction of the corresponding sample that is greater than or equal to it, and this applies to the maximum as well. The main problem of resorting to this approach is its computational cost, though it is generally the most principled possibility in situations where it is unclear whether the result reporting procedure is optimistically biased.

Comparing classification results

If our goal is to compare two classifiers based on results obtained on the same test data, we will have two samples of Bernoulli variables, one per classifier, where each variable takes the value 1 if the corresponding example was classified correctly and 0 if not. The accuracy estimate for each classifier is the mean of its respective sample of Bernoulli variables. The statement we want to make, however, is not about the difference in accuracy estimates but about the difference in true accuracy of those classifiers.

The naive way to address this problem would be to treat it as a two sample t -test. The null hypothesis would then be that the means of the distributions the two samples came from – the classifiers' true accuracies – were the same. The first reason this is inappropriate is that the two samples are not independent; moreover, they are paired and the two classifiers can be viewed as repeated measures for each example. This suggests a second possibility, a paired t -test of the sample of differences between Bernoulli variables for each example. This is feasible as long as the number of examples is enough to invoke the Central Limit Theorem but still not ideal, given that the differences in scores can only take three values (-1 , 0 or 1). A more appropriate solution is the Wilcoxon signed-rank test, a nonparametric version of the paired t -test. An excellent reference for testing in repeated measure situations (and testing in general) is Lowry, (1999).

Applications of pattern classification

We have been treating a classifier-based analysis as a sequence of stages with choices at each stage. One choice that we have not yet directly addressed is perhaps the most important: the choice of the initial scientific question to be answered. Our focus so far has been on the use of classifiers for *pattern discrimination*. Here, the basic question being addressed is simply whether fMRI data carry information about a

given variable of interest. Determining the answer to this question is of practical interest beyond answering neuroscience questions, for instance in clinical diagnosis or lie detection (Davatzikos et al., 2005). We can view traditional univariate methods as answering this question on a voxel by voxel or ROI by ROI basis. Classifiers bring to the question an increase in sensitivity of detection, both by pooling information across many voxels but also by relaxing constraints on spatial contiguity or the need for voxel responses to be similar.

In addition to the question of whether fMRI data carry class information (pattern discrimination), there are at least two other basic questions to which classifiers can be applied, namely 'where or when is class information represented' (pattern localization) and 'how is class information encoded' or 'how does its encoding relate to known relationships between stimuli' (pattern characterization). This section examines the extent to which these questions can be tackled using the machinery of linear classifiers and feature selection, as well as the limitations and caveats. For the rest of the section we will assume features are voxels, as it will simplify discussion.

Pattern localization

Once it has been established that class information is present in a dataset, one may go on to ask *where* in the brain this information resides. A natural approach for addressing this is to ask which voxels contribute most strongly and reliably to the classifier's success. There are two aspects to this: determining which voxels are being selected and also how their weight affects the classifier prediction.

Which voxels are being selected

When cross-validation is being used, a particular voxel selection method can and very likely will rank different voxels at the top from fold to fold. It may come as a surprise that, for many datasets, the overlap between the set of voxels selected in all folds versus the set selected in an individual fold is fairly small (between $\frac{1}{10}$ and $\frac{1}{3}$ is typical) even if selecting a relatively large fraction of the total number of voxels (Pereira, 2007). This is not a problem, intrinsically, it just reflects the fact that whatever voxels are selected contain *sufficient* information. There might be redundancy, in that many voxels encode the same things and a few get picked in each fold, and there might be a small *necessary* subset, present in all folds. One possible strategy for finding the maximal subset of voxels containing information would be to start from the ones present in all or almost all folds and use them as seeds in a clustering algorithm that would find voxels with similar response profiles across examples.

A related issue is whether the information carried by voxels is the same in the sets identified by various selection methods. There are two ways of considering this issue: one is to look at overlap in selected sets, now between methods, and the other to consider whether a given classifier using different voxel sets performs differently (in particular, whether it makes different prediction errors). For problems with two classes, most selection criteria will find similar voxels, as there aren't many ways that a voxel could respond differently to the two classes. In multiclass problems results suggest there are different types of voxel behaviour and the various voxel selection methods described earlier can pick on different types (Pereira, 2007). Hence, if we wish to use classifier prediction success as an indication that the voxel set it uses contains informative voxels, it is wise to consider different selection methods; it is also a good idea to determine whether there is any overlap between the voxel sets selected in all folds by each method.

Which voxels drive the classifier

Given a set of voxels, a different question is which of those voxels most strongly affect classifier predictions. The use of linear classifiers makes it relatively simple to answer this question, as each voxel affects the decision only through the magnitude of its weight, and not through the interaction with other voxels. Appendix A explains how to

get these weights for the various linear classifiers, as well as for nearest neighbour classifiers with particular distance measures (although in this case the weights are test example specific, rather than the same for any example), and also shows what a typical image of voxel weights looks like. Note that you can take this further by multiplying the weights by the average example of each class, giving you class-specific images.

A practical issue if using cross-validation is that, again, we have a different set of voxels and thus a different classifier in each fold. A heuristic approach is to group all the voxels that were selected in at least one of the folds and run the cross-validation again on that set; if the accuracy is comparable to what was obtained with the original cross-validation, the classifier weights can be averaged across folds and considered. Alternatively, if the accuracy using all the voxels available is high—more common in problems with few classes—we can just consider the weights over all those voxels.

One point to be aware of concerning linear classifiers is that the regularization for LR or linear SVMs will cause the weights on correlated voxels to shrink proportionally to the number of correlated voxels. This is beneficial when the correlation is mostly driven by a noise source and is the reason why, in general, both types of classifier outperform GNB if all voxels are used (though not if coupled with voxel selection) (Pereira, 2007). However, it is quite possible that there will be voxels that contribute to the decision but whose weights are smaller than those of other voxels contributing voxels.

Note that, in addition to asking *where* class information is present, we can also ask *when* classes are represented—temporal as opposed to spatial localization. This is discussed less commonly in the literature

and appears in two different guises. The first is to either select a feature space corresponding to voxels at all time points in a trial or simply restrict the time points data is extracted from, as in Mitchell et al., (2004) or Mourao-Miranda et al., (2007). A more sophisticated approach is used in Polyn et al., (2005), where classifiers are trained on data where classes are cleanly separated between blocks and then deployed as ‘sensors’ of the emergence of class-specific patterns in event-related test data.

From global to local classifiers

The issues outlined make localization of information to specific voxels less straightforward than one might wish and, furthermore, the use of linear classifiers means that no nonlinear relationships between voxels can be learned. There is, however, an alternative approach to answering the localization question, ‘information-based functional brain mapping’ (Kriegeskorte et al., 2006), which can ameliorate these problems and also address characterization questions to some extent.

A general version of this idea can be summarized as training classifiers on many small voxel sets which, put together, cover the whole brain. One natural candidate for this is to train a distinct classifier for each voxel, using only the voxels spatially adjacent to it. This is akin to shining a ‘searchlight’ on every voxel neighbourhood in succession, and is often referred to in the fMRI community as training ‘searchlight classifiers’; the idea of training local classifiers—and possibly combining them into a global one—has appeared under various other guises in the machine learning literature (see, for instance, RBF classifiers Schoelkopf et al., 1997). Common variants for

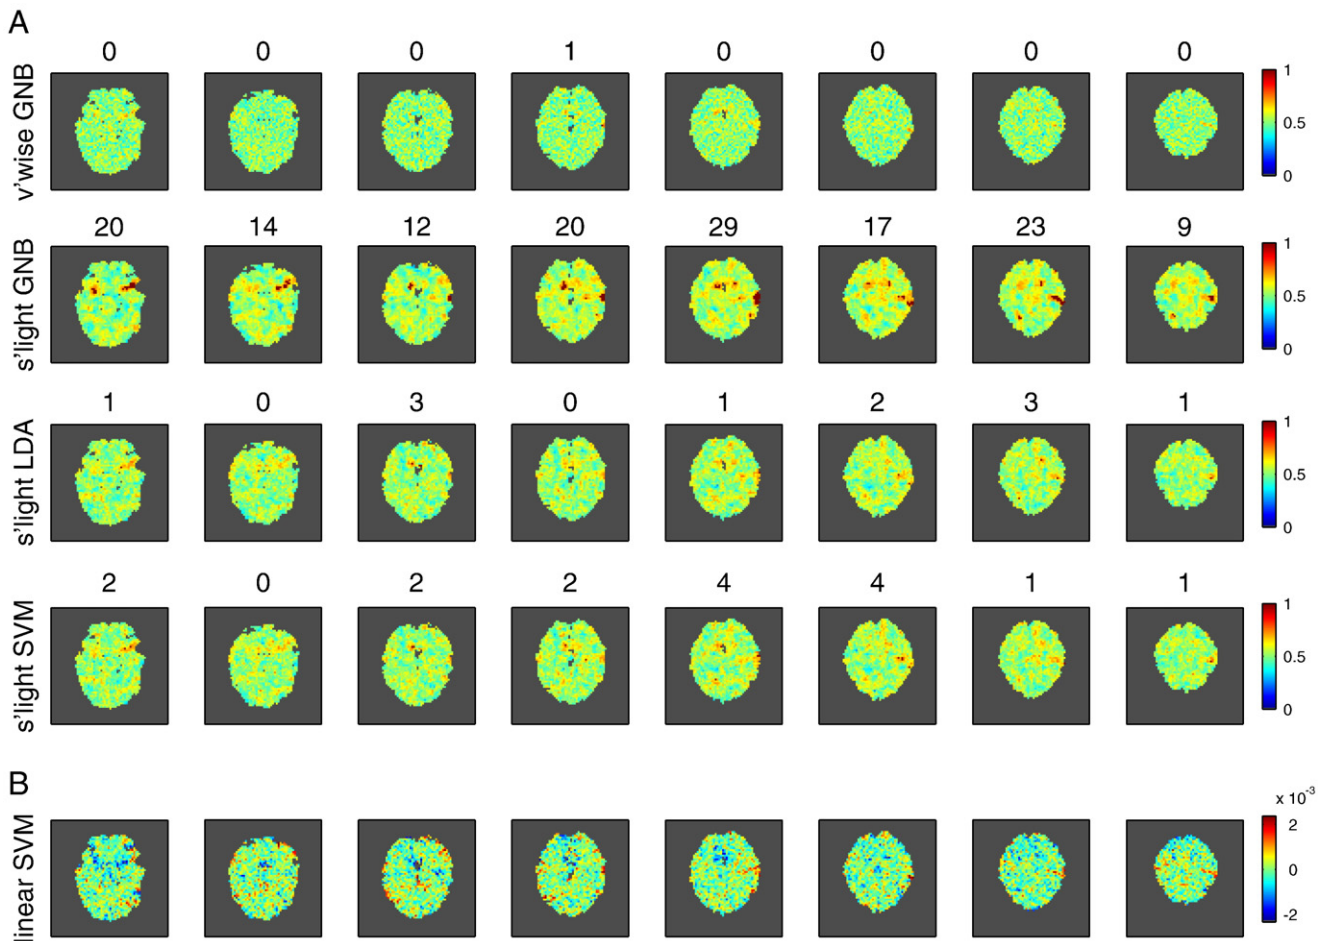


Fig. 4. Top: comparison of accuracy maps obtained with 6-fold cross-validation. Each row contains eight slices from an accuracy map, inferior to superior, top of each slice is posterior. The classifiers used in each row are single voxel GNB and radius 1 neighbourhood searchlight GNB, LDA and linear SVM. On top of each slice is the number of voxels for which accuracy was deemed significant using FDR $q=0.01$, which are highlighted in dark red. Bottom: classifier weights for a linear SVM classifier trained on the entire image.

the classifier are versions of Gaussian classifiers (GNB, LDA) or linear SVM; the original paper used the Mahalanobis distance, which is closely related to the multivariate Gaussian distribution used in LDA. The result is a brain map where, for every voxel, we have the cross-validation accuracy of a classifier trained on the neighbourhood surrounding it (typically, the 26 neighbours in 3D, although a wider radius could be considered). This accuracy can be converted into a p -value using either the analytical or permutation methods described in Evaluating results section. In either case, we will obtain a p -value map that can be thresholded using a multiple comparison correction criterion, to identify voxel neighbourhoods of significant accuracy.

We again have a choice of which classifier to use, although in this case the tradeoffs are different. Given the relatively small number of voxels in a neighbourhood, it is now feasible to train more complex classifiers. LDA is the first obvious choice, in that a covariance matrix of voxels captures relationships between them, but one could also attempt a nonlinear classifier if more complex relationships are hypothesized. This will possibly be more advantageous if there are more than two classes, although with the caveat that it is unlikely that activation distinguishing *all* the classes from each other can be found in a small region.

We provide a comparison of single voxel GNB and searchlight GNB, LDA and linear SVM in Fig. 4 for our illustrative study. Whereas there are reassuring similarities in the location of voxels deemed significantly informative, there are also interesting differences between the maps which we discuss in Appendix A. For reference, the figure also provides the weights of a linear SVM classifier on each voxel.

Pattern characterization

Beyond asking whether and where class information is represented, classifiers can also be used to answer the question of *how* classes are represented in the brain. It is perhaps here that classifiers and other machine learning methods have the greatest untapped potential.

The work on information mapping in the previous section can also be considered from the point of view of pattern characterization. Each neighbourhood classifier attempts to find relationships between the voxels in a small set that contain information. As said earlier, given the small number of voxels, we have the possibility of training nonlinear classifiers to learn fairly complex relationships. But just how complex should one expect these relationships to be? One extreme in the range of answers is the situation in Kamitani and Tong, (2005), where voxels are barely more sensitive to one class of stimulus than another, and thus do not reflect an encoding of those stimuli. There is no advantage in training a classifier that does more than weigh a 'vote' of each voxel to pool information. The other extreme is what Hanson et al., (2004) calls a combinatorial code: the information is in the complex combination of voxel activities in different classes. In between lie situations where voxels have very selective responses (e.g. FFA or PPA in Kanwisher, 2003) or voxels are selective for groups of classes (e.g. semantic category classes that correspond to animate things versus those that do not).

Thus far, we chose to describe the work done resorting to linear classifiers or information mapping, both because it is simple and a good starting point to address characterization questions. There are, however, several recent papers that tackle this kind of question in more sophisticated ways that elude a unified view, and hence we will cover them only briefly. The common points between them are the wish to correspond relationships between stimuli with the relationships between activation patterns in fMRI data elicited by those stimuli, as captured by classifiers. In the first camp, Hanson et al., (2004) shows that a neural network trained to distinguish examples of eight semantic category classes learns a representation on its hidden layer that captures a plausible taxonomy of those classes. O'Toole et al., (2005), on the same dataset, shows that the ability to distinguish the

category stimuli correlates with the ability of the classifier to distinguish the corresponding classes. Botvinick and Bylsma, (2005) shows that the similarity matrix of phonetic stimuli corresponds to the similarity matrix of patterns of activation while processing them and Aguirre, (2007) relates perceptual similarity of stimuli to the structure of their neural representation. Hutchinson et al., (2006) introduces hidden process models, a technique to identify multiple cognitive processes present simultaneously in the brain with overlapping spatial activation, taking advantage of domain information and class labels to constrain their causal ordering.

Finally, two recent studies leverage domain knowledge to predict fMRI activation for new stimuli, something which can be used to classify by comparing the predicted activation to the activation of stimuli in the training set. Kay et al., (2008) does this for visual stimuli, using a model of the responses of neuron populations in the visual cortex and predicting the fMRI signal from the model populations for each voxel. Mitchell et al., (2008) learns a model that predicts fMRI activation for new words, based on their similarity in a space of semantic features to words for which fMRI activation is known.

Conclusions

In this paper we have described the various stages in a machine learning classifier analysis of fMRI data. Aside from discussing the choices available at each analysis stage, their interactions and the practical factors conditioning them, we explored the use of this kind of analysis to answer three types of scientific question. These are 'is there information about a variable of interest' (pattern discrimination), 'where is the information' (pattern localization) and 'how is that information encoded' (pattern characterization). We purposefully focused on linear classifiers and the discrimination and localization questions, as these are both better understood and the ones most readers are likely to want to answer first. Space constraints meant that some important practical considerations had to be left out of the main text. Various appendices are provided in a longer version of this paper available at <http://www.princeton.edu/~matthewb> or <http://www.cs.cmu.edu/~fpereira>.

MATLAB software for performing a classifier analysis (the Princeton Multi-Voxel Pattern Analysis toolbox) can be found at <http://www.cs.bmb.princeton.edu/mvpa>, together with tutorials and a very responsive online community. Alternatively, there is a separately developed Python version of the toolbox (PyMVPA) available at http://www.py_mvpa.org/.

What we have tried to do in the present article is to lay out a roadmap, identifying the key decision points involved in classifier-based imaging analyses, and tracing out the available routes among them. Like any map, this one will require continual updating, as the terrain itself evolves. If the past is any model, then frequent and often dramatic changes in classifier-based techniques can be expected over coming years, especially in nascent areas such as pattern characterization. Regardless of the changes to come, however, it is clear that classifier-based approaches have already evolved to the point where it is possible to begin establishing a principled consensus concerning basic aspects of methodology, similar to the one that has coalesced around longer-established univariate analysis techniques.

Conflicts of interest

The authors declare that there are no conflicts of interest. The funding agencies listed had no involvement in the writing of this report or its submission to publication.

Acknowledgments

Tom Mitchell was supported by the W.M. Keck Foundation. Matthew Botvinick and Francisco Pereira were supported by the National Institute of Neurological Disease and Stroke (NINDS), grant

number NS053366. We are very grateful to the anonymous reviewers for their valuable feedback and to Greg Detre, Niko Kriegeskorte and Marieke Mur for enlightening discussions and helpful comments.

Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, at doi:10.1016/j.neuroimage.2008.11.007.

References

- Aguirre, G.K., 2007. Continuous carry-over designs for fmri. *NeuroImage* 35, 1480–1494.
- Botvinick, M.M., Bylsma, L.M., 2005. Imaging phonological representations and their similarity structure with pattern-analytic fmri. *NeuroImage* 26 (suppl. 1).
- Brown, Lawrence D., Cai, T. Tony, Dasgupta, Anirban, 1999. Interval estimation for a binomial proportion. *Stat. Sci.* 16, 101–133.
- Calhoun, V., Adali, T., Hansen, L., Larsen, J., Pekar, J., 2003. Ica of functional mri data: an overview. 4th International Symposium on Independent Component Analysis and Blind Signal Separation (ICA 2003).
- Carlson, T.A., Schrater, P., He, S., 2003. Patterns of activity in the categorical representations of objects. *J. Cogn. Neurosci.* 5 (15), 704–717.
- Davatzikos, C., Ruparel, K., Fan, Y., Shen, D.G., Acharyya, M., Loughhead, J.W., Gur, R.C., Langleben, D.D., 2005. Classifying spatial patterns of brain activity with machine learning methods: application to lie detection. *NeuroImage* 28 (3), 663–668.
- Friston, K.J., Ashburner, J., Kiebel, S.J., Nichols, T.E., Penny, W.D. (Eds.), 2007. *Statistical Parametric Mapping: The Analysis of Functional Brain Images*. Academic Press.
- Genovese, C.R., Lazar, N.A., Nichols, T.E., 2002. Thresholding of statistical maps in functional neuroimaging using the false discovery rate. *NeuroImage* 15, 870–878.
- Golland, P., Fischl, B., 2003. Permutation tests for classification: Towards statistical significance in image-based studies. *LNCS. Proceedings of IPMI: International Conference on Information Processing and Medical Imaging*, 2732. Springer, pp. 330–341.
- Guyon, I., Elisseeff, A., 2003. An introduction to variable and feature selection. *J. Mach. Learn. Res.* 3, 1157–1182.
- Hansen, L.K., et al., 1999. Generalizable patterns in neuroimaging: how many principal components? *NeuroImage* 9, 534–544.
- Hanson, S.J., Halchenko, Y.O., 2008. Brain reading using full brain support vector machines for object recognition: there is no face identification area. *Neural Comput.* 20 (2), 486–503 February.
- Hanson, S.J., Matsuka, T., Haxby, J.V., 2004. Combinatorial codes in ventral temporal lobe for object recognition: Haxby(2001) revisited: is there a face area? *NeuroImage* 23.
- Hastie, T., Tibshirani, R., Friedman, J.H., 2001. *The Elements of Statistical Learning: Data Mining, Inference and Prediction*. Springer-Verlag.
- Haxby, J.V., Gobbini, M.I., Furey, M.L., Ishai, A., Schouten, J.L., Pietrini, P., 2001. Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science* 293 (28 September), 2425–2430.
- Haynes, J., Rees, G., 2005. Predicting the stream of consciousness from activity in human visual cortex. *Current Biology* 16, 1301–1307.
- Haynes, J., Rees, G., 2006. Decoding mental states from brain activity in humans. *Nat. Rev. Neurosci.* 7, 523–534.
- Hutchinson, R., Mitchell, T.M., Rustandi, I., 2006. Hidden process models. *International Conference on Machine Learning*.
- Kamitani, Y., Tong, F., 2005. Decoding the visual and subjective contents of the human brain. *Nat. Neurosci.* 8, 679–685.
- Kanwisher, N., 2003. The ventral visual object pathway in humans: evidence from fmri. In: Chalupa, L., Werner, J. (Eds.), *The Visual Neurosciences*. MIT Press.
- Kay, K.N., Naselaris, T., Prenger, R.J., Gallant, J.L., 2008. Identifying natural images from human brain activity. *Nature* 452, 352–355.
- Kohavi, Ron, 1995. A study of cross-validation and bootstrap for accuracy estimation and model selection. *International Joint Conference on Artificial Intelligence*.
- Kriegeskorte, N., Goebel, R., Bandettini, P., 2006. Information-based functional brain mapping. *Proc. Natl. Acad. Sci. U. S. A.* 103, 3863–3868.
- Langford, J., 2005. Tutorial on practical prediction theory for classification. *J. Mach. Learn. Res.* 6, 273–306.
- Lowry, Richard, 1999. 2007 Concepts and Applications of Inferential Statistics. <http://faculty.vassar.edu/lowry/webtext.html>.
- McIntosh, A.R., Lobaugh, N.J., 2004. Partial least squares analysis of neuroimaging data—applications and advances. *NeuroImage* 23, 250–263.
- Mitchell, T.M., 1997. *Machine Learning*. McGraw Hill, New York.
- Mitchell, T.M., Hutchinson, R., Niculescu, R.S., Pereira, F., Wang, X., Just, M., Newman, S., 2004. Learning to decode cognitive states from brain images. *Mach. Learn.* 57, 145–175.
- Mitchell, T.M., Shinkareva, S.V., Carlson, A., Chang, K., Malave, V.L., Mason, R.A., Just, M.A., 2008. Predicting human brain activity associated with the meanings of nouns. *Science* 320, 1191–1195.
- Mourao-Miranda, J., Friston, K., Brammer, M., 2007. Dynamic discrimination analysis: a spatialtemporal svm. *NeuroImage* 36, 88–99.
- Norman, K., Polyn, S.M., Detre, G., Haxby, J.V., 2006. Beyond mind-reading: multi-voxel pattern analysis of fmri data. *Trends Cogn. Sci.* 10.
- OToole, A.J., Jiang, F., Abdi, H., Haxby, J.V., 2005. Partially distributed representations of objects and faces in ventral temporal cortex. *J. Cogn. Neurosci.* 17 (4), 580–590.
- OToole, A.J., Jiang, F., Abdi, H., Penard, N., Dunlop, J.P., Parent, M.A., 2007. Theoretical, statistical, and practical perspectives on pattern-based classification approaches to functional neuroimaging analysis. *J. Cogn. Neurosci.* 19, 1735–1752.
- Pereira, F. (2007). *Beyond Brain Blobs: Machine Learning Classifiers as Instruments for Analyzing Functional Magnetic Resonance Imaging Data*. PhD thesis, Computer Science Department, Carnegie Mellon University, Pittsburgh.
- Pereira, F., Gordon, G., 2006. The support vector decomposition machine. *International Conference on Machine Learning*.
- Polyn, S.M., Natu, V.S., Cohen, J.D., Norman, K.A., 2005. Category-specific cortical activity precedes recall during memory search. *Science* 310, 1963–1966.
- Schoelkopf, B., Sung, K., Burges, C., Girosi, F., Niyogi, P., Poggio, T., Vapnik, V., 1997. Comparing support vector machines with gaussian kernels to radial basis function classifiers. *IEEE Trans. Signal Process.* 45 (11), 2758–2765.
- SPM, (2000). *Statistical parametric mapping*. Technical report, Functional Imaging Laboratory, UCL (<http://www.fil.ion.ucl.ac.uk/spm>).
- Strother, S.C., 2006. Evaluating fmri preprocessing pipelines. *IEEE Eng. Med. Biol. Mag.* 25 (2), 27–41.

FURTHER-READING

- Detre, G.J. (2006). *Defining voxelsets using rich psychological models*. Master's thesis, Psychology Department, Princeton University, Princeton NJ.
- Dietterich, T.G., Bakiri, G., 1995. Solving multiclass learning problems via error-correcting output codes. *J. Artif. Intell. Res.* 2, 263–286.
- Good, Phillip, 2005. *Permutation, Parametric and Bootstrap Tests of Hypotheses*. Springer, New York.
- Joachims, T., 1999. Making large-scale svm learning practical. In: Scholkopf, B., Burges, C., Smola, A. (Eds.), *Advances in Kernel Methods—Support Vector Learning*. MIT Press.
- Ledoit, O., Wolf, M., 2003. Improved estimation of the covariance matrix of stock returns with an application to portfolio selection. *J. Empir. Finance* 10 (5), 365–411.
- Strother, S.C., Anderson, J., Hansen, L.K., Kjems, U., Kustra, R., Siditis, J., Frutiger, S., Muley, S., LaConte, S., Rottenberg, D., 2002. The quantitative evaluation of functional neuroimaging experiments: The npairs data analysis framework. *NeuroImage* 15, 615–624.
- Wasserman, Larry, 2004. *All of Statistics*. Springer, New York.