ARTICLE IN PRESS

YNIMG-10275; No. of pages: 9; 4C: 2, 3, 4, 5, 6, 7

NeuroImage xxx (2013) xxx–xxx

Contents lists available at SciVerse ScienceDirect

NeuroImage

journal homepage: www.elsevier.com/locate/ynimg

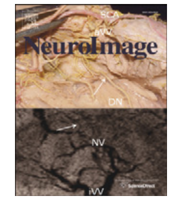

Comments and Controversies

Searchlight analysis: Promise, pitfalls, and potential

Joset A. Etzel [a,*], Jeffrey M. Zacks [b], Todd S. Braver [a]

[a] Cognitive Control & Psychopathology Lab, Washington University in St. Louis, 1 Brookings Drive, St. Louis, MO 63130, USA
[b] Dynamic Cognition Laboratory, Washington University in St. Louis, 1 Brookings Drive, St. Louis, MO 63130, USA

A R T I C L E I N F O

Article history:
Accepted 19 March 2013
Available online xxxx

Keywords:
Functional magnetic resonance imaging (fMRI)
Multivariate pattern analysis (MVPA)
Searchlight analysis
Support vector machines (SVM)
Information mapping
Classification

A B S T R A C T

Multivariate pattern analysis (MVPA) is an increasingly popular approach for characterizing the information present in neural activity as measured by fMRI. For neuroimaging researchers, the searchlight technique serves as the most intuitively appealing means of implementing MVPA with fMRI data. However, searchlight approaches carry with them a number of special concerns and limitations that can lead to serious interpretation errors in practice, such as misidentifying a cluster as informative, or failing to detect truly informative voxels. Here we describe how such distorted results can occur, using both schematic illustrations and examples from actual fMRI datasets. We recommend that confirmatory and sensitivity tests, such as the ones prescribed here, should be considered a necessary stage of searchlight analysis interpretation, and that their adoption will allow the full potential of searchlight analysis to be realized.

Introduction

Multivariate pattern analysis (MVPA) of functional MRI (fMRI) data has grown steadily since its beginnings in 2001(Haxby, 2012). Following Raizada and Kriegeskorte (2010), we illustrate the growth of the literature by showing the citation rate for several key MVPA papers in Fig. 1. Interest in MVPA spans disciplines. Advances have arisen from synergistic interactions with the machine learning community, which has developed new methods for addressing fMRI datasets and questions, as seen in the proliferation of relevant articles (e.g. Cuingnet et al., 2011; Mitchell et al., 2004; Van De Ville and Lee, 2012) and dedicated conference workshops (e.g. the International Conference on Pattern Recognition, NIPS, Cosyne, etc.). Interest in the cognitive neuroscience applications of MVPA is just as great (e.g. Heinzle et al., 2012; Tong and Pratte, 2012; Yang et al., 2012). The growing popularity of MVPA within neuroimaging has been driven by multiple factors, including: a) suggestions that it provides greater sensitivity and specificity than mass-univariate analyses with generally complementary results (Haynes and Rees, 2005; Jimura and Poldrack, 2012; Kamitani and Tong, 2005); b) the possibility of designing tests to address hypotheses which cannot be addressed with mass-univariate methods (e.g. Knops et al., 2009; Quadflieg et al., 2011; Stokes et al., 2009); and c) the intuitive appeal of a method which incorporates the signal from multiple voxels at once.

Searchlight analysis (also called information mapping) is an MVPA method introduced as a technique for identifying locally informative areas with greater power and flexibility than mass-univariate analyses (Kriegeskorte and Bandettini, 2007a; Kriegeskorte et al., 2006). Searchlight approaches are relatively unique, in that they were developed specifically for fMRI analysis, addressing both the common localization goal (many fMRI studies aim to identify small brain areas) and the spatial structure of the BOLD signal (adjacent voxels tend to have similar activation timecourses). Searchlight analysis produces maps by measuring the information in small spherical subsets ("searchlights") centered on every voxel; the map value for each voxel thus derives from the information present in its searchlight, not the voxel individually. Note that the word "information" is not used here in its formal sense (as in the field of information theory), but rather following its conventional use in the MVPA application literature. Specifically, we use the word "information" to indicate that the activity in a group of voxels varies consistently with experimental condition: a highly informative voxel cluster can be used to identify experimental condition more accurately than a weakly informative one.

Appealing aspects of searchlight analysis include its whole-brain approach (i.e., a priori region specification is not needed), the ability to pool over subject-specific activation patterns, and its minimization of the extremes of the curse of dimensionality associated with whole-brain MVPA (the "curse" refers to computational difficulties which can occur when there are more voxels than examples, see (Clarke et al., 2008; Jain et al., 2000); it is minimized in searchlight analysis since relatively few voxels are typically included in each searchlight). Additionally, searchlight analysis produces a whole-brain results map that is superficially similar in appearance to the whole-brain significance

* Corresponding author.
E-mail addresses: jetzel@artsci.wustl.edu (J.A. Etzel), jzacks@wustl.edu (J.M. Zacks), tbraver@artsci.wustl.edu (T.S. Braver).

1053-8119/$ – see front matter © 2013 Published by Elsevier Inc.
http://dx.doi.org/10.1016/j.neuroimage.2013.03.041

Please cite this article as: Etzel, J.A., et al., Searchlight analysis: Promise, pitfalls, and potential, NeuroImage (2013), http://dx.doi.org/10.1016/j.neuroimage.2013.03.041

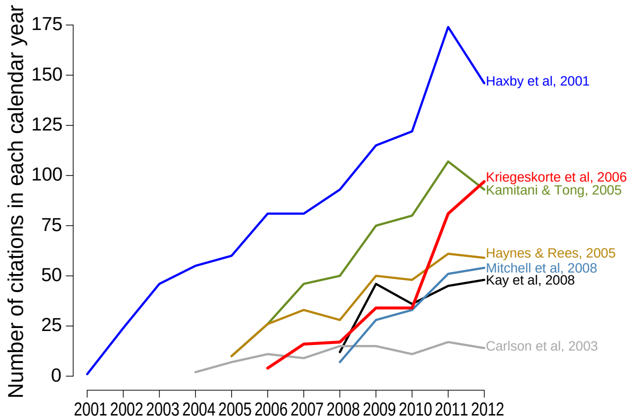

Fig. 1. Pattern-information fMRI is still a rapidly growing field, particularly searchlight analysis (note the rapid increase in papers citing Kriegeskorte et al., 2006). This figure follows Fig. 2 in Raizada and Kriegeskorte (2010), but uses the actual citation counts after 2008. The number of citations for each paper and year was obtained via Scopus (www.scopus.com) on 9 January 2013. (Carlson et al., 2003; Haxby et al., 2001; Haynes and Rees, 2005; Kamitani and Tong, 2006; Kay et al., 2008; Mitchell et al., 2008).

maps produced by more familiar mass-univariate analyses (based on the general linear model); thus, searchlight analysis results are potentially easier to interpret. These appealing aspects, plus promising early results, have led to a rapid increase in the number of studies using searchlight analyses (note the rapid rise in citations for Kriegeskorte et al., 2006 in Fig. 1, particularly in the last few years). Its acceptance as a standard approach is reflected in its inclusion in recent MVPA review and methodology articles (e.g. Bandettini, 2009; Mourao-Miranda et al., 2006; Raizada and Kriegeskorte, 2010; Tong and Pratte, 2012), as well as in the most prominent MVPA software packages (BrainVoyager QX 2.0, the Princeton MVPA Toolbox, PyMVPA).

Reflecting its potential and appeal, variations of the searchlight technique have been developed. In the spatial domain, it has been extended to circular subsets on cortical surfaces (Chen et al., 2011; Oosterhof et al., 2010, 2011), rather than the original volumetric spheres. Efforts have also been made to extend the technique to incorporate the temporal domain (Fogelson et al., 2011; Rao et al., 2011). The first searchlight analyses used the Mahalanobis distance as the similarity measure for information mapping, but a widely adopted variation is to use machine learning algorithms, often support vector machines (SVMs), instead (Haynes et al., 2007; Kriegeskorte and Bandettini, 2007b). In these approaches, generalization accuracy of the classifier is used as a proxy for information content. Group analysis is usually performed by combining individual subject's maps with a binomial or *t*-test at each voxel (with the null hypothesis that the group classification accuracy is at chance level), creating maps of voxels with significant searchlights. Here we primarily consider classification-based searchlight analysis, but much of the discussion applies regardless of the precise implementation.

Searchlight analysis is a powerful and attractive tool for understanding neuroimaging data. However, it has particular characteristics and limitations that can lead to serious interpretation errors in practice, and so we recommend that straightforward confirmatory and sensitivity tests (analogous to post-hoc tests after an ANOVA), such as the ones described here, be considered a standard part of the searchlight analysis procedure. In the following sections we describe two assumptions that often implicitly underlie the interpretation of searchlight analysis results. Unfortunately, as we illustrate, these assumptions do not always hold, and so may lead to distorted results. We then describe how confirmatory follow-up tests can be used to guard against particularly harmful distortions, using two hypotheses common in cognitive studies as illustrations. This manuscript is accompanied by Supplemental Information containing examples (with code) and technical details.

Assumption 1. Information is detected consistently.

A fundamental aspect of fMRI is that information is not distributed uniformly across voxels but rather has a three-dimensional structure: some groups of voxels (e.g. those corresponding to a specific anatomical region) are more informative for a particular task than other groups of the same size. Additionally, neuroimaging data contains information at multiple spatial frequencies (Kriegeskorte et al., 2010; Op de Beeck, 2010). For example, consider a cued finger-tapping task. The finger area of the primary motor cortex will be highly informative at a very small spatial frequency while the premotor and somatosensory cortices may be equally informative, but at a larger spatial frequency. The difference can be imagined as the size of box required to enclose the minimum set of voxels capable of task classification: a larger box is necessary to enclose the pattern in premotor or somatosensory cortices than to enclose the pattern in the primary motor cortex.

The distribution of information is relevant for searchlight analysis because interpretation of any particular map depends on whether the information can be detected equally across spatial frequencies. In a simulation designed with equal power in all spatial frequency bands, Kriegeskorte et al. (2006) showed that detection did not require a close match between the size of the searchlight and the informative area: a 4 mm radius consistently performed well. When this finding holds, it simplifies searchlight analysis interpretation: the peak areas of the map are the most informative voxels. However, if information is not present and detected equally at all spatial frequencies, then searchlight analysis results will depend fairly strongly upon the searchlight size; moreover, no single searchlight radius will be universally optimal or sufficient.

Additionally, although the Mahalanobis distance may be consistently sensitive to information across spatial frequency bands (Kriegeskorte et al., 2006), this property does not hold for all information measures used with searchlight analysis, especially the linear SVM. Training a linear SVM algorithm results in a set of weights; its decision function is a weighted linear combination of the voxels (Norman et al., 2006). Two properties of the linear SVM are particularly relevant when used in searchlight analysis: (1) It is sometimes able to correctly classify when the searchlight contains a small minority of highly informative voxels (intermixed with a majority of uninformative voxels), and conversely, (2) It is sometimes able to correctly classify when the searchlight contains a large number of weakly informative voxels.

Highly-informative voxels can be detected even when very rare

Since, as described above, linear SVMs are relatively resistant to the curse of dimensionality (Jain et al., 2000), they can sometimes classify a dataset accurately even when only a tiny minority of the voxels are informative. The degree to which this occurs varies depending on dataset properties, but it happens often enough to be relevant in practice. For instance, Supplemental Example 4 shows that introducing just five informative voxels from an actual fMRI dataset into a group of two hundred random (uninformative) voxels is sufficient to shift the median accuracy of an SVM from chance to 0.6. For an extreme example, a dataset containing a single highly informative voxel and 200 random voxels is accurately classified in Supplemental Example 5. Searchlight analysis generally includes fewer than 200 voxels in each searchlight, increasing the likelihood that searchlights containing a single or only a few informative voxels will be detected (see the “Detection of rare informative voxels” section of the Supplemental Information for further discussion).

This behavior can cause distortions in a searchlight map. To illustrate, suppose that a cluster of five highly informative voxels (capable of significant classification whenever included in a searchlight) is surrounded by hundreds of truly uninformative voxels. Any searchlight

Please cite this article as: Etzel, J.A., et al., Searchlight analysis: Promise, pitfalls, and potential, NeuroImage (2013), http://dx.doi.org/10.1016/j.neuroimage.2013.03.041