NeuroImage 162 (2017) 45–55

Contents lists available at ScienceDirect

NeuroImage

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

Effects of spatial fMRI resolution on the classification of naturalistic movies

H. Mandelkow*, J.A. de Zwart, J.H. Duyn

Advanced MRI, LFMI, NINDS, National Institutes of Health, Bethesda, MD, USA

ARTICLE INFO

Keywords:
fMRI
High-resolution
High-field
7T
Multivariate pattern analysis
MVPA
Naturalistic stimuli
Movies
Eye tracking

ABSTRACT

Studies involving multivariate pattern analysis (MVPA) of BOLD fMRI data generally attribute the success of the information-theoretic approach to BOLD signal contrast on the fine spatial scale of millimeters facilitating the *classification* or *decoding* of perceptual stimuli. However, to date MVPA studies that have actually explored fMRI resolutions at less than 2 mm voxel size are rare and limited to small sets of unnatural stimuli (like visual gratings) as well as specific sub-regions of the brain, notably the primary somatosensory cortices. To investigate what spatial scale best supports high information extraction under more general conditions this study combined naturalistic movie stimuli with high-resolution fMRI at 7 T and linear discriminant analysis (LDA) of global and local BOLD signal patterns.

Contrary to predictions, LDA and similar classifiers reached a maximum in classification accuracy (CA) at a smoothed resolution close to 3 mm, well above the 1.2 mm voxel size of the fMRI acquisition. Maximal CAs around 90% were contingent upon global fMRI signal patterns comprising 4 k–16 k of the most reactive voxels distributed sparsely throughout the occipital and ventro-temporal cortices. A Searchlight analysis of local fMRI patterns largely confirmed the global results, but also revealed a small subset of brain regions in early visual cortex showing limited increases in CA with higher resolution. Principal component analysis of the global and local fMRI signal patterns suggested that reproducible neuronal contributions were spatially auto-correlated and smooth, while other components of higher spatial frequency were likely related to physiological noise and responsible for the reduced CA at higher resolution. Systematic differences between experiments and subjects suggested that higher CA was significantly correlated with more consistent behavior revealed by eye tracking. Thus, the optimal resolution of fMRI data for MVPA was mainly limited by physiological noise of high spatial frequency as well as behavioral (in-)consistency.

1. Introduction

In the field of functional magnetic resonance imaging (fMRI), data-driven machine-learning classification methods, often referred to as multivariate pattern analysis (MVPA), have gained popularity as they promise to overcome certain limitations of traditional univariate fMRI analysis (Kriegeskorte et al., 2006; Norman et al., 2006; Pereira et al., 2009). Unlike the well-established statistical parametric mapping (SPM) technique (Friston et al., 1995), these alternative, information-theoretic methods do not depend on an explicit general linear model (GLM) of the BOLD signal, nor do they require strong a-priori assumptions about the relevant stimulus features and the hemodynamic response function (HRF) of the cortex. Based on multivariate statistics, MVPA methods like Linear Discriminant Analysis (LDA) can more efficiently exploit spatial correlations of the fMRI signal in order to gain sensitivity for the detection of small effects. fMRI data analysis by means of non-parametric, model-free machine-learning methods is particularly interesting in conjunction with naturalistic stimuli like photographs or movies that evoke complex perceptual processes, for which adequate fMRI signal models are lacking despite some unique efforts in the field (Horikawa et al., 2013; Nishimoto et al., 2011). These advantages make MVPA methods particularly useful for the kind of single-subject and single-trial analyses required for diagnostic applications, neuro-feedback (BCI) and unrepeatable cognitive processes like learning.

Many MVPA studies have loosely attributed the success of multivariate classification methods to fMRI contrast on a “fine” spatial scale (Alink et al., 2013; Emmerling et al., 2016; Guntupalli et al., 2016; Haxby et al., 2014). Indeed high-resolution (HR) fMRI studies in humans and animals have demonstrated fMRI contrast on a spatial scale of <1 mm (Menon and Goodyear, 1999; Olman et al., 2012; Yacoub et al., 2008). Such information would typically be lost to partial-volume effects and low statistical power in conventional fMRI analysis with low-resolution

* Corresponding author.
E-mail address: Hendrik.Mandelkow@nih.gov (H. Mandelkow).

http://dx.doi.org/10.1016/j.neuroimage.2017.08.053
Received 7 April 2017; Received in revised form 14 August 2017; Accepted 20 August 2017
Available online 24 August 2017
1053-8119/

Abbreviations

BOLD	blood oxygenation level dependent (fMRI)
BW	(readout) bandwidth (MRI parameter)
CA	classification accuracy
CNR	contrast-to-noise ratio
FA	flip angle (MRI parameter)
fMRI	functional MRI
fROI	functional ROI (= voxels/features)
FSL	FMRIB Software Library
FWHM	full width at half maximum
GLM	general linear model
GNB	Gaussian Naïve Bayes (classifier)
GRAPPA	(MRI parallel imaging technique)
HR	high-resolution (fMRI)
HRF	hemodynamic response function
LDA	Linear Discriminant Analysis (classifier)
LR	low-resolution (fMRI)
Mcflirt	FSL motion correction software
MRI	magnetic resonance imaging
MVPA	multivariate pattern analysis
NIH	National Institutes of Health
NN	nearest-neighbour (classifier)
NM	nearest-mean (classifier)
PC	principal component
PCA	principal component analysis
ROI	region of interest
SL	searchlight
SNR	signal-to-noise ratio
SVD	singular value decomposition
SVM	support vector machine
TE	echo time (MRI parameter)
TR	(volume) repetition time (MRI)

(LR) acquisitions (>2 mm) and additional smoothing (>3–8 mm FWHM) aimed at detecting significant deviations in the regional *mean* signal of voxel clusters (Kriegeskorte et al., 2006). MVPA methods, by contrast, are designed to take advantage of a spatially diverse (high-dimensional) fMRI signal (of high resolution), given sufficient SNR to make activation patterns distinct and reproducible (Chaimow et al., 2017). However, experimental evidence in support of this premise is scarce in the literature:

A majority of MVPA studies, even those ostensibly concerned with fMRI signal patterns "on a fine spatial scale", have relied on acquisitions far from the resolution limits of fMRI, especially at high field (≥3 T) (Alink et al., 2013; Guntupalli et al., 2016). Sacrificing fMRI resolution in the inevitable trade-off for SNR, TR and (whole-brain) coverage makes perfect sense in the context of traditional fMRI analysis, which relies on spatial smoothing to justify statistical assumptions and to overcome the limitations of anatomical co-registration procedures. The few studies that have actually applied MVPA to HR fMRI data were limited to small sets of artificial stimuli like oriented gratings as well as small cortical subregions like V1 (Emmerling et al., 2016; Gardumi et al., 2016; Sengupta et al., 2017; Swisher et al., 2010). Such HR fMRI experiments have successfully detected patterns of cortical orientation columns only 0.7 mm in diameter. However, MVPA studies that used a voxel size much larger than these cortical structures have demonstrated an equal or better ability to discriminate similar orientation stimuli (Haynes and Rees, 2005; Kamitani and Tong, 2005). This apparent conundrum triggered an ongoing debate about "fMRI hyperacuity", or mechanisms by which the fMRI signal from cortical columns might transcend to the spatial scale of voxels (Chaimow et al., 2017; Kamitani and Sawahata, 2010; Kriegeskorte et al., 2010; Op de Beeck, 2010), although similar information on different spatial scales could also be supported by independent physiological processes.

This paper is not a contribution to the hyperacuity debate regarding the discrimination of stimuli (oriented gratings) specifically tailored to elicit predictable changes in the neuronal response on a sub-millimeter scale. Instead we address the question whether or not high-resolution fMRI at 7 T in practice increases the discriminability of a wide range of naturalistic (visual) stimuli by MVPA classification methods. Some MVPA studies that have used naturalistic stimuli like movies indicate that smoothing may increase CA contrary to expectations (Haxby et al., 2011). But fMRI resolutions of <2 mm in voxel size were not explored. More recent studies have actually compared MVPA at resolutions as high as 0.8–1.1 mm but were limited to primary visual and auditory brain regions as well as small sets of highly tailored and very similar stimuli (gratings, vowels) (Gardumi et al., 2016; Sengupta et al., 2017; Swisher et al., 2010) – quite the opposite of our naturalistic stimuli expected to elicit strong fMRI contrast across a range of spatial scales. Even so, aforementioned studies tentatively corroborate a maximum in classification accuracy at a moderate fMRI resolution of 2–3 mm, although results were somewhat mixed and exhibited a substantial amount of variability above all.

In principle, high-field MRI systems (≥7 T) show promise for functional brain imaging in humans as they facilitate fMRI scans with higher spatial and temporal resolution thanks to both increased SNR and functional (T_2^*-weighted) BOLD contrast. It has been argued, however, that the co-amplification of physiological noise and the lack of precise inter-subject alignment methods limit the gains of higher field strength in many practical applications (Krüger and Glover, 2001; Triantafyllou et al., 2005). Focused on the tSNR of single voxels in resting-state fMRI these studies neglect the spatially correlated fMRI signal evoked by stimulation, which is the basis of MVPA. Ultimately, the spatial bandwidth and relative amplitude of BOLD signal and noise components must determine the optimal resolution for information extraction by SPM or MVPA. In other words, the available spatial bandwidth of the BOLD signal must overcompensate inevitable losses in instrumental SNR as well as any interference due to "physiological noise" in order to achieve a net gain in information as a result of higher fMRI resolution.

In light of these mixed results, the purpose of this study is to determine whether or not MVPA profits from HR fMRI at 7 T in practice – specifically whether the accuracy of common classifiers like LDA increases with fMRI resolution. In preparation for this study, we recently published a comparison of common classification algorithms demonstrating that comparisons between data sets are largely independent of the choice of classifier (Mandelkow et al., 2016). However, results will likely have some dependence on the type of stimuli as well as the targeted brain region. Aiming for results that would generalize across a wide range of (visual) stimuli, we chose short clips from (naturalistic) action movies (L. Wachowski and A. Wachowski, 1999), which are popular in the field for a number of reasons: Unlike abstract stimuli like moving dots or gratings, movies of human actions and environments have ecological validity for human subjects and are certain to evoke a broad range of natural cognitive processes. Such processes are of great interest in cognitive neuroscience, but challenging to analyze by traditional (GLM-based) fMRI methods, because investigators are forced to speculate on the (many) relevant stimulus features and their relative contributions to the evoked BOLD signal (Huth et al., 2012; Naselaris et al., 2015). From a practical point of view, movies are easy to deliver with precise timing in the MRI scanner and they are known to evoke a strong and reproducible BOLD fMRI response in large parts of the brain, especially if they are engaging, i.e. attracting a subject's attention (Golland et al., 2007; Hasson et al., 2008; Jääskeläinen et al., 2008).