ARTICLE IN PRESS
YNIMG-12813; No. of pages: 13; 4C: 2, 7, 8, 9, 10
NeuroImage xxx (2015) xxx–xxx

Contents lists available at ScienceDirect

NeuroImage

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

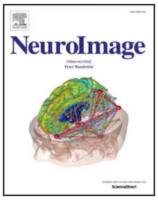

Q1 Reliability of dissimilarity measures for multi-voxel pattern analysis

Q2 Alexander Walther [a,b,*], Hamed Nili [c], Naveed Ejaz [b], Arjen Alink [a], Nikolaus Kriegeskorte [a], Jörn Diedrichsen [b]

[a] *MRC Cognition and Brain Sciences Unit, 15 Chaucer Road, CB2 7EF, Cambridge, United Kingdom*
[b] *Institute of Cognitive Neuroscience, University College London, Alexandra House, 17 Queen Square, London WC1N 3AR, United Kingdom*
[c] *Department of Experimental Psychology, University of Oxford, South Parks Road, OX1 3UD, Oxford, United Kingdom*

ARTICLE INFO

Article history:
Received 10 February 2015
Accepted 7 December 2015
Available online xxxx

Keywords:
fMRI
Multi-voxel pattern analysis
Representational similarity analysis
Crossvalidation
Linear discriminant
Noise normalization
Classification
Decoding
Machine learning

ABSTRACT

Representational similarity analysis of activation patterns has become an increasingly important tool for studying brain representations. The dissimilarity between two patterns is commonly quantified by the correlation distance or the accuracy of a linear classifier. However, there are many different ways to measure pattern dissimilarity and little is known about their relative reliability. Here, we compare the reliability of three classes of dissimilarity measure: classification accuracy, Euclidean/Mahalanobis distance, and Pearson correlation distance. Using simulations and four real functional magnetic resonance imaging (fMRI) datasets, we demonstrate that continuous dissimilarity measures are substantially more reliable than the classification accuracy. The difference in reliability can be explained by two characteristics of classifiers: discretization and susceptibility of the discriminant function to shifts of the pattern ensemble between runs. Reliability can be further improved through multivariate noise normalization for all measures. Finally, unlike conventional distance measures, crossvalidated distances provide unbiased estimates of pattern dissimilarity on a ratio scale, thus providing an interpretable zero point. Overall, our results indicate that the crossvalidated Mahalanobis distance is preferable to both the classification accuracy and the correlation distance for characterizing representational geometries.

Introduction

It has become increasingly popular to analyze functional magnetic resonance imaging (fMRI) data using multi-voxel pattern analysis (MVPA). In MVPA, activation patterns are analyzed using either classification (Cox and Savoy, 2003; Haxby et al., 2001) or representational similarity analysis (RSA, Kriegeskorte et al., 2008). Both approaches quantitatively measure the dissimilarity of fMRI response patterns for pairs of conditions. All possible pairwise dissimilarity values of an experiment can be assembled in a pairwise decoding accuracy matrix or representational dissimilarity matrix (RDM).

One important decision in RSA is the choice of dissimilarity measure. Popular dissimilarity measures are the percentage of correct pairwise classifications (accuracy) and continuous distance measures, such as the Pearson correlation distance, the Euclidean distance, and the Mahalanobis distance. In this paper we provide a careful evaluation of the reliability of these dissimilarity measures, i.e. how reliable a measure is over replications of the experiment.

In evaluating reliability, it is important to consider the inferential aim of the analysis. One hypothesis that a researcher may want to test is that the patterns associated with conditions A and B are more similar than those associated with conditions C and D. This hypothesis concerns only the ranks of the dissimilarities. A more specific hypothesis would be that the dissimilarity between the patterns for conditions A and B is twice as large as the dissimilarity between the patterns for C and D. Here it is necessary that the dissimilarity measure have a meaningful zero point, with zero indicating that the two patterns are not different. However, distances, by definition, are non-negative and always larger than zero if estimated from noisy data. Thus, even if the true patterns are not different, the estimated distance will be larger than zero. The noise creates a positive bias, which will rise with the noise level. As we will show in the results, the bias can be removed by crossvalidation (Allefeld and Haynes, 2014; Nili et al., 2014; Kriegeskorte et al., 2007). Crossvalidated distance estimator are unbiased, i.e. their expected value equals the true distance and is zero if the two patterns are not different (see Crossvalidation section). As a consequence, crossvalidated Q3 distance estimators enable us to interpret ratios between distances.

In this paper, we compare the reliability of the Euclidean distance, the Mahalanobis distance, and the correlation distance and study the influence of univariate and multivariate noise normalization on RDM reliability. We also consider crossvalidated versions of the Mahalanobis distance (including the linear-discriminant t value; Nili et al., 2014; Kriegeskorte et al., 2007). Finally, we compare continuous distance measures to classification accuracies from linear discriminant analysis

---

* Corresponding author at: MRC Cognition and Brain Sciences Unit, 15 Chaucer Road, Cambridge, CB2 7EF, United Kingdom. Tel.: +44 1223 355294.
E-mail addresses: alexander.walther@mrc-cbu.cam.ac.uk (A. Walther), hamed.nili@psy.ox.ac.uk (H. Nili), n.ejaz@ucl.ac.uk (N. Ejaz), arjen.alink@mrc-cbu.cam.ac.uk (A. Alink), nikolaus.kriegeskorte@mrc-cbu.cam.ac.uk (N. Kriegeskorte), j.diedrichsen@ucl.ac.uk (J. Diedrichsen).

http://dx.doi.org/10.1016/j.neuroimage.2015.12.012
1053-8119/

Please cite this article as: Walther, A., et al., Reliability of dissimilarity measures for multi-voxel pattern analysis, NeuroImage (2015), http://dx.doi.org/10.1016/j.neuroimage.2015.12.012

(LDA) and support vector machines (SVM). Overall, our results strongly suggest the use of continuous crossvalidated distance estimators with multivariate noise normalization to measure brain representational dissimilarities.

Materials and methods

The Euclidean distance

In RSA, we want to calculate the distance between the activation patterns b_k and b_j, corresponding to two of $k = 1,\ldots,K$ conditions. An activation pattern usually consist of the regression coefficients from a general linear model (GLM), which represent the response of the voxels $p = 1,\ldots,P$ to condition k. The Euclidean distance between two patterns in a P-dimensional voxel space, with the activity of each voxel forming a separate dimension, is defined analogously to the familiar distance in two dimensions. The squared Euclidean distance d^2 between the two row vectors $\mathbf{b}_k$ and $\mathbf{b}_j$ is:

$$d^2_{Euclidean}(\mathbf{b}_k,\mathbf{b}_j) = \|\mathbf{b}_j-\mathbf{b}_k\|^2 = (\mathbf{b}_j-\mathbf{b}_k)(\mathbf{b}_j-\mathbf{b}_k)^T = \mathbf{cBB}^T\mathbf{c}^T \qquad (1)$$

where the last term represents a compact form obtained by assembling the activation patterns into a K × P Matrix **B** and applying a 1 × K contrast vector **c**, which contains zeros except for $c_j = 1$ and $c_k = -1$.

To visualize the pattern distances, imagine each pattern as a vector extending from the origin to point $\mathbf{b}_k$, where the origin of the pattern space is usually determined by the implicit baseline estimate of the GLM. The Euclidean distance between the endpoints of two vectors is independent of the origin (Figs. 1A,B). This might be advantageous if the baseline was not reliably estimated or if it cannot be meaningfully defined.

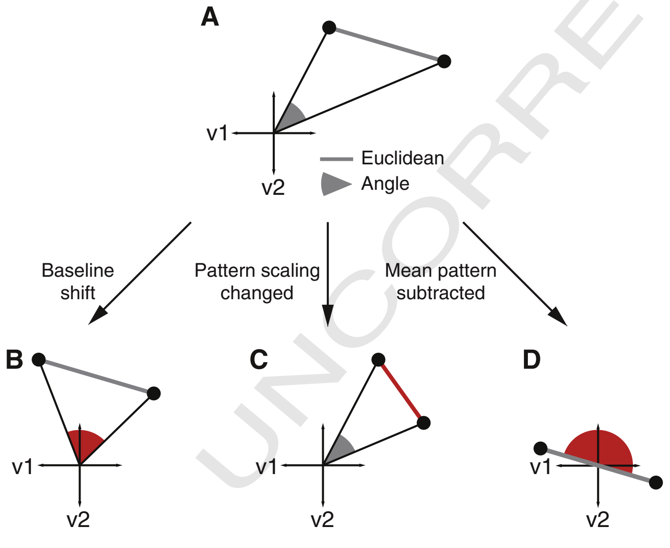

Fig. 1. Euclidean and angle-based distance in MVPA. (A) An fMRI pattern space laid out by two voxels (v1 and v2; note that the typical pattern space will often have >50 dimensions). Two pattern vectors extend from the origin. The Euclidean distance is the distance between the patterns. The cosine distance (as well as the Pearson correlation distance) measures pattern dissimilarity as a function of the angle enclosed by the vectors. (B) Shifts of the origin (i.e. the fMRI baseline) of the pattern space influence the angle (red) between the two vectors and hence the correlation distance, but not the Euclidean distance (gray). (C) Changes in the length of the two vectors (multiplicative scaling) influence the Euclidean distance (red) between the two vectors, but not the angle (gray). (D) The mean pattern of the two conditions has been subtracted (cocktail blank removal). The two vectors now extend in opposite directions from the origin, causing the cosine of the angle (red) and the correlation to become −1.

The Pearson correlation distance

Another measure of the similarity of $\mathbf{b}_k$ and $\mathbf{b}_j$ is their Pearson correlation r. The correlation is related to a slightly simpler measure, which can be more easily understood graphically: the cosine of the angle between the vectors (Fig. 1A). The cosine can be obtained by normalizing $\mathbf{b}_k$ and $\mathbf{b}_j$ by their respective L2-norms and subsequently calculating their inner product. We can then obtain a distance measure (known as cosine distance) by taking the complement:

$$d_{Cosine}(\mathbf{b}_k,\mathbf{b}_j) = 1-\frac{\langle \mathbf{b}_k,\mathbf{b}_j\rangle}{\|\mathbf{b}_k\|\|\mathbf{b}_j\|} = 1-\cos(\angle \mathbf{b}_k,\mathbf{b}_j). \qquad (2)$$

The inner product detects congruent trends between $\mathbf{b}_k$ and $\mathbf{b}_j$ (i.e. when $b_{p,k}$ tends to be high, $b_{p,j}$ tends to be high as well, and vice versa). The normalization makes the cosine distance, unlike the Euclidean distance, invariant to changes in scaling (or length) of b (Fig. 1C).

The correlation distance is equivalent to the cosine distance after subtracting the mean value from each voxel pattern. If $\bar{b}$ is the voxel mean and 1 is a 1 × P row vector of ones, the correlation distance is defined as:

$$\tilde{\mathbf{b}}_k = \mathbf{b}_k-\bar{b}_k 1 \qquad \tilde{\mathbf{b}}_j = \mathbf{b}_j-\bar{b}_j 1\; d_{Correlation}(\mathbf{b}_k,\mathbf{b}_j) = 1-\frac{\langle \tilde{\mathbf{b}}_k,\tilde{\mathbf{b}}_j\rangle}{\|\tilde{\mathbf{b}}_k\|\|\tilde{\mathbf{b}}_j\|} = 1-\cos\left(\angle \tilde{\mathbf{b}}_k,\tilde{\mathbf{b}}_j\right) \qquad (3)$$

The cosine and correlation distance are zero if two normalized patterns are identical. In the cosine similarity, only vector length is divisively normalized. In the correlation distance, the mean is first subtracted before divisive length normalization, making it invariant to both changes in the mean and variance of b_k across voxels. Importantly, both the cosine and correlation distance depend on the implicit baseline estimate of the GLM (Fig. 1B). Therefore, shifts in the origin will affect the overall distance structure.

The effect of mean pattern subtraction (cocktail-blank removal)

Before submitting the patterns to MVPA, it is common practice to subtract the mean pattern, i.e. the mean across conditions for each voxel, from each response pattern (Misaki et al., 2010; Op de Beeck, 2010; Pietrini et al., 2004; Williams et al., 2008, 2007). This normalization step is sometimes called "cocktail-blank removal". Removal of the mean pattern has a very different effect from removing the *mean value* (i.e. the mean of each condition, averaged across voxels, Eq. (3)). Mean pattern subtraction effectively moves the origin of the pattern space to lie in the mean pattern of all conditions (Fig. 1D). The reasoning behind this normalization step is that the response patterns may share a common component, which will increase all correlations and hence decrease the correlation distance. Mean pattern subtraction removes the influence of this common response pattern. However, the change in origin will cause unrelated patterns to be negatively correlated (Garrido et al., 2013; Diedrichsen et al., 2011). In the extreme case of only two conditions, the angle between them will always be 180 degrees and the cosine of the angle (and also the correlation) will be −1 (Fig. 1D). This can change the representational structure substantially, even when only considering the ranks of the distances. Unlike the correlation distance, the Euclidean distance is unaffected by mean pattern subtraction, as it does not depend on the origin of the coordinate system (Fig 1D).

Please cite this article as: Walther, A., et al., Reliability of dissimilarity measures for multi-voxel pattern analysis, NeuroImage (2015), http://dx.doi.org/10.1016/j.neuroimage.2015.12.012