NeuroImage 186 (2019) 155–163

Contents lists available at ScienceDirect

NeuroImage

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

Stability of representational geometry across a wide range of fMRI activity levels

Spencer A. Arbuckle [a], Atsushi Yokoi [b,c], J. Andrew Pruszynski [a,d,e,f], Jörn Diedrichsen [a,g,h,*]

[a] Brain and Mind Institute, Western University, Canada
[b] Graduate School of Frontier Biosciences, Osaka University, Japan
[c] Center for Information and Neural Networks (CiNet), National Institute of Information and Communications Technology, Osaka, Japan
[d] Department of Physiology and Pharmacology, Western University, Canada
[e] Department of Psychology, Western University, Canada
[f] Robarts Research Institute, Western University, Canada
[g] Department of Statistical and Actuarial Sciences, Western University, Canada
[h] Department of Computer Science, Western University, Canada

ARTICLE INFO

Keywords:
Multi-voxel pattern analysis
fMRI
Representational similarity analysis
Representational geometry

ABSTRACT

Fine-grained activity patterns, as measured with functional magnetic resonance imaging (fMRI), are thought to reflect underlying neural representations. Multivariate analysis techniques, such as representational similarity analysis (RSA), can be used to test models of brain representation by quantifying the representational geometry (the collection of pair-wise dissimilarities between activity patterns). One important caveat, however, is that non-linearities in the coupling between neural activity and the fMRI signal may lead to significant distortions in the representational geometry estimated from fMRI activity patterns. Here we tested the stability of representational dissimilarity measures in primary sensory-motor (S1 and M1) and early visual regions (V1/V2) across a large range of activation levels. Participants were visually cued with different letters to perform single finger presses with one of the 5 fingers at a rate of 0.3–2.6 Hz. For each stimulation frequency, we quantified the difference between the 5 activity patterns in M1, S1, and V1/V2. We found that the representational geometry remained relatively stable, even though the average activity increased over a large dynamic range. These results indicate that the representational geometry of fMRI activity patterns can be reliably assessed, largely independent of the average activity in the region. This has important methodological implications for RSA and other multivariate analysis approaches that use the representational geometry to make inferences about brain representations.

1. Introduction

Multivariate analysis of activity patterns has profoundly changed functional magnetic resonance imaging (fMRI) data analysis. Traditional fMRI studies have examined differences in overall activity levels in extended brain regions. In this approach, local fine-grained patterns of activity are removed by smoothing, as they are typically not consistent across individuals. However, it was realized that one could decode the experimental condition from activity patterns within individuals, even if the average activity is the same between conditions (Haxby et al., 2001). Decodability is often interpreted as evidence that the region *represents* something about the underlying distinction between conditions (Haxby et al., 2014), i.e. that another area can potentially read out information about the distinction (de-Wit et al., 2016; deCharms and Zador, 2000). Extending this idea, the degree to which different pairs of activity patterns are dissimilar may tell us something about the structure of the underlying neural population code. For example, a region involved in object recognition should show large dissimilarities between activity patterns associated with objects from different categories, but smaller dissimilarities between objects of the same category (Kriegeskorte et al., 2008b). The relationship between all pair-wise dissimilarities defines what we call the *representational geometry*. Representational similarity analysis (RSA, Kriegeskorte et al., 2008a), pattern component modelling (PCM, Diedrichsen et al., 2018) and encoding models (Naselaris et al., 2011) all analyze this representational geometry to test between models of brain representations (Diedrichsen and Kriegeskorte, 2017).

When testing representational models with fMRI data, we base our analysis on the Blood Oxygenation Level Dependent (BOLD) signal. To

* Corresponding author. Brain and Mind Institute, Western University, Canada.
E-mail address: jdiedric@uwo.ca (J. Diedrichsen).

https://doi.org/10.1016/j.neuroimage.2018.11.002
Received 17 February 2018; Received in revised form 30 October 2018; Accepted 1 November 2018
Available online 2 November 2018

what degree can we make inferences about neural representations using this indirect measure of neural activity? There are a number of reasons why representational fMRI analysis may be limited (see discussion). One important problem, which is the main focus of this paper, is that the measured representational geometry may depend strongly on the overall activity in a region. This is of concern as we often make comparisons across regions, participants, or attentional states with different activity levels. Although RSA is in theory independent of average activity, it is possible that the patterns and their dissimilarities distort with increasing activation. The relationship between the pair-wise pattern dissimilarities is only guaranteed to remain the same if all voxel activities in a region scale (up to a multiplicative constant) according to the same function. Thus, even though there is evidence that the relationship between neural activity and the average BOLD is fairly linear in M1 (Siero et al., 2013) and V1 (Boynton et al., 1996; Heeger et al., 2000), these findings do not guarantee that the representational geometry would also scale in an orderly fashion.

There are a number of potential mechanisms that could lead to substantial distortions of the representational geometry when measured with fMRI. For example, the spatial point-spread function of BOLD may differ across activity levels. Recent work using optical imaging suggests that although vasodilation of arterioles is relatively coupled with the pre- and post-synaptic activity of the surrounding neural tissues (O'Herron et al., 2016), the authors observed many instances of vasodilatory responses in the absence of local changes in neural activity. Notably, they propose that this decoupling may change with increasing neural activity. We interpret this result as cause for concern, as this effect could cause an increasing spread of the BOLD signal to neighboring voxels as activity increases. Because such spread would affect different activity patterns differently, it could lead to severe distortions of the representational geometry, which would make it difficult to draw conclusions about the representational content of neural population codes using fMRI data. Therefore, it is important to empirically test if the representational geometry remains stable across a wide range of overall activity levels.

We investigated this question with an experiment that allowed us to assess patterns in sensory-motor regions (S1 and M1) and in primary and secondary visual cortices (V1 and V2) using RSA. Ejaz et al. (2015) demonstrated a stable representational geometry across humans for individual finger movements in M1 and S1, in which the thumb had the most distinct activity pattern and neighboring fingers showed higher similarities than non-neighboring fingers. Similarly, it has been shown that letters (Miyawaki et al., 2008) as well as colors (Brouwer and Heeger, 2009), can be decoded from activity in visual cortices. We therefore cued finger presses with colored letter cues presented on a screen. We chose a specific letter and color for each finger, such that the perceptual similarities between stimuli would be different from the dissimilarities of motor actions. To increase the overall activity in both regions, we varied the letter flashing and finger pressing frequency between 0.3 Hz and 2.6 Hz. For individual finger presses on our isometric device (see methods), a rate of 2.6 Hz is close to the upper performance limit.

In interpreting the results, it is important to distinguish between changes in the representational geometry that arise from the fact that the neural activity patterns change qualitatively with higher stimulation frequency, and changes that arise through distortions when measuring the representational geometry using BOLD. We designed our task assuming that the underlying neural activity patterns would scale by a single multiplicative factor with increasing speed, allowing us to identify distortions arising in the neural-to-BOLD coupling.

We note that the data have been already included in a recent publication from our lab (Diedrichsen et al., 2018) as an example to highlight the technical aspects of fitting different types of PCM models. Although not the focus of the paper, the results already indicate that, on a group level, the representational geometry in M1 remains relatively unchanged across movement speeds. However, because possible distortion may be idiosyncratic across individuals, we here re-investigate this issue with dedicated analyses within the framework of RSA.

2. Methods & procedures

2.1. Participants

We measured cortical activity patterns in 5 female and 3 male participants (mean Age = 25.5 (2.41) years). All participants were self-reported right handers (mean Edinburgh questionnaire laterality quotient = 91.25 (7.82)), and made individuated finger presses of the right hand. Motor cortex data from these participants were used as example data in a method article describing non-linear pattern component models (Diedrichsen et al., 2018).

2.2. Apparatus and stimuli

The motor behavior was monitored by a keyboard-like device. The device had a key for each finger of the right hand, with a force transducer (Honeywell-FS series, dynamic range = 0–16N, resolution <0.02N) mounted under each key. A bevel on each key ensured fingertip placement across participants was consistent. Forces were recorded at a sampling rate of 200 Hz.

To evoke visual responses, we presented colored letter stimuli on the screen at the same frequency as the finger presses. The aim was to pick letters and colors that would induce a dissimilarity structure that would differ considerably from that of the fingers. Therefore, we chose similar letters and colors for fingers which evoke different activity profiles (see Fig. 1a for finger-letter-color pairs). The letters were presented centrally and peripherally (see Fig. 1b). The size of the letters on the screen were 8×10 cm, subtending a visual angle of approximately 70°. The screen background was black. Participants were instructed to maintain visual fixation on a gray cross presented centrally on the screen. Cues were presented centered on the fixation cue. Five lines were presented in the lower third of the screen, one for each of the five fingers. The locations of these lines were dynamically updated to indicate the real-time force applied to each key of the fingerboard device.

2.3. Behavioral task

In the MR scanner, participants completed a paced finger pressing task. Each trial lasted for 8 s and was divided into three phases (see Fig. 1b). In the announce phase (1 s), participants saw a visual cue indicating which finger of their right hand they were to press in the

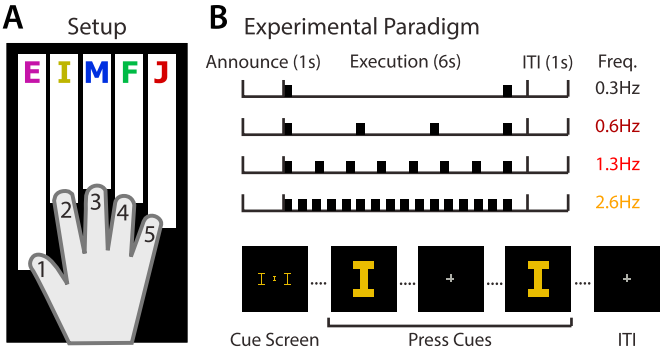

Fig. 1. Experimental paradigm. **A)** Participants made short, isometric presses of an individual finger onto a keyboard while in an MR scanner. Each finger press was cued with a unique color-letter combination. **B)** A cue at the start of each trial (1s) instructed participants which finger they would press. Participants then executed presses when prompted by a larger cue presentation. The cues flashed either 2, 4, 8, or 16 presses in 6 s (0.3, 0.6, 1.3, 2.6 Hz). A 1s inter-trial-interval (ITI) separated each trial. Random periods of rest were interleaved between trials in each block. This design yielded 20 conditions (5 fingers/letters x 4 frequencies).

current trial. During the following execution phase (6 s), participants made 2, 4, 8, or 16 isometric presses of the finger (0.3, 0.6, 1.3, and 2.6 Hz pressing frequency). Each press was paced by a visual metronome, which flashed the letter cue from the announce phase for 100 ms each. The first and the last press always occurred at the beginning and end of the execution phase, with the intermediate presses being cued at a constant rate (see Fig. 1b). Following this, there was an inter-trial interval of 1 s before the next trial started. Throughout the entire experiment, participants were instructed to refrain from moving their wrist or fingers of either hand when not instructed to do so.

There were 20 possible conditions (5 fingers/letters x 4 pressing frequencies). The task was divided into 8 runs of 40 trials each, with two repeats per condition. Trial order within each run was randomized. Seven periods of rest (13 s) were randomly interspersed between trials in each run. Each run lasted 411 s. Participants were instructed to produce a minimum of 2 N, but not to exceed a maximum of 4N with each finger press. Minimum and maximum force thresholds were visually presented on screen above the finger force lines. Between presses, the force applied to each key needed to be below 0.75N before another press could be registered. The fixation cross turned white for each correct press. Participants trained on this task prior to the scanning portion of the experiment to ensure stable performance.

2.4. fMRI data acquisition

Functional images were acquired using a Siemens Magnetom 7T MRI scanner with a 32-channel head coil at Western University (London, Ontario, Canada). Volumes were acquired using an interleaved, multi-band slice acquisition (TR = 1000 ms, 44 slices, 1.4 mm isotropic voxels, no gap between slices, in-plane acceleration factor = 3, multi-band factor = 4). The first three images of each functional run were discarded to allow magnetization to reach equilibrium. The slices covered the dorsal aspects of the cerebrum, encompassing M1 through to V1. A T1-weighted anatomical scan (3D MPRAGE sequence, TR = 6000 ms, 0.75 mm isotropic voxels, 208 vol) was also acquired at the start of the scan. Fieldmaps were collected at the end of the imaging session.

2.5. Preprocessing and first-level model

First-level fMRI analyses were conducted with SPM12 (http://www.fil.ion.ucl.ac.uk/spm/). Functional images were realigned to correct for motion across runs. Within this process, we utilized a B_0 fieldmap to correct for magnetic field inhomogeneities. Due to the short TR, no slice timing corrections were applied. The functional data was co-registered to the individual anatomical scan, but no normalization was applied.

The pre-processed images were analyzed with a general linear model (GLM) with separate task regressors for each condition (20 regressors) for each run. Each regressor was a boxcar function that was on for 6 s of the trial duration and off otherwise. These regressors were then convolved with a hemodynamic response function with a peak onset of 4.5 s and a post-stimulus undershoot minimum at 11 s.

We used the SPM FAST autocorrelation model in conjunction with restricted-maximum likelihood (ReML) estimation to estimate the long-range temporal dependencies in the functional timeseries. This relies on a minor modification of the standard SPM analysis routine. SPM proceeds in a two-step estimation process. First, SPM estimates an ordinary least square regression to collect sufficient statistics on the voxels in order to select voxels that are used for autocovariance estimation using ReML. By default, the inclusion of voxels is decided using an omnibus F-test (any difference in any condition vs. zero) with a threshold set to 0.001. This means voxels that are probably "involved" in the task of interest will be included. Using the FAST option to estimate the temporal covariance effectively attenuates low-temporal frequencies. This means it is unnecessary to apply a separate high-pass filter. Without additional high-pass filtering, however, very few voxels will be significant in the omnibus F-test in the first step because the data is dominated by low-frequency noise (which is now only removed in the second stage with the FAST option). Therefore, we modified the estimation procedure to specify separate masks that defined which voxels are included in the first and second steps of the analysis. Here, we were interested in optimal inferences for gray-matter. Therefore, the coefficients of this model were estimated from all gray-matter voxels using a gray-matter mask. We found that on several data sets, this analysis procedure improves the reliability of activity pattern estimates as compared to the standard high-pass filtering and subsequent temporal autocorrelation correction with FAST.

2.6. ROI definitions

We used Freesurfer software (Dale et al., 1999) to extract the white-gray matter and pial surfaces from each participant's anatomical image. These surfaces were inflated to a sphere and aligned using sulcal depth and curvature information to the Freesurfer average atlas (fsaverage, Fischl et al., 1999). Following alignment, both hemispheres in each participant were resampled into a 163,842 vertex grid. This allowed us to reference similar areas of the cortical surface in each participant by selecting the corresponding vertex on the group atlas.

Anatomical regions of interest (ROI) were defined using a procedure established in previous work (Wiestler and Diedrichsen, 2013; Ejaz et al., 2015). All ROIs were defined using a probabilistic cytoarchtectonic atlas (Fischl et al., 2008) projected onto the common group surface. For M1 and S1, we constrained the resulting ROIs to the hand and arm region by choosing the area of the cyctoarchteconically defined strip 2 cm above and below the hand knob (Yousry et al., 1997). To avoid cross-contamination between M1 and S1 activities along the central sulcus, voxels with more than 25% of their volume originating from the opposite side of the central sulcus were excluded. The primary and secondary visual cortices (V1/V2) were grouped as one ROI. The group area was then projected onto the individual volume using the individual surface reconstruction.

2.7. Multivariate fMRI analysis

Multi-voxel analyses were conducted within each ROI (M1, S1, and V1/V2), using the RSA (Nili et al., 2014) and PCM toolboxes (Diedrichsen et al., 2018). For each ROI, we extracted the beta-weights from the first-level GLM for each condition in each imaging run. These beta-weights were then spatially pre-whitened using multivariate noise-normalization to suppress correlated noise across voxels (Walther et al., 2016). The mean pattern was not removed from each run to preserve information about activity from baseline in each voxel.

We then calculated the squared cross-validated Mahalanobis distance (crossnobis; Walther et al., 2016; Diedrichsen et al., 2016) between activity patterns:

$$d^2(x_y, x_z) = (\mathbf{x}_y - \mathbf{x}_z)_A^T \boldsymbol{\Sigma}^{-1} (\mathbf{x}_y - \mathbf{x}_z)_B \quad (1)$$

where $(\mathbf{x}_y - \mathbf{x}_z)_A$ corresponds to the difference between the activity patterns of conditions y and z in run A, and Σ refers to the voxel-wise noise covariance matrix (Walther et al., 2016). We repeated this procedure over all possible leave-one-run-out crossvalidation folds and then averaged the resulting dissimilarities across folds. This procedure leads to an unbiased distance estimate, which on average will be zero if there is no reliable difference between the two patterns. This also means that the crossnobis estimator can become negative. The large advantage for this measure, however, is that zero is meaningfully defined and hence ratios between distances can be interpreted in a meaningful way.

The representational geometry is characterized by the dissimilarities between all possible pairs of condition activity patterns, that can be collected into a representational dissimilarity matrix (RDM). The RDM is a (number of conditions x number of conditions) symmetric matrix, with zeros along the diagonal. Dissimilarities were calculated for the left

hemisphere M1 and S1 ROIs (contralateral to the side of finger movements). For V1/V2 ROIs, we first calculated dissimilarities for each hemisphere, then averaged the dissimilarities across hemispheres within each participant. We used classical multi-dimensional scaling (eigenvalue decomposition) to visualize a low-dimensional projection of the representational geometry (Diedrichsen et al., 2018).

2.8. Stability of the representational geometry across stimulation-frequencies.

To assess the stability of the representational geometry across frequency conditions, we correlated the RDMs across all 6 possible frequencies pairs within each participant. For each frequency condition (j), we had 10 pairwise dissimilarities ($d_{i,j}$, where i corresponds to one of the 10 dissimilarities for frequency condition j). We calculated a Pearson correlation without subtracting the mean across the 10 dissimilarities first, as zero is a fixed and meaningful value for the unbiased crossnobis distance (see section 2.7). The RDM correlation between frequency conditions j and k then becomes

$$r_{j,k} = \frac{\sum_{i=1}^{10} d_{i,j} d_{i,k}}{\sqrt{\sum_{i=1}^{10} d_{i,j}^2 d_{i,k}^2}} \quad (2)$$

Because the RDM is a symmetric matrix, we correlated only the dissimilarities in the lower triangular of each RDM (excluding the zero diagonal values).

To compare the correlation values to a meaningful noise ceiling (see section 2.9), we split the data into odd and even runs and calculated crossnobis distances for each partition separately. We then calculated the correlation either between odd runs for frequency condition j and even runs for frequency condition k, or the other way around. The two correlations were then averaged for each participant and cross-frequency pair.

2.9. Reliability of representational geometries

Even if the representational geometry was perfectly stable across frequencies, the resultant correlations would not be 1 given the noise in our measurements. Therefore, to interpret these cross-frequency RDM correlations meaningfully, we used the reliability of each RDM to estimate a noise-ceiling for each of the 6-possible cross-frequency pairs. We measured the split-half reliability of the RDM at each frequency (r_j, r_k) using the same procedure used to calculate the cross-frequency correlations, but this time correlating the RDMs for odd and even runs within frequencies. If the true activity patterns for frequency condition j and k were identical, the expected cross-frequency correlation would be

$$E(r_{j,k}) = \sqrt{r_j^* r_k}. \quad (3)$$

This prediction therefore provides an appropriate noise ceiling for the measured cross-frequency correlations. We then scaled the measured cross-frequency correlations to their corresponding noise ceilings, such that a value of 1 indicates the RDMs of the two frequency conditions were highly similar after correcting for measurement noise. Due to the cross-validated calculation of the cross-frequency correlations, we can encounter scaled correlations >1 (where the cross-frequency correlation is larger than expected). We tested if the rescaled cross-frequency correlations were significantly lower than 1 using a one-tailed sign-test. Deviations significantly lower than 1 indicated that the cross-frequency RDM correlations were lower than expected given the reliability of each RDM.

2.10. Bayesian analysis to quantify the evidence for no RDM distortion

Given that we are also interested in quantifying the evidence for the Null-hypothesis (no distortions), we conducted a Bayesian analysis. For this, we first needed to determine what size of deviation from perfect stability would matter for model comparison, i.e. we needed to determine the effect size of the alternative hypothesis. As an example for a difficult model comparison problem, we used the muscle model and natural statistics model RDMs from Ejaz et al. (2015). Both models specify how finger movements are represented in the sensory-motor cortex. The natural statistics model hypothesizes that fingers that frequently move together evoke cortical activity patterns that are more similar. In contrast, the muscle model hypothesizes that finger movements that engage similar muscles would have a high overlap. Comparing between these models is difficult because the predicted distances from each model are highly correlated ($r = 0.9$).

As a slightly easier model comparison, we used the contrast between the somatotopic model and the natural statistics model (Ejaz et al., 2015). The somatotopic model hypothesizes that cortical activity patterns for single finger movements are arranged in an orderly fashion along the central sulcus, with some overlap between neighboring fingers. The RDMs of the somatotopic and natural statistics model were only moderately correlated ($r = 0.68$).

We simulated 1000 RDMs under each of the three models, then distorted the distances for each simulated RDM by increasing or decreasing each distance by a specific percentage of the true value. At each level of simulated distortion, we calculated the average correlation with the true model. We also determined whether the resulting RDM was closer to the true model or to the competing model with a Pearson correlation with a fixed intercept. Each misclassification was counted as a *model confusion*. Model RDMs that are more distinct would result in lower confusion rates at the same levels of distortion.

Next, we evaluated the probability of the observed cross-frequency RDM correlations from each ROI under the Null-hypothesis (no distortion, corrected $r = 1$), and various levels of the alternative hypothesis (some distortion, corrected $r < 1$). As a distribution of the differences, we assumed a Gaussian distribution with a mean of 1 and the empirical standard deviation (σ) of the corresponding cross-frequency correlations:

$$H_0 \sim N(1, \sigma) \quad (4)$$

For the alternative models, we shifted the Gaussian to be centered on the average correlation for each distortion level from our simulation (r_D):

$$H_D \sim N(r_D, \sigma) \quad (5)$$

To obtain a Bayes factor, we averaged the empirical differences in correlations for each participant across all non-neighboring frequency pairs, and then evaluated the probability of this group data under each of the alternative models. A Bayes factor ≥ 3 is considered to be positive evidence for the hypothesis in question, and factors ≥ 20 indicate strong positive evidence (Raftery and Kass, 1995). This approach allowed us to determine the distortion levels for which we observe positive evidence in favor of the alternative model (i.e. the data is probable under this level of distortion) and the levels at which we have positive evidence in favor of the null model (i.e. the data is more likely under the null model).

3. Results

We measured cortical BOLD activity patterns as participants saw digits flashed repeatedly on the screen (Fig. 1b), and made short, isometric presses of each finger of the right hand (Fig. 1a). We systematically increased the stimulation and pressing frequencies (Fig. 1b) to increase the overall activity in the visual and motor regions. Our main question was whether the representational geometry (the collection of relative dissimilarities between different conditions) would remain relatively stable across a large range of overall activity. Behaviorally, participants were able to follow the pacing and to perform the task with relatively high accuracy and matched forces (see Table 1).

Table 1. Mean and (in parentheses) between-subject standard error of behavioral measure of the finger pressing task. The pressing frequency is

reported in Hertz (Hz), and forces in Newtons (N). Participants (n = 8) were able to approximately match the instructed frequency and keep the pressing forces relatively stable.

Fig. 2a shows a surface representation of the activity patterns in the left hemisphere hand region of primary motor and sensory cortices (M1 and S1) from one participant (see supplemental figures for V1/V2 activity maps from the same participant, and maps from another participant). The overlapping nature of the activity patterns for the different fingers is clearly visible, as well as an overall gradient with the thumb activating more ventral and the little finger more dorsal aspects of the hand region. The overall intensity of the activity increased with increasing pressing frequency, but the spatial distribution of the activity associated with each finger movement appeared to remain stable. Overall, we observed an approximately 5-fold activity increase in M1 and S1 (Fig. 2b), with less activity evoked in visual cortices (V1/V2).

Our main interest was whether intensity-dependent variations in neural-to-BOLD coupling would lead to distortions of the representational geometry. Therefore, to assess this, we calculated the dissimilarity between pairs of activity patterns for each condition of the same frequency (i.e. the representational geometry), for each ROI in each participant. We then examined how stable this geometry remained despite the nearly 5-fold increase in overall activity. As a first step, we visualized the group average representational geometry in M1 (S1 representational geometries are very similar to those in M1: see Ejaz et al., 2015) and V1/V2 using a multi-dimensional scaling plot.

For M1 (Fig. 3a), we observed the expected representational geometry with the thumb having the most unique pattern and the other fingers being arranged according to their neighborhood relationship (Ejaz et al., 2015). As pressing frequency increased, this arrangement scaled up and substantially moved away from resting baseline (Fig. 3a, cross), but the overall geometry remained the same. This stability can also be appreciated when visualizing the 10 pairwise dissimilarities between fingers for each frequency (Fig. 4a, e).

In the visual cortices (V1/V2), we observed a distinct arrangement of the conditions, with the representational geometry relating to the letters and colors presented for each finger. Similarly to M1 and S1, this structure scaled up with increasing stimulation frequency (Fig. 3b). However, the lowest stimulation frequency was not very successful in eliciting either average activity or very reliable activity patterns (see below).

To quantify the stability of the representational geometry across different levels of activity, we correlated the dissimilarities in M1, S1, and V1/V2 (Fig. 4a, e, & i) across frequencies within each participant. The average cross-frequency correlations (Pearson correlation without intercept-see equation (2)) was r = 0.92 in M1 (Fig. 4b) and r = 0.94 in S1 (Fig. 4f). In V1/V2 (Fig. 4j), cross-frequency correlations were lower when they involved the lowest stimulation frequency (average of all cross-frequency correlations was 0.69), but increased with stimulation frequency (the average cross-frequency correlation between the RDMs of the two highest frequencies was 0.95).

Given measurement noise in the data, however, the cross-frequency correlations are expected to be < 1 even if they are perfectly stable. Therefore, to quantify the stability of the representational geometry, we calculated a noise-ceiling (the expected correlation if the true patterns were identical across frequencies-see section 2.9) for each cross-frequency pair. We first determined the split-half reliabilities of the RDMs within each frequency. In M1 and S1, the average split-half reliabilities across participants (Fig. 4c,g) was high (r > 0.88) for all pressing frequencies. In V1/V2 (Fig. 4k), the reliabilities of the dissimilarities measured at the slowest stimulation frequencies were lower, likely due to low levels of evoked activity, but increased comparably to reliabilities measured for M1 and S1 at higher frequencies (average split-half reliabilities across participants in V1/V2 for all stimulation frequencies was 0.78). We then estimated the noise-ceilings for each cross-frequency pair by calculating the geometric means of pairs of split-half reliabilities for different frequencies within each participant (see section 2.9).

Comparison between the measured and the expected (noise ceiling) correlations confirmed that the representational geometries remained as stable as could be expected based on the level of measurement noise across a broad range of overall activities. The right-most column in Fig. 4 shows the measured cross-frequency correlations as a ratio of their respective noise-ceilings. Values < 1 indicate that the measured correlations were lower than expected given an estimate of measurement noise for that cross-frequency pair, whereas values > 1 indicate the opposite. In V1/V2, one-tailed sign-tests indicated the measured cross-frequency correlations did not significantly differ from their estimated noise ceilings (p-values evaluated without corrections for multiple comparisons). In M1, only the measured correlations for the lowest and third highest frequency (Fig. 4d, pair 1 vs. 3) deviated significantly from the expected correlations ($p = 0.035$). In S1, the only significant deviations were for frequency condition pair 2 vs. 4 (Fig. 4h, $p = 0.004$). Although these deviations in M1 and S1 are statistically significant, the magnitude of these deviations were minor (average rescaled cross-frequency correlations in M1 and S1 = 0.97). More importantly, the correlations between the RDMs measured at the lowest and highest activity level were not significantly different from the noise ceiling estimates.

To assess whether these relatively minor distortions would be of practical relevance, we assessed their influence on model inference through simulations. For different levels of distortions, we determined the model confusion rate for a difficult model comparison problem (model RDMs are correlated with r = 0.9, see section 2.10), and a moderate model comparison problem (r = 0.68). As can be seen in Fig. 5a, the model confusion rate increased with increasing levels of RDM distortion. However, at the observed distortion rates (dashed lines in Fig. 5a), the confusion rates were below 12%, even for the difficult model inference.

To quantify the statistical evidence that our data shows no distortion (the Null-hypothesis), we calculated the Bayes Factor (BF, see section 2.10) of the Null against distortions of various sizes. Because we were interested in distortions that would arise due to intensity-dependent variations in neural-to-BOLD coupling, we analyzed the empirical differences between 1 and the cross-frequency correlations (corrected to their respective noise-ceiling) for all non-neighboring cross-frequency pairs (averaged across pairs within each participant). Fig. 5b shows the log Bayes factor for various levels of distortion for the three ROI. We obtained positive evidence ($|BF| \geq 3$) in favor of the null at a distortion of r = 0.95 for M1 and for r = 0.93 for S1, which were associated with a model confusion rate of 5–17% and 9–19%. For V1/V2, we excluded the lowest stimulation frequency, as we could not measure the representational structure reliably here. We observed positive evidence against a higher distortion level (r = 0.91) corresponding to a confusion range of

Table 1
Finger pressing behavior during MRI data acquisition.

Cued freq.	Pressing freq. (Hz)	Force (N)	Pressing freq. (Hz)	Force (N)	Pressing freq. (Hz)	Force (N)	Pressing freq. (Hz)	Force (N)
	0.3 Hz		0.6 Hz		1.3 Hz		2.6 Hz	
Thumb	0.33 (0.02)	3.04 (0.08)	0.66 (0.03)	3.12 (0.05)	1.32 (0.04)	3.18 (0.07)	2.65 (0.10)	3.13 (0.07)
Index	0.32 (0.03)	2.82 (0.07)	0.65 (0.04)	2.91 (0.07)	1.33 (0.02)	2.99 (0.05)	2.67 (0.21)	2.97 (0.07)
Middle	0.34 (0.02)	3.62 (0.16)	0.67 (0.02)	3.53 (0.14)	1.34 (0.03)	3.52 (0.14)	2.70 (0.14)	3.39 (0.11)
Fourth	0.33 (0.04)	2.94 (0.08)	0.67 (0.04)	2.99 (0.05)	1.34 (0.03)	3.07 (0.05)	2.63 (0.24)	2.98 (0.07)
Little	0.33 (0.02)	2.99 (0.06)	0.65 (0.04)	2.99 (0.05)	1.32 (0.05)	3.11 (0.04)	2.66 (0.21)	2.97 (0.05)

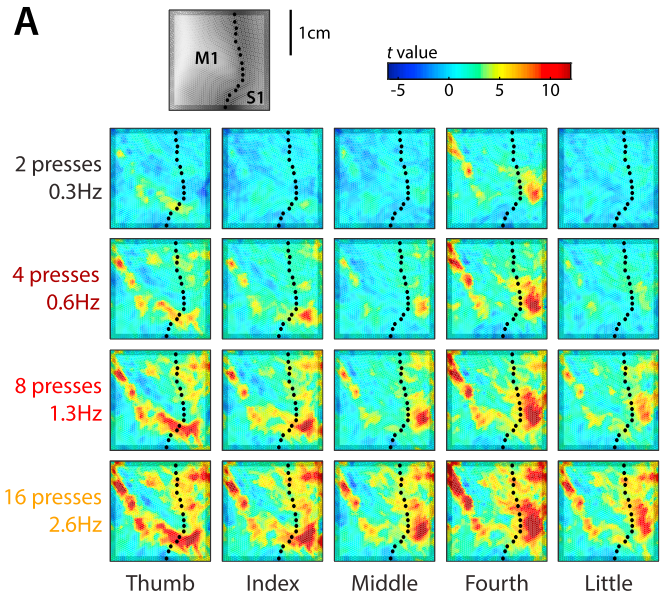

Fig. 2. Scaling of activity patterns. **A**) BOLD activity patterns from the hand area of the primary sensorimotor cortices of an example participant projected onto a flat, surface reconstruction of their cortex. Dotted lines indicate the fundus of the central sulcus. The top insert reflects sulcal depth (darker colors reflect larger depths) and denotes location of M1 and S1. Color maps reflect t-values of activity against rest. Each column corresponds to one finger (thumb to little), and each row one pressing frequency (0.3–2.6 Hz). The activity increases with increasing pressing frequency. **B**) Activity (beta-coefficients) in contralateral M1, S1, and bilateral V1/V2 as a function of pressing/stimulation frequency. Data from each participant are plotted. The group average for each region are plotted in bold. Error bars reflect s.e.m.

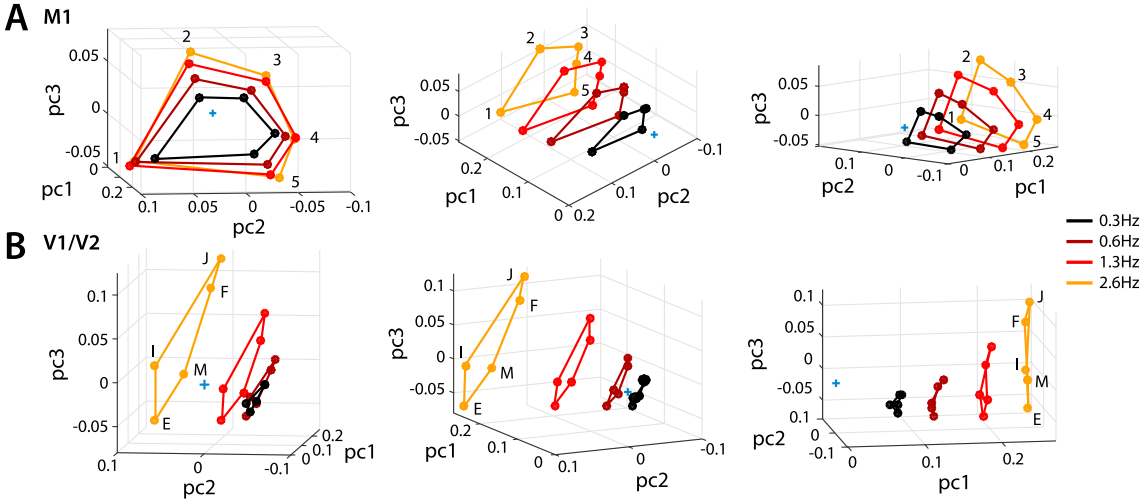

Fig. 3. Multidimensional scaling of group average representational geometries for M1 (**A**) and bilateral V1/V2 (**B**). Colors correspond to pressing/stimulation frequency. Numbers indicate fingers (1 = thumb, 5 = little finger). Letters correspond to the flashed letter cue (1 = E, 2 = I, 3 = M, 4 = F, 5 = J). Note that the conditions are connected differently for M1 and V1/V2. Panels in the same row present different views of the same space, defined by the first three eigenvectors. The blue cross indicates baseline.

approximately 12–28%. Evidence in favor of the Null hypothesis for these high distortion levels indicates the true distortion is substantially lower. Indeed, the average of the non-neighboring scaled cross-frequency correlations in M1, S1, and V1/V2 were close to 1 (M1: r = 0.98, S1: r = 0.97, V1/V2: r = 0.97). The evidence for the alternative hypothesis at these levels of distortion was weak in M1 (|BF| = 2.8) and positive in S1 (|BF| = 9). The associated model confusion rates were 0–6% in M1 and 2–12% in S1, respectively. We did not observe positive evidence for RDM distortions at any level in V1/V2.

Together, these results demonstrate that the relationship between crossnobis dissimilarities remains relatively stable across a broad range of overall activity in sensorimotor cortices, and to a lesser extent in and primary and secondary visual cortices.

4. Discussion

Here we critically investigated whether there are measurable distortions of the representational geometry as measured by BOLD as overall neural activity increases. We tested this assumption by stimulating both sensory-motor and visual regions at increasing frequencies. We assumed that on the neural level, each repeated event should elicit approximately the same activity pattern. Therefore, the temporally integrated patterns should scale in an orderly fashion across frequencies. Importantly, this does not imply that neural activity would increase linearly with the number of events. Indeed, previous findings have shown that the neural response to subsequent finger taps is strongly attenuated in M1 (Hermes et al., 2012). However, as pointed out in the introduction, this non-linearity between behavioral and average BOLD does not provide insight into whether the representational geometry would distort.

Indeed, we found that in M1 and S1 the representational geometry scales in a relatively orderly manner, even though the local activity increased over a large dynamic range – likely close to the achievable maximum for this paradigm. The drop in correlation across pressing frequencies (as compared to the noise ceiling derived from within-frequency consistency) was minor (average drop = 0.03). It also needs to be kept in mind that some of the observed distortions may have been

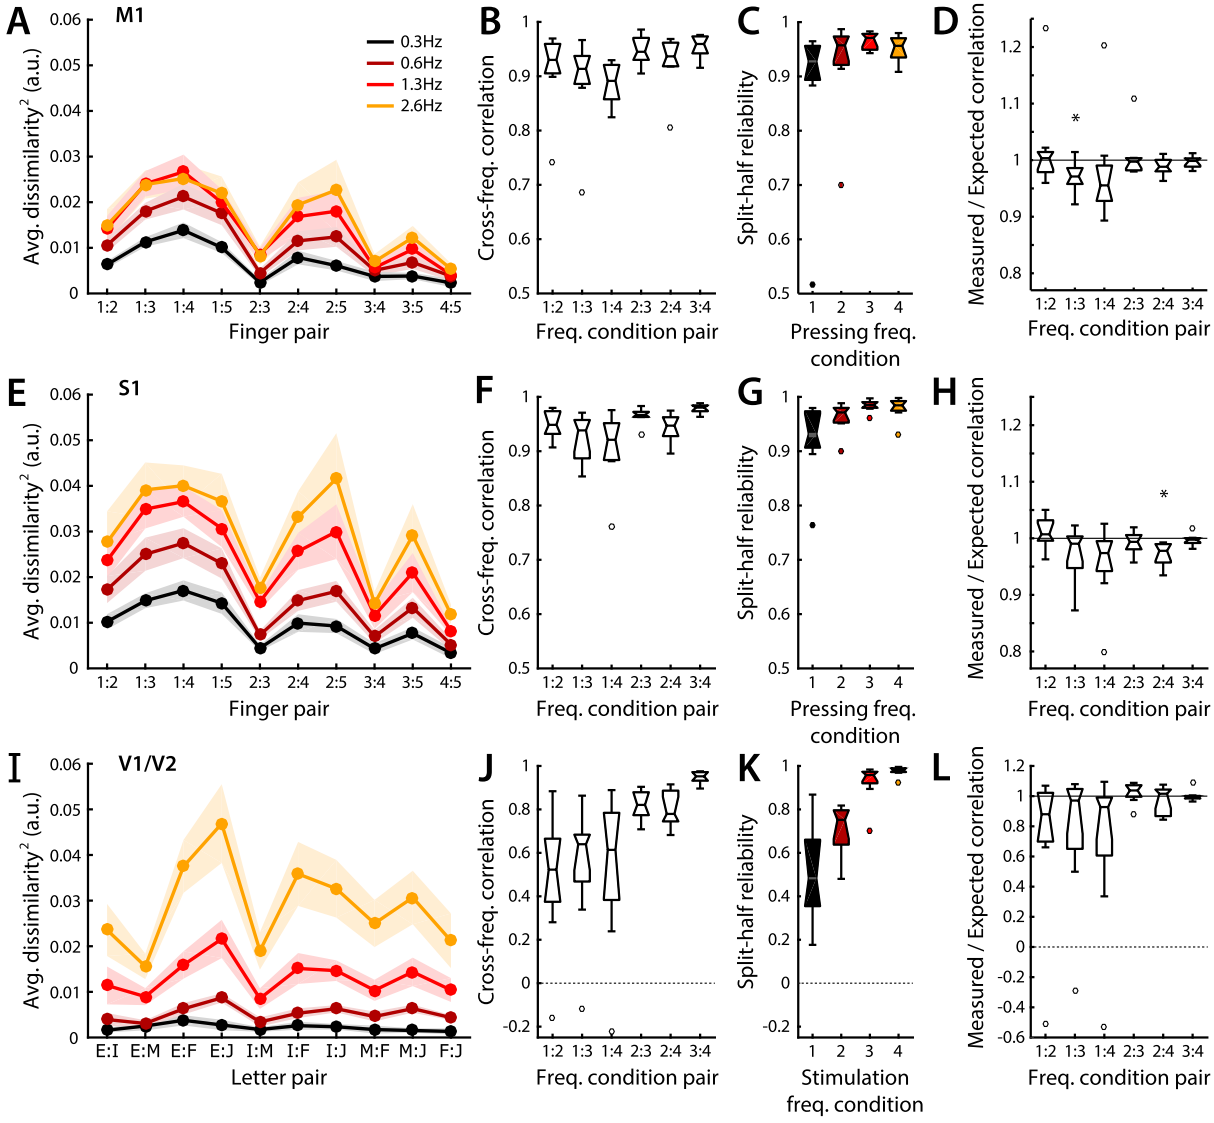

Fig. 4. Stability of representational geometry across stimulation frequencies. **A**) Average dissimilarity between all possible pairs of the five activity patterns measured at each frequency in M1. Colors indicate pressing/stimulation frequency, and shaded regions reflect s.e.m. **B**) Cross-frequency correlations (equation (2)) between dissimilarities depicted in **A**. The box plot extends from the 25th to 75th percentiles, and whiskers extend to the full range of the data. The line in the center of each box denotes the median. Outliers (circles) are points more than 1.5 times the box length away from the median. **C**) Within participant split-half reliabilities (Pearson correlation with a forced intercept) of the dissimilarities depicted in **A**. **D**) Measured cross-frequency correlations from **B** scaled by their respective expected correlations (noise ceiling, see equation (3)). Values < 1 indicate the observed correlation (stability) is lower than what would be expected given the internal reliability of each RDM. Deviations from 1 were evaluated with one-tailed signed rank tests. Asterisks indicate significant deviations ($p < 0.05$). The outlier data in M1 does not drive this result. **E-H**) Results for S1. **I-L**) Results for V1/V2. Note that the representational geometry in visual regions is different from the one found in M1/S1, reflecting the finger-to-letter assignment. Due to low stimulation intensity, visual regions have lower reliabilities for low stimulation frequencies.

due to real changes in the representational geometry of the neural activity patterns - after all our assumption that the neural representational geometry would be completely stable may not be true. Thus, our results provide an upper bound for distortions that can be attributed the neural-to-BOLD coupling. Our ability to identify the correct representational model at this level of distortion remained good, even for a difficult model comparison.

Our results in V1/V2 were slightly weaker. While the measured distortions were at a similar level as in M1/S1, the variability of our RDM estimates were much higher in than in sensorimotor regions, especially for lower stimulation frequencies. This likely reflects limitations in our task design, which succeeded in driving the overall activity level in M1 and S1 across a large range, but was not optimal to elicit maximal activity in visual regions.

Overall, our findings provide an important extension of previous studies that report linear coupling between neural activity and BOLD responses in M1/S1 (Siero et al., 2013) and V1 (Boynton et al., 1996; Heeger et al., 2000). Showing that BOLD signal and neural activity are linearly coupled *on average* does not guarantee that the fine-grained activity patterns would also retain their representational geometry. For this, every individual voxel would have to obey the same scaling function. Testing this commonly-held assumption, we provide here the strongest empirical evidence to date that the dissimilarities of multivariate fMRI activity patterns are stable across a wide range of activation levels. Our findings are broadly consistent with a PCM-style analysis of the same data (Diedrichsen et al., 2018), which already showed the stability of the representational structure on the group level in M1. The current paper extends these results by quantifying the stability on the individual level, in multiple brain regions, and by estimating the effect the distortions may have on model comparison.

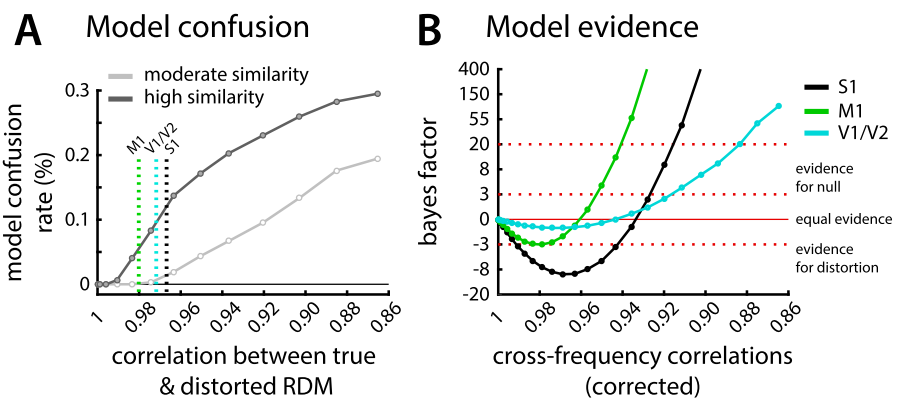

Fig. 5. The effect of RDM distortions on model comparisons. **A**) Simulated model confusion rates between two highly correlated (dark gray, r = 0.9) and two moderately correlated (light gray, r = 0.68) model RDMs (see section 2.10). Dotted lines reflect the average observed cross-frequency correlations (scaled to noise-ceilings) from each of the three ROIs. RDMs are more commonly misclassified as distortion increases (shown here as a decrease in the correlation between the true and distorted RDM). The magnitude of misclassification depends on how similar the two model RDMs are. **B**) Bayes factors (BF) for evidence of the null (no distortion) over the alternative (distortion) models for increasing levels of RDM distortion (x-axis). Negative BF indicate evidence in favor of the alternative model, which means that the cross-frequency RDM correlations are more probable under the corresponding level of distortion. In contrast, positive BF indicate evidence in favor of the null model, suggesting that the cross-frequency RDM correlations exhibit a lower level of distortion than the corresponding level tested. The dashed red lines indicate the thresholds for positive evidence ($|BF| \geq 3$) and strong evidence ($|BF| \geq 20$).

These findings have important implications for multivariate fMRI analyses. In general, most papers rely explicitly or implicitly on the assumption that the representational geometry as measured with fMRI veridically reflects the representational geometry of the underlying neural population code. While this strong assumption may still be violated for other reasons (see below), we at least show here that the representational geometry can be meaningfully compared across a large range of overall activation levels. Thus, under appropriate conditions, representational geometries can be compared across different regions, individuals, or patient populations. Furthermore, from our data it appears justified to compare RDMs between different attentional states or different levels of learning, even if these differ in their average activity. Given our results, changes in the representational geometry across these conditions are likely due to real changes in the underlying neural population code, rather than due to distortions induced by neural-to-BOLD coupling.

For RSA, these findings also have implications for the choice of a dissimilarity measure. Not all measures make equally strong assumptions about the relation between the underlying neural representational geometry and the one measured with fMRI. For example, a common practice in RSA is to evaluate rank-correlation between measured and predicted RDMs. Arguably, this approach makes interpretations more robust against minor distortions in the measurement process (Kriegeskorte et al., 2008; Nili et al., 2014). Our results indicate that the exact ratio-relationship between dissimilarity measures can be meaningfully interpreted across a large range of average activation states. For this to be true, we of course need to use a dissimilarity measure that provides a meaningful zero point unbiased by noise – a condition met by the crossnobis estimator (Diedrichsen et al., 2016; Kriegeskorte and Diedrichsen, 2016; Walther et al., 2016). The additional information in the exact ratio-relationships allow for more powerful inference about the underlying representations (Diedrichsen and Kriegeskorte, 2017).

The stability of the representational geometry is also good news for alternative approaches that test representational models. PCM (Diedrichsen et al., 2011; Diedrichsen et al., 2018) and encoding models (Naselaris et al., 2011) make inferences about the underlying representational geometry in a very similar way to RSA (Diedrichsen and Kriegeskorte, 2017). Therefore, PCM and encoding approaches are subject to similar assumptions as RSA – and our findings generalize, such that inferences using these two methods will also be stable across activity levels.

Do these results suggest that one can make inferences about the representational geometry of the underlying neural population code from fMRI measures? While our results are reassuring in some aspects, there are two important caveats that we did not address in the current paper. First, fMRI samples neural activity with dramatic spatial averaging, even when using sub-millimeter resolution. Representations that exist at a finer spatial scale in the neural population code will be under-represented in BOLD activity patterns, while representations at a large spatial scale will be over-represented (see Kriegeskorte and Diedrichsen, 2016). Therefore, the representational geometry of BOLD activity patterns may differ systematically from the underlying neural code.

Secondly, the physiological processes underlying the BOLD signal and underlying extracellular neural recording are fundamentally different: While the BOLD signal reflects to a large degree the metabolically expensive processes of ion transport after excitatory postsynaptic potentials (Attwell and Laughlin, 2001; Harris et al., 2012), extracellular recordings reflect neural spiking. Crudely stated, therefore, BOLD reflects more the input to a region, while neural extracellular recordings reflect the output. Additionally, most extracellular recordings are biased towards large output neurons, as these provide the clearest extracellular signal (Firmin et al., 2014; Harris et al., 2016), whereas the BOLD signal indiscriminately averages metabolic activity. These important caveats need to be kept in mind when drawing parallels between representational analysis of extracellular recordings and BOLD signal.

5. Conclusion

One common assumption in multivariate fMRI analyses is that the relationship between activity patterns can be meaningfully interpreted. Intensity-dependent variations in neural-to-BOLD coupling, however, could lead to substantial distortions of multivariate measures when overall activity increases. Our results demonstrate that, across a broad range of overall activation states in M1 and S1 the ratio-relationships between pair-wise dissimilarities remain stable. This suggests that it is viable to leverage more powerful techniques, such as the use of cross-validated dissimilarities and likelihood-based RSA (Diedrichsen et al., 2016; Diedrichsen and Kriegeskorte, 2017), for model comparison. The finding also applies to other multivariate techniques that analyze the relationship of BOLD activity pattern.

Acknowledgements

The work was supported by a NSERC discovery grant (RGPIN-2016-04890) to JD. Functional imaging was supported partly by a Platform

Support Grant from Brain Canada and the Canada First Research Excellence Fund (BrainsCAN). SA is supported by a doctoral scholarship from NSERC (PGSD3-519263-2018). We thank Nikolaus Kriegeskorte, Timothy Verstynen, and Patrick Beukma for their helpful comments on earlier drafts of the paper.

Appendix A. Supplementary data

Supplementary data to this article can be found online at https://doi.org/10.1016/j.neuroimage.2018.11.002.

References

Attwell, D., Laughlin, S.B., 2001. An energy budget for signaling in the grey matter of the brain. J. Cerebr. Blood Flow Metabol. 21 (10), 1133–1145. https://doi.org/10.1097/00004647-200110000-00001.

Boynton, G.M., Engel, S.A., Glover, G.H., Heeger, D.J., 1996. Linear systems analysis of functional magnetic resonance imaging in human V1. J. Neurosci. 16 (13), 4207–4221. Retrieved from. http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?cmd=Retrieve&db=PubMed&dopt=Citation&list_uids=8753882.

Brouwer, G.J., Heeger, D.J., 2009. Decoding and reconstructing color from responses in human visual cortex. J. Neurosci. 29 (44), 13992–14003. https://doi.org/10.1523/JNEUROSCI.3577-09.2009.

Dale, A.M., Fischl, B., Sereno, M.I., 1999. Cortical surface-based analysis. I. Segmentation and surface reconstruction. Neuroimage 9 (2), 179–194. https://doi.org/10.1006/nimg.1998.0395.

de-Wit, L., Alexander, D., Ekroll, V., Wagemans, J., 2016. Is neuroimaging measuring information in the brain? Psychonomic Bull. Rev. 23 (5), 1415–1428. https://doi.org/10.3758/s13423-016-1002-0.

deCharms, R.C., Zador, A., 2000. Neural representation and the cortical code. Annu. Rev. Neurosci. 23 (1), 613–647. https://doi.org/10.1146/annurev.neuro.23.1.613.

Diedrichsen, J., Kriegeskorte, N., 2017. Representational models: a common framework for understanding encoding, pattern-component, and representational-similarity analysis. PLoS Comput. Biol. 13 (4), e1005508. https://doi.org/10.1371/journal.pcbi.1005508.

Diedrichsen, J., Ridgway, G.R., Friston, K.J., Wiestler, T., 2011. Comparing the similarity and spatial structure of neural representations: a pattern-component model. Neuroimage 55 (4), 1665–1678. https://doi.org/10.1016/J.NEUROIMAGE.2011.01.044.

Diedrichsen, J., Provost, S., Zareamoghaddam, H., 2016. On the Distribution of Cross-validated Mahalanobis Distances. Retrieved from. http://arxiv.org/abs/1607.01371.

Diedrichsen, J., Yokoi, A., Arbuckle, S.A., 2018. Pattern component modeling: a flexible approach for understanding the representational structure of brain activity patterns. Neuroimage 180, 119–133. https://doi.org/10.1016/J.NEUROIMAGE.2017.08.051.

Ejaz, N., Hamada, M., Diedrichsen, J., 2015. Hand use predicts the structure of representations in sensorimotor cortex. Nat. Neurosci. 18 (7). https://doi.org/10.1038/nn.4038.

Firmin, L., Field, P., Maier, M.A., Kraskov, A., Kirkwood, P.A., Nakajima, K., et al., 2014. Axon diameters and conduction velocities in the macaque pyramidal tract. J. Neurophysiol. 112 (6), 1229–1240. https://doi.org/10.1152/jn.00720.2013.

Fischl, B., Sereno, M.I., Tootell, R.B., Dale, A.M., 1999. High-resolution intersubject averaging and a coordinate system for the cortical surface. Hum. Brain Mapp. 8 (4), 272–284. Retrieved from. http://www.ncbi.nlm.nih.gov/pubmed/10619420.

Fischl, B., Rajendran, N., Busa, E., Augustinack, J., Hinds, O., Yeo, B.T.T., et al., 2008. Cortical folding patterns and predicting cytoarchitecture. Cerebr. Cortex 18 (8), 1973–1980. https://doi.org/10.1093/cercor/bhm225.

Harris, J.J., Jolivet, R., Attwell, D., 2012. Synaptic energy use and supply. Neuron 75 (5), 762–777. https://doi.org/10.1016/j.neuron.2012.08.019.

Harris, K.D., Quiroga, R.Q., Freeman, J., Smith, S.L., 2016. Improving data quality in neuronal population recordings. Nat. Neurosci. 19 (9), 1165–1174. https://doi.org/10.1038/nn.4365.

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 (5539), 2425–2430. https://doi.org/10.1126/science.1063736.

Haxby, J.V., Connolly, A.C., Guntupalli, J.S., 2014. Decoding neural representational spaces using multivariate pattern analysis. Annu. Rev. Neurosci. 37 (1), 435–456. https://doi.org/10.1146/annurev-neuro-062012-170325.

Heeger, D.J., Huk, A.C., Geisler, W.S., Albrecht, D.G., 2000. Spikes versus BOLD: what does neuroimaging tell us about neuronal activity? Nat. Neurosci. 3 (7), 631–633. https://doi.org/10.1038/76572.

Hermes, D., Siero, J.C.W., Aarnoutse, E.J., Leijten, F.S.S., Petridou, N., Ramsey, N.F., 2012. Dissociation between neuronal activity in sensorimotor cortex and hand movement revealed as a function of movement rate. J. Neurosci. 32 (28), 9736–9744. https://doi.org/10.1523/JNEUROSCI.0357-12.2012.

Kriegeskorte, N., Diedrichsen, J., 2016. Inferring brain-computational mechanisms with models of activity measurements. Philos. Trans. R. Soc. Lond. Ser. B Biol. Sci. 371 (1705), 20160278. https://doi.org/10.1098/rstb.2016.0278.

Kriegeskorte, N., Mur, M., Bandettini, P., 2008a. Representational similarity analysis - connecting the branches of systems neuroscience. Front. Syst. Neurosci. 2, 4. https://doi.org/10.3389/neuro.06.004.2008.

Kriegeskorte, N., Mur, M., Ruff, D.A., Kiani, R., Bodurka, J., Esteky, H., et al., 2008b. Matching categorical object representations in inferior temporal cortex of man and monkey. Neuron 60 (6), 1126–1141. https://doi.org/10.1016/j.neuron.2008.10.043.

Miyawaki, Y., Uchida, H., Yamashita, O., Sato, M., Morito, Y., Tanabe, H.C., et al., 2008. Visual image reconstruction from human brain activity using a combination of multiscale local image decoders. Neuron 60 (5), 915–929. https://doi.org/10.1016/J.NEURON.2008.11.004.

Naselaris, T., Kay, K.N., Nishimoto, S., Gallant, J.L., 2011. Encoding and decoding in fMRI. Neuroimage 56 (2), 400–410. https://doi.org/10.1016/j.neuroimage.2010.07.073.

Nili, H., Wingfield, C., Walther, A., Su, L., Marslen-Wilson, W., Kriegeskorte, N., 2014. A toolbox for representational similarity analysis. PLoS Comput. Biol. 10 (4), e1003553. https://doi.org/10.1371/journal.pcbi.1003553.

O'Herron, P., Chhatbar, P., Levy, M., Shen, Z., Schramm, A., Lu, Z., Kara, P., 2016. Neural correlates of single-vessel haemodynamic responses in vivo. Nature 534, 378–394. https://doi.org/doi:10.1038/nature17965.

Raftery, Adrian E., Kass, Robert E., 1995. Bayes factors. J. Am. Stat. Assoc. 90 (430), 773–795. Retrieved from. http://www.stat.washington.edu/people/raftery/Research/PDF/kass1995.pdf.

Siero, J.C., Hermes, D., Hoogduin, H., Luijten, P.R., Petridou, N., Ramsey, N.F., 2013. BOLD consistently matches electrophysiology in human sensorimotor cortex at increasing movement rates: a combined 7T fMRI and ECoG study on neurovascular coupling. J. Cerebr. Blood Flow Metabol. 3397 (10), 1448–1456. https://doi.org/10.1038/jcbfm.2013.97.

Walther, A., Nili, H., Ejaz, N., Alink, A., Kriegeskorte, N., Diedrichsen, J., 2016. Reliability of dissimilarity measures for multi-voxel pattern analysis. Neuroimage 137, 188–200. https://doi.org/10.1016/j.neuroimage.2015.12.012.

Wiestler, T., Diedrichsen, J., 2013. Skill learning strengthens cortical representations of motor sequences. ELife 2013 (2). https://doi.org/10.7554/eLife.00801.

Yousry, T.A., Schmid, U.D., Alkadhi, H., Schmidt, D., Peraud, A., Buettner, A., Winkler, P., 1997. Localization of the motor hand area to a knob on the precentral gyrus. A new landmark. Brain 120 (1), 141–157. https://doi.org/10.1093/brain/120.1.141.