

Received 10 March 2019; revised 12 July 2019; accepted 15 August 2019. Date of publication 20 August 2019; date of current version 30 August 2019.

Digital Object Identifier 10.1109/JTEHM.2019.2936348

Locally Linear Embedding and fMRI Feature Selection in Psychiatric Classification

GAGAN SIDHU 

Department of Computing Science, 1-337 Athabasca Hall, University of Alberta, Edmonton, AB T6G 2E8, Canada
This paper has supplementary downloadable material available at <http://ieeexplore.ieee.org>, provided by the author.
(e-mail: gagan@g-a.ca)

ABSTRACT

Background:

Functional magnetic resonance imaging (fMRI) provides non-invasive measures of neuronal activity using an endogenous Blood Oxygenation-Level Dependent (BOLD) contrast. This article introduces a nonlinear dimensionality reduction (Locally Linear Embedding) to extract informative measures of the underlying neuronal activity from BOLD time-series. The method is validated using the Leave-One-Out-Cross-Validation (LOOCV) accuracy of classifying psychiatric diagnoses using resting-state and task-related fMRI.

Methods:

Locally Linear Embedding of BOLD time-series (into each voxel's respective tensor) was used to optimise feature selection. This uses Gauß' Principle of Least Constraint to conserve quantities over both space and time. This conservation was assessed using LOOCV to greedily select time points in an incremental fashion on training data that was categorised in terms of psychiatric diagnoses.

Findings:

The embedded fMRI gave highly diagnostic performances (> 80%) on eleven publicly-available datasets containing healthy controls and patients with either Schizophrenia, Attention-Deficit Hyperactivity Disorder (ADHD), or Autism Spectrum Disorder (ASD). Furthermore, unlike the original fMRI data before or after using Principal Component Analysis (PCA) for artefact reduction, the embedded fMRI furnished significantly better than chance classification (defined as the majority class proportion) on ten of eleven datasets.

Interpretation:

Locally Linear Embedding appears to be a useful feature extraction procedure that retains important information about patterns of brain activity distinguishing among psychiatric cohorts.

INDEX TERMS Nonlinear, dimensionality reduction, image processing, machine learning, kernel methods, optimization, least squares, neurophysiology, evidence-based medicine, (computer-assisted) diagnosis, fMRI, method of image charges, integration, oscillations, theorema egregium.

I. INTRODUCTION

Over a century ago, Charles Darwin alluded to an experimental paradigm that involved direct observation of the brain's physical mechanisms (nervous matter) [14], where such observations [55] would serve as the physical basis for dichotomising species-specific behaviour. In the early 90s, an indirect and non-invasive measurement of mental activity over uniformly-spaced time points became possible through functional Magnetic Resonance Imaging (fMRI), which allows paramagnetic deoxyhemoglobin to act as an endogenous Blood Oxygenation-Level Dependent (BOLD) contrast [49]. This article pursues Darwin's proposed experimental paradigm by Locally Linear Embedding (LLE) [54] the BOLD time-series to produce precise summaries of

cerebral activity that may optimise the classification of different brain states, such as mental disorders.

Logothetis *et al* found that an increase in invasively-measured neural activity directly and monotonically reflects local BOLD signal increases and, for short stimulus presentations, there is a linear relationship between BOLD and neural responses [39]. This suggests the unobservable neural activity is spatially-localised in anatomical space. Locally Linear Embedding of fMRI data in space and time can, in principle [2], [20], [24] (see Appendix IV), summarise these *local* measurements of neuronal mass activity [40]—via the notion of *analytic capacity* [45]—to disclose information about *global* (i.e., whole-brain) activity patterns.

II. METHODS: TESTING LOCALLY LINEAR EMBEDDING (LLE) WITH CROSS VALIDATION

The discriminatory power of Locally Linear Embedding was compared to the original fMRI, both before and after applying Principal Component Analysis (PCA) [31] for artefact reduction [39]. This comparison was performed on eleven datasets containing cohorts with different mental disorders, using a combination of Leave One Out Cross Validation (LOOCV) on the training set, with greedy feature selection based on Fisher discriminability [17], [18]. The purpose of using LOOCV and feature selection is to find the time points that discriminate patients from controls on the respective dataset. The feature selection step initially starts with an empty candidate set of time points and proceeds to select the time point with the highest discriminatory power [63]. Then, time points that improve discrimination in conjunction with those already in the candidate set are added to this set in incremental fashion; this process is terminated when there are no more time points that can be added to the candidate set to improve discriminability. Note that the selection of time points is based upon cross validation and does not induce any biased sampling.

To illustrate the patterns that best discriminate between groups, a paired two-sample t-test between the patient and control groups is performed to both threshold and identify the statistically-significant differences ($p < 0.05$ uncorrected) in space (at the time points identified by the greedy feature selection). Per Mill's Methods of Induction (Method of difference), the functional differences (depicted by the statistically significantly-different regions) at the respective time point are therefore a necessary part of the cause of the phenomena that distinguish the subject groups [46], which in this case pertain to a neuropsychiatric disorder. Neuropsychiatric disorders are diagnosed using clinical assessments that include: evaluating the background demographics, collecting first and third party observations, and a structured psychiatric interview with the subject [53]. In detail:

Every subject's fMRI time-series is treated as a four-dimensional array $\mathbf{X} \in \mathbb{R}^{L \times W \times H \times T}$ with $V = LWH$ T -dimensional voxel waveforms $\mathbf{x}_i \in \mathbb{R}^T$ for $i = 1, \dots, V$. Assume each subject scan \mathbf{X}_i is associated with a binary-valued class label y_i representing the diagnosis and that, for any subject scan, every voxel waveform $\mathbf{x} \in \mathbb{R}^T$ is generated by a vector $\mathbf{z} \in \mathbb{R}^d$ corresponding to a point on the manifold. Our approach to fMRI-based diagnosis involves two stages:

fMRI reconstruction takes the subject's fMRI as input and outputs a reconstructed fMRI that is more informative than the original. Formally, this reconstruction is a mapping $\mathcal{L} : \mathbb{R}^{L \times W \times H \times T} \rightarrow \mathbb{R}^{L \times W \times H \times d}$.

In other words, Locally Linear Embedding reduces a time-series of length T to a smaller number of spatial modes of dimensionality d ; these modes contain all the information used for the subsequent step.

Classification builds a classifier that takes the subject's reconstructed fMRI as input, and outputs a class label

$y_i \in \{0, 1\}$. The classifier is therefore a mapping $\mathcal{C} : \mathbb{R}^{L \times W \times H \times d} \rightarrow \{0, 1\}$.

The reconstructed, or reduced, fMRI data produced from step 1 is hereon referred to as \mathbf{Z} . All reconstructions initially vectorise the fMRI data to produce a two-dimensional array $\mathbf{X} \in \mathbb{R}^{V \times T}$, and conclude by reshaping the resulting two-dimensional reconstruction $\mathbf{Z} \in \mathbb{R}^{V \times d}$ into a four-dimensional array $\mathbf{Z} \in \mathbb{R}^{L \times W \times H \times d}$.

a) *Principal Component Analysis (PCA)* [5], [31]: reconstructs \mathbf{X} by finding an orthogonal rotation that minimises the reconstruction cost

$$\min_{\mathbf{B}=[\mathbf{b}_1, \dots, \mathbf{b}_d]} \sum_{i=1}^V \|\mathbf{x}_i - \bar{\mathbf{X}} - \mathbf{B}\mathbf{z}_i\|_2^2$$

where $\bar{\mathbf{x}} \in \mathbb{R}^T$ is the mean over all voxel waveforms, and $\mathbf{z}_i = \mathbf{B}^T(\mathbf{x}_i - \bar{\mathbf{X}}) \in \mathbb{R}^d$. To find the optimal \mathbf{B} , compute the right-hand matrix for the singular value decomposition (SVD) of $\mathbf{X} = \mathbf{U}\mathbf{D}\mathbf{B}^T$, which contains the T right-singular vectors of \mathbf{X} [26]. It follows that $\mathbf{Z} = \mathbf{X}\mathbf{B} \in \mathbb{R}^{V \times d}$ is the rotated matrix that minimises the reconstruction cost of the subject's fMRI, where every column of \mathbf{Z} is a *principal component*. It is assumed the first d principal components (columns) of \mathbf{Z} capture the "systematic structure," where the confounding factors are relegated to the remaining $T - d$ principal components, which produces the two-dimensional reconstructed fMRI $\mathbf{Z} \in \mathbb{R}^{V \times d}$.

Using PCA's "systematic structure" for distinguishing humans with different neurological disorders has been met with caution [16], largely because PCA's application to fMRI has some subjective components [1]. The main limitation behind the PCA reconstruction is that it assumes the lower-dimensional manifold is a linear subspace. We demonstrate that introducing the Cauchy stress tensor [9] on the Cartesian space with the Pythagorean distance metric enables three-dimensional measurements over time, thereby revealing the local (group) action in the physical system.

b) *Locally Linear Embedding (LLE)* [54], [56]: reconstructs \mathbf{X} by constructing the Cauchy stress tensor [9] at every voxel i for $i = 1, \dots, V$, which is achieved by minimising the reconstruction cost of its waveform \mathbf{x}_i in terms of its spatially-adjacent neighbours:

$$\min \|\mathbf{x}_i - \sum_{j \in \mathcal{N}(i)} w_{i,j} \mathbf{x}_j\|_2^2, \quad \text{where} \quad \sum_{j \in \mathcal{N}(i)} w_{i,j} = 1 \quad (1)$$

where the neighbourhood set $\mathcal{N}(i)$ for voxel i is the complement of its K spatial neighbours on the surface of the sphere with radius r , and $\mathbf{w}_i = [w_{i,1}, \dots, w_{i,|\mathcal{N}(i)|}] \in \mathbb{R}^{|\mathcal{N}(i)|}$ are the reconstruction weights containing the spatially-invariant geometric properties of the Cauchy stress tensor at voxel i .

To determine the *analytic capacity* [45] at every voxel location i , LLE first subtracts the K spatial patterns of the voxels on the boundary of the sphere centred around voxel i to determine the separation distance from the origin of the tensor at the respective voxel. Then, it computes the

local (symmetric) spatiotemporal covariance matrix:

$$\begin{aligned}
 \mathbf{G}_i &= \mathbf{C}_i^T \mathbf{C}_i \\
 &= [(\mathbf{x}_j - \mathbf{x}_i), \dots, (\mathbf{x}_{j+|\mathcal{N}(i)|} - \mathbf{x}_i)]^T [(\mathbf{x}_j - \mathbf{x}_i), \dots, (\mathbf{x}_{j+|\mathcal{N}(i)|} - \mathbf{x}_i)] \\
 &\quad + \xi \mathbf{I}_{|\mathcal{N}(i)|} \quad (2)
 \end{aligned}$$

where $\xi \mathbf{I}_{|\mathcal{N}(i)|}$ is a non-negative regularisation term to enforce positive-definiteness (for this study, $\xi = 0$). LLE calculates the reconstruction weights by finding the unique minimum-norm solution [50] to the constrained least-squares problem defined by:

$$\mathbf{G}_i \mathbf{w}_i = \mathbf{1}_{|\mathcal{N}(i)|} \iff \mathbf{G}_i^+ \mathbf{1}_{|\mathcal{N}(i)|} = \mathbf{w}_i \quad (3)$$

where $\mathbf{1}_{|\mathcal{N}(i)|} \in \mathbb{R}^{|\mathcal{N}(i)|}$ is a vector of ones and the j^{th} element of $\mathbf{w}_{i,j}$ can be thought of as the average height of a curve representing the *mean transit time of the indicator* [34] of voxel j from voxel i over the duration of the scan. Since \mathbf{G}_i represents the squared distance of the surface forces from voxel i , the reconstruction weights \mathbf{w}_i are Lebesgue measures [36] summarising the *analytic capacity*, or spatially-invariant geometry [25], of the space-filling curve [27], where the constraint ensures that the areas between the imaginary surface (acting as the origin that divides the body) and curves (defined by the stress vectors) are 1 in each of the $|\mathcal{N}(i)|$ directions. In practice the weights can be brittle [65] due to any number of reasons. Modified Locally Linear Embedding (MLLE) therefore computes the $1 \leq s_i \leq K$ linearly-independent (orthogonal) vectors¹ $\mathbf{Q}_i \in \mathbb{R}^{K \times K}$ of \mathbf{G}_i using the eigendecomposition $\mathbf{G}_i = \mathbf{Q}_i^T \mathbf{A} \mathbf{Q}_i$, thereby allowing the definition of multiple weight vectors for each voxel. Assuming the columns (eigenvectors) $[\mathbf{q}_1, \dots, \mathbf{q}_K] = \mathbf{Q}_i$ are sorted in descending order of their respective eigenvalues $\lambda_1^{(i)}, \dots, \lambda_K^{(i)}$, MLLE uses the first s_i columns to compute multiple local weight vectors for a single voxel:

$$\begin{aligned}
 \mathbf{w}_i^{(\ell)} &= (1 - \alpha_i) \mathbf{w}_i + \mathbf{Q}_i \mathbf{H}_i^{(\ell)} \quad \text{where: } \alpha_i = \frac{1}{\sqrt{s_i}} \|\mathbf{Q}_i^T \mathbf{1}_K\|_2^2 \\
 \mathbf{H}_i^{(\ell)} &\in \mathbb{R}^{s_i}, \quad \mathbf{H}_i = \mathbf{I} - 2\mathbf{h}\mathbf{h}^T \in \mathbb{R}^{s_i \times s_i}, \quad \mathbf{h} \in \mathbb{R}^{s_i}, \quad \text{and} \\
 s_i &= \max_{\ell} \left\{ \ell \leq K - d, \frac{\sum_{p=K-\ell+1}^K \lambda_p^{(i)}}{\sum_{p=1}^{K-\ell} \lambda_p^{(i)}} < \eta \right\}
 \end{aligned}$$

where $\mathbf{h} = \frac{\mathbf{h}_0}{\|\mathbf{h}_0\|}$ if $\mathbf{h}_0 = \alpha_i \mathbf{1}_{s_i} - \mathbf{Q}_i^T \mathbf{1}_K \neq \mathbf{0} \in \mathbb{R}^{s_i}$ (else $\mathbf{h} = \mathbf{h}_0 = \mathbf{0} \in \mathbb{R}^{s_i}$), $\eta = \rho_{\lceil V/2 \rceil}$, $\hat{\rho}_i = \frac{\sum_{p=d+1}^K \lambda_p^{(i)}}{\sum_{p=1}^d \lambda_p^{(i)}}$, $\rho = \text{sort}(\hat{\rho}, \text{ascending})$, and $\hat{\rho}, \rho \in \mathbb{R}^V$. Since every measure's invariant properties are determined in a square-integrable space [58], LLE performs a global least-squares optimisation based on Gauß' Principle of Least Constraint [23] to calculate the vectors $\mathbf{z}_1, \dots, \mathbf{z}_V$ corresponding to points on

¹When using MLLE it is possible for $d > K$, thus the optimal number of weight vectors s_i for each voxel i is determined by setting $d = 1$ so that $K - 1 \leq s_i \leq K$ —i.e., s_i is set to span as large of a basis as possible. After this step the desired dimensionality d is then input to the eigensolver.

the manifold:

$$\begin{aligned}
 \mathbb{E}[\mathbf{Z}] &= \min_{\mathbf{Z}=[\mathbf{z}_1, \dots, \mathbf{z}_V]} \sum_{i=1}^V \sum_{\ell=1}^{s_i} \|\mathbf{z}_i - \sum_{j \in \mathcal{N}(i)} w_{i,j}^{(\ell)} \mathbf{z}_j\|_2^2, \\
 &\quad \text{such that } \mathbf{Z}\mathbf{Z}^T = \mathbf{I} \quad (4)
 \end{aligned}$$

where $d \leq T$ is the dimensionality parameter selected by the user and $\mathbf{Z} \in \mathbb{R}^{V \times d}$ is the two-dimensional reconstructed fMRI. The global optimisation therefore calculates the points on the manifold [57], [59] that act as the four-dimensional orthogonal basis that best retains the geometry of the stress vectors. These represent the second order invariant properties of each voxel's Cauchy stress tensor. A detailed explanation of this optimisation is provided below.

Define $\hat{\mathbf{W}}_i \in \mathbb{R}^{V \times s_i}$ as the local sparse adjacency matrix, where:

$$\begin{aligned}
 \hat{\mathbf{W}}_i(\mathcal{N}(i), :) &= \mathbf{w}_i, \quad \hat{\mathbf{W}}_i(i, :) = -\mathbf{1}_{s_i}^T, \quad \text{and} \\
 \hat{\mathbf{W}}_i(j, :) &= 0, \quad \forall j \notin \{\mathcal{N}(i) \cup i\}
 \end{aligned}$$

The optimisation in Equation 4 can be written as a minimisation of the expected reconstruction cost, or error:

$$\begin{aligned}
 \mathbb{E}[\mathbf{Z}] &= \sum_{i=1}^V \|\mathbf{Z} \hat{\mathbf{W}}_i\|_2^2 = \text{trace}(\mathbf{Z} \sum_{i=1}^V \hat{\mathbf{W}}_i \hat{\mathbf{W}}_i^T \mathbf{Z}^T) = \text{trace}(\mathbf{Z} \Phi \mathbf{Z}^T) \quad (5)
 \end{aligned}$$

where $\hat{\mathbf{W}}_i \hat{\mathbf{W}}_i^T$ is the orthogonal projection for voxel i , and $\Phi = \sum_{i=1}^V \hat{\mathbf{W}}_i \hat{\mathbf{W}}_i^T$ is the sparse, symmetric and positive-definite *alignment matrix*, and therefore admits the eigendecomposition [42]:

$$\mathbf{Z} \Phi \mathbf{Z}^T = \mathbf{Z} \mathbf{\Lambda} \mathbf{Z}^T \quad (6)$$

where $\mathbf{\Lambda} \in \mathbb{R}^{(d+1) \times (d+1)}$ is the diagonal matrix containing the $d + 1$ smallest eigenvalues of $\mathbf{Z} \Phi \mathbf{Z}^T$, and $\mathbf{Z} \in \mathbb{R}^{V \times (d+1)}$ are the corresponding eigenvectors; Rayleigh's variational principle [12], [52] enables calculation of these bottom $(d + 1)$ eigenvectors. Each eigenvector represents a degree of freedom in space and time, where the $(d + 1)$ eigenvector is the global unit vector that fills three-dimensional space. The global unit vector is discarded to enforce the constraint that the manifolds have mean zero.

Note: To avoid degenerate solutions, LLE requires the manifolds to be centred around the origin in *both* space and time—i.e., $\sum_i^V \mathbf{z}_{i,:} = \mathbf{0} \in \mathbb{R}^d$ and $\sum_i^d \mathbf{z}_{:,i} = \mathbf{0} \in \mathbb{R}^V$ —and also have outer products with unit covariance—i.e., $\mathbf{Z}\mathbf{Z}^T = \mathbf{I}$. Centring the manifolds about the origin ensures they are of the same scale, which is superficially similar to the common practice of signal, or count rate, normalisation [51]. The unit covariance constraint imposes the requirement that the reconstruction errors of the extracted manifolds are measured on the same scale.

TABLE 1. Results.

Dataset	Partition	Specificity				Sensitivity				Precision				Accuracy											
		Chance	Original	LLE	PCA	Chance	Original	LLE	PCA	Chance	Original	LLE	PCA	Chance	Original	LLE	PCA								
Beijing (Peking_3)	Training 100%	73.3% ± 15.8%	80%	± 14.3%	66.7%	± 16.9%	0%	53.3% ± 17.9%	93.3%	± 8.9%	80%	± 14.3%	0%	66.7% ± 16.9%	82.4%	± 13.6%	70.6%	± 16.3%	50%	± 17.9%	63.3%	± 17.2%	86.7% ± 12.2%	73.3% ± 15.8%	
	Holdout 100%	75%	87.5%	25%	0%	100%	50%	25%	0%	66.7%	66.7%	14%	67%	83.3%	75%	25%									
COBRE	Training 100%	57.7% ± 13.6%	92.3%	± 7.3%	96.2%	± 5.3%	0%	60%	± 13.4%	88%	± 8.9%	52%	± 13.7%	0%	57.7% ± 13.6%	91.7%	± 7.6%	92.9%	± 7.1%	51%	± 13.7%	54.9%	± 13.7%	90.2% ± 8.2%	74.5% ± 6.1%
	Holdout 100%	60%	100%	100%	0%	70%	100%	60%	0%	63.6%	100%	100%	50%	65%	100%	80%									
Mind Research Network (MRN)	Training 100%	N/A	88.2%	± 11.5%	100%	0%	N/A	84.6%	± 12.9%	23.1%	± 15.1%	0%	N/A	84.6%	± 12.9%	100%	56.7%	± 17.7%	N/A	86.7% ± 12.2%	66.7%	± 16.9%			
	Holdout 100%	N/A	67%	100%	0%	N/A	71%	14%	0%	N/A	63%	100%	56%	N/A	68.8%	62.5%									
Stanford	Training 100%	73.3% ± 15.8%	80%	± 14.3%	100%	0%	66.7% ± 16.9%	86.7%	± 12.2%	33.3%	± 16.9%	0%	71.4% ± 16.2%	81.3%	± 14%	100%	50%	± 17.9%	70%	± 16.4%	83.3% ± 13.3%	66.7%	± 16.9%		
	Holdout 100%	80%	80%	100%	0%	40%	60%	0%	0%	67%	75%	0%	50%	60%	70%	50%									
University of Michigan (UM_2)	Training 100%	69.2% ± 20.2%	100%	0%	42.9% ± 21.7%	71.4%	± 19.8%	14.3%	± 15.3%	0%	42.9% ± 21.7%	100%	65%	± 20.9%	60%	± 21.5%	90% ± 13.1%	70%	± 20.1%						
	Holdout 100%	66.6%	100%	83.3%	0%	25%	50%	0%	0%	20%	100%	0%	60%	30%	80%	50%									

TABLE 2. Dataset summary.

Dataset	Experimental Design	Pulse Sequence	Patient Disorder	Source	Training Data		Holdout Data		fMRI Dimensionality ($L \times W \times H \times T$)	Parameters		Age Range [min,mean,max]			
					total patients	controls	total patients	controls		LLE (r,d)	PCA (d)				
Beijing (Peking_3)	resting-state	EPI	ADHD	ADHD200	30	15	15	12	4	8	57 × 68 × 42 × 236	(2,236)	236	[11, 13.24, 16]	
	resting-state	EPI	Schizophrenia	COBRE	51	25	26	20	10	10	53 × 62 × 52 × 150	(2,150)	150	[18, 37.45, 65]	
Mind Research Network	block-design	EPI	Schizophrenia	MCIC	30	13	17	16	7	9	57 × 68 × 38 × 177	(2,177)	177	[18, 34.15, 60]	
	resting-state	Spiral-IN/OUT	EPI	Autism	ABIDE	30	15	15	10	5	5	63 × 75 × 42 × 180	(2,60)	60	[7.52, 9.95, 12.93]
University of Michigan (UM_2)	resting-state	Spiral-IN	EPI	Autism	ABIDE	20	7	13	10	4	6	57 × 68 × 63 × 300	(2,222)	222	[13.1, 16.26, 28.8]

A. FEATURE & PARAMETER SELECTION

Hyper-parameter Selection involves selecting the algorithm parameters that will be used to generate the reconstructed fMRI from LLE and PCA, respectively. Both PCA and LLE require one to specify the number of reconstruction dimensions d , where LLE’s additional hyper-parameter, $K = (1 + 2r)^3 - 1$, selects the neighbouring voxels whose coordinates are on the boundary of a sphere with radius $r \in \mathbb{Z}_+$, $r \geq 1$ that is centred around the respective voxel. For both PCA and LLE, the number of reconstruction dimensions is bounded by the number of time points. In the case of PCA, the optimal d is often chosen by the proportion of variance captured by the first d eigenvectors, expressed as $\frac{\sum_{i=1}^d \lambda_i}{\sum_{j=1}^T \lambda_j}$. For LLE, however, there is no analogous interpretation because the d eigenvectors are uniformly spaced “time points” on the *world line* [47]. Thus cross-validation procedures are implemented (see following subsection for details) on the training data to systematically select the best hyper-parameters, iterating over $d \in \{1, \dots, T\}$, where d is generated on a log-scale from 1 to T .

SEQUENTIAL FORWARD SELECTION (SFS)² [63]

is a nonparametric method for measurement (feature) selection that starts with an empty “candidate,” or near-optimal, set of “time points”. The method first finds the “time point,” defined as the second moment [6], with the highest classification accuracy and adds the corresponding image volume to the set. The method then finds an *additional* “time point” that *strictly* improves the classification accuracy in *conjunction* with the volume(s) whose “time points” are already in the “candidate” set, and terminates when no such volume can be found. Thus, SFS produces a near-optimal set of “time points” (volumes) that distinguish the subject groups with high classification accuracy, where this set is determined by

²Implemented using MATLAB’s (R2013a) *sequentialfs* function

adding “time points” to the near optimal set in a one-by-one fashion. It follows that the near-optimal set of “time points” represent points in time during the scan that enable discrimination of patients from controls with high classification accuracy. Note that SFS is only used on the training data (see Section II-C).

B. CLASSIFIER

A slight abuse of notation is introduced by redefining variable $\mathbf{z} \in \mathbb{R}^{V_c}$ as the one-dimensional representation of the c diagnostic volumes from the respective subject’s reconstructed fMRI \mathbf{Z} produced in the previous step. Here, it is assumed there are c diagnostic time points and that \mathcal{Z} is a random variable containing the collection of cohorts’ reconstructed fMRI.

FISHER’S LINEAR DISCRIMINANT (LDA)³ [17], [18]

is a linear classification rule [26] that assumes both the patient and control class densities (at each voxel location) can be represented as multivariate Gaussians in three-dimensional space, each with some intrinsic curvature [21], [22]. Each class density is expressed as

$$f_k(\mathbf{z}) = \frac{1}{(2\pi)^{\frac{V_c}{2}} |\boldsymbol{\Sigma}_k|^{\frac{1}{2}}} e^{-\frac{1}{2}(\mathbf{z}-\boldsymbol{\mu}_k)^\top \boldsymbol{\Sigma}_k^{-1}(\mathbf{z}-\boldsymbol{\mu}_k)}$$

where it is assumed the k classes have a common covariance matrix—i.e., $\boldsymbol{\Sigma}_k = \boldsymbol{\Sigma}$, $\forall k$. This assumption allows the log ratio between the posterior distribution of each class to form a decision boundary that lies between patients (class 1) and controls (class 0), written as $P(Y = 0 | \mathcal{Z} = \mathbf{z}) = P(Y = 1 | \mathcal{Z} = \mathbf{z})$, which is linear in \mathbf{z} . LDA therefore calculates the (V_c)-dimensional hyperplane that best discriminates, or separates, the diagnostic volumes that have been determined using

³Implemented using MATLAB’s (R2013a) *classify* function with the ‘diaglinear’ argument to estimate the positive diagonal covariance matrix

SFS with LOOCV on the respective dataset, for each class' reconstructed fMRI. It follows these diagnostic volumes contain spatial locations with sufficiently different patient and control class densities, where statistically-significantly different regions possess the requisite margin between the class densities such that they are perceptible [51].

C. EVALUATION CRITERIA

Each dataset was split such that the proportion of patients in the training partition was near-equal to the proportion of controls, and the holdout dataset contained at least 10 cohorts. In all but two cases, the aforementioned criteria could not be met because the datasets were unbalanced. For these situations, the training dataset was constructed such that it was a representative sample of the overall data, with the remaining cohorts being assigned to the holdout dataset. The holdout set therefore contained group proportions that could deviate from the training set. The cohorts used to define the training and holdout partition for each dataset are provided in Table 3 of the Supplementary Information.

Performance was compared to the original fMRI, both before and after applying PCA for artefact reduction [39], and chance, which is defined as the proportion of the majority class on the respective dataset. The performance of the reconstruction methods are evaluated using a combination of Leave-One-Out Cross Validation (LOOCV) and holdout data classification performance [26]. LOOCV is used on the training data to find both the best reconstruction parameter d , and the diagnostic volumes that produce the highest accuracy for this d . The holdout data classification accuracies use the parameters found from performing LOOCV on the training data. Note: the cohorts in the holdout set are *never* involved in determining the optimal hyper-parameter d , or the diagnostic volumes produced from this d .

III. RESULTS

A. PERFORMANCE & VISUALISATION

Tables 1 and 2 demonstrate that Locally Linear Embedded fMRI can distinguish various mental disorders from healthy controls with high discriminatory power ($> 80\%$); results for six additional resting-state datasets are provided in Table 2 of the Supplementary Information. The datasets contained healthy controls and patients with either Schizophrenia, Attention-Deficit Hyperactivity Disorder (ADHD), or Autism Spectrum Disorder (ASD). Ten of these datasets contained cohorts in the resting-state, with the remaining containing schizophrenic patients and healthy controls performing the Sternberg Item Recognition Paradigm (SIRP) task [61]. For a description of the data sources and technical details, please consult the appendix.

Given that performance on the training partition uses Leave One Out Cross Validation (LOOCV) to individually predict each subject's diagnosis, which is analogous to performing n Bernoulli trials [4] (where n is the number of cohorts),

the training data's performance metrics can be interpreted as the mean of successes over n binomially-distributed observations. To calculate the error of these estimates, we follow Laplace's approach of employing a normal distribution to estimate the error of binomially-distributed observations [33]. Given that discrimination performance on the holdout partition is determined in a one-time fashion, variance estimates are not applicable.

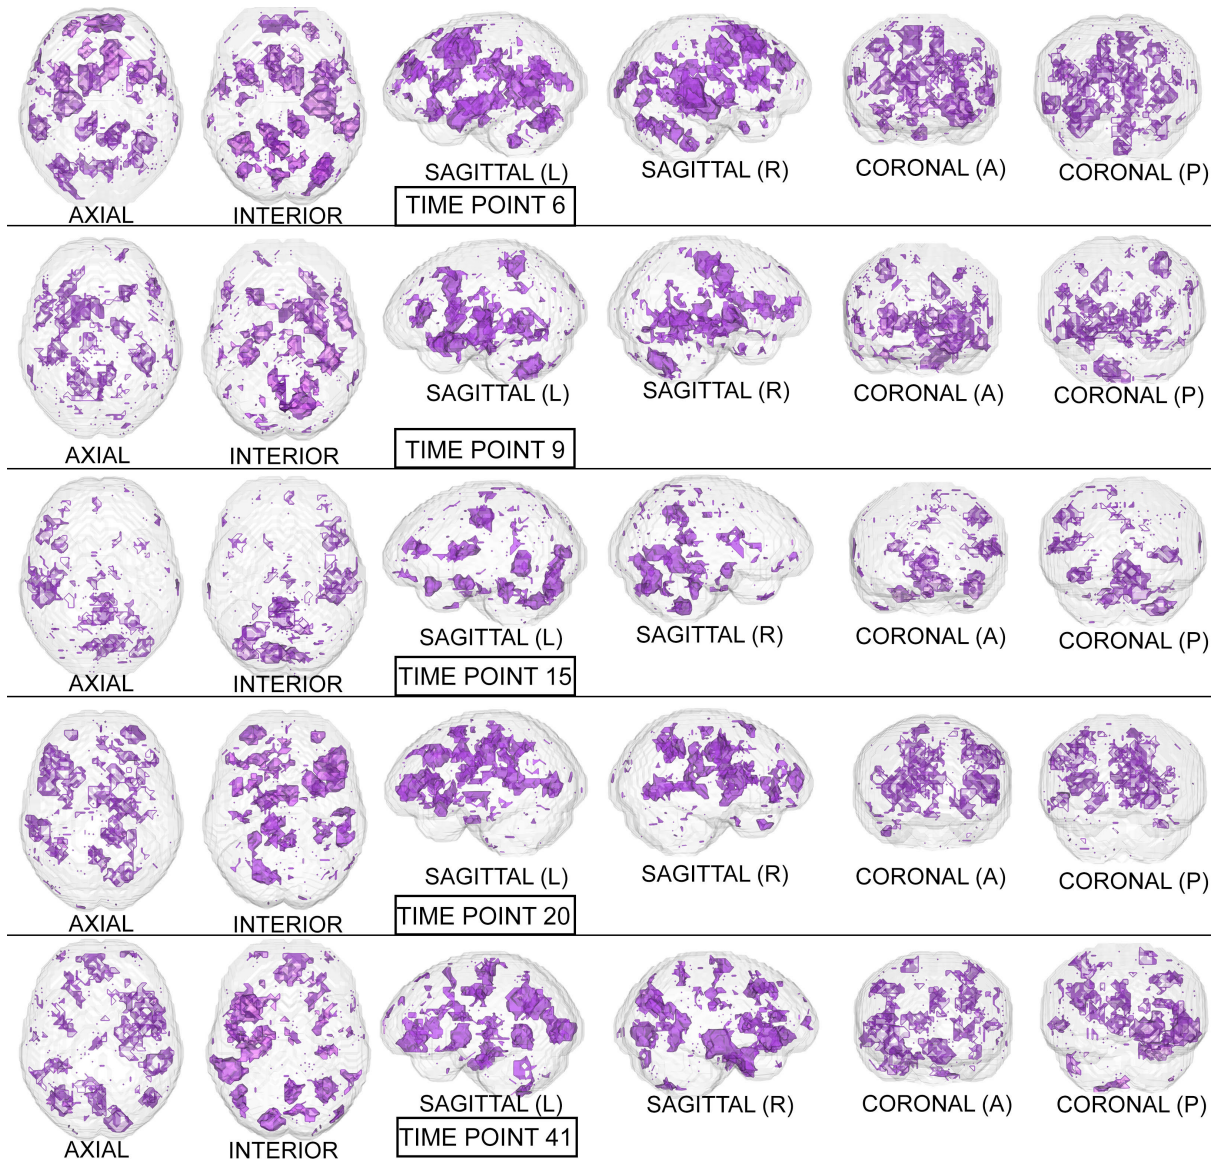
The Harvard-Oxford Subcortical/Cortical and Cerebellum atlases [30] are used to identify the statistically-significant differences between patients and controls in the diagnostic volumes for the respective dataset, shown in Figures 1, 2, and 3 (Figures 1, 2, 3, 4, 5, and 6 for datasets in the Supplementary Information). Each dataset's figure contains six different views of the statistically-significant differences at the time reflected by the respective time point(s); the coloured voxels at these time points denote statistically significant ($p < 0.05$ uncorrected) physical differences between the patient and control groups, where these groups include cohorts from *both* the training and holdout partitions. The proportion of significantly different voxels in each region, calculated by dividing the number of significantly different voxels by the total number of voxels in the respective region defined by the atlas, are provided in Tables 4 and 5 of the Supplementary Information.

B. NEUROBIOLOGICAL INTERPRETATION

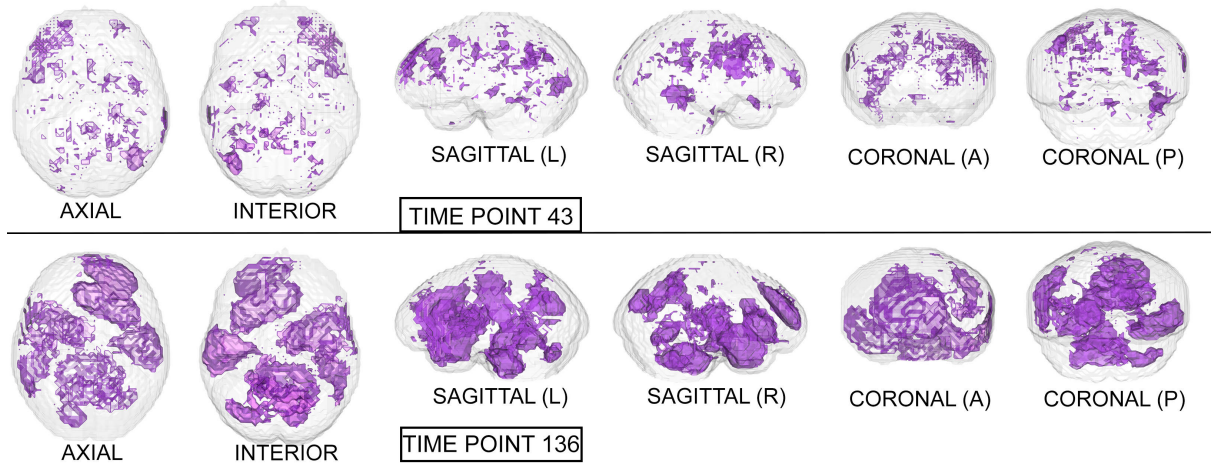
1) SCHIZOPHRENIA

David Ingvar & Göran Franzén found that healthy controls exhibit increased flows in the prefrontal regions and decreased flow in the post central regions, and schizophrenic patients exhibited the reversed pattern, with low flows prefrontally and high flows postcentrally [28]. Furthermore, they noticed that a lower flow in the premotor and frontal regions was associated with symptoms of indifference, activity and autism, and a higher postcentral flow over the temporo-occipito-parietal regions was associated with disturbed cognition [28]. Inspecting the statistical maps for the volumes in Figure 1, the significantly different areas seem to further substantiate Ingvar & Franzén's observations, as there are significant differences in the various temporo-occipital-parietal regions in all of the volumes.

The Sternberg Item Recognition Paradigm (SIRP) task evaluates cohorts' short-term, or working, memory. Each of the seven tasks during the scan involves the subject memorising a set of objects, followed by presentation of a new object whose membership in this set is identified by a 'yes' or 'no.' The tasks therefore evaluate the hypothesised information processing differences between schizophrenics and healthy controls [16]. It has been shown that the prefrontal and medial temporal regions are involved in encoding information, and it is believed the interactions between these regions are central to retrieval of stored information [60]. Figure 1b, especially volume 136, illustrates significant differences in the areas associated with the prefrontal and medial temporal regions;




(a) Center of Biomedical Research Excellence (COBRE). Cohorts were at rest for the duration of the scan.



(b) Mind Research Network (MRN). Cohorts performed the Sternberg Item Recognition Paradigm (SIRP) task.

FIGURE 1. Statistical maps illustrating the individual differences in mental activity (schizophrenic patients versus healthy controls) for the discriminative time points determined on the training partition.



Beijing (Peking_3) University. Cohorts were at rest for the duration of the scan.

FIGURE 2. Statistical maps illustrating the individual differences in mental activity (ADHD patients versus healthy controls) for the discriminative time points determined on the training partition.

this is further supported by the fact that the accuracy on the holdout data rose from 68.8% to 75% when using only volume 136. Thus, it is possible that schizophrenia indeed affects the physical mechanisms associated with retrieving stored information, as these mechanisms are central to the SIRP task. The reconstruction method therefore successfully reveals physical differences associated with task performance between patients and controls, which are different from the resting-state differences for the same subject groups.

2) ATTENTION-DEFICIT HYPERACTIVITY DISORDER (ADHD)

Similar to the goals of Ingvar & Franzén [28], previous work used PET scans to compare the regional cerebral blood flow of children with Attention-Deficit Hyperactivity Disorder (ADHD) to healthy controls, where it was found that the disorder was associated with hypoperfusion in the striatal and posterior periventricular regions [41]; these results provide biological evidence that is consistent with the canonical model for ADHD as a fronto-striatal deficient disorder. Figure 2 shows significant differences in the various occipital, striatal, cerebellar and ventral regions of the brain.

3) AUTISM SPECTRUM DISORDER (ASD)

Similar to Ingvar & Franzén's observations that schizophrenic patients and healthy controls had normal hemisphere flows [28], studies using PET to compare the regional cerebral blood flow of Autism Spectrum Disorder (ASD) patients to healthy controls observed normal metabolism and blood flow. Hypoperfusion in the temporal lobes, centred in the associative auditory and adjacent multimodal cortex [66], was observed in autistic children. Furthermore, this temporal hypoperfusion was individually identifiable in 75% of autistic children [66]. Figure 3 illustrates statistical differences in many areas of the temporal lobe.

In summary, Locally Linear Embedding appears to have conserved spatiotemporal patterns in resting-state fMRI (and task-related responses) that are consistent with literature on regionally-specific abnormalities of cerebral activity in the psychiatric conditions used to assess classification performance.

IV. DISCUSSION

The multidisciplinary nature of this work undoubtedly introduces difficulty when discussing its motivations, which is the deployment of this methodology in a clinical setting. Such a goal imposes some conditions. First, while the results support deployment, the methodology must be further evaluated by trained clinicians well-versed in the etymology of the disease under investigation. More importantly, however, experimental designs are compulsory when discovering biological markers for disease; that is, patients must be subject to the same stimulus presentation at the same time during the scan in order to homogenise comparisons. With respect to resting-state fMRI, it is felt that a consensus is required to glean neurobiological insight that can generalise, which makes further discussion more appropriate for future work.

This exposition focused on Locally Linear Embedding as a promising and effective form of dimensionality reduction as a pre-processing step for the analysis of fMRI time-series. The approaches and aims of this form of pre-processing share a close relationship with other approaches in imaging neuroscience. These approaches include Independent Component Analysis, the use of Support Vector Machines (and regression) to classification problems (and prediction), and algorithms based on adaptive smoothing. In future work, it will be interesting to explore the formal connections between other approaches and assess their relative sensitivity in the context of the classification problems considered above.

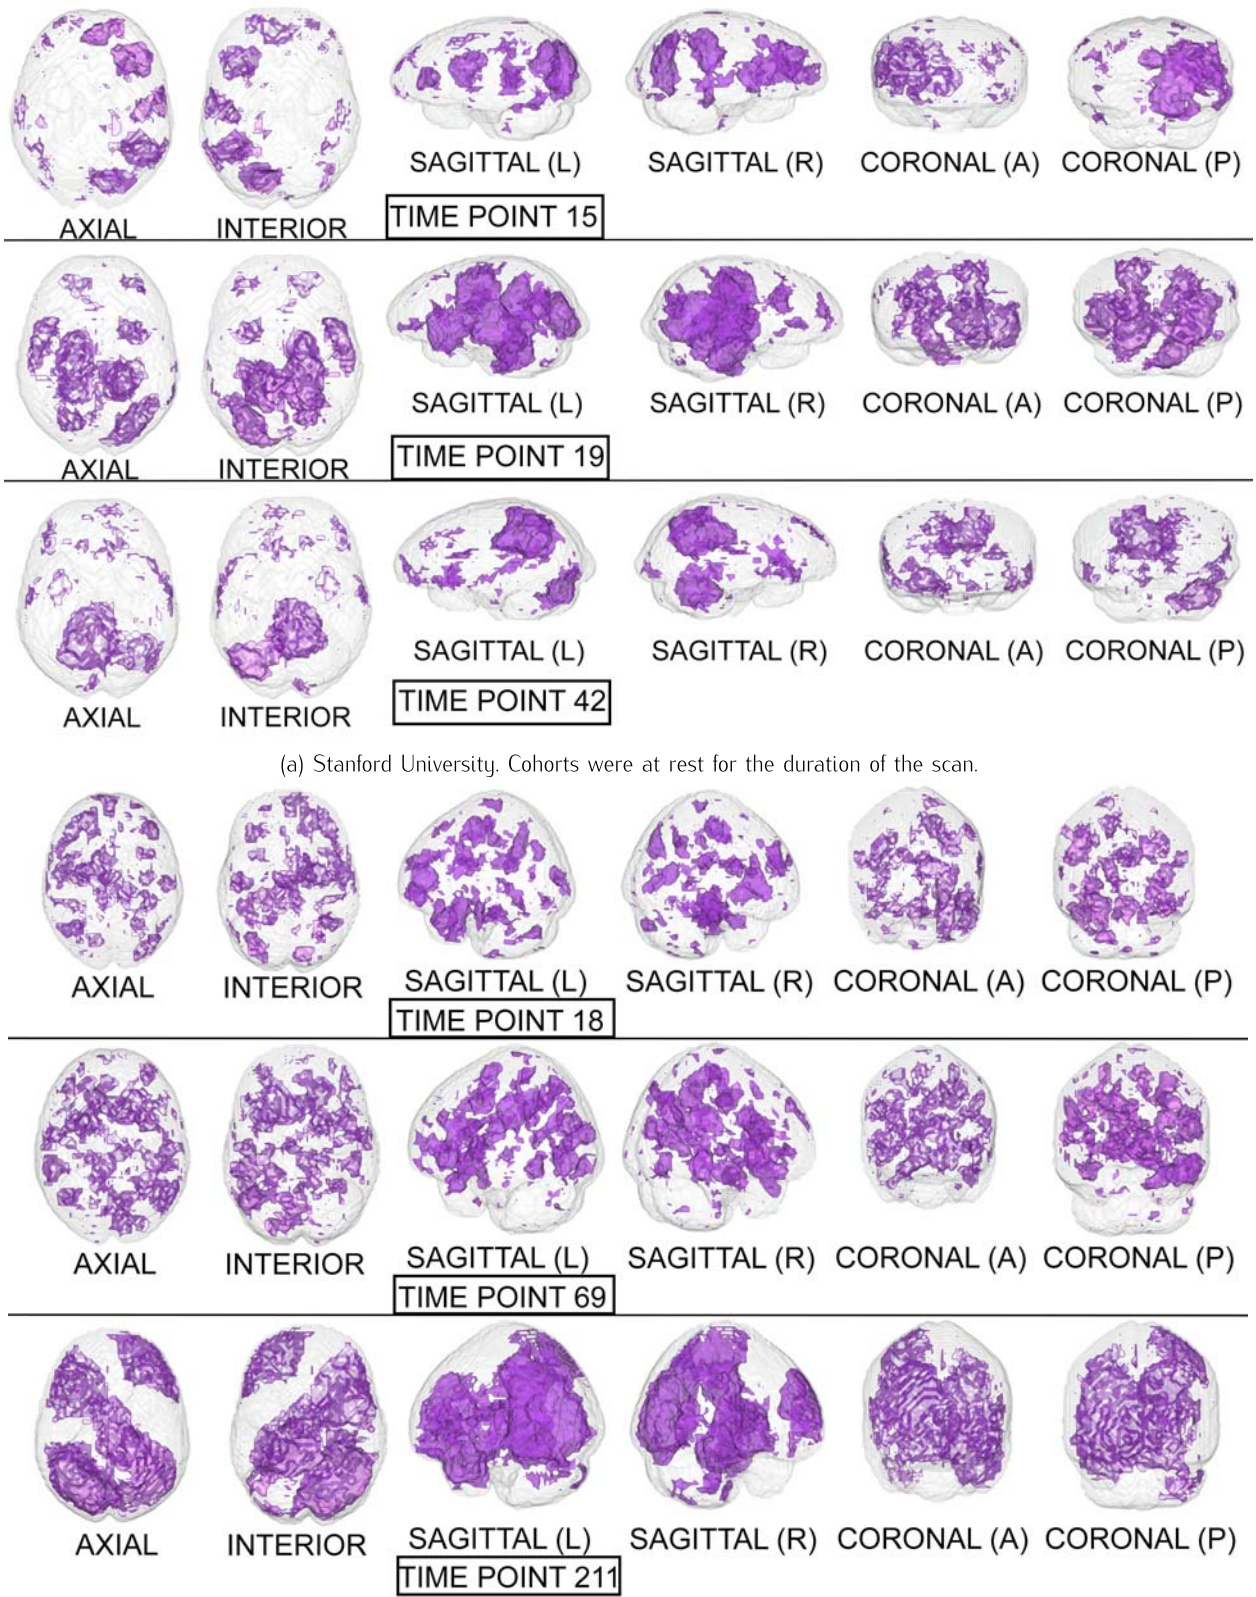
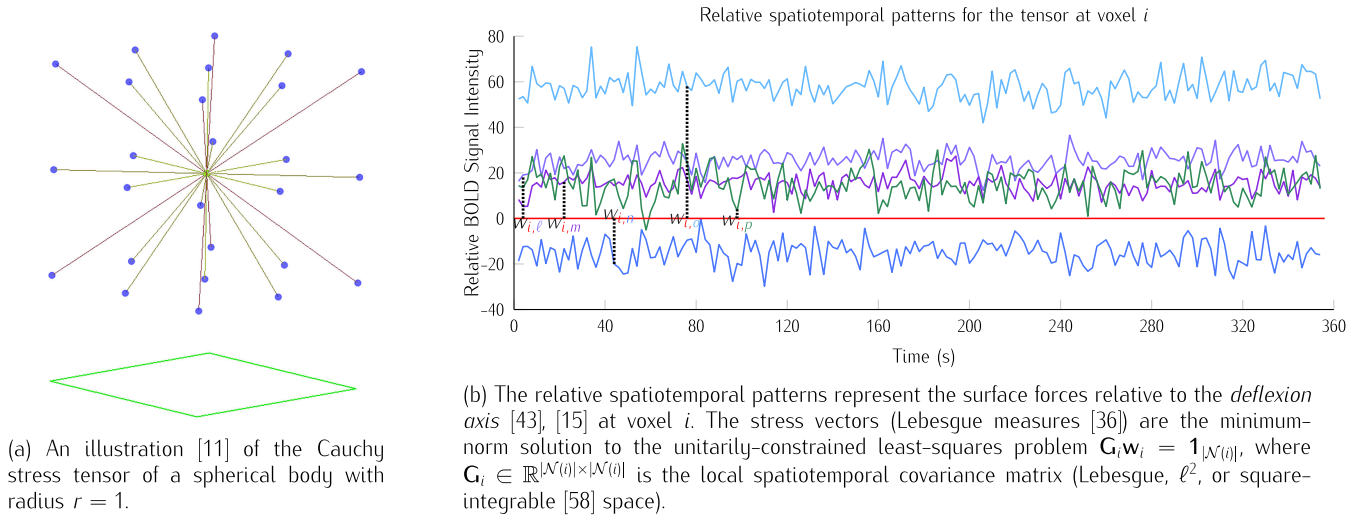


FIGURE 3. The statistical maps illustrating the individual differences in mental activity (ASD patients versus healthy controls) for the discriminative time points determined on the training partition.



(a) An illustration [11] of the Cauchy stress tensor of a spherical body with radius $r = 1$.

(b) The relative spatiotemporal patterns represent the surface forces relative to the *deflexion axis* [43], [15] at voxel i . The stress vectors (Lebesgue measures [36]) are the minimum-norm solution to the unitarily-constrained least-squares problem $\mathbf{G}_i \mathbf{w}_i = \mathbf{1}_{|\mathcal{N}(i)|}$, where $\mathbf{G}_i \in \mathbb{R}^{|\mathcal{N}(i)| \times |\mathcal{N}(i)|}$ is the local spatiotemporal covariance matrix (Lebesgue, ℓ^2 , or square-integrable [58] space).

FIGURE 4. The local geometry of the Cauchy stress tensor and its relative spatiotemporal patterns on topology [7] \mathcal{X} , which is defined on the Cartesian space with the Pythagorean distance metric.

One hundred and forty-eight years after Darwin ascertained that mental activity invokes physical mechanisms in the brain [14], the brain of man and ant alike are among the marvellous collection of atoms in the world.

**APPENDIX
LOCALLY LINEAR EMBEDDING**

fMRI contain T uniformly-spaced time points, where time point $t \in \{1, \dots, T\}$ describes a three-dimensional space comprised of $V = L \times W \times H$ voxels (volumetric elements); this is the *global* description. Each voxel is a volumetric measurement of the brain’s physical mechanisms in both space and time— i.e., every voxel represents a measurement at spatial coordinates (x, y, z) at time t , which can also be expressed as a four-tuple (x, y, z, t) . The BOLD contrast over T uniformly-spaced time points can be viewed as a random variable, or measurable function, [6] that is caused by local neuronal fluctuations. The terms global and local are only valid under a rigorous definition of three-dimensional space. A topological space requires specification of the relationships between points in a set and their respective open sets, which

are defined as the sets that generalise the concept of an open interval in the real line while also providing a rigorous definition for nearness of points in this space [7].





FIGURE 5. Illustrating the reconstruction of the waveform at voxel i , located in the left cerebral cortex of a schizophrenic patient performing the Sternberg Item Recognition Paradigm (SIRP), as a linear combination of the Lebesgue measures [36] $w_{i,1}, \dots, w_{i,|\mathcal{N}(i)|}$ (defined on the Cauchy stress tensor) and the spatial patterns of the voxels on the boundary of the sphere, where spatial distance is defined using the Pythagorean distance metric; in this example, $r = 2$. For comprehension, only the five nearest voxel waveforms are shown.

discarded spatial properties of the data in the neighbourhood selection step [44], thereby failing to preserve the inherent spatial configuration. In contrast, specifying the topological space \mathcal{X} allows LLE to treat each data point's open set as its neighbourhood, thereby forgoing the neighbourhood selection step entirely. It follows that the benefit of LLE depends on the suitability of the distance metric used in determining the neighbourhood of each data point. Figures 4 and 5 suggest that the physical reality represented by topological space \mathcal{X} is a collection of locally-linear subspaces whose properties are subject to change at any time.

LLE uses the *local* description to construct the Cauchy stress tensor [9] at every voxel, which proceeds as follows: Initially, the spatial pattern of each voxel i , given by $\mathbf{x}_i \in \mathbb{R}^T$, is subtracted from both itself and the spatial patterns of its neighbours $\mathcal{N}(i)$, thereby allowing its “zero waveform” (represented as $\mathbf{0} \in \mathbb{R}^T$) to serve as the imaginary plane, or *deflexion axis* [15], [43], that divides the spherical body; the subtracted spatial patterns, given by (vector space) $\mathbf{C}_i = [\mathbf{x}_j - \mathbf{x}_i, \dots, \mathbf{x}_{j+|\mathcal{N}(i)} - \mathbf{x}_i] \in \mathbb{R}^{T \times |\mathcal{N}(i)|}$, represent the distances (over time) from the respective voxel's spatial location.

We then use the inner-product $\langle (\mathbf{x}_j - \mathbf{x}_i), (\mathbf{x}_j - \mathbf{x}_i) \rangle \in \mathbb{R}$ for $j \in \mathcal{N}(i)$ to compute the squared distances, which results in the real-symmetric positive-definite spatiotemporal covariance matrix $\mathbf{G}_i = \mathbf{C}_i^T \mathbf{C}_i \in \mathbb{R}^{|\mathcal{N}(i)| \times |\mathcal{N}(i)|}$ for every voxel i . Since the contact forces are inversely proportional to the squared differences represented in elements of \mathbf{G}_i [48], we use the Moore-Penrose inverse [50] to solve for the minimum-norm solution $(\mathbf{G}_i^+ \mathbf{1}_{|\mathcal{N}(i)|}) = \mathbf{w}_i$ to the constrained least squares problem $\mathbf{G}_i \mathbf{w}_i = \mathbf{1}_{|\mathcal{N}(i)|}$. Here, $\mathbf{1}_{|\mathcal{N}(i)|} \in \mathbb{R}^{|\mathcal{N}(i)|}$ is the unit-length direction vector in each of the $|\mathcal{N}(i)|$ directions, and weights $\mathbf{w}_i \in \mathbb{R}^{|\mathcal{N}(i)|}$ represent the stress vectors originating from voxel i . By the spectral theorem, there exists an orthogonal basis $\mathbf{Q}_i \in \mathbb{R}^{d \times |\mathcal{N}(i)|}$ such that $\mathbf{G}_i = \mathbf{Q}_i^T \mathbf{A} \mathbf{Q}_i$ and $\mathbf{Q}_i^T \mathbf{Q}_i = \mathbf{I} \in \mathbb{R}^{|\mathcal{N}(i)| \times |\mathcal{N}(i)|}$ [62]; thus the stress vectors represented by \mathbf{w}_i are shift, rotation, and translation invariant in the space defined on basis \mathbf{Q}_i . The substantial overlap between spatially-adjacent voxels' tensors precludes independently calculating orthogonal bases $\mathbf{Q}_1, \dots, \mathbf{Q}_V$, as this defines each tensor's basis on separate vector spaces.

The stress vector \mathbf{w}_i for each voxel's Cauchy stress tensor, \mathbf{Q}_i , represents the spatially-invariant properties of the three-dimensional forces applied over the duration of the scan in the respective subspace. LLE retains these spatially-local invariant properties by constructing the (global) adjacency matrix $\mathbf{W} \in \mathbb{R}^{V \times V}$ such that $\mathbf{W}_{i,j} \neq 0 \iff j \in \mathcal{N}(i) \iff \mathbf{W}_{j,i} \neq 0 \iff i \in \mathcal{N}(j)$; by construction \mathbf{W} contains every voxel's stress vectors, which were determined using the Cauchy stress tensor. Given that the stress vectors' invariant properties are determined using squared distances, LLE then computes the squared distances between every voxel in the (global) normalised Laplacian matrix [10] $(\mathbf{I} - \mathbf{W}) \in \mathbb{R}^{V \times V}$. This is expressed by $\mathbf{M} = (\mathbf{I} - \mathbf{W})^T (\mathbf{I} - \mathbf{W}) \mathbf{1}_V$, where $(\mathbf{I} - \mathbf{W}) \mathbf{1}_V = \mathbf{0} \in \mathbb{R}^V$ [35]. By construction, \mathbf{M} is

symmetric and positive-definite, which means there exists an orthogonal basis $\mathbf{Z} \in \mathbb{R}^{V \times (d+1)}$ such that $(\mathbf{I} - \mathbf{W})^T (\mathbf{I} - \mathbf{W}) \mathbf{1}_V \equiv \mathbf{Z} \mathbf{\Lambda} \mathbf{Z}^T$ [42], where $\mathbf{1}_V$ represents the *global* unit vector in three-dimensional space.⁴ Since the Cauchy stress tensor is of second order, LLE uses Rayleigh's variational principle [12], [52] to calculate the resonance frequencies that best preserve the geometry of the *deflexion axis* at every voxel's stress tensor. These are given by the bottom $(d+1)$ eigenvectors, each of which represent one degree of freedom in space and time. LLE therefore uses a *global* optimisation to embed the relative measurements of every voxel's Cauchy stress tensor [9] in a global coordinate system $\mathbf{Z}^{V \times (d+1)}$ that conserves quantities over *both* space and time [19], where this coordinate system contains a mechanical system [32] in static equilibrium [8]. Applying LLE to fMRI data can therefore be viewed as using Carl Friedrich Gauß' Principle of Least Constraint [23] to determine the true motion of the mechanical system defined on the topology \mathcal{X} , where the Cauchy stress tensor [9] allows preservation of the intrinsic local Gaussian curvature [21], [22] in space and time.

ACKNOWLEDGEMENTS

This work is dedicated to the late Sam Roweis (1972-2010), Donald MacCrimmon MacKay (1922-1987) and his son David John Cameron MacKay (1967-2016). Additionally, many thanks to colleagues Ruitong Huang & Dr. Csaba Szepesvári, Joshua T. Vogelstein, Dr. Vincent D. Calhoun, Dr. Neil Lawrence, Dr. Bert Vogelstein, Dr. Michael Milham, Dr. Karl J. Friston, Dr. Klaas Enno Stephan, Dr. Nikos K. Logothetis, Dr. Christof Koch, and Dr. Yifan Hu of AT&T Labs' Information Visualisation department for his assistance with GraphViz, which was used to depict the Cauchy Stress Tensor in three dimensions.

REFERENCES

- [1] A. H. Andersen, D. M. Gash, and M. J. Avison, “Principal component analysis of the dynamic response measured by fMRI: A generalized linear systems framework,” *Magn. Reson. Imag.*, vol. 17, no. 6, pp. 795–815, 1999.
- [2] R. C. Archibald, “Time as a fourth dimension,” *Bull. Amer. Math. Soc.*, vol. 20, no. 8, pp. 409–412, 1914.
- [3] F. Bacon, *Francisci De Verulamio, Summi Angliae Cancellarii, Instauratio Magna*. London, U.K.: Apud Joannem Billium, Typographum Regium, 1620.
- [4] J. Bernoulli, *Ars Conjectandi*. Basel, Switzerland: Impensis Thurnisiorum, Fraetum, 1713.
- [5] C. M. Bishop, *Pattern Recognition and Machine Learning*. New York, NY, USA: Springer, 2006.
- [6] P. Billingsley, *Probability and Measure*. Hoboken, NJ, USA: Wiley, 1979.
- [7] N. Bourbaki, *Éléments de mathématique. Topologie algébrique*. Berlin, Germany: Springer-Verlag, 2016, chs. 1–4.
- [8] C. Carathéodory, “Untersuchungen über die Grundlagen der thermodynamik,” *Mathematische Annalen*, vol. 67, no. 3, pp. 355–386, 1909.
- [9] A. L. Cauchy, “De la pression ou tension dans un corps solide,” *Exercices De Mathématiques*, vol. 2, pp. 42–56, 1827.
- [10] F. R. K. Chung, *Spectral Graph Theory*. Providence, RI, USA: American Mathematical Society, 1997.
- [11] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to Algorithms*, 2nd ed. Cambridge, MA, USA: MIT Press, 2001.

⁴ $\mathbf{1}_V$ is also the $d + 1$ eigenvector of \mathbf{Z} , which is discarded after the global optimisation.

- [12] R. Courant, "Variational methods for the solution of problems of equilibrium and vibrations," *Bull. Amer. Math. Soc.*, vol. 49, no. 1, pp. 1–23, 1943.
- [13] H. S. M. Coxeter, *Introduction to Geometry*, 2nd ed. Hoboken, NJ, USA: Wiley, 1989.
- [14] C. Darwin, *The Descent Man*. London, U.K.: Murray, 1871.
- [15] E. M. Deeley and D. M. MacKay, "Multiplication and division by electronic-analogue methods," *Nature*, vol. 163, no. 4147, p. 650, Apr. 1949.
- [16] O. Demirci *et al.*, "A review of challenges in the use of fMRI for disease classification / characterization and a projection pursuit application from a multi-site fMRI schizophrenia study," *Brain Imag. Behav.*, vol. 2, no. 3, pp. 207–226, Sep. 2008.
- [17] R. A. Fisher, "The use of multiple measurements in taxonomic problems," *Ann. Eugenics*, vol. 7, no. 2, pp. 179–188, 1936.
- [18] R. A. Fisher, "The statistical utilization of multiple measurements," *Ann. Eugenics*, vol. 8, no. 4, pp. 376–386, 1938.
- [19] K. Friston, "The free-energy principle: A unified brain theory?" *Nature Rev. Neurosci.*, vol. 11, no. 2, pp. 127–138, 2010.
- [20] P. R. Garabedian and M. Schiffer, "On existence theorems of potential theory and conformal mapping," *Ann. Math.*, vol. 52, no. 1, pp. 164–187, 1950.
- [21] C. F. Gauß, *Disquisitiones Generales Circa Superficies Curvas*. Berlin, Germany: Typis Ditericianis, 1828.
- [22] C. F. Gauß, *Theoria Motus Corporum Coelestium in Sectionibus Conicis Solem Ambientium*. Hamburg, Germany: Perthes et Besser, 1809.
- [23] C. F. Gauß, "Über ein neues allgemeines grundgesetz der mechanik," *J. Für Die Reine Angewandte Mathematik*, vol. 4, no. 4, pp. 232–235, 1829.
- [24] R. J. Gillespie, "Fifty years of the VSEPR model," *Coordination Chem. Rev.*, vol. 252, nos. 12–14, pp. 1315–1327, Jul. 2008.
- [25] J. Gregory, *Geometriae Pars Universalis*. Patavii, Italy: Typis heredum Pauli Frambotti, 1668.
- [26] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction* (Springer Series in Statistics), 2nd ed. New York, NY, USA: Springer, 2009.
- [27] D. Hilbert, "Ueber die stetige Abbildung einer Linie auf ein Flächenstück," *Mathematische Annalen*, vol. 3, no. 38, pp. 459–460, 1891.
- [28] D. H. Ingvar and G. Franzén, "Distribution of cerebral activity in chronic schizophrenia," *Lancet*, vol. 304, no. 7895, pp. 1484–1486, Dec. 1974.
- [29] D. H. Ingvar and N. A. Lassen, "Quantitative determination of regional cerebral blood-flow in man," *Lancet*, vol. 278, no. 7206, pp. 806–807, Oct. 1961.
- [30] M. Jenkinson, C. F. Beckmann, T. E. J. Behrens, M. W. Woolrich, and S. M. Smith, "FSL," *NeuroImage*, vol. 62, no. 2, pp. 782–790, Aug. 2012.
- [31] I. Jolliffe, *Principal Component Analysis*. New York, NY, USA: Springer-Verlag, 2002.
- [32] C. Koch and K. Hepp, "Quantum mechanics in the brain," *Nature*, vol. 440, no. 7084, pp. 611–612, 2006.
- [33] P. Laplace, *Théorie Analytique Des Probabilités*. Paris, France: Courcier, 1812.
- [34] N. A. Lassen, D. H. Ingvar, E. J. Potchen, and V. R. McCready, Eds., "Radioisotopic assessment of regional cerebral blood flow," *Progress in Nuclear Medicine: Neuro-Nuclear Medicine*, vol. 1. Baltimore, MD, USA: Univ. Park Press, 1972, pp. 376–409.
- [35] N. D. Lawrence, "A unifying probabilistic perspective for spectral dimensionality reduction: Insights and new models," *J. Mach. Learn. Res.*, vol. 13, pp. 1609–1638, Jun. 2012.
- [36] H. Lebesgue, "Leçons sur L'intégration et la recherche des fonctions primitives, professées au collège de France," in *Collection de Monographies Sur la Théorie des Fonctions*. Paris, France: Gauthier-Villars, 1904.
- [37] P. F. Liddle, K. J. Friston, C. D. Frith, S. R. Hirsch, T. Jones, and R. S. J. Frackowiak, "Patterns of cerebral blood flow in schizophrenia," *Brit. J. Psychiatry*, vol. 160, no. 2, pp. 179–186, 1992.
- [38] J. Locke, *An Essay Concerning Human Understanding*. Oxford, U.K.: Thomas Basset, 1690.
- [39] N. K. Logothetis, J. Pauls, M. Augath, T. Trinath, and A. Oeltermann, "Neurophysiological investigation of the basis of the fMRI signal," *Nature*, vol. 412, no. 6843, pp. 150–157, Jul. 2001.
- [40] N. K. Logothetis, "What we can do and what we cannot do with fMRI," *Nature*, vol. 453, no. 7197, pp. 869–878, Jun. 2008.
- [41] H. C. Lou, L. Henriksen, and P. B. C. Psych, "Focal cerebral dysfunction in developmental learning disabilities," *Lancet*, vol. 335, no. 8680, pp. 8–11, Jan. 1990.
- [42] D. J. C. MacKay, *Information Theory, Inference and Learning Algorithms*. Cambridge, U.K.: Cambridge Univ. Press, 2003.
- [43] D. M. MacKay, "A high-speed electronic function generator," *Nature*, vol. 159, no. 4038, pp. 406–407, Mar. 1947.
- [44] P. Mannfolk, R. Wirestam, M. Nilsson, F. Ståhlberg, and J. Olsrud, "Dimensionality reduction of fMRI time series data using locally linear embedding," *Magn. Reson. Mater. Phys., Biol. Med.*, vol. 23, nos. 5–6, pp. 327–338, Dec. 2010.
- [45] M. S. Mel'nikov, "Analytic capacity: Discrete approach and curvature of measure," *Sbornik, Math.*, vol. 186, no. 6, pp. 827–846, 1995.
- [46] J. S. Mill, *A System of Logic, Ratiocinative and Inductive: Being a Connected View of the Principles of Evidence and the Methods of Scientific Investigation*, vol. 1. London, U.K.: John W. Parker, 1843.
- [47] H. Minkowski, "Raum und zeit," *Physikalische Zeitschrift*, vol. 10, pp. 104–111, Feb. 1909.
- [48] I. Newton, *Philosophiae Naturalis Principia Mathematica*.
- [49] S. Ogawa, T. M. Lee, A. R. Kay, and D. W. Tank, "Brain magnetic resonance imaging with contrast dependent on blood oxygenation," *Proc. Nat. Acad. Sci. USA*, vol. 87, no. 24, pp. 9868–9872, 1990.
- [50] R. Penrose, "A generalized inverse for matrices," *Math. Proc. Cambridge Phil. Soc.*, vol. 51, no. 3, pp. 406–413, Jul. 1955.
- [51] M. G. Popham, E. J. Potchen, and V. R. McCready, Eds., "Numerical methods for the detection of abnormalities in radionuclide brain scans," in *Progress in Nuclear Medicine: Neuro-Nuclear Medicine*, vol. 1. Baltimore, MD, USA: Univ. Park Press, 1972, pp. 117–141.
- [52] J. W. S. Rayleigh, *The Theory of Sound*. London, U.K.: Macmillan, 1877.
- [53] L. N. Robins and J. E. Helzer, "Diagnosis and clinical assessment: The current state of psychiatric diagnosis," *Annu. Rev. Psychol.*, vol. 37, no. 1, pp. 409–432, 1986.
- [54] S. T. Roweis and L. K. Saul, "Nonlinear dimensionality reduction by locally linear embedding," *Science*, vol. 290, no. 5500, pp. 2323–2326, Dec. 2000.
- [55] C. S. Roy and C. S. Sherrington, "On the regulation of the blood-supply of the brain," *J. Physiol.*, vol. 11, nos. 1–2, pp. 85–158, Jan. 1890.
- [56] L. K. Saul and S. T. Roweis, "Think globally, fit locally: Unsupervised learning of low dimensional manifolds," *J. Mach. Learn. Res.*, vol. 4, pp. 119–155, Dec. 2003.
- [57] H. S. Seung and D. D. Lee, "The manifold ways of perception," *Science*, vol. 290, no. 5500, pp. 2268–2269, 2000.
- [58] E. Schmidt, "Über die auflösung linearer gleichungen mit unendlich vielen unbekanntem," *Rendiconti Circolo Matematico Palermo*, vol. 25, no. 1, pp. 53–77, 1908.
- [59] C. S. Sherrington, *Man on His Nature—Gifford Lectures*. Edinburgh, U.K.: Edinburgh Univ. Press, 1951.
- [60] J. S. Simons and H. J. Spiers, "Prefrontal and medial temporal lobe interactions in long-term memory," *Nature Rev. Neurosci.*, vol. 4, no. 8, pp. 637–648, Aug. 2003.
- [61] S. Sternberg, "High-speed scanning in human memory," *Science*, vol. 153, no. 3736, pp. 652–654, 1966.
- [62] J. J. Sylvester, "A demonstration of the theorem that every homogeneous quadratic polynomial is reducible by real orthogonal substitutions to the form of a sum of positive and negative squares," *London, Edinburgh, Dublin Philos. Mag. J. Sci.*, vol. 4, no. 23, pp. 138–142, 1852.
- [63] A. W. Whitney, "A direct method of nonparametric measurement selection," *IEEE Trans. Comput.*, vol. C-20, no. 9, pp. 1100–1103, Sep. 1971.
- [64] J. B. W. Williams *et al.*, "The structured clinical interview for DSM-III-R (SCID): II. Multisite test-retest reliability," *Arch. Gen. Psychiatry*, vol. 49, no. 8, pp. 630–636, Aug. 1992.
- [65] Z. Zhang and J. Wang, "MLLE: Modified locally linear embedding using multiple weights," *Adv. Neural Inf. Process. Syst.*, vol. 19, no. 19, pp. 1593–1600, Dec. 2007.
- [66] M. Zilbovicius *et al.*, "Temporal lobe dysfunction in childhood autism: A PET study," *Amer. J. Psychiatry*, vol. 157, no. 12, pp. 1988–1993, 2000.