Please cite the Published Version

Additional Information: This is an open access article published in NeuroImage, by Elsevier.

A tutorial and tool for exploring feature similarity gradients with MRI data

Claude J. Bajada [a,b,c,*], Lucas Q. Costa Campos [c,d], Svenja Caspers [c,e,f], Richard Muscat [a], Geoff J.M. Parker [g,h], Matthew A. Lambon Ralph [i], Lauren L. Cloutman [b], Nelson J. Trujillo-Barreto [b,**]

[a] Department of Physiology and Biochemistry, Faculty of Medicine and Surgery, The University of Malta, Malta
[b] Division of Neuroscience & Experimental Psychology, School of Biological Sciences, The University of Manchester, UK
[c] Institute of Neuroscience and Medicine (INM-1), Research Centre Jülich, 52425, Jülich, Germany
[d] Institute of Complex Systems and Institute for Advanced Simulation (ICS-2/IAS-2), Research Centre Jülich, 52425, Jülich, Germany
[e] Institute for Anatomy I, Medical Faculty, Heinrich-Heine-University Duesseldorf, 40221, Duesseldorf, Germany
[f] JARA-BRAIN, Jülich-Aachen Research Alliance, 52425, Jülich, Germany
[g] Centre for Medical Image Computing, Department of Computer Science, and Queen Square MS Centre, Department of Neuroinflammation, UCL Institute of Neurology, University College London, UK
[h] Bioxydyn Limited, Manchester, UK
[i] MRC Cognition and Brain Sciences Unit, University of Cambridge, Cambridge, UK

ARTICLE INFO

ABSTRACT

There has been an increasing interest in examining organisational principles of the cerebral cortex (and subcortical regions) using different MRI features such as structural or functional connectivity. Despite the widespread interest, introductory tutorials on the underlying technique targeted for the novice neuroimager are sparse in the literature.

Articles that investigate various "neural gradients" (for example based on region studied "cortical gradients," "cerebellar gradients," "hippocampal gradients" etc … or feature of interest "functional gradients," "cytoarchitectural gradients," "myeloarchitectural gradients" etc …) have increased in popularity. Thus, we believe that it is opportune to discuss what is generally meant by "gradient analysis". We introduce basics concepts in graph theory, such as graphs themselves, the degree matrix, and the adjacency matrix. We discuss how one can think about gradients of feature similarity (the similarity between timeseries in fMRI, or streamline in tractography) using graph theory and we extend this to explore such gradients across the whole MRI scale; from the voxel level to the whole brain level. We proceed to introduce a measure for quantifying the level of similarity in regions of interest. We propose the term "the Vogt-Bailey index" for such quantification to pay homage to our history as a brain mapping community.

We run through the techniques on sample datasets including a brain MRI as an example of the application of the techniques on real data and we provide several appendices that expand upon details. To maximise intuition, the appendices contain a didactic example describing how one could use these techniques to solve a particularly pernicious problem that one may encounter at a wedding. Accompanying the article is a tool, available in both MATLAB and Python, that enables readers to perform the analysis described in this article on their own data.

We refer readers to the graphical abstract as an overview of the analysis pipeline presented in this work.

1. Introduction

Every discrete point in the brain (modelled as a voxel or vertex in the context of MR imaging) has several co-existing features such as the cytological architecture, the functional signature, the receptor density etc … Parcellation is one method of describing neural features and their similarities. The technique groups area of the brain that have similar features together. One of the most recognisable names in modern neuroscience is that of Korbinian Brodmann and his cytoarchitectonic parcellations of the cortex from the early 20th century. Despite his

* Corresponding author. Department of Physiology and Biochemistry, Faculty of Medicine and Surgery, The University of Malta, Msida, Malta
** Corresponding author. Division of Neuroscience and Experimental Psychology, School of Biological Sciences, Faculty of Biology, Medicine and Health, University of Manchester, Academic Health Science Centre, Brunswick Street, Manchester, M13 9PL, UK
E-mail address: claude.bajada@um.edu.mt (C.J. Bajada).

modern fame, Brodmann was not the most ardent proponent of parcellation. That honour arguably goes to Oskar and Cecile Vogt, who were the true parents of modern parcellation, and Brodmann's mentors.

Brodmann's regions attempted to define areas of local cytoarchitectural homogeneity in the cortex. Unfortunately, the convenience of using such areas in neuroimaging studies comes at a high price. First, Brodmann areas were defined histologically and do not necessarily correspond to gross anatomical landmarks that are visible on MRI (Zilles and Amunts, 2010). Second, Brodmann's maps certainly are not the final word on what constitutes a cytoarchitectural brain parcel. His contemporaries (von Economo and Koskinas, 1925) as well as current researchers (Amunts et al., 2005; Bludau et al., 2014; Caspers et al., 2013, 2006; Rottschy et al., 2007; Scheperjans et al., 2008) are still investigating and refining cytoarchitectonic parcellation. Third, cytoarchitecture is not the only feature with which one can parcellate the cortex. Myeloarchitecture, for instance, has been used since the times of Brodmann (Geyer and Turner, 2013; Nieuwenhuys et al., 2015; Vogt and Vogt, 1911) and recently a vast array of neuroimaging features has been used (Glasser et al., 2016). Fourth, the boundaries (or borders) between areas characterised by any particular feature may sometimes be sharp while other times they can be blurred (Bailey and von Bonin, 1951; Brodmann, 1909; Glasser et al., 2016). Finally, different "distinct" areas may also share some relationship with other areas and thus show a non-random pattern across the whole cortex; for example functional areas involved in resting state networks (Damoiseaux et al., 2006).

In the 1950s, Percival Bailey and Gerhardt von Bonin proposed another, competing conceptualisation of cortical organisation. Vogt and Bailey can be, *prima facie,* thought of as expounding opposing points of views. While the Vogts, championed cortical parcellation, Bailey and von Bonin (1951) argue that the isocortex (or neocortex – the six layered cortex) is much more similar throughout its extent than it is different. They go so far as to state that:

> "The drawing of sharp areal boundaries, on the basis of many structural peculiarities of varying distinctiveness and significance, is the fundamental defect of most maps and has been carried to absurd lengths by the Vogt school." (p. 189)

They elaborate by stating:

> "Anybody can see, to give an example, the difference between Brodmann's areas 17 and 18. But the differences between his 18 and 19 are quite tenuous and very difficult to recognize. To draw a map on which these three areas are given three different markings - such as dots, cross-hatchings, and broken lines - is to create an entirely misleading impression. Useful as such maps are for the description of corticocortical connections, they do not translate accurately cytoarchitectonic data." (p. viii)

An interest in similar issues of parcellations versus gradual transitions between areas has reappeared in the modern neuroscience literature. Indeed, Brodmann himself asserted that some regions demonstrated transitionary zones (Brodmann, 1909). While the feature of interest has moved on from cytoarchitecture to fMRI time series analysis or diffusion MRI tractography (more generally within a context of network analysis, connectomes and connectivity based parcellations), some themes of the early debates have lived on. For example: when, and to what extent, is clustering the cortex into distinct parcels appropriate? And, as has been explored by (Mesulam, 2008, 1998), what are the interareal relationships between cortical territories?

In 2004, a novel approach, based on spectral graph theory, appeared in the literature to investigate changes in cortical connectivity patterns across the brain using diffusion MRI tractography (Johansen-Berg et al., 2004). Similar approaches have become a popular tool for parcellating the cerebral cortex using both diffusion and functional MRI (Cloutman et al., 2012; Craddock et al., 2012; Devlin et al., 2006; Eickhoff et al., 2015; O'Donnell et al., 2013, 2006). Recently, similar techniques have

also been used to explore interareal connectivity pattern changes (such as structural connectivity through tractography or functional connectivity) as one traverses the cerebral cortex; so-called feature similarity gradients (Bajada et al., 2017; Cerliani et al., 2012; Haak et al., 2018; Jackson et al., 2017; Margulies et al., 2016). Indeed, Margulies et al. (2016) have shown that, under certain constraints that likely depend on the construction of the similarity matrix, the primary rs-fMRI feature gradient reflects the interareal relationships outlined by Mesulam (2008) and elaborated by Buckner et al. (2013) where this gradient has modality selective and modality general cortices on either end. While this concept may not be fully understood by a novice reader, we hope that by the end of this article (and particularly after reading the supplementary text) the interpretation of such a statement will be obvious.

These techniques are appealing to the neuroimaging community since they have the potential to provide a flexible, unified framework for understanding similarities of neural structure or function across the brain (c.f. Paquola et al., 2020, 2019; Vázquez-Rodríguez et al., 2019 for examples of how flexible these concepts can be used across multiple modalities). In this article we further extend this framework by introducing a way to measure how sharply defined each area is, showing the full spectrum of possibilities between the ideas of the Vogts, and those of Bailey and von Bonin; The Vogt-Bailey index.

We use the historical context to highlight the importance of having a way of thinking about cortical organisation through "feature gradients" – e.g. fMRI, tractography, cytoarchtectonic etc … that bridges the gap between old debates. Specifically, we will apply the tools described here to help settle a discussion started in the middle of the 20th century. We have also made available a pair of tools written in MATLAB and Python implementing the algorithms outlined in this work, thus making it possible for the interested reader to calculate the VB index using their own data (https://doi.org/10.5281/zenodo.3609459, https://github.com/VBIndex/) (Da Costa Campos and Bajada, 2020). It is also possible to install the recommended production version of the software using Python's *de facto* package manager, *pip*, with "*pip install vb_toolbox*". Once this is done, the software *vb_tool* will be available for use. For usage details, we refer to the full documentation of the software's GitHub repository (https://github.com/VBIndex/py_vb_toolbox).

We note that various groups have released their own "gradient analysis" pipelines and toolboxes including the early "gradient pipeline" by Margulies et al. (2016), connectotopic mapping focused on regional modes of connectivity changes and their statistical tests by Haak et al. (2018), LittleBrain focusing on Cerebellar gradient by Guell et al. (2019), and BrainSpace a recent all-purpose gradient toolbox by Vos de Wael et al. (2019). All these workflows and toolboxes have minor differences in certain choices that are described below. Some also include the ability to perform statistical testing on gradient maps. The software presented in this article creates similar, but not necessarily identical, 'gradient maps' as the other software packages available, in addition it is the only package to-date, that allows calculation of the VB index (as described in section 5.1).

In the rest of this article we describe and explore the details of the steps needed to extract feature similarity gradients and the VB Index from data. We discuss methods of measuring similarity between brain regions, why it is useful to think of these resultant similarity measurements as a mathematical graph, and how to further process the graph to obtain the desired gradient maps. In this article we restrict our discussion to a technique based heavily upon Laplacian Eigenmaps (c.f. Belkin and Niyogi, 2003, 2002). In general, the problem of finding meaningful structures and geometric descriptions of such data is usually stated as some sort of nonlinear dimensionality reduction. Although several forms of dimensionality reductions for constructing cortical gradients (or subcortical, cerebellar etc.) have been used in the literature (Coifman and Lafon, 2006; Haak et al., 2018; Johansen-Berg et al., 2004; Margulies et al., 2016), they are similar in spirit to the Laplacian Eigen-mapping reviewed here. We refer the interested reader to the relevant literature and hope that the tutorial presented here will serve as a useful

introduction to understand the principles behind those related approaches.

2. What is a graph?

Most algorithms for feature gradient analyses emerge from the mathematical discipline of spectral graph theory. This is an approach to studying properties of graphs by computing the eigenvalues and eigenvectors of matrices that summarise the graph. While it would be lengthy to go into a detailed explanation of eigenvalues and eigenvectors in this text, we hope that their use in the context described will become clearer in later sections.

A graph is a mathematical structure that defines relationships between various objects. For example, the structure in Fig. 1 is a graph that defines the relationship between four objects. Each object is called a node.

The nodes could be thought of as voxels (or surface vertices) in the cortex or as cortical regions of interest.

The lines that link the nodes are called edges. The edges can be binary or have a weight associated with them (creating a weighted graph). Within neuroimaging, the edges are almost always undirected meaning that if node a connects to node b, the opposite is also true.

Some basic concepts are needed in order to proceed. The adjacency matrix is a square matrix (i.e., the same number of rows and columns) where every row and every column represent a single node, and the elements in the matrix represent the relationships between the row node and the column node. For the unweighted graph in Fig. 1, the adjacency matrix is

$$A = \begin{pmatrix} 0 & 1 & 1 & 1 \\ 1 & 0 & 1 & 0 \\ 1 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{pmatrix}.$$

The rows and columns are ordered from 1 to 4 such that entries of 1s in columns 2, 3 and 4 of row 1 means that node a (row 1) is connected to nodes b, c and d (columns 2, 3 and 4).

The degree matrix is a diagonal matrix where the entries along the diagonal represent the degree of each node, that is, the number of nodes that are connected (adjacent) to that node. For example, node a has a degree of 3 because it is adjacent to three nodes (b, c and d). The degree matrix D can be computed as the row/column wise sum of the adjacency matrix. For the graph in Fig. 1, the degree matrix is

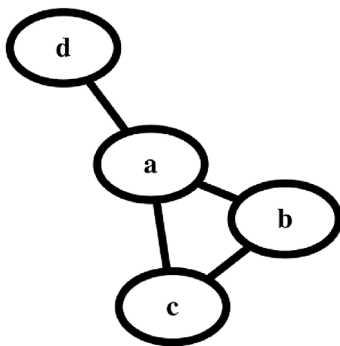$$D = \begin{pmatrix} 3 & 0 & 0 & 0 \\ 0 & 2 & 0 & 0 \\ 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}.$$

Fig. 1. A representation of a graph with 4 nodes. Every node can be considered to be a voxel or a region of interest. The edges between the nodes represent their relationships; these can either be structural connections or a measure of similarity (affinity) between the nodes.

The Laplacian is defined as the degree matrix minus the adjacency matrix

$$L = D - A.$$

While the exact meaning of the Laplacian may be difficult to intuit for many readers, we hope that the use of it in Section 4 will give the readers some intuition. At this point it is useful to note that in many applications, including in neuroimaging one can define a weighted graph, where each edge connecting the nodes carry different weights. A high weight, for instance, could mean that two nodes are strongly connected, while a low weight would indicate the nodes are not as strongly connected. One can now define a weighted version of the adjacency matrix, which can be used to describe a weighted graph. In the general case, the weighted adjacency matrix can be defined as

$$W = \begin{pmatrix} w_{11} & w_{12} & \cdots & w_{1n} \\ w_{21} & w_{22} & \cdots & w_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ w_{n1} & w_{n2} & \cdots & w_{nn} \end{pmatrix}.$$

We will see later that this matrix can be associated to the concept of a similarity or affinity matrix. In the example above, the weighted adjacency matrix can be written replacing the 1s in the adjacency matrix with the corresponding weights. Weighted versions of the Degree and Laplacian matrices can be defined in the same way as before. Henceforward, unless otherwise specified, we will refer to the Laplacian, degree and adjacency matrix as their generalised weighted versions.

2.1. Graphs in neuroimaging

For neuroimaging purposes, a graph can be one of two distinct types. The most conceptually straightforward way of creating a graph of the brain is to consider its structural connections. For example, the nodes in the graph of Fig. 1 can represent brain areas (e.g. cerebellum, brainstem, etc …) and the edges can represent the tracts that connect those brain areas. In other words, these graphs can be obtained through tractography and assuming that each voxel (or region of interest) is a node and that every tract is an edge connecting two nodes. These graphs we refer to as *direct graphs* since the edges are the direct connections between each node. One can think of these as friendship networks where an edge between two individuals is determined on whether they are friends or not.

Another approach to constructing a graph is to compute a measure of similarity between a feature of interest of one region and the same feature of another. For structural data, the whole output of a tractography algorithm (the tractogram of a single voxel) can be considered to be a feature of a voxel and is compared pairwise to obtain a similarity matrix between voxels (Bajada et al., 2017; Cerliani et al., 2012; Devlin et al., 2006; Johansen-Berg et al., 2004). For functional data, one can use the fMRI time series. In this case, the nodes are still voxels, but the edges are weights of how similar one voxel's feature is to another's. If we consider that two people are nodes in a network their edge weight would be determined by how similar two individuals are, based on individual features (e.g. dress sense, job, etc …). We call this a *feature similarity graph*. For simplicity, this is the type of graph that we will be discussing in the rest of this article. It is important to note that much work in the fMRI literature performs the similarity computation not on the features themselves (e.g. the time-series), but on a "functional connectivity" matrix (c.f. Margulies et al., 2016). In terms of the "people network" proposed above, if we assume that the "functional connectivity" gives us information about the "level of friendship" between two individuals, then the similarity matrix of this last approach indicates the similarity between each individual's friendship network.

3. The similarity, affinity or adjacency matrix

The adjacency matrix is a simple mathematical representation of a graph that describes the structure of the connectivity in the graph, that is,

whether nodes are connected or not. A more detailed description is provided by using a weighted adjacency matrix. The question is then how to define the weights, which in turn depends on what kind of graph we want to describe.

In neuroimaging the weights can be defined in terms of a *similarity metric* describing to what extent a feature of one voxel, or vertex (e.g. an fMRI time series or a set of streamlines) is similar to every other voxel in the region of interest. This is done across all voxels (or vertices). We will refer to such a weighted adjacency matrix as the *similarity* or *affinity* matrix.

Choosing a similarity metric is extremely important since it will affect any clustering that may be done on the data. By far the most popular similarity measure between two voxels is the Pearson's correlation coefficient, which can be interpreted as a centred and normalised dot product (Cerliani et al., 2012; Craddock et al., 2012; Devlin et al., 2006; Johansen-Berg et al., 2004; Klein et al., 2007; Zhang et al., 2014); see next section for a discussion of the dot product and other similarity measures. In order to understand the utility of correlation as a similarity metric, a short description of its precursors is given below (cf O'Connor, 2012 for an intuitive review). We will then introduce a slight modification that we employ in our adjoining code.

One caveat of the algorithms used for these analyses is that the adjacency matrix must be non-negative. This is not automatically true for most similarity measures, including those introduced in section 4.1. Hence, some technique must be used to ensure that the corresponding similarity matrix only contains non-negative weights (Haak et al., 2018; Johansen-Berg et al., 2004; Von Luxburg, 2007). For example, Johansen-Berg et al. (2004) proposed to add a scalar constant to the similarity matrix to ensure that all values are positive, others have only kept positive values at some threshold (Margulies et al., 2016).

Once a similarity matrix has been created, it can be used to represent the graph that all computations are carried out on. A final consideration regarding the similarity matrix is whether the full set of similarities should be used or if the similarity matrix should be thresholded in some way (Von Luxburg, 2007). For example, all weights below some arbitrary value ε could be set to zero; the remainder of the weights can be retained or binarised (this will be the same as using a simple adjacency matrix). Another approach to limiting the neighbourhood is to restrict the weights to the k-nearest neighbours. Advantages of both these data reduction approaches are that they remove noisy weights and they sparsify the matrix, leading to faster and cheaper computations. For example, in fMRI voxels may have a very low correlation (weight) not because of any intrinsic functional connectivity, but because of noise.

3.1. Similarity measures

The most basic way to measure the similarity of two datasets (thought of as vectors) is the dot product of the two vectors (consider an fMRI time series or a three-dimensional image of a tract density map (or tractogram) that is flattened into one long vector).

Geometrically, the dot product of two vectors is the projection of one vector onto the other. There are many equivalent ways to calculate the dot product. For this paper, the one offers the most insight is

$$dot(\boldsymbol{x}, \boldsymbol{y}) = \sqrt{\sum_i x_i^2} \sqrt{\sum_i y_i^2} \cos(\boldsymbol{x}, \boldsymbol{y}).$$

In this form, the dot product has two components: the cosine of the angle between the two datasets (treated as vectors) and their magnitudes. This means that magnitude and angular similarity (as measured by the cosine of an angle) are confounded. In order to solve this problem, one can normalise the dot product by dividing by the magnitude of each dataset and that leaves us with the cosine function.

One problem with the cosine similarity is that it is sensitive to relative shifts in the data between samples (such as can occur in fMRI time series due to absolute signal differences that are of no interest). The most

common way to create a shift invariant similarity is to mean centre the data and then compute the cosine similarity, which is the sample Pearson's correlation coefficient:

$$corr(\boldsymbol{x}, \boldsymbol{y}) = \cos((\boldsymbol{x} - \bar{\boldsymbol{x}}), (\boldsymbol{y} - \bar{\boldsymbol{y}})),$$

where $\bar{\boldsymbol{x}}$ and $\bar{\boldsymbol{y}}$ are constant vectors the size as $\boldsymbol{x}$ and $\boldsymbol{y}$, where each element is the mean of $\boldsymbol{x}$ and $\boldsymbol{y}$, respectively. Being shift invariant is an appealing property of the correlation coefficient and is especially useful to compare variables that have different means. Examples of works in the literature that use the cosine similarity can be found in the following articles (Bajada et al., 2017; Hong et al., 2019; Jackson et al., 2020, 2017; Margulies et al., 2016).

One should remember that the cosine is a sinusoidal function. As a result, a cosine similarity, or a correlation, value of 0.5 does not have the neat interpretation that the angle between the two datasets is 45°. An easy solution to this is to calculate the angle between the two data sets by using the inverse cosine function (the arccosine), normalise by 90° (or π/2):

$$normAngle = \frac{cos^{-1}(\cos((\boldsymbol{x} - \bar{\boldsymbol{x}}), (\boldsymbol{y} - \bar{\boldsymbol{y}})))}{90}.$$

The above formula will measure a normalised "angular distance" between two datasets bound between 0 and 2. We can thus define the quantity, as follows

$$AngSim = 1 - normAngle.$$

This measure returns a value that has an almost identical interpretation to the correlation coefficient (or cosine similarity) but has the nice property that a value of $+0.5$ implies that the two datasets are half way between orthogonal and colinear while a value of -0.5 implies that the two datasets are half way between orthogonal and anti-colinear. Examples of works in the literature that use a normalised angle include (Larivière et al., 2020; Vos de Wael et al., 2018).

4. The spectral transformation and the graph laplacian

Once a similarity (also affinity or adjacency) matrix is computed, we have all the information that describes the relationships between individual nodes. Our next step is to embed our data into a low dimensional space (for the moment a one-dimensional line) where the nodes distances from each other and the centre of the space reflect the internodal affinity.

While we refer readers to the supplementary material (Appendix C) for an informal discussion of the problem, the process can be formulated as the solution of an optimisation problem where a suitable cost function $U(\boldsymbol{x})$ is minimised (described by Leskovec et al., 2014)

$$\widehat{\boldsymbol{x}} = \underset{\boldsymbol{x}}{\operatorname{argmin}}\{U(\boldsymbol{x})\}.$$

Shortly, $\widehat{\boldsymbol{x}}$ is the vector that minimises the cost function $U(\boldsymbol{x})$. Such cost function can be written as a weighted sum of squared internode (Euclidean) distances across all connected nodes,

$$U(\boldsymbol{x}) = \sum_{(i,j) \in E} W_{ij}(x_i - x_j)^2.$$

One can think of the weights W_{ij} as a measure of the relationship between two nodes, for example, cortical vertices. Using graph theory language, one can think of the above situation as defining a weighted graph where the value of each node represents the location of each individual and the weight of the edge connecting two nodes represent their relationship.

The minimisation of the above cost function means that long distances between pairs of vertices with a high relationship value (i.e., high weight) is penalised. As a result, a pair of vertices with a high weight will be placed close to each other, while a pair of vertices with a low weight

will be placed far apart. The aim is to find a positioning where the sum of costs associated with all pairs of vertices is at its minimum.

The above problem, however, is not well-posed. First, it has a trivial solution which is to place all vertices at the same location (x is a constant vector), which produces zero cost. While this satisfies the minimisation of our cost function, it is not useful since it tells us nothing about the relationships between vertices. Second, if $\hat{x}$ is a solution then any shifting or re-scaling of $\hat{x}$ by a constant value c (i.e., $\hat{y} = \hat{x} + c$ or $\hat{y} = c\hat{x}$, respectively) will also be a solution because the resulting cost function is invariant to shifting or re-scaling. Therefore, in order to circumvent trivial and non-unique solutions some constraints are required. The simplest constraint is that the cost function must be minimised subject to (s.t.) the magnitude of the position vector x being equal to 1. Mathematically this is written as

$$\hat{x} = \underset{x}{\mathrm{argmin}}\{U(x)\} \quad s.t. \quad x^T x = 1.$$

Note that this constraint does not solve the problem of a constant solution completely since a constant vector, can still produce a zero cost and also satisfy the constraint. We will see later that since this solution is known in advance, one can easily account for it after optimisation. In general, the constrained minimisation problem can be solved using the method of Lagrange multipliers (cf Hagen and Kahng, 1992). While a detailed explanation of Lagrangian multipliers is beyond the scope of this text, the modification makes the computation easier to solve while maintaining its accuracy. The idea of this method is simple: We incorporate the constraints into the cost function itself. Thus, we rewrite the problem as

$$\hat{x} = \underset{x}{\mathrm{argmin}}\{U(x) + \lambda(1 - x^T x)\}.$$

Now, any putative solution that does not conform to the constraint imposed will be penalised. This penalty is dictated by the weight λ, often called the Lagrange multiplier. We proceed to solve the problem in the following way. First, rewrite the optimisation problem as follows (see Appendix A for details):

$$\hat{x} = \underset{x}{\mathrm{argmin}}\{\tilde{U}(x)\} = \underset{x}{\mathrm{argmin}}\{x^T L x + \lambda(1 - x^T x)\}.$$

The minimisation can now proceed in the usual way by taking the first derivative of the modified cost function $\tilde{U}(x)$ and equating to zero (extremum condition)

$$\partial \frac{\tilde{U}(x)}{\partial x} = 2Lx - 2\lambda x = 0,$$

and hence

$$Lx = \lambda x.$$

The last expression defines a *standard eigenvalue problem* for the Laplacian that can be solved using standard numerical libraries (e.g., the MATLAB function *eig*). The pair $(\lambda, \hat{x})$ is called an eigenpair, with λ called the eigenvalue, and $\hat{x}$ the eigenvector. Eigenvalues and eigenvectors are useful in a broad range of applications, with the interpretation of these pairs depending on the context in which they are used.

In this paper, we will focus on the eigenvectors of the Laplacian, as they contain the information which we will use to create our gradients. However, the eigenvalues also encode important information. Given that the Laplacian matrix is positive semi-definite, the smallest eigenvalue is zero and its associated eigenvector is a constant, thus, we will focus upon the second smallest eigenvalue, which is termed the algebraic connectivity of a graph (Fiedler, 1973).

Fiedler (1973) showed that the magnitude of the algebraic connectivity reflects how well connected the overall graph is, i.e., the larger the algebraic connectivity is, the more difficult it is to cut a graph into independent components. If the algebraic connectivity is zero it means that the graph is not connected; i.e. there are at least two graph partitions. In other words, if a graph has at least two hard clusters (i.e. it is two completely disconnected subgraphs), the algebraic connectivity will be zero. The more connected a graph gets, the higher the algebraic connectivity becomes. This intuition will be revisited in section 5.3.

The fact that the first eigenvalue is zero directly dictates that its associated eigenvector does not carry any useful information regarding the relative position of the nodes. Hence, the optimal solution is encoded in the eigenvector associated with the second smallest eigenvalue. This is called the Fiedler vector after the mathematician who first described this solution in the context of graph partitioning (Fiedler, 1973).

At this stage, it is worth noting that the described solution to the problem is biased in the sense that nodes with high degree will dominate the minimisation since the corresponding row (or column) of the Laplacian matrix is dominant. This means that nodes with a high number of neighbours (i.e., high degree) will tend to be grouped together irrespective of their similarity. This bias can be compensated for by using a modified constraint $x^T D x = 1$ so that our optimisation problem is transformed to

$$\hat{x} = \underset{x}{\mathrm{argmin}}\{x^T L x + \lambda(1 - x^T D x)\}.$$

The new constraint means that nodes are penalised (i.e., they are assigned a higher cost) according to their degree (Johansen-Berg et al., 2004). By following the same mathematical derivation, as above, the associated eigenvalue problem is then

$$Lx = \lambda D x.$$

This is known as the *generalised eigenvalue problem* for matrices L and D, which can also be solved numerically using standard toolboxes. Commonly, toolboxes that are able to solve the standard eigenvalue problem can also be used to solve the generalised problem. This is the case, for instance, for MATLAB's and Scipy's *eig* function.

Often, the Laplacian matrix is used in normalised form (i.e., normalised with respect to the nodes degree), so that its diagonal elements are all one. However, it can be demonstrated that normalising the Laplacian is equivalent to changing the constraint of the minimisation problem in some way and therefore one must be clear of how a given normalisation affects the solution. Several versions of the normalised Laplacian have been used in the literature. In Appendix B we describe the symmetric normalised Laplacian and the random walk normalised Laplacian.

5. Reordering, eigenmaps, and the Vogt-Bailey Index

If we think of brain voxels, or cortical surface vertices, as nodes with associated features (such as an fMRI time series, or a tractogram etc …) and the relationships between these features as edges on a graph, we previously described that the second smallest eigenvector describes the location (coordinate) of each node in one dimension, a line, which is dictated by each nodes relationship (affinity) to each-other. Using the location as a heatmap value becomes a way to visualise those relationships on the brain (the so-called macroscale gradients). Further, the components of eigenvectors denote the coordinates of the node in a space containing as many dimensions as there are eigenvectors (it is not restricted to a single dimension). Hence the eigenvector with the second smallest eigenvalue would give coordinates of the nodes on a line, the second and third eigenvectors would give the coordinates on a plane and so on. For this, more complicated visualisations are needed. One may also present the higher dimensional gradients independently, but one must always remember that the second gradient is influenced by the first and the third by the previous two etc.

Further, the algebraic connectivity indicates the sharpness of the best split (or cluster) in the region of interest. If a searchlight VB index analysis is performed on local neighbourhoods (see section 5.3), we can

investigate mesoscopic gradients (or transitions between areal borders).

5.1. Spectral reordering

The simplest approach to mapping the embedded location onto the brain is by assigning each voxel or vertex a heatmap value that corresponds to their rank order in the embedding. This approach allows for investigating the general pattern of changes in features across the brain, but, being rank ordered, does not provide any details about the actual feature distance between vertices (c.f. Bajada et al., 2017 for an example; c.f. Johansen-Berg et al., 2004 for the original spectral reordering paper in the literature where it was used for parcellation).

5.2. Eigenmaps

Laplacian eigenmaps (Belkin and Niyogi, 2003, 2002) are closely related to spectral reordering. Use of eigenmaps has been introduced to the neuroscience literature (Cerliani et al., 2012; Haak et al., 2018). In one dimension, the approach uses the coordinate points given by the primary eigenvector of the Laplacian as the intensity of the voxel of interest.

This approach can be particularly advantageous since one can explore the relationships between voxels in more than one dimension. Indeed, coordinates of the similarity or eigen-space can be mapped into a colour palette and the resultant colour map value can be mapped onto the brain space by assigning that value to the corresponding voxel (c.f. Bajada et al., 2019). This means that one can only map as many dimensions as the dimensions of the colour palette (in our case the 3-dimensional RGB colour palette).

Reordering and eigenmaps give us maps of optimal embedding of voxels in a low dimensional space. Effectively, voxels (or ROIs) with a similar value, have a greater affinity to one another. This establishes the large-scale organisational gradients of the cortex. It also gives some indication as to whether there are sharp discontinuities across that gradient but a focus on the eigenvectors alone fails to quantify the extent of discontinuity in cortical intra-areal relationships. The eigenvalues provide a solution.

5.3. The algebraic connectivity and the Vogt-Bailey Index

The algebraic connectivity of a graph is an indicator of how "well connected" that graph is. It is the second smallest eigenvalue of the Laplacian matrix (see section 4).

Once normalised to be bounded between zero and one (by dividing by the mean of all eigenvalues save for the first, which is the maximum value a graph with an affinity matrix one ones would have), the algebraic connectivity can be used as an indicator that a particular neural region has at least one sharp delineation or comprises only graded differences. This allows for a quantification of the historical issue about the degree of interareal transitions present in the cortex. While the Vogts primarily argued for clearly demarcated brain areas, Brodmann, in his 1909 monograph clearly stated that some areal cytoarchitectonic boundaries were graded. In the extreme, Bailey and von Bonin (1951) argued for an effectively graded cortex (with some minor exceptions). We thus propose the term "Vogt-Bailey Index" to describe the normalised algebraic connectivity of the graph Laplacian when used to describe the extent of feature similarity in a neuroscientific context.

Such an approach can be done across the entire cortex to give a single value for the "gradedness" of the whole cortex, across predefined clusters (such as the resting state networks) or to give a value per region of interest that quantifies how similar features in the region are. We note, however, that this value alone tells us little since the value will be affected by smoothing (which exist in MRI signals). It can, however, be used as a relative measure where one can compare regions across the same brain or between different subjects (see section 5.3. for notes on statistical analysis).

Finally, one can use a vertex-wise searchlight to calculate the Vogt-Bailey (VB) Index across the entire cortex. Using this approach, a neighbourhood of adjoining cortical voxels, or vertices (as is assumed in the adjoining code) is calculated and the (normalised) algebraic connectivity of its affinity graph is calculated. The calculated VB index gives a value of how similar a feature (fMRI, tractography, or others) in the centre of the searchlight is to its nearest neighbours. The final result is effectively a cortical edge detection algorithm (see Figs. 2 and 3 for intuitive examples) where boundaries between parcels should emerge
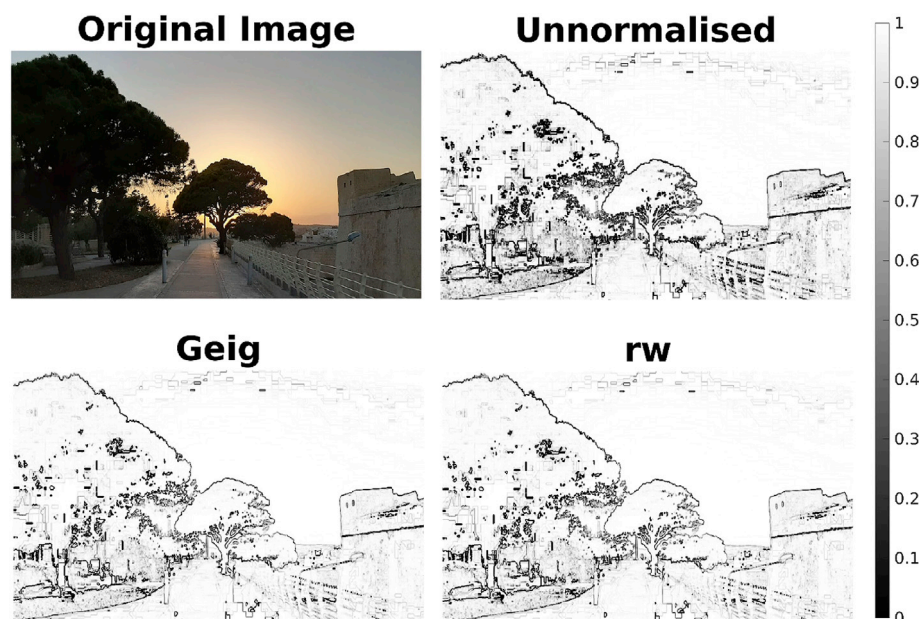
Fig. 2. The VB Index applied to a photograph (top left) using all three normalisations of tha Laplacian matrix. The colormap ranges from black (0) where there are sharp transitions to white (1) where there is homogenous structure in the image.
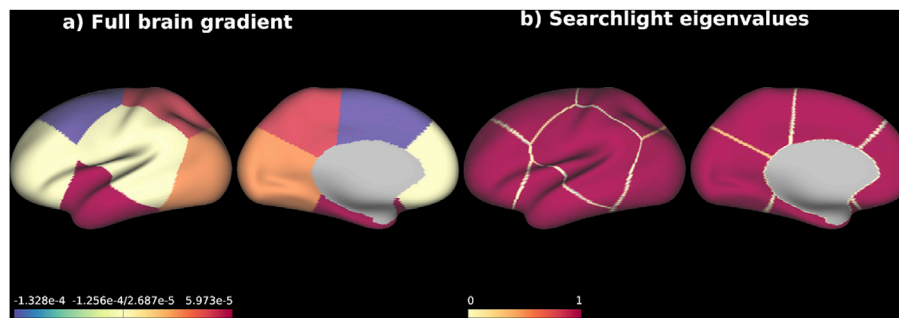
Fig. 3. a) The gradient map on synthetic data showing the expected pattern with values within parcels being extremely similar but different across parcels. b) The VB Index on a cortical surface highlighting the arbitrary parcels. All results show the default generalised eigenvalue problem solution.

naturally and their relative sharpness should be apparent. Of course, smoothing effects and voxel/vertex size will limit the resolution that one can expect. Indeed, such notions in MRI analysis are not new. The idea of a measure of regional homogeneity (ReHo) has been present since the early days of fMRI (Jiang and Zuo, 2016; Zang et al., 2004). Further, the approach has similarities to the observer independent method for microstructural parcellations (Schleicher et al., 1999). Our approach simply fits these ideas of regional homogeneities and boundaries into a flexible and more general framework that does not restrict either the metric for similarity that is used (such as the method for microstructural parcellation) or the feature of interest (such as in both ReHo and the observationally independent approach for microstructural parcellation).

In summary, the VB Index is the proposed term for the normalised algebraic connectivity of the graph Laplacian when used to describe the extent of feature similarity in a neuroscientific context. The adjoining software can produce three "types" of VB Indices: 1) the full brain analysis which also computes a whole brain gradient single VB Index for the whole brain; 2) clustered analysis that computes a gradient and VB Index per region of interest, and 3) the searchlight VB Index which computes a VB Index per vertex based on the neighbourhood data of directly adjacent vertices. The size of the cluster (going from the nearest connected 5 or 6 neighbours to the full brain) is relevant in the interpretation of the VB Index. As Fielder (1973) showed in the original paper, the value of algebraic connectivity provides a measure of how difficult it is to split the graph (higher value indicating more "connectedness", i.e. more difficult to split). If it is zero, then it indicates that there is at least one complete split in the graph. Hence as the cluster gets larger there is more of a chance that the graph will be easier to split into two, that is what the VB Index measures. In the case of the searchlight, since we are only looking at 5 or 6 connected neighbours, the interpretation is more straightforward: is there an edge near that vertex?

6. Notes on statistical analysis

This article focused on creating a conceptual understanding of the large-scale gradients and the quantification of boundary edges using MRI data.

With respect to the statistical treatment of large-scale feature gradients, a literature is emerging that discusses various statistical approaches to use on gradient maps. The reader is encouraged to explore the articles in this section for current approaches on gradient statistics (Haak et al., 2018; Hong et al., 2019; Langs et al., 2015; Tian and Zalesky, 2018; Xu et al., 2019).

Regarding the statistical analysis of the VB Index maps, future work is needed to disentangle the effects of MR noise and inherent smoothness from real gradations in feature similarity. A research avenue for noise removal includes the generation of null models with similar noise and baseline smoothness as the underlying MRI data (c.f. Gordon et al., 2016; Tian and Zalesky, 2018).

The VB toolbox is a freely available, open source, project under the terms of a GPL licence, we hope that with interest growing in the field of Gradient analysis that the toolbox will grow to also incorporate various statistical approaches for making inferences on both gradient maps and the VB Index.

7. Experiments

7.1. Photography example: searchlight VB index

Before applying the VB Index to the rather abstract notion of function MRI, we have provided a MATLAB script within the respective version of the toolbox that implements the VB Index on a colour photograph (http s://github.com/VBIndex/matlab_vb_toolbox/tree/master/vb_index_int uitive_example). Every pixel within the 2D photo can be thought of as a vertex within a brain surface. The functional data is represented by the hue, saturation and their brightness value of the pixel. Performing the VB Index searchlight operation on the photograph, as described in Section 5.3. results in a quantification of boundaries of the image (See Fig. 2). Readers are invited to explore this script with other images made freely available or try it out using their own photographs.

7.2. A neuroimaging example: simulated MRI data example

Following the validation of the technique on colour photographs, the performance of the technique was evaluated on synthetic MRI data where the ground truth is known. To this end a cortical surface from the HCP dataset was arbitrarily split into 6 contiguous parcels. The vertices within the same parcel were assigned identical time-series, which differed across different parcels. The analysis was carried out using version 1.1.0 of the python *vb_tool*. Fig. 3 show the results of applying the proposed method to the simulated data. As expected, the full brain gradient shows a piece-wise constant pattern, reflecting the similarity structures between parcels (unknown) as described above (Fig. 3a). Consistent with this result, the vertex-wise VB Index shows a pattern where the edges between parcels are highlighted (Fig. 3b).

7.3. A neuroimaging example: human connectome project data example

The adjoining toolbox was run on twenty-four (24) individuals (12 F) from the human connectome project database. The calculations were carried out on two separate rs-fMRI runs per participant across both hemispheres. The dataset was pre-processed by the HCP using the minimal processing pipeline (Glasser et al., 2013). The data collection was approved by the Washington University Institutional Review Board (IRB) and further approval for processing the data was obtained by the University of Malta's University Research Ethics Committee.

The data were processed according to the procedures outlined in the above text using version 1.1.0 of the python *vb_tool*; all calculations used the generalised eigenvalue problem for computations.

First, the whole brain affinity matrix was computed for all 24 subjects
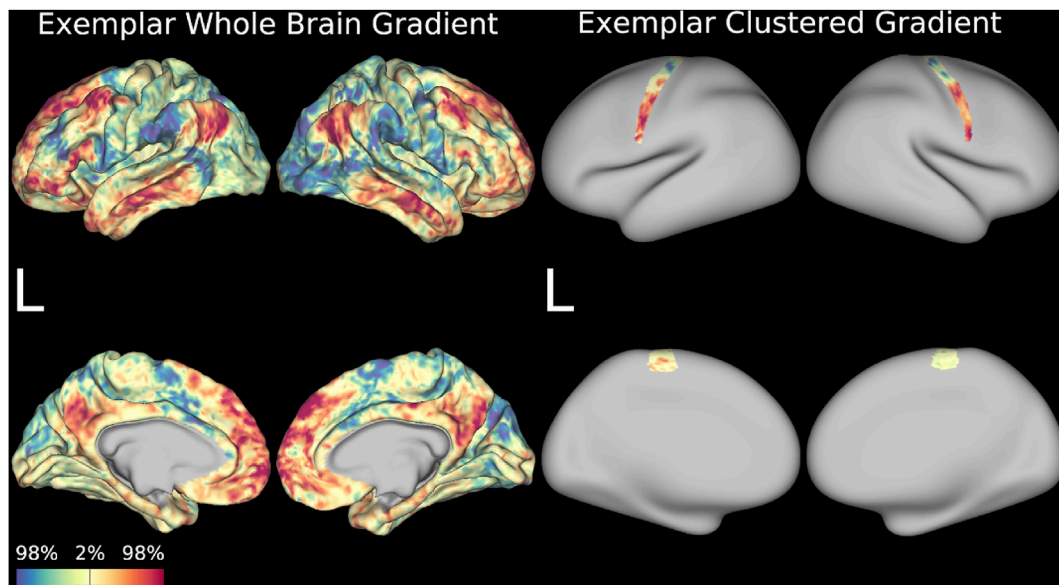
Fig. 4. An exemplar of the principal similarity gradient across the whole cortex based on rs-fMRI as a feature (left). An exemplar of the principal gradient computed on a pre-clustered cortex (using the HCP Multimodal Parcellation, right).
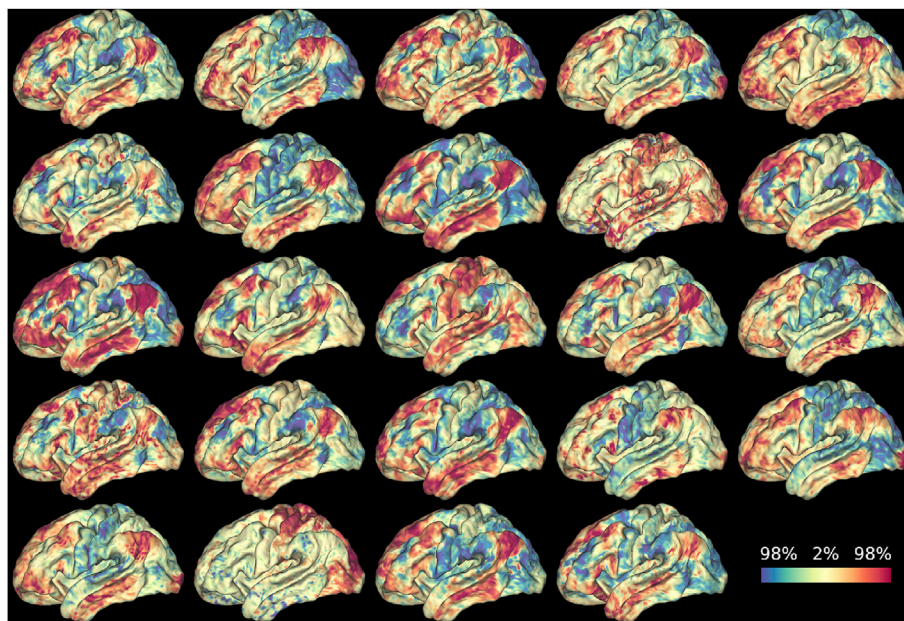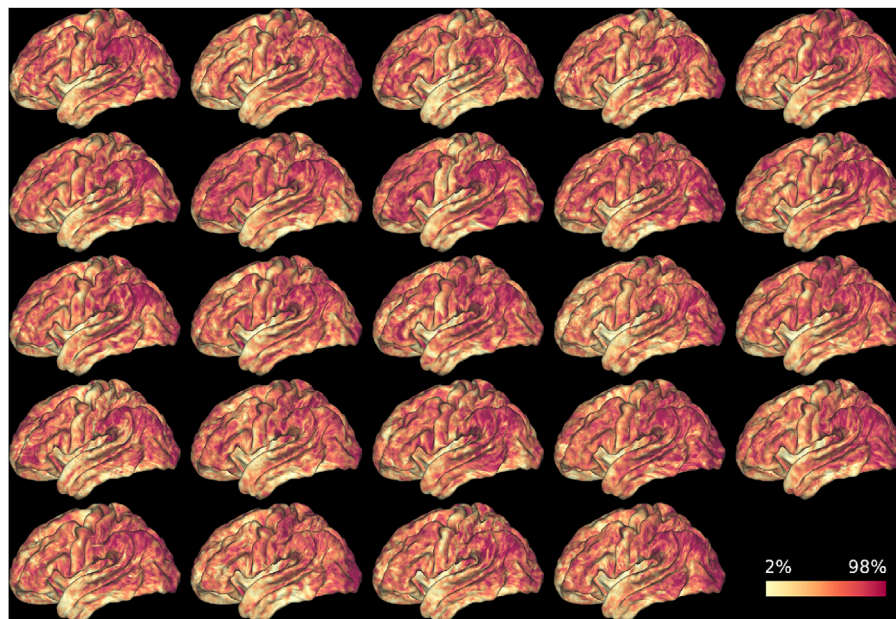

Fig. 5. The principle similarity gradient across 24 individuals on a single run. The image of the same participants on a second fMRI run can be found in the supplementary material.

(per run) and an eigendecomposition of its Laplacian was computed. This resulted in Fig. 4 (left) which shows an exemplar of the primary large scale inter-areal (feature) gradient of the whole cortex in both hemispheres of a single subject. Fig. 5 shows all 24 subjects on the lateral surface of the left hemisphere for a single rs-fMRI run.

We note that the principle gradients differ somewhat from previously published work (c.f. Margulies et al., 2016). This may be due to multiple differences in parameter choices that were made (and discussed briefly in section 2.1). In short, the principle gradient is highly dependent upon the properties of the affinity matrix; the most pertinent difference between our affinity matrix and that of Margulies et al. (2016) was the latter's retention of only the top 10% of functional connections and subsequent re-computation of a cosine similarity while our approach (c.f. Jackson et al., 2020, 2017 for a slightly modified example) accepts all positive correlations that were then transformed to a normalised angular distance. Thresholding plays an important role in the interpretation of the results. A high threshold (such as retaining only the top 10% of connections) will only consider the similarity of "well-connected" vertices giving no weight to moderately and poorly connected ones. Our approach, which only eliminates negative weights, takes these connections into account but would also be more sensitive to "spurious" connections.

Second, the toolbox was used to calculate the primary gradients and their associated VB Index for the data parcellated using the Multimodal HCP parcellation (Glasser et al., 2016). An exemplar of these results can be found in Fig. 3 (right), where the principle gradient is computed in each parcel. Associated with these parcels are the VB Indices per parcel (see Fig. 7 right and a further discussion below).

Fig. 6. The VB Index computed across 24 individuals on a single run. The image of the same participants on a second fMRI run can be found in the supplementary material.
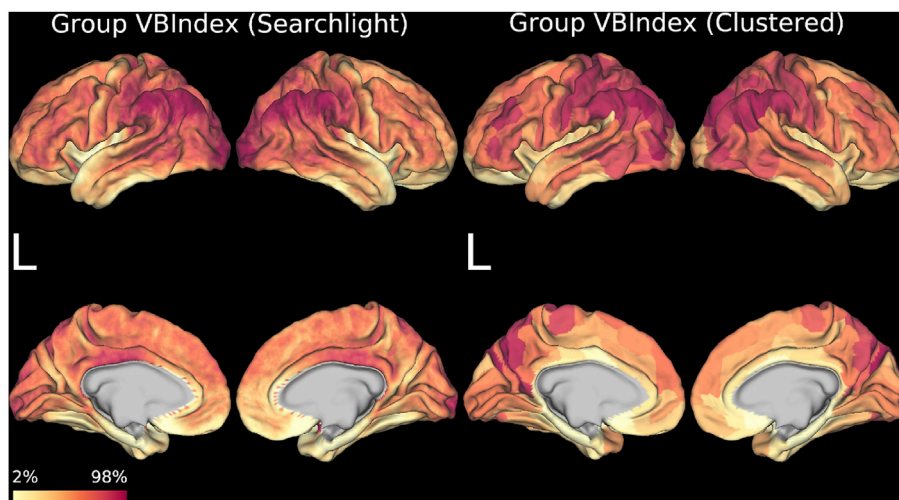
Fig. 7. An average (on a single run) of 24 individual's searchlight, local neighbourhood, whole brain VB-index identifying regions of relatively sharper borders across the cortex.

Finally, the vertex-wise searchlight VB Index was computed on all participants. This approach highlights, in a data-driven fashion, the feature edges and boundaries across the cortex. Fig. 6 shows the searchlight VB Index across all 24 participants while Fig. 7 (left) shows the mean vertex-wise VB Index. One can also see similar patterns between the group vertex-wise VB Index (Fig. 4 left) and the group VB Index computed on clusters (Fig. 7 right).

The full set of results can be found on the HCP BALSA database (https ://balsa.wustl.edu/study/show/kND1N).

8. Conclusion

The idea of gradations in neural features has been present since at least the beginning of the twentieth century and has gained traction in the neuroimaging community throughout the early twenty-first. This paper has outlined the general concepts and mathematical intuition behind the spectral transformation and has introduced the related

techniques of spectral reordering, Laplacian eigenmaps and clustering. As an accompaniment to this paper, MATLAB and Python tools that performs the different spectral transformations discussed in Section 4 are available. Depending on the size of the data, the technique can take up a considerable amount of RAM and computation time, however, at standard mesh sampling our attached code can run a full brain gradient analysis (using HCP 32K surfaces) on a standard desktop or notebook with 32 GB of RAM.

While the described framework can be used to reason about relationships between neural features, there are plenty of unanswered questions. The first important issue regards the choice of similarity measurement. Although all the metrics discussed above have been used to some extent, a systematic comparison along with guidance regarding which metric to use in different circumstances is needed.

In summary, it is hoped that this article and accompanying tools will be used as a guide to researchers interested in performing anatomical investigations using neural features and their interareal relationships in

the brain.

Declaration of competing interest

GJMP is a shareholder and director of Bioxydyn, a company with an interest in imaging services and products.

CRediT authorship contribution statement

Claude J. Bajada: Conceptualization, Data curation, Formal analysis, Investigation, Methodology, Resources, Software, Validation, Visualization, Writing - original draft, Writing - review & editing. **Lucas Q. Costa Campos:** Data curation, Formal analysis, Software, Validation, Writing - review & editing. **Svenja Caspers:** Funding acquisition, Resources, Supervision, Writing - review & editing. **Richard Muscat:** Resources, Supervision, Writing - review & editing. **Geoff J.M. Parker:** Conceptualization, Methodology, Supervision, Writing - review & editing. **Matthew A. Lambon Ralph:** Funding acquisition, Conceptualization, Supervision, Writing - review & editing. **Lauren L. Cloutman:** Funding acquisition, Conceptualization, Supervision, Writing - review & editing. **Nelson J. Trujillo-Barreto:** Formal analysis, Conceptualization, Methodology, Supervision, Writing - original draft, Writing - review & editing.

Acknowledgments

Appendix A. Supplementary data

Supplementary data to this article can be found online at https://doi.org/10.1016/j.neuroimage.2020.117140.

References

Amunts, K., Kedo, O., Kindler, M., Pieperhoff, P., Mohlberg, H., Shah, N.J., Habel, U., Schneider, F., Zilles, K., 2005. Cytoarchitectonic mapping of the human amygdala, hippocampal region and entorhinal cortex: intersubject variability and probability maps. Anat. Embryol. 210, 343–352. https://doi.org/10.1007/s00429-005-0025-5.

Bailey, P., von Bonin, G., 1951. The Isocortex of Man, Illinois Monographs in the Medical Sciences. University of Illinois Press, Urbana.

Bajada, C.J., Jackson, R.L., Haroon, H.A., Azadbakht, H., Parker, G.J.M., Lambon Ralph, M.A., Cloutman, L.L., 2017. A graded tractographic parcellation of the temporal lobe. Neuroimage. https://doi.org/10.1016/j.neuroimage.2017.04.016.

Bajada, C.J., Trujillo-Barreto, N.J., Parker, G.J.M., Cloutman, L.L., Lambon Ralph, M.A., 2019. A structural connectivity convergence zone in the ventral and anterior temporal lobes: data-driven evidence from structural imaging. Cortex 120, 298–307. https://doi.org/10.1016/j.cortex.2019.06.014.

Belkin, M., Niyogi, P., 2002. Laplacian eigenmaps and spectral techniques for embedding and clustering. Adv. Neural Inf. Process. Syst. 14, 585–591. Vols 1 and 2 14.

Belkin, M., Niyogi, P., 2003. Laplacian eigenmaps for dimensionality reduction and data representation. Neural Comput. 15, 1373–1396. https://doi.org/10.1162/089976603321780317.

Bludau, S., Eickhoff, S.B., Mohlberg, H., Caspers, S., Laird, A.R., Fox, P.T., Schleicher, A., Zilles, K., Amunts, K., 2014. Cytoarchitecture, probability maps and functions of the human frontal pole. Neuroimage 93 (Pt 2), 260–275. https://doi.org/10.1016/j.neuroimage.2013.05.052.

Brodmann, K., 1909. In: Brodmann's Localisation in the Cerebral Cortex, (Electronic bk.). Springer, New York, NY]. Springer, c2006.

Buckner, R.L., Krienen, F.M., 2013. The evolution of distributed association networks in the human brain. Trends Cognit. Sci. 17, 648–665. https://doi.org/10.1016/j.tics.2013.09.017.

Caspers, J., Zilles, K., Eickhoff, S.B., Schleicher, A., Mohlberg, H., Amunts, K., 2013. Cytoarchitectonical analysis and probabilistic mapping of two extrastriate areas of the human posterior fusiform gyrus. Brain Struct. Funct. 218, 511–526. https://doi.org/10.1007/s00429-012-0411-8.

Caspers, S., Geyer, S., Schleicher, A., Mohlberg, H., Amunts, K., Zilles, K., 2006. The human inferior parietal cortex: cytoarchitectonic parcellation and interindividual variability. Neuroimage 33, 430–448. https://doi.org/10.1016/j.neuroimage.2006.06.054.

Cerliani, L., Thomas, R.M., Jbabdi, S., Siero, J.C.W., Nanetti, L., Crippa, A., Gazzola, V., D'Arceuil, H., Keysers, C., 2012. Probabilistic tractography recovers a rostrocaudal trajectory of connectivity variability in the human insular cortex. Hum. Brain Mapp. 33, 2005–2034. https://doi.org/10.1002/hbm.21338.

Cloutman, L.L., Binney, R.J., Drakesmith, M., Parker, G.J.M., Lambon Ralph, M.A., 2012. The variation of function across the human insula mirrors its patterns of structural connectivity: evidence from in vivo probabilistic tractography. Neuroimage 59, 3514–3521. https://doi.org/10.1016/j.neuroimage.2011.11.016.

Coifman, R.R., Lafon, S., 2006. Diffusion maps. Appl. Comput. Harmon. Anal. 21, 5–30. https://doi.org/10.1016/j.acha.2006.04.006.

Craddock, R.C., James, G.A., Holtzheimer, P.E., Hu, X.P., Mayberg, H.S., 2012. A whole brain fMRI atlas generated via spatially constrained spectral clustering. Hum. Brain Mapp. 33, 1914–1928. https://doi.org/10.1002/hbm.21333.

Da Costa Campos, L., Bajada, C.J., 2020. VBIndex/py_vb_toolbox 1.0.0. Zenodo. https://doi.org/10.5281/zenodo.3609460.

Damoiseaux, J.S., Rombouts, S.A.R.B., Barkhof, F., Scheltens, P., Stam, C.J., Smith, S.M., Beckmann, C.F., 2006. Consistent resting-state networks across healthy subjects. Proc. Natl. Acad. Sci. U.S.A. 103, 13848–13853. https://doi.org/10.1073/pnas.0601417103.

Devlin, J.T., Sillery, E.L., Hall, D.A., Hobden, P., Behrens, T.E.J., Nunes, R.G., Clare, S., Matthews, P.M., Moore, D.R., Johansen-Berg, H., 2006. Reliable identification of the auditory thalamus using multi-modal structural analyses. Neuroimage 30, 1112–1120. https://doi.org/10.1016/j.neuroimage.2005.11.025.

Eickhoff, S.B., Thirion, B., Varoquaux, G., Bzdok, D., 2015. Connectivity-based parcellation: critique and implications. Hum. Brain Mapp. 36, 4771–4792. https://doi.org/10.1002/hbm.22933.

Geyer, S., Turner, R. (Eds.), 2013. Microstructural Parcellation of the Human Cerebral Cortex: from Brodmann's Post-Mortem Map to in Vivo Mapping with High-Field Magnetic Resonance Imaging, Illustrated. Springer Science & Business Media.

Glasser, M.F., Coalson, T.S., Robinson, E.C., Hacker, C.D., Harwell, J., Yacoub, E., Ugurbil, K., Andersson, J., Beckmann, C.F., Jenkinson, M., Smith, S.M., Van Essen, D.C., 2016. A multi-modal parcellation of human cerebral cortex. Nature 536, 171–178. https://doi.org/10.1038/nature18933.

Glasser, M.F., Sotiropoulos, S.N., Wilson, J.A., Coalson, T.S., Fischl, B., Andersson, J.L., Xu, J., Jbabdi, S., Webster, M., Polimeni, J.R., Van Essen, D.C., Jenkinson, M., WU-Minn HCP Consortium, 2013. The minimal preprocessing pipelines for the Human Connectome Project. Neuroimage 80, 105–124. https://doi.org/10.1016/j.neuroimage.2013.04.127.

Gordon, E.M., Laumann, T.O., Adeyemo, B., Huckins, J.F., Kelley, W.M., Petersen, S.E., 2016. Generation and evaluation of a cortical area parcellation from resting-state correlations. Cerebr. Cortex 26, 288–303. https://doi.org/10.1093/cercor/bhu239.

Guell, X., Goncalves, M., Kaczmarzyk, J.R., Gabrieli, J.D.E., Schmahmann, J.D., Ghosh, S.S., 2019. LittleBrain: a gradient-based tool for the topographical interpretation of cerebellar neuroimaging findings. PloS One 14, e0210028. https://doi.org/10.1371/journal.pone.0210028.

Haak, K.V., Marquand, A.F., Beckmann, C.F., 2018. Connectopic mapping with resting-state fMRI. Neuroimage 170, 83–94.

Hagen, L., Kahng, A.B., 1992. New spectral methods for ratio cut partitioning and clustering. IEEE Trans. Comput. Aided Des. Integrated Circ. Syst. 11, 1074–1085. https://doi.org/10.1109/43.159993.

Hong, S.-J., Vos de Wael, R., Bethlehem, R.A.I., Lariviere, S., Paquola, C., Valk, S.L., Milham, M.P., Di Martino, A., Margulies, D.S., Smallwood, J., Bernhardt, B.C., 2019. Atypical functional connectome hierarchy in autism. Nat. Commun. 10, 1022. https://doi.org/10.1038/s41467-019-08944-1.

Jackson, R.L., Bajada, C.J., Lambon Ralph, M.A., Cloutman, L.L., 2020. The graded change in connectivity across the ventromedial prefrontal cortex reveals distinct subregions. Cerebr. Cortex 30, 165–180. https://doi.org/10.1093/cercor/bhz079.

Jackson, R.L., Bajada, C.J., Rice, G.E., Cloutman, L.L., Lambon Ralph, M.A., 2017. An emergent functional parcellation of the temporal cortex. Neuroimage. https://doi.org/10.1016/j.neuroimage.2017.04.024.

Jiang, L., Zuo, X.-N., 2016. Regional homogeneity: a multimodal, multiscale neuroimaging marker of the human connectome. Neuroscientist 22, 486–505. https://doi.org/10.1177/1073858415595004.

Johansen-Berg, H., Behrens, T.E.J., Robson, M.D., Drobnjak, I., Rushworth, M.F.S., Brady, J.M., Smith, S.M., Higham, D.J., Matthews, P.M., 2004. Changes in connectivity profiles define functionally distinct regions in human medial frontal cortex. Proc. Natl. Acad. Sci. U.S.A. 101, 13335–13340. https://doi.org/10.1073/pnas.0403743101.

Klein, J.C., Behrens, T.E.J., Robson, M.D., Mackay, C.E., Higham, D.J., Johansen-Berg, H., 2007. Connectivity-based parcellation of human cortex using diffusion MRI: establishing reproducibility, validity and observer independence in BA 44/45 and SMA/pre-SMA. Neuroimage 34, 204–211. https://doi.org/10.1016/j.neuroimage.2006.08.022.

Langs, G., Golland, P., Ghosh, S.S., 2015. Predicting activation across individuals with resting-state functional connectivity based multi-atlas label fusion. Med. Imag. Comput. Comput. Assist. Interv. 9350, 313–320. https://doi.org/10.1007/978-3-319-24571-3_38.

Larivière, S., Vos de Wael, R., Hong, S.-J., Paquola, C., Tavakol, S., Lowe, A.J., Schrader, D.V., Bernhardt, B.C., 2020. Multiscale structure-function gradients in the neonatal connectome. Cerebr. Cortex 30, 47–58. https://doi.org/10.1093/cercor/bhz069.

Leskovec, A., Rajaraman, A., Ullman, J.D., 2014. Mining of Massive Datasets. Cambridge University Press, Cambridge. https://doi.org/10.1017/CBO9781139058452.

Margulies, D.S., Ghosh, S.S., Goulas, A., Falkiewicz, M., Huntenburg, J.M., Langs, G., Bezgin, G., Eickhoff, S.B., Castellanos, F.X., Petrides, M., Jefferies, E., Smallwood, J., 2016. Situating the default-mode network along a principal gradient of macroscale cortical organization. Proc. Natl. Acad. Sci. U.S.A. 113, 12574–12579. https://doi.org/10.1073/pnas.1608282113.

Mesulam, M., 2008. Representation, inference, and transcendent encoding in neurocognitive networks of the human brain. Ann. Neurol. 64, 367–378. https://doi.org/10.1002/ana.21534.

Mesulam, M.M., 1998. From sensation to cognition. Brain 121 (Pt 6), 1013–1052. https://doi.org/10.1093/brain/121.6.1013.

Nieuwenhuys, R., Broere, C.A.J., Cerliani, L., 2015. A new myeloarchitectonic map of the human neocortex based on data from the Vogt-Vogt school. Brain Struct. Funct. 220, 2551–2573. https://doi.org/10.1007/s00429-014-0806-9.

O'Connor, B., 2012. Cosine similarity, Pearson correlation, and OLS coefficients [WWW document]. AI and social science. URL. http://brenocon.com/blog/2012/03/cosine-similarity-pearson-correlation-and-ols-coefficients/. accessed 6.1.16.

O'Donnell, L.J., Golby, A.J., Westin, C.-F., 2013. Fiber clustering versus the parcellation-based connectome. Neuroimage 80, 283–289. https://doi.org/10.1016/j.neuroimage.2013.04.066.

O'Donnell, L.J., Kubicki, M., Shenton, M.E., Dreusicke, M.H., Grimson, W.E.L., Westin, C.F., 2006. A method for clustering white matter fiber tracts. AJNR Am. J. Neuroradiol. 27, 1032–1036.

Paquola, C., Seidlitz, J., Benkarim, O., Royer, J., Klimes, P., Bethlehem, R.A.I., Lariviere, S., Vos de Wael, R., Hall, J.A., Frauscher, B., Smallwood, J., Bernhardt, B.C., 2020. The cortical wiring scheme of hierarchical information processing. BioRxiv. https://doi.org/10.1101/2020.01.08.899583.

Paquola, C., Vos De Wael, R., Wagstyl, K., Bethlehem, R.A.I., Hong, S.-J., Seidlitz, J., Bullmore, E.T., Evans, A.C., Misic, B., Margulies, D.S., Smallwood, J., Bernhardt, B.C., 2019. Microstructural and functional gradients are increasingly dissociated in transmodal cortices. PLoS Biol. 17, e3000284 https://doi.org/10.1371/journal.pbio.3000284.

Rottschy, C., Eickhoff, S.B., Schleicher, A., Mohlberg, H., Kujovic, M., Zilles, K., Amunts, K., 2007. Ventral visual cortex in humans: cytoarchitectonic mapping of two extrastriate areas. Hum. Brain Mapp. 28, 1045–1059. https://doi.org/10.1002/hbm.20348.

Scheperjans, F., Hermann, K., Eickhoff, S.B., Amunts, K., Schleicher, A., Zilles, K., 2008. Observer-independent cytoarchitectonic mapping of the human superior parietal cortex. Cerebr. Cortex 18, 846–867. https://doi.org/10.1093/cercor/bhm116.

Schleicher, A., Amunts, K., Geyer, S., Morosan, P., Zilles, K., 1999. Observer-independent method for microstructural parcellation of cerebral cortex: a quantitative approach to cytoarchitectonics. Neuroimage 9, 165–177. https://doi.org/10.1006/nimg.1998.0385.

Tian, Y., Zalesky, A., 2018. Characterizing the functional connectivity diversity of the insula cortex: subregions, diversity curves and behavior. Neuroimage 183, 716–733. https://doi.org/10.1016/j.neuroimage.2018.08.055.

Vázquez-Rodríguez, B., Suárez, L.E., Markello, R.D., Shafiei, G., Paquola, C., Hagmann, P., van den Heuvel, M.P., Bernhardt, B.C., Spreng, R.N., Misic, B., 2019. Gradients of structure-function tethering across neocortex. Proc. Natl. Acad. Sci. U.S.A. 116, 21219–21227. https://doi.org/10.1073/pnas.1903403116.

Vogt, C., Vogt, O., 1911. Nouvelle contribution à l'étude de la myéloarchitecture de l'écorce cérébrale. Presented at the XX. Congres des médecins aliénistes et neurologistes de France (Brüssel).

Von Luxburg, U., 2007. A tutorial on spectral clustering. Stat. Comput. 17, 395–416. https://doi.org/10.1007/s11222-007-9033-z.

von Economo, C., Koskinas, G.N., 1925. Die Cytoarchitektonik der Hirnrinde des erwachsenen Menschen. Atlas mit 112 mikrophotographischen Tafeln inbesonderer Mappe.

Vos de Wael, R., Benkarim, O., Paquola, C., Lariviere, S., Royer, J., Tavakol, S., Xu, T., Hong, S.-J., Valk, S.L., Misic, B., Milham, M.P., Margulies, D.S., Smallwood, J., Bernhardt, B.C., 2019. BrainSpace: a toolbox for the analysis of macroscale gradients in neuroimaging and connectomics datasets. BioRxiv. https://doi.org/10.1101/761460.

Vos de Wael, R., Larivière, S., Caldairou, B., Hong, S.-J., Margulies, D.S., Jefferies, E., Bernasconi, A., Smallwood, J., Bernasconi, N., Bernhardt, B.C., 2018. Anatomical and microstructural determinants of hippocampal subfield functional connectome embedding. Proc. Natl. Acad. Sci. U.S.A. 115, 10154–10159. https://doi.org/10.1073/pnas.1803667115.

Xu, T., Nenning, K.-H., Schwartz, E., Hong, S.-J., Vogelstein, J.T., Fair, D.A., Schroeder, C.E., Margulies, D.S., Smallwood, J., Milham, M.P., Langs, G., 2019. Cross-species functional alignment reveals evolutionary hierarchy within the connectome. BioRxiv. https://doi.org/10.1101/692616.

Zang, Y., Jiang, T., Lu, Y., He, Y., Tian, L., 2004. Regional homogeneity approach to fMRI data analysis. Neuroimage 22, 394–400. https://doi.org/10.1016/j.neuroimage.2003.12.030.

Zhang, Yaqin, Fan, L., Zhang, Yu, Wang, J., Zhu, M., Zhang, Yuanchao, Yu, C., Jiang, T., 2014. Connectivity-based parcellation of the human posteromedial cortex. Cerebr. Cortex 24, 719–727. https://doi.org/10.1093/cercor/bhs353.

Zilles, K., Amunts, K., 2010. Centenary of Brodmann's map–conception and fate. Nat. Rev. Neurosci. 11, 139–145. https://doi.org/10.1038/nrn2776.