Meta-analytic decoding of the cortical gradient of functional connectivity

Julio A. Peraza[1], Taylor Salo[2], Michael C. Riedel[3], Katherine L. Bottenhorn[4], Jean-Baptiste Poline[5], Jérôme Dockès[5], James D. Kent[6], Jessica E. Bartley[1], Jessica S. Flannery[7], Lauren D. Hill-Bowen[8], Rosario Pintos Lobo[8], Ranjita Poudel[9], Kimberly L. Ray[6], Jennifer L. Robinson[10], Robert W. Laird[1], Matthew T. Sutherland[8], Alejandro de la Vega[6], Angela R. Laird[1]

[1] Department of Physics, Florida International University, Miami, FL, USA
[2] Department of Psychiatry, University of Pennsylvania, Philadelphia, PA, USA
[3] LTI Engineering and Software, Quebec City, QC, Canada
[4] Department of Population and Public Health Sciences, University of Southern California, Los Angeles, CA, USA
[5] Department of Neurology and Neurosurgery, McGill University, Montreal, QC, Canada
[6] Department of Psychology, University of Texas at Austin, Austin, TX, USA
[7] Department of Psychology and Neuroscience, University of North Carolina, Chapel Hill, NC, USA
[8] Department of Psychology, Florida International University, Miami, FL, USA
[9] Department of Health Outcomes and Behavior, Moffitt Cancer Center, Tampa, FL, USA
[10] Department of Psychology, Auburn University, Auburn, AL, USA

Keywords: fMRI, gradients, functional connectivity, functional decoding, meta-analysis, meta-analytic decoding

Supplemental Information for this study can be found underline{here}.

Abstract

Macroscale gradients have emerged as a central principle for understanding functional brain organization. Previous studies have demonstrated that a principal gradient of connectivity in the human brain exists, with unimodal primary sensorimotor regions situated at one end, and transmodal regions associated with the default mode network and representative of abstract functioning at the other. The functional significance and interpretation of macroscale gradients remains a central topic of discussion in the neuroimaging community, with some studies demonstrating that gradients may be described using meta-analytic functional decoding techniques. However, additional methodological development is necessary to more fully leverage available meta-analytic methods and resources and quantitatively evaluate their relative performance. Here, we conducted a comprehensive series of analyses to investigate and improve the framework of data-driven, meta-analytic methods, thereby establishing a principled approach for gradient segmentation and functional decoding. We found that a small number of segments determined by a K-means segmentation approach and an LDA-based meta-analysis combined with the NeuroQuery database was the optimal combination of methods for decoding functional connectivity gradients. Taken together, the current work aims to provide recommendations on best practices, along with flexible methods, for gradient-based functional decoding of fMRI data.

Corresponding Author: **Julio A. Peraza**, Florida International University, Miami, FL, USA | Email: jperaza@fiu.edu

Introduction

Recent methodological developments using nonlinear dimensionality reduction techniques on functional magnetic resonance imaging (fMRI) datasets, such as diffusion map embedding (Coifman et al., 2005), have allowed the mapping of vertex-level connectomes to low dimensional manifolds known as "macroscale gradients" (Vos de Wael et al., 2020). This representation of connectomes has complemented discrete parcellation methods by capturing a continuous transition between different functional systems, offering potential insight into hierarchical information processing in the brain. In a seminal publication, Margulies and colleagues demonstrated that a functional spectrum of the principal gradient of connectivity exists in the human brain, with unimodal primary sensorimotor regions situated at one end, and transmodal regions associated with the default mode network and representative of abstract functioning at the other end (Margulies et al., 2016). This work provides a foundation for understanding macroscale gradients within the framework of functional hierarchies. Since then, gradient methods have been used to describe the continuous axis of functional specialization in regions such as the hippocampus (Vos de Wael et al., 2018) and insula (Royer et al., 2020; Tian and Zalesky, 2018). Gradient methods have also provided novel insight into disrupted connectivity among clinical populations (Hong et al., 2019), the organization of neonatal connectomes (Larivière et al., 2020a), and the description of functional connectivity changes after structural alterations (e.g., focal lesions) (Bayrak et al., 2019). However, the functional significance and behavioral interpretation of macroscale gradients is still a central topic of discussion in the community.

Prior work has demonstrated that the functional significance of gradients may be described using meta-analytic functional decoding techniques (Margulies et al., 2016; Paquola et al., 2019). Using the Neurosynth database (Yarkoni et al., 2011), meta-analytic decoding of a connectivity gradient has previously been performed in voxel space by first segmenting the gradient spectrum (i.e., from high- to low-intensity voxels) into five-percentile increments, binarizing each of the 20 resultant maps, and decoding each map individually (Margulies et al., 2016). This segmented decoding approach allows characterization of the full gradient spectrum, thereby addressing limitations of standard meta-analytic decoding, which would yield associations to only the highest loading vertices and leave lower loading vertices unexplained. Although this gradient decoding strategy has been used in multiple prior studies (Margulies et al., 2016; Paquola et al., 2019), additional methodological development is necessary to more fully evaluate these procedures, including the segmentation approach and decoding strategy.

First, existing segmentation procedures are arbitrarily determined with equidistant intervals (i.e., five-percentile increments) (Caciagli et al., 2022; Cross et al., 2021; Larivière et al., 2020b; Margulies et al., 2016; Murphy et al., 2018; Paquola et al., 2019) and could be based on a more data-driven approach. In some cases, the segmental boundary may fall in a high-density area (i.e., a point in the gradient space with a large number of vertices in its vicinity), thus splitting a highly interconnected group of vertices into two different segments. As a result, vertices at the boundary of the gradient's coordinate system may show a low confidence value. Although the vertices are close in gradient space to vertices of their corresponding segment (i.e., highly interconnected), they are as close to the neighboring vertices from the boundary of the subsequent segment.

Second, multiple different meta-analytic approaches for functional decoding are available as potential strategies for decoding connectivity gradients. Previously, gradient decoding has been conducted with

term-based meta-analytic maps using 24 terms from Neurosynth, covering a limited range of cognitive functions (Paquola et al., 2019). Although such maps often approximate results of manual meta-analyses of the same domain and produce accurate classifications of very specific stimuli, such an approach can be accompanied by a high level of redundancy and ambiguity between terms (de la Vega et al., 2016). Further, given the heterogeneous terminology across the literature, term-based meta-analyses from Neurosynth may map terms of the same construct (e.g., calculation, arithmetic, addition, and computation) to different maps (Dockès et al., 2020); thus, decoding based on term-based maps may yield association to only a few terms of a whole construct.

Alternatively, gradient decoding has used topic-based meta-analytic maps generated using a Latent Dirichlet allocation (LDA) model of publication abstracts in Neurosynth (Margulies et al., 2016). In this case, the "*v3-topics-50*" version of the database was used, which includes a set of 50 topics extracted from LDA of Neurosynth abstracts as of September 2014. Standard topic approaches provide a broad collection of meta-analytic maps associated with not just one but with different terms (Poldrack et al., 2012), reducing the level of redundancy and ambiguity of term-based meta-analysis (de la Vega et al., 2016). In addition, the topic model is likely to group similar terms under the same construct and thus under the same meta-analytic map, thus overcoming the limitation of term-based meta-analysis with similar terms. However, the topics that a standard LDA produces are not constrained by neural data because the model works only on the text of publications (Rubin et al., 2017). In addition, the spatial components of the maps associated with each topic may be widely distributed and not reflect associations with clearly localized brain regions. Recently, a new decoding framework was introduced based on Generalized Correspondence LDA (GC-LDA), an extension to the LDA model (Rubin et al., 2017). The GC-LDA model adds spatial and semantic constraints to the data which allows the model to generate topic maps with a maximized correspondence between cognition and brain activity. However, such models have not yet been applied to decode functional connectivity gradients and to date, no formal comparison across decoding strategies has been investigated.

The overall objective of the current study was to investigate and improve the framework of data-driven methods for decoding the principal gradient of functional connectivity, thereby promoting best practices for understanding its underlying mechanisms. We comprehensively examined and evaluated different methods to address the above limitations and establish a principled approach for gradient segmentation and meta-analytic decoding. To this end, we utilized the resting-state fMRI (rs-fMRI) group-average dense connectome from the Human Connectome Project (HCP) S1200 data release (Smith et al., 2013) to identify the principal gradient of functional connectivity. We evaluated three different segmentation approaches: (i) percentile-based (PCT), (ii) segmentation based on a 1D *k*-means clustering approach (KMeans), and (iii) segmentation based on the Kernel Density Estimation (KDE) curve of the gradient axis. We assessed six different decoding strategies that used two meta-analytic databases (i.e., Neurosynth (Yarkoni et al., 2011) and NeuroQuery (Dockès et al., 2020)) and three methods to produce meta-analytic maps (i.e., term-based, LDA-based, and GC-LDA-based decoding). First, we evaluated the three segmentation approaches for 30 different segmentation solutions using the silhouette scores. Then, we assessed performance of the 18 meta-analytic segmentation/decoding strategies using multiple metrics, including correlation profile (CP), two different semantic similarity measures, and signal-to-noise ratio (SNR). Finally, we propose visualization approaches that may help on the presentation of decoded gradients and thus their overall functional interpretation.

We hypothesized that a data-driven segmentation approach (e.g., KMeans or KDE) is preferable to an equidistant segmentation based on an arbitrary percentile. In addition, we predicted that a continuous decoding approach based on a topic model (e.g., LDA ot GC-LDA) would produce a broad set of meta-analytic maps with high association to each segment of the gradient, providing a reliable functional description and aiding overall interpretability. Further, we expected that the group of terms provided by such topic-based decoders would yield a high information content and SNR in comparison with other strategies. Taken together, the current work aims to provide recommendations on best practices, along with flexible methods for gradient-based functional decoding of fMRI data.

Materials and Methods

Figure 1. **Workflow for Evaluating Meta-Analytic Functional Decoding of the Functional Connectivity Gradient.** HCP S1200 resting-state fMRI data were used (**Step 1.1**) to generate functional connectivity matrices (**Step 1.2**) and compute the affinity matrix (**Step 1.3**). Diffusion map embedding was applied to identify the principal gradient of functional connectivity (**Step 1.4**). Whole-brain gradient maps were segmented to split the gradient spectrum into a finite number of brain maps. Three different segmentation approaches were evaluated: percentile-based (PCT), *k*-means (KMeans), and KDE (**Step 2.1**). Individual segments were transformed into "activation" brain maps for decoding (**Step 2.2**). The three segmentation approaches were evaluated using the silhouette scores. Six different meta-analytic decoding strategies were implemented in surface space, derived from three sets of meta-analytic maps (i.e., term-based (Term), LDA, and GC-LDA) and two databases (i.e., Neurosynth and NeuroQuery) (**Step 3.1**). Results from the decoding strategies were evaluated using four performance metrics, assessed by comparing correlation profiles (**Step 4.1**), semantic similarity metrics (i.e., information content and TFIDF), (**Step 4.2**) and signal-to-noise ratio (SNR) (**Step 4.3**). In addition, we remove the non-functional terms from the model. Finally, we select the strategy with better performance across metrics for visualization (**Step 5**).

Our analysis plan was preregistered at https://osf.io/5q29z. **Figure 1** provides an overview of our methodological approach. First, we performed gradient decomposition of the group-average dense connectome from HCP resting-state fMRI data. Second, we evaluated three different segmentation approaches to split the gradient spectrum into a finite number of brain maps. Third, we implemented six different decoding strategies and performed an optimization test to identify the segment size to split the gradient for each strategy, for a total of 18 decoding strategies when combined with each segmentation approach. Fourth, we evaluated the different decoding strategies using multiple metrics to compare relative performance. Fifth, and as a final step, we propose a figure for reporting the result of the strategies that showed better performance across metrics. Subsequent sections describe these five steps in more detail.

Analysis Step 1: Functional Connectivity Gradient

Connectivity Data

We utilized the rs-fMRI group-average dense connectome from the Human Connectome Project (HCP) S1200 data release (Smith et al., 2013; Uğurbil et al., 2013; Van Essen et al., 2013, 2012). The rs-fMRI experiment consisted of four runs of approximately 14.5 minutes (1,200 frames) each, for a total of 4,800 frames of resting-state data per participant. 1,003 subjects that completed all four rfMRI runs (i.e., their resting-state data comprising 4,800 frames) were included in the analysis. Each run was preprocessed using HCP pipelines (Fischl, 2012; Glasser et al., 2013; Jenkinson et al., 2012, 2002), and denoised using ICA-FIX (Feinberg et al., 2010; Glasser et al., 2016; Salimi-Khorshidi et al., 2014). The data were temporally demeaned, and its variance was normalized according to (Beckmann and Smith, 2004). The cerebral cortex between subjects was coregistered using the Multimodal Surface Matching algorithm (MSMAll). Then, group-PCA from the 1,003 subjects, performed by MELODIC's Incremental Group-PCA (MIGP), produced 4,500 weighted spatial eigenvectors. These eigenvectors were renormalized, eigenvalue-reweighted, and correlated to produce the dense connectome ($59,412 \times 59,412$ z-transformed functional connectivity matrix) (**Fig. 1, Step 1.1**). Further details on this approach can be found at www.humanconnectome.org/storage/app/media/documentation/s1200/HCP1200-DenseConnectome+PTN+Appendix-July2017.pdf. Following Margulies et al. (Margulies et al., 2016), an inverse Fisher transform was applied to the z-transformed correlation to scale the values between -1 and 1 (**Fig. 1, Step 1.2**).

Gradient Decomposition

Diffusion embedding (Coifman et al., 2005) was applied to the group-averaged connectivity matrix using the mapalign Python package (Langs et al., 2015), capturing spatial gradients in macroscale resting-state functional organization. The dimensionality reduction algorithm was run using diffusion maps embedding. The procedure is detailed in Vos de Wael et al., 2020, and includes the following steps: (i) the input matrix was made sparse by retaining the top 10% of weighted connections per row, and a cosine similarity matrix was computed to capture similarity in connectivity matrices between vertices (**Fig. 1, Step 1.3**); (ii) diffusion map embedding was applied to identify principal gradient components explaining connectome variance (n_components x n_vertices: 5 x 59,412); and (iii) the amount of variance explained by each gradient was calculated (**Fig. 1, Step 1.4**). We utilized the gradient (i.e., eigenvector) with the highest variance explained (i.e., eigenvalue), also called the

principal gradient, for further analysis. Although previous work has also utilized others lower variance gradients, here we focus only on the principal gradient, which is an indicator of global connectivity of the graph (Lioi et al., 2021) and has been observed to be highly consistent across resting-state dimensionality decomposition studies (Caciagli et al., 2022; Hong et al., 2019; Margulies et al., 2016; Paquola et al., 2019).

Analysis Step 2: Segmentation and Gradient Maps

Once diffusion embedding was applied to decompose the connectivity matrix into whole-brain gradient connectivity maps, the transition from high to low-intensity vertices was examined by segmenting the gradient into individual maps in preparation for meta-analytic decoding.

Segmentation of the Cortical Gradient

We evaluated three different segmentation approaches: (i) percentile-based segmentation (PCT), (ii) segmentation based on a 1D k-means clustering approach (KMeans), and (iii) segmentation based on the Kernel Density Estimation (KDE) curve of the gradient axis (**Fig. 1, Step 2.1**). For percentile-based segmentation, we applied the previously implemented method (Margulies et al., 2016) in which the whole-brain gradient is segmented into percentile increments. Then, we applied two additional, data-driven segmentation approaches to minimize the effect of hard, arbitrary boundaries. The second segmentation approach relied on 1D k-means clustering, which has previously been used to define clusters of functional connectivity matrices to establish a brain-wide parcellation (Yeo et al., 2011). The third segmentation approach was inspired by methods of cytoarchitectonic borders detection using gray level index (Bludau et al., 2014). This method identifies local minima of a Kernel Density Estimation (KDE) curve of the gradient axis, which are points with minimal numbers of vertices in their vicinity. The number of segments for the KDE method was modified by tuning the width parameter used to generate the KDE curve.

To compare the three segmentation approaches (i.e., PCT, KMeans, and KDE), silhouette measures were estimated with the python library `scikit-learn` (Abraham et al., 2014); the silhouette measure is useful to estimate the confidence of each spatial location in cortical parcellations (i.e., "clusters") (Yeo et al., 2011). Silhouette scores were used to generate confidence maps, which allows whole-brain visualization across different segmentation approaches. For each of the segmentation approaches, $n = 30$ different segmentations of the principal gradient were generated, corresponding to numbers of segments ranging from $k = 3 - 32$, for a total of 90 different confidence maps. Mean silhouette scores were computed to determine relative performance across PCT, KMeans, and KDE segmentations, with the highest mean silhouette score representing the "best" relative performance.

From Gradient Scores to Brain Maps

Following segmentation, segmented gradient maps were transformed to brain activation maps prior to meta-analytic decoding. In the original approach (Margulies et al., 2016), the gradient maps were binarized; however, binarization removes the continuous representation of connectivity within each segment. Although the segments are still hierarchically organized, binarization eliminates the

within-segment organization, which may remove important information and thus distort functional decoding results.

For the current study, we omitted the binarization step and leveraged the within-segment continuous axis to generate activation maps. First, we defined a peak activation vertex for each segmentation approach. For PCT we chose the median point within subsegments, while for KMeans we selected the center of inertia as provided by the k-means algorithm. For KDE, on the other hand, we thought that the peak activation was better represented by the local maximum. The extreme value theorem guarantees the existence of a maximum between two minima for a continuous function in a closed interval. The local-maximum vertex represented the closest in space on average to all vertices in the corresponding segment. Finally, vertices located above the peak activation point in each sub-segments were translated relative to that maximum, and the final values were translated relative to zero to avoid the introduction of negative correlation in the decoding algorithm (**Fig. 1, Step 2.2**). These new values for each vertex captures their level of relevance in its corresponding segment.

It is worth noting that utilizing continuous maps on surface space required the functional decoding approach to be performed in surface space. Margulies et al.'s approach did not warrant consideration of this issue because it made use of binary maps in place of continuous maps (Margulies et al., 2016), which may be accurately transformed from surface to volume space using a nearest vertex approach (Wu et al., 2018). However, surface-to-volume mapping is ill-defined for continuous maps because not every MNI voxel has a corresponding `fsaverage` vertex (Markello et al., 2022; Wu et al., 2018).

Analysis Step 3: Meta-Analytic Functional Decoding

Meta-analytic functional decoding provides a quantitative estimation of behavioral profiles associated with brain regions, or networks (Amft et al., 2015; Bzdok et al., 2013b, 2013a; Cieslik et al., 2013; Laird et al., 2009; Nickl-Jockschat et al., 2015; Poldrack, 2011; Rottschy et al., 2013; Smith et al., 2009), from neuroimaging data by leveraging large-scale meta-analytic databases (e.g., Neurosynth (Yarkoni et al., 2011) and NeuroQuery (Dockès et al., 2020)). Margulies and others pioneered meta-analytic functional decoding of the connectivity gradient in volume space using the Neurosynth database (Margulies et al., 2016). However, multiple different meta-analytic decoding approaches are available as potential strategies for decoding gradients. In the current study, we assessed six different decoding strategies that used two meta-analytic databases (i.e., Neurosynth and NeuroQuery) and three sets of meta-analytic maps (i.e., term-based, LDA-based, and GC-LDA-based decoding) (**Fig. 1, Step 3.1**).

Meta-Analytic Databases

In neuroimaging, meta-analytic databases can be found in three different formats: (1) coordinate-based databases with automatically annotated studies, such as Neurosynth (Yarkoni et al., 2011) and NeuroQuery (Dockès et al., 2020), (2) coordinate-based databases with manually annotated studies, such as BrainMap (Laird et al., 2011b, 2011a, 2005), and (3) image-based databases such as NeuroVault (Gorgolewski et al., 2015).

Considering the trade-off between breadth and depth in the goal to develop a sufficiently large meta-analytic database that covers a broad range of mental functions (Varoquaux et al., 2018), in this work, we used large coordinate-based databases whose studies were automatically annotated, which

provide a wide set of cognitive states/tasks and activation coordinates. We used the popular Neurosynth database (Yarkoni et al., 2011), which contains activation coordinates that were automatically extracted from 14,371 published fMRI studies, as well as associated 3,228 semantic terms from the corresponding article abstracts. In addition, we used the recently released NeuroQuery database (Dockès et al., 2020), which contains activation coordinates that were automatically extracted from 13,459 published fMRI studies, as well as associated 6,308 terms or phrases related to neuroscience extracted from the articles' abstract, body, author-provided keywords, and title. We downloaded the database and features from the Neurosynth version 7 data release (github.com/neurosynth/neurosynth-data) and the NeuroQuery version 1 data release (github.com/neuroquery/neuroquery_data) using the fetch_neurosynth and fetch_neuroquery functions, respectively, using the NiMARE (Neuroimaging Meta-Analysis Research Environment) Python package (Salo et al., 2022a, 2022b). In both databases, each document text was represented with the term-frequency inverse document frequency (TFIDF) for each term of the vocabulary in the entire corpus of fMRI studies. The databases were converted to a NiMARE object using `convert_neurosynth_to_dataset` (Salo et al., 2022a). Then we downloaded article abstracts from PubMed using the PubMed IDs with the NiMARE's function `download_abstracts` to create a corpus to train topic-based models.

Meta-Analytic Maps

Next, two sets of meta-analytic maps with associated terms were generated from the Neurosynth and NeuroQuery datasets. Meta-analytic maps are generated from these databases using coordinate-based meta-analysis (CBMA). Resultant meta-analytic maps are associated with a vocabulary of concepts related to cognitive states or tasks, which are then used for functional decoding. Each set of maps was generated using three different meta-analytic methods: term-based, LDA-based, and GC-LDA-based approaches. Term-based (Paquola et al., 2019) and LDA-based maps (Margulies et al., 2016) have previously been used for gradient decoding; the GC-LDA method (Rubin et al., 2017) is more recent and has not yet been applied to decode connectivity gradients.

Term-Based Meta-Analytic Maps

Term-based meta-analysis (Yarkoni et al., 2011) uses database-specific vocabulary terms, their frequency of occurrence across articles, and the activation coordinates of these articles in a dataset to create meta-analytic maps. As mentioned above, each term in the database is represented by the TFIDF across articles, which is used to identify studies where the term was frequently used. To create the meta-analytic sample of studies for a term, we applied a frequency threshold of 0.001 to eliminate studies that used the term incidentally. Study-specific maps were generated by convolving the peak activation with a binary sphere of 10mm radius centered at the voxel of each activation coordinate, then combining the map from each activation coordinate by taking the maximum value for each voxel. Next, meta-analytic maps were generated by fitting the meta estimator `MKDAChi2` class from NiMARE's CBMA module to the resulting model activation maps from the meta-analytic sample. The χ^2 estimator differs from the typical multilevel kernel density analysis (MKDA) (Kober et al., 2008; Wager et al., 2009, 2007) as it uses voxel-wise chi-squared tests to generate statistical maps showing the dependency between term and activation. Finally, for each database, we generated one meta-analytic map for each

term from the vocabulary of each database; that is, 3,228 meta-analytic maps from Neurosynth and 6,145 meta-analytic maps from NeuroQuery.

LDA-Based Meta-Analytic Maps

The topic modeling technique, LDA, has been applied to text from published fMRI articles to identify latent structure and produce semantically coherent neuroimaging topics (Poldrack et al., 2012), revealing patterns consistent with previous meta-analyses (Poldrack et al., 2012; Rubin et al., 2017). LDA-based meta-analytic maps have been used for the functional decoding of brain networks (de la Vega et al., 2016; Wang et al., 2020) and individual segments of functional connectivity gradients (Margulies et al., 2016) as an alternative to overcome drawbacks associated with term-based meta-analytic methods.

In the current study, we performed topic-based meta-analyses using a LDA model on article abstracts from Neurosynth and on article abstracts, bodies, author-provided keywords, and titles from NeuroQuery. We generated 200 different topics with two sets of probability distributions: (i) the probability of a word given topic *p(word|topic)* and (ii) the probability of a topic given article *p(topic|article)*. To generate topic-wise meta-analytic maps we conducted a meta-analysis for each topic, where spatial mappings were indirectly computed via the documents' topic loadings. Similar to term-based meta-analysis, we applied a frequency threshold of 0.05 to the probability *p(topic|article)* to eliminate articles that used the topic incidentally and created the meta-analytic sample of studies for a topic. Then, meta-analytic maps were generated by fitting the meta estimator `MKDAChi2` class from NiMARE's CBMA module to the resulting model activation maps of each article from a topic. Finally, for each database, we generated one meta-analytic map for each topic with the collection of words given by the top probability values of *p(word|topic)*; that is, 200 meta-analytic maps from Neurosynth and 200 meta-analytic maps from NeuroQuery.

GC-LDA-Based Meta-Analytic Maps

Recently, an extension to the LDA model, the GC-LDA, was introduced by Rubin et al. (Rubin et al., 2017). The GC-LDA model adds spatial and semantic constraints to the data, which allows the model to generate topic maps with a maximized correspondence between cognition and brain activity. This method, in addition to the two probability distributions from LDA (i.e., *p(word|topic)* and *p(topic|article)*), produces an additional probability distribution: (iii) the probability of each topic given a voxel: *p(topic|voxel)*. The default implementation of the GC-LDA decoding in NiMARE makes use of a dot product continuous decoding approach: $p(word|image) = \tau_t \cdot p(word|topic)$, where $p(word|image)$ is the vector of term/word weight associated with an input image *I* (e.g., unthresholded and standardized segment map) and $\tau_t = p(topic|voxel) \cdot I(voxel)$ is the topic weight vector, $I(voxel)$ is a vector with *z*-score value for each masked voxel of the input image. The term $p(word|image)$ gives the most likely word from the top associated topics for a given unthresholded statistical map of interest. However, to keep the GC-LDA decoding strategy comparable to the other two strategies in terms of producing meta-analytic maps with high correlation across a segment of the gradient axis, we used the topic-wise meta-analytic maps given by the *p(topic|voxel)* distribution for functional decoding. Thus, similar to the LDA approach, the final step of the GC-LDA approach was to generate, for each database, one meta-analytic map for each topic with the collection of words given by

the top probability values of *p(topic|voxel)*; that is, 200 meta-analytic maps from Neurosynth and 200 meta-analytic maps from NeuroQuery.

Continuous Decoding Approach

Ultimately, these procedures yielded six different decoding strategies, which were derived from three different sets of meta-analytic maps (i.e., term-based, LDA-based, GC-LDA-based) and two different databases (i.e., Neurosynth and NeuroQuery). We then applied a continuous decoding approach to decode the segmented gradient maps from the three different segmentation approaches. Continuous decoding refers to the decoding of unthresholded statistical maps (e.g., unthresholded and standardized segment maps from the spectrum of the gradient) and can be accomplished either by generating correlation coefficients between unthresholded meta-analytic maps and unthresholded statistical maps or by computing the dot product between the two maps. We used the first approach, which is implemented in NiMARE's `CorrelationDecoder` class object (**Fig. 1, Step 3.2**). **Table 1** summarizes the six different strategies used in this study.

Table 1: Characteristics of Different Strategies for Decoding Connectivity Gradients.

Method to Generate Meta-Analytic Maps	No. of Maps
Term + NS + `CorrelationDecoder`	3,228
LDA + NS + `CorrelationDecoder`	200
GC-LDA + NS + `CorrelationDecoder`	200
Term + NQ + `CorrelationDecoder`	6,145
LDA + NQ + `CorrelationDecoder`	200
GCLDA + NQ + `CorrelationDecoder`	200

Note. Term, term-based meta-analysis. LDA, Latent Dirichlet allocation-based meta-analysis. GCLDA, Generalized Correspondence LDA-based meta-analysis. NS, Neurosynth. NQ, NeuroQuery.

It is worth noting that although NiMARE could natively support any kind of data masker for both volumetric and surface data, it relies on Nilearn's masker (Abraham et al., 2014), which does not currently support the transformation from surface to volume space. Therefore, we implemented a surface version of `CorrelationDecoder` to perform the functional decoding on surface space, given that gradient methods are natively implemented on the cortical surface (Vos de Wael et al., 2020). First, we transformed meta-analytic maps across decoding strategies from the MNI152 space to the standard MNI fsLR 32K 2-mm mesh surface space of the HCP, using the `mni152_to_fslr` function from the neuromaps' transforms modules (Markello et al., 2022). Then for each strategy, similar to the `CorrelationDecoder` approach, we generated correlation coefficients between vertex-level unthresholded meta-analytic maps and unthresholded, gradient segment maps. As described in Step

2.2 (**Fig. 1, Step 2.2**), $n = 30$ different segmentations of the principal gradient were generated, corresponding to numbers of segments ranging from $k = 3 - 32$.

Evaluating Significance of Correlation Values

Note that `CorrelationDecoder` does not actually test if null hypotheses of the correlations are driven by the specific topography of the map of interest. To address this issue, permutation tests have been implemented in past studies (Alexander-Bloch et al., 2013a, 2013b; Vandekar et al., 2015), which consists in comparing the empirical correlation amongst two spatial maps of interest to a set of null correlations. In this work, we used a permutation test to evaluate the null hypothesis that the meta-analytic map and gradient map were not significantly different from each other. The null correlations were generated by correlating the maps of interest with a group of null maps. We leveraged non-parametric spatial permutation model also known as "spin permutation" model (Alexander-Bloch et al., 2018; Baum et al., 2020; Cornblath et al., 2020; Váša et al., 2018; Vázquez-Rodríguez et al., 2019). We generated 1000 null maps with identical spatial autocorrelation to the maps of interest, in surface space, by randomly rotating the spherical projection of the meta-analytic maps. In particular, we applied the Váša method (Váša et al., 2018) using the `gen_spinsamples` function from the Python netneurotools's stats module (Markello and Misic, 2021). This method produces a lower false positive rate compared to naive, parameterized, and most non-parameterized data models; it also generates permutations without duplicates in the parcels assignments, and without rotating the medial wall (Markello and Misic, 2021). Non-parametric p-values were estimated from the null distribution of correlation values by calculating the fraction of null maps that generated a correlation equal to or greater than the correlation of interest. Finally, false discovery rate (FDR) correction was performed with the Benjamini-Hochberg procedure (Benjamini and Hochberg, 1995). We considered the correlation between meta-analytic maps and gradient maps to be significant for FDR-corrected p-value < 0.05.

Analysis Step 4: Performance of Decoding Strategies

Next, we evaluated the six different decoding strategies for each segmentation approach, for a total of 18 methods. We utilized multiple different metrics to compare the relative performance and utility across decoding strategies, as summarized in **Table 2**. Our goal was to identify a decoding strategy with high association to the gradient maps, along with meaningful and informative terms associated with coherent cognitive functions. We anticipated that an optimal decoding strategy would be one that would exhibit the highest performance in its correlation profile and the highest score in one or more of other metrics. Subsequent sections describe each metric in detail.

Table 2: Summary of Decoding Performance Metrics.

Metric	Summary	Database	References
CP	Top correlation for each segment of the gradient	NeuroQuery Neurosynth	(Caciagli et al., 2022)

IC	Negative log-likelihood of the probability of occurrence (*p*) of a term (*t*) in a document (*d*).	NeuroQuery Neurosynth	(Pesquita et al., 2009)
TFIDF	Inverse proportion of the frequency of the word in a particular document to the percentage of documents the word appears in. TFIDF = tf x idf	NeuroQuery Neurosynth	(Aizawa, 2003)
SNR	SNR = # functional / # non-functional	Expert Annotation	(Bottenhorn et al., 2018)

Note. CP, correlation profile. IC, information content. TFIDF, term-frequency inverse document frequency. SNR, signal-to-noise ratio.

Meta-Analytic Correlation Profile

A successful gradient decoding framework should be able to produce meta-analytic maps with a high correlation across segments of the gradient (**Fig. 1, Step 4.1**). It is common practice in continuous decoding approaches using correlation coefficients to select the term/topic maps that show the highest correlation with the input map and use the associated labels to describe the input map. We selected the meta-analytic map with maximum correlation value for each segment for further analysis. We defined a correlation profile for the principal gradient as a two-dimensional representation of the gradient axis, where the x-axis corresponds to the coordinates in gradient space and the y-axis corresponds to the top correlation for each segment of the gradient. The correlation profile metric was inspired by previous work (Caciagli et al., 2022). Note that the correlations could be high by design given the large number of meta-analytics maps that are evaluated. In addition, the different number of meta-analytic maps per strategy prevents us from comparing the correlation values that are different from the maximum.

Semantic Similarity

A successful decoding framework should also be able to classify a map of interest with meaningful and informative terms. To this end, we used two different semantic similarity node-based metrics (**Fig. 1, Step 4.2**) to evaluate properties of terms (Pesquita et al., 2009).

First, we examined the information content (IC), which measures how specific and informative a term is, based on its frequency of annotation. The IC of term *t* in document *d* is defined as the negative log-likelihood of the probability of occurrence (*p*) of a term in a document:

$$IC(t, d) = -log\ p(t, d)$$

Higher values of IC are associated with less probable terms and thus more specific concepts (e.g., pain), while lower values correspond with high probable terms and thus more general/abstract concepts (e.g., fmri, resonance) (Zhu and Iglesias, 2017). We utilized the combined corpora of article abstracts from Neurosynth and article abstracts, bodies, author-provided keywords, and titles from NeuroQuery to calculate the probability of occurrence:

$$p(t, d) = \frac{f(t,d)}{N(d)},$$

where $f(t, d)$ is the number of times that the term t appears in document d, and $N(d)$ is the total number of terms from the vocabulary in document d:

$$N(d) = \sum_{t' \in d} f(t', d)$$

Next, the IC value of a term associated with a meta-analytic map was determined by the average across the documents that were included in the meta-analysis. In the case of topic-based meta-analytic maps, the IC was calculated as the sum of individual ICs from the top words in the topic.

Second, we computed an additional metric that accounts for the popularity of a term, thus establishing a balance between the specificity (i.e., IC) and the popularity (e.g., frequency of annotation) of a term (Aizawa, 2003). This metric is known as the term frequency-inverse document frequency (TFIDF), which is defined as the proportion of the frequency of the word in a particular document (term frequency: *tf*) to the percentage of documents the word appears in (inverse document frequency: *idf*):

$$TFIDF(t, d) = tf \times idf = -f(t, d) \times [log(\frac{1+d(t,d)}{1+D(d)}) + 1]$$

where *tf* is the number of times that the term t appears in the abstract of document d $f(t, d)$, and $idf(t, d)$ is the negative log-likelihood of the proportion of the number of documents where the term appears to the total number of documents in the corpus D. The constant "1" was added to the denominator and denominator to prevent zero division, and another "1" was added to the logarithm to avoid ignoring terms with zero idf. The `TfidfTransformer` class from the python library `scikit-learn` (Abraham et al., 2014) was used to calculate the TFIDF values. TFIDF determines how relevant a given word is in a particular document. Words that are common in a single or a small group of documents tend to have higher TFIDF numbers than common words, such as articles and prepositions. The relevance of the term in a meta-analytic map can be computed, similarly to the IC, by averaging across the studies used in the meta-analysis. Moreover, the TFIDF of a topic-based decoder, was calculated as the sum of individual TFIDF from the top words in the topic.

Signal-to-Noise Ratio Profile

Finally, in a successful gradient decoding framework, the top terms that label a particular segment of the gradient axis should be associated with clear, coherent, and useful cognitive terms, which facilitates meaningful functional interpretations. However, Neurosynth terms are automatically extracted from papers and not all extracted terms contain helpful information. That is, terms such as "*memory*" or "*attention*" carry important functional meaning and are thus more useful than "*brain*" or "*association*". Thus, in a typical decoding workflow, researchers often manually filter meta-analytic terms from Neurosynth and decode maps using terms they determine to be most meaningful, a process that is commonly carried out by a single individual. Here, we used both Neurosynth and NeuroQuery terms to calculate the signal-to-noise ratio (SNR) for each decoding framework (**Fig. 1, Step 4.3**).

For Neurosynth terms, we implemented a crowdsourced approach in which manual annotations of meta-analytic terms were provided by 10 neuroimaging researchers of varying backgrounds and career stages. Each term was classified according to four categories: "*Anatomical*", "*Functional*", "*Clinical*", and "*Non-specific*"; these categories were initially described by Bottenhorn et al. (Bottenhorn et al., 2018).

"*Anatomical*" terms describe regions, networks, or locations in the brain (e.g., "cortex", "acc", "mpfc"). "*Functional*" terms describe mental processes, tasks, or behaviors (e.g., "hand movements", "visual"). "*Clinical*" terms describe participant-related characteristics, including clinical diagnoses and symptoms (e.g., "adhd", "alzheimer"), as well as age and gender. "*Non-specific*" terms provide no useful information about the study, such as general terms found in articles (e.g., "useful", "influences") or methodological terms (e.g., "correlation", "high resolution"). For NeuroQuery, we also used the expert annotation for terms that were present in Neurosynth. The rest of the terms we annotated based on the classification provided by NeuroQuery, which is available on GitHub (www.github.com/neuroquery/neuroquery_data/blob/main/data/data-neuroquery_version-1_termcategories.csv).

We created a frequentist probability distribution using the annotation to automatically classify each term: *p(category|word)*. We classified a term using the category with probability greater than 0.5. For topic-based decoding, we used the term classification distribution to classify topics. In this case the probability for topics were calculated as the the sum per categories of the *p(category|word)* weighted by the *p(word|topic)*. If the categories failed to reach the 0.5 threshold, the topic was classified as "*Non-specific*". As demonstrated by Bottenhorn et al., SNR for a decoding framework is defined as the proportion of functional (i.e., "*Functional*") to non-functional (i.e., "*Anatomical*", "*Clinical*", and "*Non-specific*") terms (Bottenhorn et al., 2018). The SNR for each segment solution of the connectivity gradients was calculated from the classification of the individual meta-analytic maps associated with each segment of the spectrum. Greater SNR was expected for a more informative functional characterization of the connectivity gradient.

Analysis Step 5: Visualization of Decoded Gradient Maps

The final step of meta-analytic decoding is to report findings and generate an appealing and effective visualization. Across the literature, four different visualization strategies are commonly employed. First, radar plots (Bartley et al., 2019; Chang et al., 2013; Chase et al., 2020; de la Vega et al., 2016) are often used to illustrate the top meta-analytic terms or topics associated with a decoded map. Second, word clouds (Fan et al., 2022; Laubach et al., 2018; Riedel et al., 2019; Rubin et al., 2017) are also frequently used as a visual representation of term frequency. We combined both plots in a single figure to show a hierarchical view of the decoding results across gradient maps (**Fig. 1, Step 5**). Prior to generating the visualizations, we filtered out the meta-analytic terms/topics without coherent cognitive function association using our expert annotations (i.e., "*Non-specific*" terms were excluded), and those with correlation not statistically significant.

Results

Analysis Step 1: Functional Connectivity Gradient

To achieve the goals of this paper, we used the rs-fMRI group-average dense connectome from the Human Connectome Project (HCP) S1200 data release. The z-values from the dense connectome were inverted to correlation coefficients, thresholded, and converted to cosine similarity values. These matrices consist of 91,282 X 91,282 grayordinates, including cortical structures (59,412 vertices) in surface space and subcortical structures (31,870 voxels) in volume space. Diffusion embedding was applied to the cosine similarity affinity matrix using the `mapalign` repository. We used the gradient

with the highest variance explained, also known as the principal gradient, for further analysis. Only the cortical regions were considered for this analysis. We added 5,572 vertices belonging to the medial wall for a total of 64,984 vertices (32,492 per hemisphere), only for visualization purposes. **Fig. 2** shows the connectivity matrix, affinity matrices, and the principal gradient. We observed that the principal gradient of the HCP S1200 connectome (**Fig. 2c**) was very similar to the one obtained by Margulies et al. (Margulies et al., 2016) on the connectome from the HCP S900 data release (**Fig. 2d**).

Figure 2. **Cortical Gradient of Functional Connectivity**. A) Z-scored functional connectivity matrix of the dense connectome HCP S1200 data release. B) Affinity matrix. C) Principal gradient observed in the current paper. D) Principal gradient observed from HCP S900 data release by Margulies et. al. (Margulies et al., 2016).

Analysis Step 2: Segmentation and Gradient Maps

Our first aim was to evaluate three different segmentation approaches: (i) percentile-based segmentation, (ii) segmentation based on a 1D k-means clustering approach, and (iii) segmentation based on the Kernel Density Estimation (KDE) curve of the gradient axis. For each segmentation approach, $n = 30$ different segmentations of the principal gradient were generated, corresponding to numbers of segments ranging from $k = 3 - 32$, for a total of 90 maps. **Fig. 3** shows the segmentation for 3, 17, and 32 segment solutions. **Fig. 3a** shows the position of the boundaries as determined by the different segmentation algorithms, with the peak activation point (red dot), which corresponded to the median point, center of inertia, and local maximum for PCT, KMeans, and KDE, respectively. For a small number of segments, PCT produced one segment covering a small range at the left end and two segments with similar sizes covering the rest of the spectrum. Similarly, for a large segment solution, most segments yielded by PCT covered a relatively small range in the left end of the spectrum. KMeans showed an evenly distributed range for all segment solutions. On the other hand, the segment assignment produced by KDE across segment solutions showed the segments close to both ends covering a larger range than the segment in the middle of the spectrum. Next, we transformed the segments into "activation" maps by translating the values of vertices located above the peak, ensuring that the peak represented the maximum value in the map. Finally, the resulting segment's map was translated from zero to prevent negative correlations. The new vertex-wise values quantify the level of relevance of vertices in their corresponding segment. These maps were the target of the decoder algorithms. **Fig. 3b** illustrates the resultant continuous gradient maps by segments.

Figure 3. **Segmentation of the Principal Gradient**. Segmentation of the principal gradient with three different approaches: PCT, KMeans, and KDE. 30 different segmentations were generated, ranging from a solution of 3 to 32 segments. A) Three distinctive and representative solutions are provided for 3, 17, and 32 segments. The boundaries are determined by the segmentation algorithm, where the red point represents the peak "activation" that corresponds to the median point, center of inertia, and local maximum for PCT, KMeans, and KDE, respectively. B) The continuous vertex-level values quantify the level of relevance of a vertex in its corresponding segments in a scale from 0 to the maximum value in the segment as determined by the peak "activation" point. See **Fig. S3** for the full set of segment solution plots. The rest of the gradient map plots can be found in the Supplementary Results.

Silhouette scores were computed for each segment solution and segmentation algorithm at the vertex level. **Fig. 4** presents a comprehensive analysis of the silhouette coefficient by depicting it in three distinct ways: a) as a distribution, b) in terms of cluster imbalance, and c) through confidence maps. **Fig. 4a** illustrates the distribution of the silhouette coefficient for three different clustering algorithms (i.e., PCT, KMeans, and KDE) and for three distinctive segment solutions. The results show that with a small number of segments (e.g., 3 segments), both PCT and KDE algorithms have a small tail toward negative silhouette coefficients, indicating vertices assigned to the wrong segment. However, with an increase in the number of segments, we see a smaller negative tail indicating better performance for these methods. Interestingly, the KDE method shows a higher concentration of vertices closer to one, indicating a better clustering solution. However, their distribution is still widespread toward negative values, suggesting that many vertices were wrongly assigned. Overall, the KMeans algorithm shows the most consistent result across segment solutions, with a distribution having the highest median value and the lowest number of vertices with negative silhouette coefficients.

Fig. 4b presents the silhouette coefficient across segments within each segment solution. For a small number of segments (e.g., 3 segments), the PCT solution yielded segments at both ends of the spectrum with all positive coefficients; nevertheless, it wrongly assigned a significant number of vertices to the middle segment. KDE yielded a small fraction of segments incorrectly assigned in the first and last segments. Conversely, KMeans shows the most precise assignment with all vertices with positive

silhouette coefficients. Increasing the segment solution in PCT improves the segment assignment, but we can still find negative silhouette coefficients, especially in segments in the middle of the spectrum. KDE performed poorly across the spectrum. In general, the highest balance in assignment across segment solutions was produced by KMeans. Lastly, **Fig. 4c** illustrates which regions of the cortical surface were most affected by the segmentation approach. Although PCT and KDE show a large area with confidence values higher than 0.8, they also show a large area with negative values. The KMeans approach consistently resulted in better segment assignment across the cortical surface.

Figure 4. **Silhouette Coefficients, Cluster Imbalance, and Confidence Maps**. A) Distributions of vertex-wise silhouette coefficient are provided for three different and representative segment solutions (e.g., 3, 17, and 32 segments); the values range from -1 to 1, where +1 indicates that the vertex was assigned to the correct cluster, a value of 0 indicates that the vertex is on or very close to the decision boundary between two neighboring clusters, and -1 indicates the vertex might have been assigned to the wrong cluster. B) Cluster imbalance is plotted as the silhouette coefficient density for each segment solution and segmentation approach per segment, to visualize the density of vertices that were assigned to the correct (i.e., positive values) or the wrong (i.e., negative values) segments within each segment. Values in the x-axis are sorted according to their values in the y-axis. C) Confidence maps are represented by the silhouette coefficient of each vertex with respect to its assigned network. Regions close to the boundaries between networks were less confident of their assignment. The best value is 1 and the worst value is -1; values near 0 indicate overlapping clusters. The distribution, cluster imbalance and confidence maps for the rest of the segment solutions can be found in the Supplementary Results.

Mean silhouette scores were computed to determine relative performance across percentile-based, k-means-based, and KDE-based segmentations, with the highest mean silhouette score representing better relative performance. As illustrated in **Fig. 5**, PCT exhibited a constant trend in the mean silhouette coefficient across different segment solutions. Conversely, KDE showed satisfactory performance for both low and high numbers of segments, but demonstrated the worst performance of

the three approaches for segment solutions ranging from 6-30. On the other hand, KMeans provided the best and most consistent segment assignments for the 30 segmentation solutions, with the highest values observed between 3-7 segments. Generally, all segmentation approaches yielded better results for five segment solutions. In summary, these findings highlight the importance of choosing the most appropriate clustering method based on the number of segments required and the desired accuracy.

Figure 5. **Mean Silhouette Scores**. Mean silhouette coefficient plotted by segment solution across segmentation approaches, to determine relative performance, with the highest mean silhouette score representing better relative performance, in which more vertices were assigned to the correct segment.

Analysis Step 3: Meta-Analytic Functional Decoding

Our second aim was to assess six different decoding strategies, which used two meta-analytic databases (i.e., Neurosynth and NeuroQuery) and three sets of meta-analytic maps (i.e., term-based, LDA-based, and GC-LDA-based decoding. The mata-analytics maps that resulted from each strategy can be found in the Supplementary Results. **Fig. 6** presents the distributions of correlation values across the three different segment solutions (i.e., PCT, KMeans, KDE), for a total of 18 decoding/segmentation strategies tested. For a small number of segments, LDA-based showed a wider distribution than term-based strategies, whereas GC-LDA showed the smallest spread of correlation values of all strategies, across segment solutions and subsegments. With the increase in the segment solution, the width of the distributions for all strategies was narrower, given that the number of high correlation values became smaller. Moreover, for a small number of segments, term-based strategies produced the highest median for the left end, while LDA-based strategies showed slightly higher medians overall for the middle subsegments and to the right end of the spectrum. Overall, term-based decoders showed the slightly highest medians with values close to zero, while LDA showed medians slightly below and over zero in some cases. GC-LDA strategies, on the other hand, showed the most extreme values, with a high proportion of outliers compared to the other strategies. Regarding the segmentation approach, the spread of the distribution of PCT and KMeans were very similar, while the KDE sometimes showed a widespread and other narrow spread compared to the other segmentations. The distribution of NS and NQ were very similar in most cases.

Figure 6. Correlation Coefficient Distributions. CorrelationDecoder estimates across models (i.e., Term, LDA, and GC-LDA), databases (i.e., Neurosynth and Neuroquery), and segmentation approaches (i.e., PCT, KMeans,

and KDE) for three distinctive segment solutions (e.g., 3, 17, and 32). For each segment solution, we plotted three subsegments, including the segments situated at both ends of the spectrum and the segment in the middle of the spectrum. The values in the distribution (y-axis) represent the correlations between meta-analytic maps and gradient maps. Vertex-wise correlation coefficients range from 1 to -1, where positive values indicate high correlation or similar maps and negative values (in blue) indicate anticorrelations or dissimilar maps. The distribution for the rest of the segment solutions can be found in the Supplementary Results.

Analysis Step 4: Performance of Decoding Strategies

Next, we assessed the performance of 18 meta-analytic decoding strategies using multiple metrics, including a correlation profile, two semantic similarity measures, and a normalized signal-to-noise ratio (SNR).

Correlation Profile

First, we assumed that a successful gradient decoding framework should be able to produce meta-analytic maps with a high correlation across segments of the gradient. Thus, we next performed a detailed analysis of the correlation profile and maximum correlation coefficient per segment for three distinct segment solutions, as shown in **Fig. 7**. The correlation profile curve (**Fig. 7a**) shows an exponential decrease in the average of maximum correlation coefficients across segment solutions for each combination of segmentation approach, meta-analytic method, and database. Then, the correlation curve reaches a plateau for a large number of segments (i.e., for segment solutions >15). The highest average correlation profile was observed in the range of 3 to 10 segments, with a maximum value of around 0.46 for 3 segments of the Term-PCT approach. The highest average of maximum correlation coefficients was produced by term-based and LDA-based decoders. In particular, the best performance was achieved through PCT segmentation and term-based decoder with Neurosynth. With regards to GC-LDA-based decoders, they showed the worst performance across segmentation solutions; however, the differences in performance as compared with the other methods were higher for a small number of segments. In general, we did not observe a larger difference between NS and NQ combinations; in some cases, NS performed slightly better (e.g., Term-PCT, LDA-PCT, GCLDA-KDE), and in others, NQ showed higher correlation values to some degree (e.g., GCLDA-PCT). With regard to segmentation approaches, PCT showed the best performance, while KMeans performed better than KDE across most segmentation solutions..

Fig. 7b represents the correlation profile within segment solutions, which exhibits a u-shaped curve, where correlation at both ends of the spectrum was higher than the rest of the subsegments, with the left end showing the highest correlation coefficient from maximum values of 0.8 for 3 segments to around 0.6 for 32 segments. Similarly, the right end shows maximum values of 0.5 and 0.4 for 3 and 32 segments, respectively. The low correlations achieved by all models in the middle of the spectrum drove the low overall performance of all decoder approaches. An increase in the segment solution also increased the number of segments in the middle of the spectrum, making it worse as indicated by the exponential decrease of the vertical plot. The decrease of the standard deviation of the mean with the increase in the segment solution also indicated an increase in a majority of segments with low performance in the overall mean per segment solution.

Figure 7. **Correlation Profile.** Correlation profile of the means within segment solution of maximum correlation coefficients across the combination of models (Term, LDA, and GC-LDA), databases (Neurosynth: NS and Neuroquery: NQ), and segmentation approach (PCT, KMeans, and KDE). In the horizontal block, three distinctive segment solutions (3, 17, and 32) are shown, containing the meta-analytic maps (the second 6 rows) that showed the highest correlation coefficient with gradient maps (top row) from a segmentation approach (columns). The horizontal plot shows the correlation profile within a segment solution, and the rectangular box at the right shows the mean and standard deviation of each model within each segment solution. The correlation profile for the rest of the segment solutions can be found in the Supplementary Results.

Semantic Similarity

Second, we assumed that a successful decoding framework should produce meaningful and informative terms. Therefore, we used two different semantic similarity node-based metrics to evaluate properties of words, such as information content (IC), which measures how specific and informative words are, based on their frequency of annotation, and term frequency-inverse document frequency (TFIDF), which accounts for the popularity of the term in the whole corpus. **Fig. 8** summarizes the results of these semantic similarity metrics across segment solutions and for three different segment solutions.

Fig. 8a presents the result of the mean information content of the top term across segment solutions. Decoding strategies that included LDA or GC-LDA yielded a collection of terms that resulted in a high information content, with LDA showing a slightly better performance than GC-LDA. Moreover, the topic-based strategies that made use of NS performed better than those that were trained with NQ. Term-based models, on the other hand, only produced a single word per map which yielded the worst information content of all pipelines.

Fig. 8b illustrates the mean TFIDF of the top terms for all segment solutions. Similarly, topic-based models produced relatively good performance, with the model that uses NQ having higher TFIDF than NS across segment solutions. Term-based decoders also showed poor performance for this metric, with mean values lower than 0.05 across segment solutions.

In Fig. 8c, we present the top words produced by the different decoding strategies for three distinctive segment solutions and three subsegments within each segmentation solution. For solutions with a small number of segments, we obtained a large number of terms that represent cognitive processes. We observed a high discrepancy between the top word between the models. However, within each model, we noticed a certain similarity in the terms, for example, "sensory", and "motor" for the left end of term-based decoders; "hand", and "sensorimotor" for LDA; and "movement" and "task" for GC-LDA. In particular, strategies where only the segmentation approach was different show the highest similarities. Similarly, for a large number of segments (e.g., 17 and 32), the words within the meta-analytic decoder continued to be very similar, and there was still a more substantial difference between modes. Interestingly, the terms at both ends for this number of segments were associated with terms with less cognitive content, for example, words like "cortex", "cortical", "magnetic", and "resonance", which do not provide much value for functional interpretation of a given map.

Figure 8. **Semantic Similarity Profile**. Average semantic similarity metrics: A) information content (IC) and b) term-frequency inverse document frequency (TFIDF), across segment solutions for the 18 different decoding strategies. C) Term associated with the meta-analytic maps of maximum correlation with the corresponding segment ID. The IC and TFIDF profile for each segment solution can be found in the Supplementary Results.

Signal-to-Noise Ratio

Finally, we assumed that in a successful gradient decoding framework, the top terms that label a particular segment of the gradient axis should be associated with clear, coherent, and useful cognitive terms to facilitate meaningful functional interpretations. For that purpose, terms were classified according to four categories: "*Anatomical*", "*Functional*", "*Clinical*", and "*Non-specific*" by a team of 10 neuroimaging researchers. In particular, we were interested in functional terms.

Fig. 9 presents the result of the classification. LDA and GC-LDA displayed the highest proportion of functional terms, with LDA-NQ producing almost 50% of functional topics. Approximately 27% of terms were classified as functional in the term-based decoder; however, these models comprised 3,153 and 6,308 terms for NS and NQ, respectively, so the total number of functional terms (i.e., 810 in NS and 1,857 in NQ) were still larger than functional LDA and GC-LDA topics.

Figure 9. **Proportion of Functional Terms**. The proportion of "*Functional*", "*Clinical*", "*Anatomical*", and "*Non-specific*" terms/topics are plotted against the 6 different models. Term, term-based meta-analysis. LDA, LDA-based meta-analysis. GC-LDA, GC-LDA-based meta-analysis. NS, Neurosynth. NQ, Neuroquery.

We next calculated a normalized signal-to-noise ratio (SNR) for each decoding strategy using the classification of the top terms with a significantly correlated map across subsegments for each segment solution. Here, the SNR for a decoding framework was defined as the normalized proportion of functional to non-functional terms. **Fig. 10** presents the results of the SNR across segment solutions and the individual classification within three distinctive segment solutions.

Fig. 10a illustrates the normalized SNR values of all decoding frameworks. Term- and LDA-based decoders showed the highest SNR values for all segment solutions, while GC-LDA decoders produced the lowest SNR across segment solutions. In terms of meta-analytic databases, NQ showed the highest SNR for each segment solution, whereas NQ combined with KDE showed the highest SNR values.

In **Fig. 10b** we present the individual classification within three distinctive segment solutions. For a small number of segments, we see a high majority of models found a functional term/topic at the end of the spectrum (i.e., 50 and 30 for the left and rightmost segments, respectively), and, for the middle segments, slightly over 50 percent of the terms were classified as functional. However, an increase in the number of segments yielded fewer functional terms at both ends of the spectrum while producing more functional terms in middle segments. In general, frameworks that used NQ were most likely associated with functional terms at both ends of the spectrum, regardless of the meta-analytic model or segmentation approach. For a large number of segments (e.g., 17 and 32), NS did not produce any top maps at the left end that were associated with functional terms, as they were classified as "*Anatomical*" or "*Non-Specific*". In particular, we noticed that when the term was classified as non-functional, term-based tended to be associated with "*Anatomical*" or "*Clinical*" terms, while LDA- and GC-LDA-based decoder were mostly associated with "*Non-specific*". Lastly, we noticed that the difference in the number of functional terms between NS and NQ tends to be less distinctive in the middle of the spectrum.

We next sought to filter the results of the top meta-analytic maps and only keep those classified as "*Functional*". After examining the final results for all the functional decoders, it was evident that the segmentation technique had minimal to no influence on the ultimate result of functional decoding, especially at both ends of the spectrum, as the word clouds of each approach appeared surprisingly similar for most of the meta-analytic decoding frameworks regardless of the segmentation algorithm used (**Fig. S4**).

Figure 10. **Normalized Signal-to-Noise Ratio**. A) SNR from classification of terms and topics across segment solutions for the 18 different decoding strategies. B) Classification of the term associated with the meta-analytic maps of maximum correlation with the corresponding segment ID; and proportion of *"Functional"*, *"Clinical"*, *"Anatomical"*, and *"Non-specific"* terms within segment solution. The classification plots for the rest of the segment solutions can be found in the Supplementary Results.

Together with the above results, we reviewed the 18 different combinations of segmentation approach, meta-analytic decoder, and database. **Table 2** presents a summary of these results with a check mark "✓" indicating frameworks that produced a good performance. We assessed performance as measured by the different metrics and concluded that KMeans segmentation combined with LDA-based decoding using the NeuroQuery database yielded the best overall gradient decoding strategy.

Table 2: Summary of Relative Performance for 18 Gradient Decoding Frameworks.

Strategy	Small number of segments (<15)					Large number of segments (>15)					Total
	SC	CP	IC	TFIDF	SNR	SC	CP	IC	TFIDF	SNR	
NS-Term-PCT		✓				✓				✓	3
NQ-Term-PCT		✓				✓				✓	3
NS-Term-KMeans	✓	✓			✓	✓				✓	5
NQ-Term-KMeans	✓	✓			✓	✓				✓	5
NS-Term-KDE		✓			✓					✓	3
NQ-Term-KDE		✓			✓					✓	3
NS-LDA-PCT		✓	✓	✓		✓		✓	✓	✓	7
NQ-LDA-PCT		✓	✓	✓	✓	✓		✓	✓	✓	8
NS-LDA-KMeans	✓	✓	✓	✓		✓		✓	✓	✓	8
NQ-LDA-KMeans	✓	✓	✓	✓	✓	✓		✓	✓	✓	9*
NS-LDA-KDE		✓	✓	✓				✓	✓		5
NQ-LDA-KDE		✓	✓	✓	✓			✓	✓	✓	7
NS-GCLDA-PCT		✓	✓	✓		✓		✓	✓		5
NQ-GCLDA-PCT		✓	✓	✓		✓		✓	✓		5
NS-GCLDA-KMeans	✓		✓	✓	✓	✓		✓	✓		7
NQ-GCLDA-KMeans	✓		✓	✓	✓	✓		✓	✓		7
NS-GCLDA-KDE			✓	✓				✓	✓		4
NQ-GCLDA-KDE			✓	✓				✓	✓		4

Analysis Step 5: Visualization of Decoded Gradient Maps

Our final aim was to generate a figure that delineates the results in a way that most effectively facilitates the interpretation of the functional gradients. To demonstrate this, we first selected the KMeans segmentation approach as it showed higher consistency in the silhouette coefficient analysis. Next, we used five segments as it showed a balanced performance across segmentation solutions, as well as one of the highest maximum correlation values on average across segmentation solutions, and also one of the highest SNR values. We selected an LDA-based model combined with NeuroQuery, as these showed high correlation values for five segments and one of the highest SNR, as shown in **Table 2**.

Figure 11 presents the functional decoding of the connectivity gradient for this best-performing framework. "Non-Functional" terms, and terms that were associated with a meta-analytic map that did not have a significant correlation (i.e., FDR-corrected p-value > 0.05) with the gradient map, were removed from the radar and word cloud plots. The first column shows the principal gradient, as the target for decoding (**Fig. 11A**). The second column presents the gradient maps that resulted from the segmentation of the gradient axis (**Fig. 11B**). The last two columns illustrate the result of the decoding algorithm for each gradient map (**Fig. 11C**). The radar plots show the top meta-analytic terms and their corresponding correlation values. The word clouds, on the other hand, provide a more inclusive representation of the terms more likely associated with the corresponding gradient maps. Overall, Figure 12 displays everything in a hierarchical view that helps to describe the gradient from top to bottom, capturing the continuous transition between different functional systems, with unimodal primary sensorimotor regions associated with "motor" and "movement" at the top, all the way down to transmodal regions associated with the default mode network and representative of abstract functioning at the bottom, as was originally described by Margulies et al. (Margulies et al., 2016).

Figure 11. **Visualization Approach for Decoding Results of the Principle Gradient of Connectivity**. A) Principal gradient of functional connectivity of the cortical surface. B) Segmentation of the principal gradient using a KMeans approach for a segment solution equal to five. C) Results of the functional decoding as shown by a radar plot and a word cloud using an LDA-based meta-analytic map generated using NeuroQuery. Radar plots present

the top ten functional terms from statistically significant meta-analytic maps associated with each segment, sorted according to their values, where the values represent correlation coefficients ranging from -1 to 1, with the maps that yielded the highest correlation plotted in the angle equal to zero degrees in polar coordinates. Note that the second and third gradient shows less than 10 terms as they were the only functional and significant terms for those gradient maps. Word cloud plots were generated using a frequency estimated by the normalized probability of the topic given a word weighted by the correlation coefficient of the corresponding maps.

Discussion

In recent years, investigation of macroscale gradients of functional connectivity have become one of the most popular methods for understanding functional brain organization. In the present study, we investigated and sought to improve the framework of data-driven methods for decoding the principal gradient of functional connectivity, thereby promoting best practices for understanding its underlying mechanisms. We evaluated 18 different decoding strategies, which consisted of three segmentation approaches combined with three meta-analytic methods and two databases. First, we found that a data-driven segmentation determined by a KMeans algorithm produced the most balanced distributions of boundaries with the highest vertex-wise and mean silhouette coefficient across segment solutions. We also determined that a small number of segments is preferred for gradients with high silhouette scores and the highest association between the subsegment and meta-analytic maps. Furthermore, we observed that an LDA-based decoder, in addition to showing a high average of top correlation scores, yielded the highest information content, TFIDF, and SNR. Finally, although we noticed minimal differences between the two meta-analytic databases tested in this work, the NeuroQuery database showed a higher SNR, provided by the richness of its vocabulary as compared to Neurosynth. The following sections discuss the implication of using a specific segmentation algorithm, meta-analytic decoder, and database for decoding connectivity gradients.

Segmentation of Connectivity Gradients

The segmentation of the connectivity gradient, the first step in the decoding framework, allows us to characterize the entire gradient spectrum by individually decoding the smaller segments as if they were independent brain maps. This step helps address limitations of standard meta-analytic decoding, which otherwise would yield associations to only the highest loading vertices of the axis, leaving lower loading vertices unexplained. In this study, we demonstrated that the number of segments and the segmentation approach affect the final results of a decoder strategy.

First, we found that the segment solution had a major influence on the final result. Indeed, for small numbers of segments, a KMeans algorithm yields the most confident distribution of boundaries as shown by the vertex-wise and mean silhouette coefficients. Furthermore, we determined that a large number of segments was detrimental to the performance of correlation decoders, as the average of the top correlation values decreases exponentially with the increase of the segment solution. Although a large number of segments does not directly affect the final result of the decoder after filtering the word to keep only functional terms, and as shown by the word cloud at both ends of the spectrum, a higher number of segments resulted in additional challenges to condense the decoding results for subsegments in the middle of the spectrum. More importantly, the decoder results for these middle segments are not as reliable as those on both ends, given their low top correlation values with meta-analytic maps. As a result, we recommend treating results from these subsegments with caution.

Prior work has made use of an arbitrary number of segments ranging from 10 to 20 (Caciagli et al., 2022; Hong et al., 2019; Margulies et al., 2016; Paquola et al., 2019; Wang et al., 2023). Even though we consider 10 segments to be an acceptable segmentation for connectivity gradients, we recommend against using a number larger than that (e.g., 20 segments) as there was a decrease in the silhouette score and a reduction in the reliability of correlation decoders. Furthermore, a large number of subsegments complicates the classification of segments in the middle of the spectrum, without adding extra benefits to the characterization of the gradient as a whole. Taken together, we recommend using a solution in the range of 3-8 segments, given the high silhouette coefficient, good relative performance for correlation decoders, and the small number of subsegments in the middle of the spectrum.

Notably, however, the selection of the segmentation algorithm only presents minor effects on the final results. In fact, the word clouds produced by all segmentation algorithms for each strategy and across segment solutions were very similar at both ends of the spectrum. However, strategies using KDE showed poor performance in their segment assignment as measured by the silhouette coefficient, and the computation of the boundaries is computationally expensive, especially for a large number of segments. As such, we recommend against using this algorithm. Previous work has used an arbitrary PCT segmentation approach, which as shown here yields a well-balanced segment assignment across solutions given the relatively high silhouette coefficients. Ultimately, we recommend using a KMeans algorithm as it yields the highest silhouette coefficient, especially for segment solutions with less than eight subsegments. Furthermore, KMeans allows the possibility of using additional components from the dimensionality reduction, which might improve segmentation of the spectrum.

Meta-Analytic Decoder

Starting from the premise that the labels associated with the meta-analytic maps accurately represent the underlying cognitive constructs, our goal was to identify the meta-analytic maps with the highest correlation value to the target map. Further, we expected that such maps should be associated with a meaningful collection of terms and thus comprise high information content, TFIDF, and SNR.

As predicted, term-based decoders performed better than LDA and GC-LDA in terms of their correlation values, which is likely a natural outcome for a model with a large number of available maps in comparison to the fixed number of maps for topic-based decoders (i.e., 200 topics maps). Although the top term-based maps accurately matched specific stimuli, their labels are known to be redundant and ambiguous given that they only consist of a single term. As a result, we hypothesized that such a model offers a limited characterization of the target maps. Here, we demonstrated the veracity of those previous assumptions by showing how the term-based decoder yielded lower information content and TFIDF. Topic-based meta-analytic models, on the other hand, generated maps labeled with numerous terms per map, thus addressing the redundancy and ambiguity of the single term from term-based decoders. In particular, LDA-based produced meta-analytic maps that yielded both a relatively high correlation value and a collection of terms that naturally improved the information content, TFIDF, and SNR. Although GC-LDA maps exhibited poor performance in terms of correlation coefficient when used in a correlation decoder framework, we think that the GC-LDA decoder as implemented in the original paper by Rubin et al. (Rubin et al., 2017), will provide a more accurate characterization of the gradient maps.

Therefore, we recommend against using GC-LDA maps in a correlation decoder framework. Although term-based decoders perform very well in terms of their correlation coefficient, we suggest using an LDA-based decoder combined with a large number of topics (e.g., between 100 and 200 topics), as it performed reasonably well in terms of their correlation coefficient and also offers a collection of terms for improved characterization of the target maps. Critically, we advise filtering out terms that do not provide functional information, given that anatomical, clinical, or non-specific labels only add noise to the classification.

In the current study, we chose to report only the top three words for topics in order to simplify the comparisons and given the fact that those three words accounted for the majority of the probability of all the words within a topic. However, we note that researchers are encouraged to include additional words from a topic if that will improve the characterization of the target map.

Meta-Analytic Database

Last but not least, the database used to train the models represents a key element of a meta-analytic decoding strategy, as the accuracy of the coordinate extraction and the heterogeneity of terms, combined with the extent of the corpus, will directly affect the performance of a decoder. As such, here we showed that NS and NQ performed similarly in terms of their correlation profile. This is likely the result of the similarity between their corpora, which consist of a similar number of overlapping documents. However, we observed that the size of the vocabulary may help improve the information content of a decoder, as well as the production of more functional terms. Therefore, we recommend using a large database with a large and rich vocabulary like NeuroQuery, or even better, a combined database that includes studies from both Neurosynth and NeuroQuery.

Limitations

The present study may be limited by several considerations. First, our results reflect different methodological approaches for investigating only the principal gradient of connectivity. That is, other gradients produced by the decomposition were discarded. While we postulate that our results will hold for other gradients, future work is required to verify this assumption. Second, we only tested the correlation decoders algorithm. Additional analyses are needed to determine if more advanced decoding strategies, such as the one implemented in a GC-LDA model or the one used by NeuroQuery, produce similar word clouds as the correlation decoder used in the present study. Correlation decoder modes are limited by design and although previous research has shown that the labels assigned to a meta-analytic map accurately reflect its cognitive state (Poldrack et al., 2012; Rubin et al., 2017; Yarkoni et al., 2011), we may identify cases where the label does not correctly describe such maps. This kind of result may be found with more frequency among labels that were classified as non-specific, clinical, or anatomical. Although the maps that were associated with these labels were filtered out, leaving only functional meta-analytic maps, we cannot be certain that one of these labels was wrongly assigned.

Conclusions

In conclusion, we provide recommendations on best practices for gradient-based functional decoding of fMRI data. We found that a small number of segments determined by a K-means segmentation

approach and an LDA-based meta-analysis combined with the NeuroQuery database was the optimal combination of methods for decoding functional connectivity gradients. Overall, we conclude that this combination of approaches, along with our recommended visualization method for reporting meta-analytic decoding findings, will successfully enhance the overall interpretability of macroscale gradients in the fMRI community.

Acknowledgments

Special thanks to the FIU Instructional & Research Computing Center (IRCC, http://ircc.fiu.edu) for providing the HPC and computing resources that contributed to the research results reported within this paper.

Ethical Statement

The Human Connectome Project provided the ethics and consent needed for study and dissemination of HCP data. This secondary data analysis was approved by the Institutional Review Board of Florida International University.

Funding Statement

Funding for this project was provided by NIH R01-MH096906 (AdlV, JBP, JDK, JD, ARL, JAP), NIH R01-DA041353 (ARL, MTS, MCR), and NIH U01-DA041156 (ARL, MTS, MCR, KLB, RPL).

Data Accessibility

Functional connectivity data were provided by the Human Connectome Project, WU-Minn Consortium (Principal Investigators: David Van Essen and Kamil Ugurbil; U54-MH091657) funded by the 16 NIH Institutes and Centers that support the NIH Blueprint for Neuroscience Research; and by the McDonnell Center for Systems Neuroscience at Washington University.

The meta-analytic databases used in this project are publicly available to download on GitHub. Neurosynth version 7 data release can be accessed at github.com/neurosynth/neurosynth-data; the NeuroQuery version 1 data release can be accessed at (github.com/neuroquery/neuroquery_data)

Code Availability

This project relied on multiple open-source Python packages, including: BrainSpace (Vos de Wael et al., 2020), Jupyter (Kluyver et al., 2016), mapalign (github.com/satra/mapalign), Matplotlib (Hunter, 2007), Netneurotools (github.com/netneurolab /netneurotools), Neuromaps (Markello et al., 2022), NiBabel (Brett et al., 2020), Nilearn (Abraham et al., 2014), NiMARE (Salo et al., 2022b; RRID:SCR_017398), NumPy (van der Walt et al., 2011), Pandas (McKinney, 2010), PtitPrince (github.com/pog87/PtitPrince), PySurfer (Waskom et al., 2020), RainCloudPlots (Allen et al., 2021), Scikit-learn (Pedregosa et al., 2011), SciPy (Virtanen et al., 2020), Seaborn (Waskom et al., 2017), SurfPlot (Gale et al., 2021), and word_cloud

(Mueller et al., 2018). We also made use of the HCP software Connectome Workbench (`wb_command` version 1.5.0, (Marcus et al., 2011)).

All code required to reproduce the analyses and figures in this paper will be made available on GitHub at https://github.com/NBCLab/gradient-decoding. High-resolution figures are available through FigShare (https://figshare.com/projects/Meta-analytic_decoding_of_the_cortical_gradient_of_functional_connectivity/172347). The decoding and segmentation modules are available for future re-use here https://github.com/JulioAPeraza/gradec (Peraza et al., 2023), which will be linked to NiMARE (https://nimare.readthedocs.io/en/latest/) as a Python module for gradient meta-analytic decoding in the future. All data and resources that resulted from this paper (e.g., connectivity gradients and trained meta-analytic decoders) is openly disseminated and made available on the Open Science Framework (OSF) at https://osf.io/xzfrt, including the links to the GitHub repository, and figures.

Competing Interests

The authors declare no competing interests.

Authors' Contributions

ARL, JAP, TS, MCR conceived and designed the project. ARL, KLB, JEB, JSF, LDHB, RPL, RP, KLR, JLR, RWL served as expert annotators for the meta-analytic terms. JAP, MCR, TS analyzed data. JAP, MCR, TS, KLB, AdlV, JBP, JD, JDK contributed scripts and pipelines. JAP, ARL wrote the paper and all authors contributed to the revisions and approved the final version.

References

Abraham, A., Pedregosa, F., Eickenberg, M., Gervais, P., Mueller, A., Kossaifi, J., Gramfort, A., Thirion, B., Varoquaux, G., 2014. Machine learning for neuroimaging with scikit-learn. Front. Neuroinformatics 8. https://doi.org/10.3389/fninf.2014.00014

Aizawa, A., 2003. An information-theoretic perspective of tf–idf measures. Inf. Process. Manag. 39, 45–65. https://doi.org/10.1016/S0306-4573(02)00021-3

Alexander-Bloch, A., Giedd, J.N., Bullmore, E., 2013a. Imaging structural co-variance between human brain regions. Nat. Rev. Neurosci. 14, 322–336. https://doi.org/10.1038/nrn3465

Alexander-Bloch, A., Raznahan, A., Bullmore, E., Giedd, J., 2013b. The Convergence of Maturational Change and Structural Covariance in Human Cortical Networks. J. Neurosci. 33, 2889–2899. https://doi.org/10.1523/JNEUROSCI.3554-12.2013

Alexander-Bloch, A.F., Shou, H., Liu, S., Satterthwaite, T.D., Glahn, D.C., Shinohara, R.T., Vandekar, S.N., Raznahan, A., 2018. On testing for spatial correspondence between maps of human brain structure and function. NeuroImage 178, 540–551. https://doi.org/10.1016/j.neuroimage.2018.05.070

Allen, M., Poggiali, D., Whitaker, K., Marshall, T.R., van Langen, J., Kievit, R.A., 2021. Raincloud plots: a multi-platform tool for robust data visualization. Wellcome Open Res. https://doi.org/10.12688/wellcomeopenres.15191.2

Amft, M., Bzdok, D., Laird, A.R., Fox, P.T., Schilbach, L., Eickhoff, S.B., 2015. Definition and characterization of an extended social-affective default network. Brain Struct. Funct. 220, 1031–1049. https://doi.org/10.1007/s00429-013-0698-0

Bartley, J.E., Riedel, M.C., Salo, T., Boeving, E.R., Bottenhorn, K.L., Bravo, E.I., Odean, R., Nazareth, A., Laird, R.W., Sutherland, M.T., Pruden, S.M., Brewe, E., Laird, A.R., 2019. Brain activity links performance in science reasoning with conceptual approach. NPJ Sci. Learn. 4, 20. https://doi.org/10.1038/s41539-019-0059-8

Baum, G.L., Cui, Z., Roalf, D.R., Ciric, R., Betzel, R.F., Larsen, B., Cieslak, M., Cook, P.A., Xia, C.H., Moore, T.M., Ruparel, K., Oathes, D.J., Alexander-Bloch, A.F., Shinohara, R.T., Raznahan, A., Gur, R.E., Gur, R.C., Bassett, D.S., Satterthwaite, T.D., 2020. Development of structure–function coupling in human brain networks during youth. Proc. Natl. Acad. Sci. 117, 771–778. https://doi.org/10.1073/pnas.1912034117

Bayrak, Ş., Khalil, A.A., Villringer, K., Fiebach, J.B., Villringer, A., Margulies, D.S., Ovadia-Caro, S., 2019. The impact of ischemic stroke on connectivity gradients. NeuroImage Clin. 24, 101947. https://doi.org/10.1016/j.nicl.2019.101947

Beckmann, C.F., Smith, S.M., 2004. Probabilistic independent component analysis for functional magnetic resonance imaging. IEEE Trans. Med. Imaging 23, 137–152. https://doi.org/10.1109/TMI.2003.822821

Benjamini, Y., Hochberg, Y., 1995. Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing. J. R. Stat. Soc. Ser. B Methodol. 57, 289–300. https://doi.org/10.1111/j.2517-6161.1995.tb02031.x

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, In-vivo Brodmann Mapping of the Human Brain 93, 260–275. https://doi.org/10.1016/j.neuroimage.2013.05.052

Bottenhorn, K.L., Salo, T., Sutherland, M.T., Laird, A.L., 2018. Quantitative comparison of functional decoding approaches across meta-analytic frameworks. INCF Neuroinformatics. https://doi.org/10.7490/f1000research.1115906.1

Brett, M., Markiewicz, C.J., Hanke, M., Côté, M.-A., Cipollini, B., McCarthy, P., Jarecka, D., Cheng, C.P.,

Halchenko, Y.O., Cottaar, M., Larson, E., Ghosh, S., Wassermann, D., Gerhard, S., Lee, G.R., Wang, H.-T., Kastman, E., Kaczmarzyk, J., Guidotti, R., Duek, O., Daniel, J., Rokem, A., Madison, C., Moloney, B., Morency, F.C., Goncalves, M., Markello, R., Riddell, C., Burns, C., Millman, J., Gramfort, A., Leppäkangas, J., Sólon, A., van den Bosch, J.J.F., Vincent, R.D., Braun, H., Subramaniam, K., Gorgolewski, K.J., Raamana, P.R., Klug, J., Nichols, B.N., Baker, E.M., Hayashi, S., Pinsard, B., Haselgrove, C., Hymers, M., Esteban, O., Koudoro, S., Pérez-García, F., Oosterhof, N.N., Amirbekian, B., Nimmo-Smith, I., Nguyen, L., Reddigari, S., St-Jean, S., Panfilov, E., Garyfallidis, E., Varoquaux, G., Legarreta, J.H., Hahn, K.S., Hinds, O.P., Fauber, B., Poline, J.-B., Stutters, J., Jordan, K., Cieslak, M., Moreno, M.E., Haenel, V., Schwartz, Y., Baratz, Z., Darwin, B.C., Thirion, B., Gauthier, C., Papadopoulos Orfanos, D., Solovey, I., Gonzalez, I., Palasubramaniam, J., Lecher, J., Leinweber, K., Raktivan, K., Calábková, M., Fischer, P., Gervais, P., Gadde, S., Ballinger, T., Roos, T., Reddam, V.R., freec84, 2020. nipy/nibabel: 3.2.1. https://doi.org/10.5281/zenodo.4295521

Bzdok, D., Laird, A.R., Zilles, K., Fox, P.T., Eickhoff, S.B., 2013a. An investigation of the structural, connectional, and functional subspecialization in the human amygdala. Hum. Brain Mapp. 34, 3247–3266. https://doi.org/10.1002/hbm.22138

Bzdok, D., Langner, R., Schilbach, L., Jakobs, O., Roski, C., Caspers, S., Laird, A.R., Fox, P.T., Zilles, K., Eickhoff, S.B., 2013b. Characterization of the temporo-parietal junction by combining data-driven parcellation, complementary connectivity analyses, and functional decoding. NeuroImage 81, 381–392. https://doi.org/10.1016/j.neuroimage.2013.05.046

Caciagli, L., Paquola, C., He, X., Vollmar, C., Centeno, M., Wandschneider, B., Braun, U., Trimmel, K., Vos, S.B., Sidhu, M.K., Thompson, P.J., Baxendale, S., Winston, G.P., Duncan, J.S., Bassett, D.S., Koepp, M.J., Bernhardt, B.C., 2022. Disorganization of language and working memory systems in frontal versus temporal lobe epilepsy. Brain awac150. https://doi.org/10.1093/brain/awac150

Chang, L.J., Yarkoni, T., Khaw, M.W., Sanfey, A.G., 2013. Decoding the role of the insula in human cognition: functional parcellation and large-scale reverse inference. Cereb. Cortex N. Y. N 1991 23, 739–749. https://doi.org/10.1093/cercor/bhs065

Chase, H.W., Grace, A.A., Fox, P.T., Phillips, M.L., Eickhoff, S.B., 2020. Functional differentiation in the human ventromedial frontal lobe: A data-driven parcellation. Hum. Brain Mapp. 41, 3266–3283. https://doi.org/10.1002/hbm.25014

Cieslik, E.C., Zilles, K., Caspers, S., Roski, C., Kellermann, T.S., Jakobs, O., Langner, R., Laird, A.R., Fox, P.T., Eickhoff, S.B., 2013. Is there "one" DLPFC in cognitive action control? Evidence for heterogeneity from co-activation-based parcellation. Cereb. Cortex N. Y. N 1991 23, 2677–2689. https://doi.org/10.1093/cercor/bhs256

Coifman, R.R., Lafon, S., Lee, A.B., Maggioni, M., Nadler, B., Warner, F., Zucker, S.W., 2005. Geometric diffusions as a tool for harmonic analysis and structure definition of data: Diffusion maps. Proc. Natl. Acad. Sci. 102, 7426–7431. https://doi.org/10.1073/pnas.0500334102

Cornblath, E.J., Ashourvan, A., Kim, J.Z., Betzel, R.F., Ciric, R., Adebimpe, A., Baum, G.L., He, X., Ruparel, K., Moore, T.M., Gur, R.C., Gur, R.E., Shinohara, R.T., Roalf, D.R., Satterthwaite, T.D., Bassett, D.S., 2020. Temporal sequences of brain activity at rest are constrained by white matter structure and modulated by cognitive demands. Commun. Biol. 3, 1–12. https://doi.org/10.1038/s42003-020-0961-x

Cross, N., Paquola, C., Pomares, F.B., Perrault, A.A., Jegou, A., Nguyen, A., Aydin, U., Bernhardt, B.C., Grova, C., Dang-Vu, T.T., 2021. Cortical gradients of functional connectivity are robust to state-dependent changes following sleep deprivation. NeuroImage 226, 117547. https://doi.org/10.1016/j.neuroimage.2020.117547

de la Vega, A., Chang, L.J., Banich, M.T., Wager, T.D., Yarkoni, T., 2016. Large-Scale Meta-Analysis of Human Medial Frontal Cortex Reveals Tripartite Functional Organization. J. Neurosci. 36,

6553–6562. https://doi.org/10.1523/JNEUROSCI.4402-15.2016

Dockès, J., Poldrack, R.A., Primet, R., Gözükan, H., Yarkoni, T., Suchanek, F., Thirion, B., Varoquaux, G., 2020. NeuroQuery, comprehensive meta-analysis of human brain mapping. eLife 9, e53385. https://doi.org/10.7554/eLife.53385

Fan, Y.-S., Xu, Y., Bayrak, Ş., Shine, J.M., Li, H., Li, L., Yang, S., Meng, Y., Valk, S.L., Chen, H., 2022. Macroscale Thalamic Functional Organization Disturbances And Underlying Core Cytoarchitecture In Early-Onset Schizophrenia. https://doi.org/10.1101/2022.05.11.489776

Feinberg, D.A., Moeller, S., Smith, S.M., Auerbach, E., Ramanna, S., Glasser, M.F., Miller, K.L., Ugurbil, K., Yacoub, E., 2010. Multiplexed Echo Planar Imaging for Sub-Second Whole Brain FMRI and Fast Diffusion Imaging. PLOS ONE 5, e15710. https://doi.org/10.1371/journal.pone.0015710

Fischl, B., 2012. FreeSurfer. NeuroImage, 20 YEARS OF fMRI 62, 774–781. https://doi.org/10.1016/j.neuroimage.2012.01.021

Gale, D.J., Vos de Wael., R., Benkarim, O., Bernhardt, B., 2021. Surfplot: Publication-ready brain surface figures. https://doi.org/10.5281/zenodo.5567926

Glasser, M.F., Smith, S.M., Marcus, D.S., Andersson, J.L.R., Auerbach, E.J., Behrens, T.E.J., Coalson, T.S., Harms, M.P., Jenkinson, M., Moeller, S., Robinson, E.C., Sotiropoulos, S.N., Xu, J., Yacoub, E., Ugurbil, K., Van Essen, D.C., 2016. The Human Connectome Project's neuroimaging approach. Nat. Neurosci. 19, 1175–1187. https://doi.org/10.1038/nn.4361

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., 2013. The minimal preprocessing pipelines for the Human Connectome Project. NeuroImage, Mapping the Connectome 80, 105–124. https://doi.org/10.1016/j.neuroimage.2013.04.127

Gorgolewski, K.J., Varoquaux, G., Rivera, G., Schwarz, Y., Ghosh, S.S., Maumet, C., Sochat, V.V., Nichols, T.E., Poldrack, R.A., Poline, J.-B., Yarkoni, T., Margulies, D.S., 2015. NeuroVault.org: a web-based repository for collecting and sharing unthresholded statistical maps of the human brain. Front. Neuroinformatics 9, 8. https://doi.org/10.3389/fninf.2015.00008

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

Hunter, J.D., 2007. Matplotlib: A 2D Graphics Environment. Comput. Sci. Eng. 9, 90–95. https://doi.org/10.1109/MCSE.2007.55

Jenkinson, M., Bannister, P., Brady, M., Smith, S., 2002. Improved Optimization for the Robust and Accurate Linear Registration and Motion Correction of Brain Images. NeuroImage 17, 825–841. https://doi.org/10.1006/nimg.2002.1132

Jenkinson, M., Beckmann, C.F., Behrens, T.E.J., Woolrich, M.W., Smith, S.M., 2012. FSL. NeuroImage, 20 YEARS OF fMRI 62, 782–790. https://doi.org/10.1016/j.neuroimage.2011.09.015

Kluyver, T., Ragan-Kelley, B., Pé, Rez, F., Granger, B., Bussonnier, M., Frederic, J., Kelley, K., Hamrick, J., Grout, J., Corlay, S., Ivanov, P., Avila, D., n, Abdalla, S., Willing, C., Team, J.D., 2016. Jupyter Notebooks – a publishing format for reproducible computational workflows. Position. Power Acad. Publ. Play. Agents Agendas 87–90. https://doi.org/10.3233/978-1-61499-649-1-87

Kober, H., Barrett, L.F., Joseph, J., Bliss-Moreau, E., Lindquist, K., Wager, T.D., 2008. Functional grouping and cortical–subcortical interactions in emotion: A meta-analysis of neuroimaging studies. NeuroImage 42, 998–1031. https://doi.org/10.1016/j.neuroimage.2008.03.059

Laird, A.R., Eickhoff, S.B., Fox, P.M., Uecker, A.M., Ray, K.L., Saenz, J.J., McKay, D.R., Bzdok, D., Laird, R.W., Robinson, J.L., Turner, J.A., Turkeltaub, P.E., Lancaster, J.L., Fox, P.T., 2011a. The BrainMap strategy for standardization, sharing, and meta-analysis of neuroimaging data. BMC Res. Notes 4, 349. https://doi.org/10.1186/1756-0500-4-349

Laird, A.R., Eickhoff, S.B., Li, K., Robin, D.A., Glahn, D.C., Fox, P.T., 2009. Investigating the functional

heterogeneity of the default mode network using coordinate-based meta-analytic modeling. J. Neurosci. Off. J. Soc. Neurosci. 29, 14496–14505. https://doi.org/10.1523/JNEUROSCI.4004-09.2009

Laird, A.R., Fox, P.M., Eickhoff, S.B., Turner, J.A., Ray, K.L., McKay, D.R., Glahn, D.C., Beckmann, C.F., Smith, S.M., Fox, P.T., 2011b. Behavioral Interpretations of Intrinsic Connectivity Networks. J. Cogn. Neurosci. 23, 4022–4037. https://doi.org/10.1162/jocn_a_00077

Laird, A.R., Lancaster, J.L., Fox, P.T., 2005. BrainMap: the social evolution of a human brain mapping database. Neuroinformatics 3, 65–78. https://doi.org/10.1385/ni:3:1:065

Langs, G., Golland, P., Ghosh, S.S., 2015. Predicting Activation Across Individuals with Resting-State Functional Connectivity Based Multi-Atlas Label Fusion, in: Navab, N., Hornegger, J., Wells, W.M., Frangi, A. (Eds.), Medical Image Computing and Computer-Assisted Intervention -- MICCAI 2015, Lecture Notes in Computer Science. Springer International Publishing, Cham, pp. 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., 2020a. Multiscale Structure–Function Gradients in the Neonatal Connectome. Cereb. Cortex 30, 47–58. https://doi.org/10.1093/cercor/bhz069

Larivière, S., Weng, Y., Vos de Wael, R., Royer, J., Frauscher, B., Wang, Z., Bernasconi, A., Bernasconi, N., Schrader, D.V., Zhang, Z., Bernhardt, B.C., 2020b. Functional connectome contractions in temporal lobe epilepsy: Microstructural underpinnings and predictors of surgical outcome. Epilepsia 61, 1221–1233. https://doi.org/10.1111/epi.16540

Laubach, M., Amarante, L.M., Swanson, K., White, S.R., 2018. What, If Anything, Is Rodent Prefrontal Cortex? eNeuro 5. https://doi.org/10.1523/ENEURO.0315-18.2018

Lioi, G., Gripon, V., Brahim, A., Rousseau, F., Farrugia, N., 2021. Gradients of connectivity as graph Fourier bases of brain activity. Netw. Neurosci. 5, 322–336. https://doi.org/10.1162/netn_a_00183

Marcus, D., Harwell, J., Olsen, T., Hodge, M., Glasser, M., Prior, F., Jenkinson, M., Laumann, T., Curtiss, S., Van Essen, D., 2011. Informatics and Data Mining Tools and Strategies for the Human Connectome Project. Front. Neuroinformatics 5. https://doi.org/10.3389/fninf.2011.00004

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. 113, 12574–12579. https://doi.org/10.1073/pnas.1608282113

Markello, R.D., Hansen, J.Y., Liu, Z.-Q., Bazinet, V., Shafiei, G., Suárez, L.E., Blostein, N., Seidlitz, J., Baillet, S., Satterthwaite, T.D., Chakravarty, M.M., Raznahan, A., Misic, B., 2022. neuromaps: structural and functional interpretation of brain maps. https://doi.org/10.1101/2022.01.06.475081

Markello, R.D., Misic, B., 2021. Comparing spatial null models for brain maps. NeuroImage 236, 118052. https://doi.org/10.1016/j.neuroimage.2021.118052

McKinney, W., 2010. Data Structures for Statistical Computing in Python. Presented at the Python in Science Conference, Austin, Texas, pp. 56–61. https://doi.org/10.25080/Majora-92bf1922-00a

Mueller, A., Fillion-Robin, J.-C., Boidol, R., Tian, F., Nechifor, P., yoonsubKim, Peter, Rampin, R., Corvellec, M., Medina, J., Dai, Y., Petrushev, B., Langner, K.M., Hong, Alessio, Ozsvald, I., vkolmakov, Jones, T., Bailey, E., Rho, V., IgorAPM, Roy, D., May, C., foobuzz, Piyush, Seong, L.K., Goey, J.V., Smith, J.S., Gus, Mai, F., 2018. amueller/word_cloud: WordCloud 1.5.0. https://doi.org/10.5281/zenodo.1322068

Murphy, C., Jefferies, E., Rueschemeyer, S.-A., Sormaz, M., Wang, H., Margulies, D.S., Smallwood, J., 2018. Distant from input: Evidence of regions within the default mode network supporting perceptually-decoupled and conceptually-guided cognition. NeuroImage 171, 393–401. https://doi.org/10.1016/j.neuroimage.2018.01.017

Nickl-Jockschat, T., Rottschy, C., Thommes, J., Schneider, F., Laird, A.R., Fox, P.T., Eickhoff, S.B., 2015. Neural networks related to dysfunctional face processing in autism spectrum disorder. Brain Struct. Funct. 220, 2355–2371. https://doi.org/10.1007/s00429-014-0791-z

Paquola, C., Wael, R.V.D., 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

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., 2011. Scikit-learn: Machine Learning in Python. J. Mach. Learn. Res. 12, 6.

Peraza, J.A., Kent, J.D., Salo, T., De La Vega, A., Laird, A.R., 2023. JulioAPeraza/gradec: 0.0.1rc3. https://doi.org/10.5281/zenodo.8161766

Pesquita, C., Faria, D., Falcão, A.O., Lord, P., Couto, F.M., 2009. Semantic Similarity in Biomedical Ontologies. PLOS Comput. Biol. 5, e1000443. https://doi.org/10.1371/journal.pcbi.1000443

Poldrack, R.A., 2011. Inferring mental states from neuroimaging data: from reverse inference to large-scale decoding. Neuron 72, 692–697. https://doi.org/10.1016/j.neuron.2011.11.001

Poldrack, R.A., Mumford, J.A., Schonberg, T., Kalar, D., Barman, B., Yarkoni, T., 2012. Discovering Relations Between Mind, Brain, and Mental Disorders Using Topic Mapping. PLOS Comput. Biol. 8, e1002707. https://doi.org/10.1371/journal.pcbi.1002707

Riedel, M.C., Salo, T., Hays, J., Turner, M.D., Sutherland, M.T., Turner, J.A., Laird, A.R., 2019. Automated, Efficient, and Accelerated Knowledge Modeling of the Cognitive Neuroimaging Literature Using the ATHENA Toolkit. Front. Neurosci. 13. https://doi.org/10.3389/fnins.2019.00494

Rottschy, C., Caspers, S., Roski, C., Reetz, K., Dogan, I., Schulz, J.B., Zilles, K., Laird, A.R., Fox, P.T., Eickhoff, S.B., 2013. Differentiated parietal connectivity of frontal regions for "what" and "where" memory. Brain Struct. Funct. 218, 1551–1567. https://doi.org/10.1007/s00429-012-0476-4

Royer, J., Paquola, C., Larivière, S., Vos de Wael, R., Tavakol, S., Lowe, A.J., Benkarim, O., Evans, A.C., Bzdok, D., Smallwood, J., Frauscher, B., Bernhardt, B.C., 2020. Myeloarchitecture gradients in the human insula: Histological underpinnings and association to intrinsic functional connectivity. NeuroImage 216, 116859. https://doi.org/10.1016/j.neuroimage.2020.116859

Rubin, T.N., Koyejo, O., Gorgolewski, K.J., Jones, M.N., Poldrack, R.A., Yarkoni, T., 2017. Decoding brain activity using a large-scale probabilistic functional-anatomical atlas of human cognition. PLOS Comput. Biol. 13, e1005649. https://doi.org/10.1371/journal.pcbi.1005649

Salimi-Khorshidi, G., Douaud, G., Beckmann, C.F., Glasser, M.F., Griffanti, L., Smith, S.M., 2014. Automatic denoising of functional MRI data: Combining independent component analysis and hierarchical fusion of classifiers. NeuroImage 90, 449–468. https://doi.org/10.1016/j.neuroimage.2013.11.046

Salo, T., Yarkoni, T., Nichols, T.E., Poline, J.-B., Bilgel, M., Bottenhorn, K.L., Jarecka, D., Kent, J.D., Kimbler, A., Nielson, D.M., Oudyk, K.M., Peraza, J.A., Pérez, A., Reeders, P.C., Yanes, J.A., Laird, A.R., 2022a. NiMARE: Neuroimaging Meta-Analysis Research Environment. NeuroLibre Jupyter Book Prepr.

Salo, T., Yarkoni, T., Nichols, T.E., Poline, J.-B., Kent, J.D., Gorgolewski, K.J., Glerean, E., Bottenhorn, K.L., Bilgel, M., Wright, J., Reeders, P., Kimbler, A., Nielson, D.N., Yanes, J.A., Pérez, A., Oudyk, K.M., Jarecka, D., Enge, A., Peraza, J.A., Laird, A.R., 2022b. neurostuff/NiMARE: 0.0.12rc5. https://doi.org/10.5281/zenodo.6603673

Smith, S.M., Beckmann, C.F., Andersson, J., Auerbach, E.J., Bijsterbosch, J., Douaud, G., Duff, E., Feinberg, D.A., Griffanti, L., Harms, M.P., Kelly, M., Laumann, T., Miller, K.L., Moeller, S., Petersen, S., Power, J., Salimi-Khorshidi, G., Snyder, A.Z., Vu, A.T., Woolrich, M.W., Xu, J.,

Yacoub, E., Uğurbil, K., Van Essen, D.C., Glasser, M.F., WU-Minn HCP Consortium, 2013. Resting-state fMRI in the Human Connectome Project. NeuroImage 80, 144–168. https://doi.org/10.1016/j.neuroimage.2013.05.039

Smith, S.M., Fox, P.T., Miller, K.L., Glahn, D.C., Fox, P.M., Mackay, C.E., Filippini, N., Watkins, K.E., Toro, R., Laird, A.R., Beckmann, C.F., 2009. Correspondence of the brain's functional architecture during activation and rest. Proc. Natl. Acad. Sci. U. S. A. 106, 13040–13045. https://doi.org/10.1073/pnas.0905267106

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

Uğurbil, K., Xu, J., Auerbach, E.J., Moeller, S., Vu, A.T., Duarte-Carvajalino, J.M., Lenglet, C., Wu, X., Schmitter, S., Van de Moortele, P.F., Strupp, J., Sapiro, G., De Martino, F., Wang, D., Harel, N., Garwood, M., Chen, L., Feinberg, D.A., Smith, S.M., Miller, K.L., Sotiropoulos, S.N., Jbabdi, S., Andersson, J.L.R., Behrens, T.E.J., Glasser, M.F., Van Essen, D.C., Yacoub, E., 2013. Pushing spatial and temporal resolution for functional and diffusion MRI in the Human Connectome Project. NeuroImage, Mapping the Connectome 80, 80–104. https://doi.org/10.1016/j.neuroimage.2013.05.012

van der Walt, S., Colbert, S.C., Varoquaux, G., 2011. The NumPy Array: A Structure for Efficient Numerical Computation. Comput. Sci. Eng. 13, 22–30. https://doi.org/10.1109/MCSE.2011.37

Van Essen, D.C., Smith, S.M., Barch, D.M., Behrens, T.E.J., Yacoub, E., Ugurbil, K., 2013. The WU-Minn Human Connectome Project: An overview. NeuroImage, Mapping the Connectome 80, 62–79. https://doi.org/10.1016/j.neuroimage.2013.05.041

Van Essen, D.C., Ugurbil, K., Auerbach, E., Barch, D., Behrens, T.E.J., Bucholz, R., Chang, A., Chen, L., Corbetta, M., Curtiss, S.W., Della Penna, S., Feinberg, D., Glasser, M.F., Harel, N., Heath, A.C., Larson-Prior, L., Marcus, D., Michalareas, G., Moeller, S., Oostenveld, R., Petersen, S.E., Prior, F., Schlaggar, B.L., Smith, S.M., Snyder, A.Z., Xu, J., Yacoub, E., 2012. The Human Connectome Project: A data acquisition perspective. NeuroImage, Connectivity 62, 2222–2231. https://doi.org/10.1016/j.neuroimage.2012.02.018

Vandekar, S.N., Shinohara, R.T., Raznahan, A., Roalf, D.R., Ross, M., DeLeo, N., Ruparel, K., Verma, R., Wolf, D.H., Gur, R.C., Gur, R.E., Satterthwaite, T.D., 2015. Topologically Dissociable Patterns of Development of the Human Cerebral Cortex. J. Neurosci. 35, 599–609. https://doi.org/10.1523/JNEUROSCI.3628-14.2015

Varoquaux, G., Schwartz, Y., Poldrack, R.A., Gauthier, B., Bzdok, D., Poline, J.-B., Thirion, B., 2018. Atlases of cognition with large-scale human brain mapping. PLOS Comput. Biol. 14, e1006565. https://doi.org/10.1371/journal.pcbi.1006565

Váša, F., Seidlitz, J., Romero-Garcia, R., Whitaker, K.J., Rosenthal, G., Vértes, P.E., Shinn, M., Alexander-Bloch, A., Fonagy, P., Dolan, R.J., Jones, P.B., Goodyer, I.M., the NSPN consortium, Sporns, O., Bullmore, E.T., 2018. Adolescent Tuning of Association Cortex in Human Structural Brain Networks. Cereb. Cortex 28, 281–294. https://doi.org/10.1093/cercor/bhx249

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

Virtanen, P., Gommers, R., Oliphant, T.E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., Walt, S.J. van der, Brett, M., Wilson, J., Millman, K.J., Mayorov, N., Nelson, A.R.J., Jones, E., Kern, R., Larson, E., Carey, C.J., Polat, İ., Feng, Y., Moore, E.W., VanderPlas, J., Laxalde, D., Perktold, J., Cimrman, R., Henriksen, I., Quintero, E.A., Harris, C.R., Archibald, A.M., Ribeiro, A.H., Pedregosa, F., Mulbregt, P. van, 2020. SciPy 1.0: fundamental algorithms for scientific computing in Python. Nat. Methods 1–12.

https://doi.org/10.1038/s41592-019-0686-2

Vos de Wael, R., Benkarim, O., Paquola, C., Lariviere, S., Royer, J., Tavakol, S., Xu, T., Hong, S.-J., Langs, G., Valk, S., Misic, B., Milham, M., Margulies, D., Smallwood, J., Bernhardt, B.C., 2020. BrainSpace: a toolbox for the analysis of macroscale gradients in neuroimaging and connectomics datasets. Commun. Biol. 3, 1–10. https://doi.org/10.1038/s42003-020-0794-7

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

Wager, T.D., Lindquist, M., Kaplan, L., 2007. Meta-analysis of functional neuroimaging data: current and future directions. Soc. Cogn. Affect. Neurosci. 2, 150–158. https://doi.org/10.1093/scan/nsm015

Wager, T.D., Lindquist, M.A., Nichols, T.E., Kober, H., Van Snellenberg, J.X., 2009. Evaluating the consistency and specificity of neuroimaging data using meta-analysis. NeuroImage, Mathematics in Brain Imaging 45, S210–S221. https://doi.org/10.1016/j.neuroimage.2008.10.061

Wang, S., Tepfer, L.J., Taren, A.A., Smith, D.V., 2020. Functional parcellation of the default mode network: a large-scale meta-analysis. Sci. Rep. 10, 16096. https://doi.org/10.1038/s41598-020-72317-8

Wang, Y., Royer, J., Park, B., Vos de Wael, R., Larivière, S., Tavakol, S., Rodriguez-Cruces, R., Paquola, C., Hong, S.-J., Margulies, D.S., Smallwood, J., Valk, S.L., Evans, A.C., Bernhardt, B.C., 2023. Long-range functional connections mirror and link microarchitectural and cognitive hierarchies in the human brain. Cereb. Cortex 33, 1782–1798. https://doi.org/10.1093/cercor/bhac172

Waskom, M., Botvinnik, O., O'Kane, D., Hobson, P., Lukauskas, S., Gemperline, D.C., Augspurger, T., Halchenko, Y., Cole, J.B., Warmenhoven, J., de Ruiter, J., Pye, C., Hoyer, S., Vanderplas, J., Villalba, S., Kunter, G., Quintero, E., Bachant, P., Martin, M., Meyer, K., Miles, A., Ram, Y., Yarkoni, T., Williams, M.L., Evans, C., Fitzgerald, C., Brian, Fonnesbeck, C., Lee, A., Qalieh, A., 2017. Mwaskom/Seaborn: V0.8.1 (September 2017). Zenodo. https://doi.org/10.5281/zenodo.883859

Waskom, M., Larson, E., Brodbeck, C., Gramfort, A., Burns, S., Luessi, M., Weidemann, C.T., Bitzer, S., LaPlante, R., Markiewicz, C., Halchenko, Y., Vliet, M. van, Engemann, D.A., diego0020, Klein, N., Ghosh, S., Dingkun.Liu, KING, J.-R., Gwilliams, L., Brett, M., Piantoni, G., 2020. nipy/PySurfer: 0.11.0. https://doi.org/10.5281/zenodo.3905195

Wu, J., Ngo, G.H., Greve, D., Li, J., He, T., Fischl, B., Eickhoff, S.B., Yeo, B.T.T., 2018. Accurate nonlinear mapping between MNI volumetric and FreeSurfer surface coordinate systems. Hum. Brain Mapp. 39, 3793–3808. https://doi.org/10.1002/hbm.24213

Yarkoni, T., Poldrack, R.A., Nichols, T.E., Van Essen, D.C., Wager, T.D., 2011. Large-scale automated synthesis of human functional neuroimaging data. Nat. Methods 8, 665–670. https://doi.org/10.1038/nmeth.1635

Yeo, B.T.T., Krienen, F.M., Sepulcre, J., Sabuncu, M.R., Lashkari, D., Hollinshead, M., Roffman, J.L., Smoller, J.W., Zöllei, L., Polimeni, J.R., Fischl, B., Liu, H., Buckner, R.L., 2011. The organization of the human cerebral cortex estimated by intrinsic functional connectivity. J. Neurophysiol. 106, 1125–1165. https://doi.org/10.1152/jn.00338.2011

Zhu, G., Iglesias, C.A., 2017. Computing Semantic Similarity of Concepts in Knowledge Graphs. IEEE Trans. Knowl. Data Eng. 29, 72–85. https://doi.org/10.1109/TKDE.2016.2610428

(A) Connectivity Matrix (B) Affinity Matrix (C) HPC S1200 Gradient (D) HPC S900 Gradient

(A) Segmentation Algorithm

Segment Solution

⋮

⋮

(B) Gradient Maps

PCT   KMeans   KDE

(A) Distributions   **(B) Cluster Imbalance**   **(C) Confidence Maps**

(B) Segmentation

(C) Functional Decoding

(A) Principal Gradient

Segment Solution	Algorithm	Segment Solution	Segment. Algorthm

Figure showing Segment Solution results for segment numbers 3–17 (left column) and 18–32 (right column), each with PCT, KMeans, and KDE algorithm rows.