Published in final edited form as:

Med Image Anal. 2016 October ; 33: 149–154. doi:10.1016/j.media.2016.06.026.

Computational neuroanatomy using brain deformations: From brain parcellation to multivariate pattern analysis and machine learning

Christos Davatzikos

Center for Biomedical Image Computing and Analytics, University of Pennsylvania, Philadelphia, PA 19104, United States

Abstract

The past 20 years have seen a mushrooming growth of the field of computational neuroanatomy. Much of this work has been enabled by the development and refinement of powerful, high-dimensional image warping methods, which have enabled detailed brain parcellation, voxel-based morphometric analyses, and multivariate pattern analyses using machine learning approaches. The evolution of these 3 types of analyses over the years has overcome many challenges. We present the evolution of our work in these 3 directions, which largely follows the evolution of this field. We discuss the progression from single-atlas, single-registration brain parcellation work to current ensemble-based parcellation; from relatively basic mass-univariate *t*-tests to optimized regional pattern analyses combining deformations and residuals; and from basic application of support vector machines to generative-discriminative formulations of multivariate pattern analyses, and to methods dealing with heterogeneity of neuroanatomical patterns. We conclude with discussion of some of the future directions and challenges.

Keywords

Computational neuroanatomy; Pattern analysis; Machine learning; Brain image analysis

1. Introduction

Lord Kelvin concisely summarized his thoughts in his address at the Institution of Civil Engineers in London, in 1883: “If you cannot measure it, then it is not science”. The field of Computational Neuroanatomy has largely followed this vision over the past two to three decades and has focused on going beyond qualitative descriptions of brain anatomy, and developing mathematical and computational frameworks and algorithms for quantitatively representing the complex anatomical patterns associated with normal brain development, aging, and with various neurologic and neuropsychiatric diseases. The roots of this field could be attributed to D’Arcy Thompson, an evolutionary biologist of the early 20th century, whose book “On Growth and Form” showed a remarkable attempt to quantify anatomical shape, considering the limited technological tools available at the time (Fig. 1). Modern

URL: https://ipp.cbica.upenn.edu/

computational neuroanatomy, perhaps having its roots in the seminal work of Miller et al. (1993), basically implements D'Arcy Thompson's general approach, albeit using increasingly advanced tools for image deformation and statistical analysis.

Deformable templates were initially used to map labels from a digital atlas to an individual's magnetic resonance imaging (MRI) scan, then later to map data from multiple individuals to a common coordinate system for voxel-based analytics of various flavors, and relatively more recently to quantify complex patterns of brain anatomy via multivariate analytics and machine learning, eventually leading to personalized diagnostic and predictive indices that get closer to being clinically useful. Herein, we review some of our work over the past 20 years, which largely followed the broader field's direction.

2. Atlas-based labeling of brain MRI

Labeling regions of interest (ROI) on brain MRIs has been of interest since the 1980's. Most brain disorders and diseases have been associated with changes in the volumes or signal (e.g. diffusion, perfusion, T2) within certain regions. MR provides information regarding a variety of tissue characteristics, such as volume, shape, connectivity, microstructural integrity, perfusion etc. Having the ability to define boundaries of various anatomical brain regions allows us to measure these anatomical and functional characteristics regionally. Early work involved very laborious and not easily reproducible manual tracing of brain ROIs, and was limited to a small number of *a priori* defined ROIs and lower numbers of subjects. Deformable neuroanatomical templates changed this, by developing computational algorithms that warped one brain MRI to another, thereby allowing one to map predefined anatomical boundaries from an anatomical template, or atlas, to any individual's MRI. Some of our initial work in this direction involved boundary-based curvature mapping, followed by elastic deformations (Davatzikos et al., 1996), based on the assumption that the geometric characteristics of cortical curvature are largely in line with well-established anatomical boundaries. Similar approaches were pursued by other investigators (e.g. Thompson and Toga, 1996). Our methods were later extended to 3D volumetric mapping, by attaching a morphological signature to each and every voxel in an MRI scan, and achieving very accurate 3D deformations by matching respective morphological signatures (Shen and Davatzikos, 2002; Ou et al., 2011). Since our certainty in the accuracy of the anatomical match achieved by these algorithms is generally variable throughout the brain, with some brain regions being less variable and more consistently identified and matched to the atlas, these methods also relied on an overarching mutual saliency function, which modulated the underlying optimization criterion and weighted different regions according to the certainty of respective matches. A variety of methods developed in parallel in the literature, paying special attention to inverse consistency of the derived transformations (Christensen and Johnson, 2001; Avants et al., 2008), using different optimization criteria and constraints on topological properties of the derived deformation fields (Vercauteren et al., 2009; Joshi and Miller, 2000), and estimating optimal population-based templates (Joshi et al., 2004).

The past 5 years have seen a remarkable transformation of these single-template, deformation-based ROI labeling methods into ensemble-based labeling approaches, which have, quite dramatically, increased the accuracy and robustness of these methods (Wang et

al., 2013). Our group has participated in this effort, and has contributed MUSE (Doshi et al., 2016), which achieved top performance in a recent MICCAI challenge (Asman, 2013). MUSE utilizes an ensemble of atlases warped to an individual MRI with multiple transformations, each emanating from a different algorithm/optimization criterion, and multiple parameters, which influence the amount of regularization and other properties of these transformations. An image-based template selection and weighting is then used to achieve consensus labeling of the individual MRI. We have found this method to be significantly more robust than single-template, single-warping methods, partially due to the noise reducing properties of ensemble-based methods, and in part due to the diversity of anatomical characteristics present in the ensemble that allow the algorithm to choose the right labels from the right brains and warps for a given location, based on the individual's anatomical characteristics (Fig. 2).

3. Voxel-based morphometic analysis (VBMA)

Parallel to the development of deformable templates for anatomical parcellation have been methods that exploit the imaging data in their full resolution, instead of the resolution of predefined ROIs. These data-driven methods are primarily suited for knowledge discovery, i.e. when there is no strong *a priori* hypothesis for the anatomical regions of interest. These approaches leveraged the inverse transformation, i.e. the transformation that maps individuals' imaging data, or derivative images extracted from them, to a stereotaxic coordinate system residing in an atlas space. Since this process alters the very data one wants to analyze, a variety of methods were developed aiming to retain as much information as possible from this transformation process. Accordingly, a series of such methods were pursued under the names of deformation-based morphometry (primarily looking at Jacobian determinants) (Chung et al., 2001), tensor-based morphometry (looking at primary directions of variation of the displacement fields around each voxel) (Thompson et al., 2000), voxel-based morphometry (looking at residuals after low-dimensional registration) (Ashburner and Friston, 2000) and others. Our group developed a method termed Regional Analysis of Volumes Embedded in Normalized Space (RAVENS) (Davatzikos et al., 2001), which was designed to precisely preserve tissue volumes while mapping them to the atlas space, thereby generating RAVENS maps whose voxel-wise values reflect regional volumetrics of various tissues. Because of the tissue preserving properties, analysis of RAVENS maps on a voxel-by-voxel, or region-by-region, basis effectively pertains to respective regional tissue volumetrics.

One of the most persistent problems with VBMA has been the need to apply some spatial filter to the data, prior to statistical analyses (such as group comparisons via mass-univariate statistical tests). This spatial smoothing is necessary for many reasons: it leads to data that follow Gaussian distributions, thereby facilitating subsequent linear statistics; and it removes noise by examining larger regions, rather than individual voxels, thereby rendering it far easier to detect group effects. However this spatial smoothing has generally been performed in a very ad hoc, or empirical way, typically involving some Gaussian filter of 6, 8 or 10 mm. One approach that formalized the concept of optimal smoothing was termed optimally discriminant voxel-based analysis (ODVBA) (Zhang and Davatzikos, 2011), and it leveraged the strengths of machine learning by applying large ensembles of regional discriminant

analysis, and subsequently integrating the results into a coherent voxel-based map that optimized our ability to detect group effects. In addition to formalizing the way in which imaging data needs to be filtered in VBMA, this approach was shown to remarkably increase our ability to detect, especially subtle, group differences.

4. From mass-univariate to multivariate pattern analysis (MVPA)

While VBMA has been a main-stream method in computational neuroanatomy for 2 decades, it has been quite limited by virtue of applying a large number of independent tests, voxel-by-voxel, in an attempt to characterize group differences. The past decade experienced a mushrooming interest in MVPA methods using machine learning, mainly for two reasons: (1) MVPA methods examine associations among different brain regions, and hence build patterns that can be quite distinctive of various pathologies, even though their individual voxel-by-voxel constituents might not be when looked at in isolation; (2) these machine learning methods were repeatedly shown to produce highly specific and sensitive personalized indices, thereby enabling the clinical use of computational neuroanatomy: a neuroanatomical (or functional) signature of a disease has little clinical value if it cannot be detected on an individual-patient basis, even if it is robustly identified in group comparisons. Initial investigations using support vector machines (SVM) were followed by application of various other methods, such as random forests, and more recently, deep learning methods.

Even though MVPA methods offer a more comprehensive way of quantifying neuroanatomical patterns, compared to ROI or VBMA approaches, they are limited by the sheer size of medical images, especially when compared to the typically available sample sizes of one or a few hundred scans. Machine learning methods require extensive training databases, due to the complexity of the multivariate relationships they explore. Various schemes have therefore emerged, attempting to select an optimal set of features, thereby reducing the dimensionality of imaging signals down to a manageable level. Many of these approaches use different flavors of forward and backward sequential feature selection and integration. In our approach of Fan et al. (2007), named COMPARE, we used these techniques along with spatial clustering methods, to partition brain images into relatively homogeneous regional clusters of voxels, prior to providing these clusters to machine learning algorithms. In a somewhat more mathematically principled way, we developed the GRASP method (Honnorat et al., 2015) for resting state fMRI (rsfMRI) images, which utilized a novel graph-based parcellation method that relies on a discrete Markov Random Field framework. The spatial connectedness of the parcels was explicitly enforced by shape priors. The shape of the parcels was adapted to underlying data through the use of functional geodesic distances. The performance of GRASP was assessed using a large developmental cohort of more than 850 subjects. Fig. 3 shows an example of this partitioning approach on rsfMRI data of the cortical surface.

These approaches to feature selection, largely driven by spatial and signal criteria (nearby voxels of similar signals are grouped together), have been complemented by statistically-derived feature extraction methods. The main premise of these approaches has been that statistically co-varying measurements (e.g. voxels or vertices) can be grouped together without significant loss of signal, and on the contrary, with potential significant gains, since

noise is averaged out. Our group, among others, has used non-negative matrix factorization (NNMF) as a way to achieve data-driven feature extraction and dimensionality reduction. NNMF has gained a great deal of attention in the computer vision field, where it has shown to effectively achieve parts-based decompositions of images into meaningful components, not otherwise obtained via commonly used PCA- and ICA-types of methods. Sotiras et al. (2015) applied this method to brain MRI images, and discovered structural brain networks that display coordinated change across individuals. Such components might reflect the influence of underlying neurodevelopmental and neurodegenerative biological processes that affect brain structure and function in coordinated ways that are manifested by statistical covariance. Fig. 4 shows an example of neuroanatomical components derived from a large study of brain aging, and which largely agree with known anatomical and functional brain units. Such data-driven, yet anatomically and functionally meaningful reductions of high-dimensional imaging data can potentially create parsimonious, yet interpretable feature sets.

The most commonly used approach to feature selection and dimensionality reduction is to first calculate a number of features, either according to prior hypotheses, or by capturing broad aspects of shape, intensity, texture and other characteristics, and then select the ones that seem most relevant. However, more holistic approaches have been developed that combine the feature extraction and classification steps into a single framework. Such approaches are often generative-discriminative, in that they aim to capture variance of the imaging signal, while simultaneously paying attention to the task of classification achieved using these features. Our group has developed similar approaches for structural (Batmanghelich et al., 2012) and functional connectivity (Eavani et al., 2015) MRI, using sparse decomposition methods that aim to describe the data by a number of components that are discriminative for the task at hand (e.g. classification of patients and controls). Generative-discriminative methods also offer additional advantages when it comes to interpretation. In particular, commonly used feature selection schemes might find a set of features that achieve good classification, but these features may be hard to interpret. For example, one might be interested in a classifier of healthy control vs. AD patients, which can be achieved by measuring cortical thickness at a number of vertices throughout the cortex. Even if good classification can be achieved, these samples do not inform us about all cortical thickness changes occurring with AD, but rather just a small subset that is sufficient for diagnosis. Mechanistic interpretations of such classification models greatly benefit from generative-discriminative methods, which in contrast try to capture as much of the underlying neuroanatomy as possible, as long as classification accuracy is not compromised. Therefore they are likely to be discriminative, but also amenable to biological and clinical interpretation.

5. MVPA and machine learning models are “black boxes”. How can I find out which aspects of brain anatomy or function are important?

Machine learning tools have often been considered as black boxes, which may achieve good diagnostic or predictive performance, but are not informative in terms of understanding disease mechanisms and phenotypes. The main question that has been largely unanswered is: “Which of the aspects of the imaging signal and features are statistically significant, in

contributing to the classification?". One approach to tackle this problem has been to perform random permutations, and measure the null distribution of each of the imaging features when a classifier tries to achieve optimal separation of randomly labeled images. Recent work has significantly expedited this otherwise very costly process, by developing analytical approximations to statistical significance maps, especially for support vector machine classifiers and in a high-dimensionality setting (Gaonkar and Davatzikos, 2013). These techniques allow us to not only report accuracy or the area under the curve (AUC) of an SVM classifier, but also to derive spatial maps of brain regions that are statistically significant in terms of achieving this diagnostic or predictive accuracy. These methods promise to render this process more interpretable and acceptable by the clinical world.

6. Current and future challenges and directions

As Niels Bohr said, "prediction is very difficult, especially about the future", so extrapolation into the future is a dangerous undertaking. However, this section briefly reviews some current challenges that are likely to receive more attention in the near future.

6.1. Personalized indices

The vast majority of the work in computational neuroanatomy has focused on group comparisons, such as differences between patients and controls, progressors and non-progressors, responders and non-responders, etc. Population-based results can inform us about disease mechanisms, but they are of little value in the clinic, unless they translate to individualized biomarkers. Achieving sufficient sensitivity and specificity on an individual patient basis has been, and will remain challenging. Alzheimer's disease for sure causes atrophy in the hippocampus, however hippocampal volume has limited value on an individual patient basis, due to inter-individual variability. Prodromal stages of disease, or diseases with more subtle imaging manifestations, pose even higher challenges for personalized medicine. Finally, numerous other brain diseases also cause hippocampal loss, thereby rendering specificity of this measure quite challenging. To a large extent, the need for individualized biomarkers has largely driven the adoption and development of machine learning tools in this field, since it allows for weak individual imaging characteristics to be integrated into multivariate patterns that carry sufficient sensitivity and specificity on an individual person basis. This line of work is therefore likely to receive increasing attention in the upcoming decade. Some of the challenges that come with it will need to be addressed. In particular, individualized indices are likely to eventually include multi-modal imaging and non-imaging data, which need to be synergistically combined. Moreover, effectively learning high-dimensional and complex imaging patterns theoretically requires thousands of training samples. Prior work has focused on much lower scale studies, in part due to limited data availability, and in part due to computational demands of these methods, including cross-validation and permutation tests that establish generalization and statistical significance of the results. Large databases of neuroimaging data become increasingly available, with examples like ADNI, Philadelphia neurodevelopmental cohort, Human Connectome project, NDAR, 1000 Functional Connectomes project, and other initiatives. Still, to pull together perhaps 10,000 scans requires multi-site integration of datasets, which

raises challenges of data harmonization across studies, which is not a solved or trivial problem.

6.2. Heterogeneity

Regardless of whether ROIs, VBMA, or MVPA are used, respective methods have largely neglected heterogeneity of neuroimaging phenotypes. In particular, all these methods seek a single pattern of difference between two groups, such as patients and controls, progressors and non-progressors, decliners and non-decliners, etc., and analogously for regression-type analyses. However, heterogeneity poses significant challenges to this type of analytics. If certain imaging characteristics are seen in some, but not all individuals, they can be lost in these analyses that seek a strongest common denominator. Brain development, aging, and diseases are known to be highly heterogeneous. Therefore, detection, classification, prognosis, and stratification into treatments, are all likely to benefit tremendously from a more refined characterization of this heterogeneity. Some semi-supervised machine learning methods have been recently developed that do not simply ask “how does group A differ from group B?”, but rather ask “in what and how many ways/patterns do groups A and B differ?”. Our group has contributed two methods, termed Chimera (Dong et al., 2016) and Hydra (Varol et al., 2016) in this direction, and is likely to look at this problem in greater detail in the upcoming years, along with other groups.

6.3. Computational neuroanatomy using consensus approaches

Consensus based approaches have been a game changer in image segmentation in the past 5 years or so, since they have been able to dramatically reduce labeling errors associated with the use of individual atlases, warping algorithms, and parameter sets, and have leveraged the power of combining multiple solutions. Several machine learning paradigms also leverage the power of integration of multiple “relatively weaker” classifiers. However, these methods still have not been formulated in the context of the various flavors of VBMA approaches. This is a direction that is likely to be fruitful, but it also poses mathematical and computational challenges. For example, if multiple atlases are used in a deformation-based morphometry framework, the multiple deformation field, each defined in a different brain’s domain, must be combined in a reasonable way. Simple approaches to this problem are certainly immediately obvious (e.g. Baloch and Davatzikos, 2009), however mathematically principled ways might be more challenging.

Acknowledgments

The author would like to thank Amanda Shacklett and Dr. Aristeidis Sotiras for help with editing the manuscript.

References

Ashburner J, Friston KJ. Voxel-based morphometry – the methods. Neuroimage. 2000; 11(6 Pt 1):805–821. [PubMed: 10860804]

Asman, A. MICCAI 2013 Challenge Workshop on Segmentation: Algorithms, Theory and Applications (SATA). MICCAI SATA 2013 competition. 2013. Available from: http://masi.vuse.vanderbilt.edu/submission/leaderboard.html

Avants BB, et al. Symmetric diffeomorphic image registration with cross–correlation: evaluating automated labeling of elderly and neurodegenerative brain. Med Image Anal. 2008; 12(1):26–41. [PubMed: 17659998]

Baloch S, Davatzikos C. Morphological appearance manifolds in computational anatomy: Groupwise registration and morphological analysis. NeuroImage. 2009; 45(1):S73–S85. [PubMed: 19061962]

Batmanghelich NK, Taskar B, Davatzikos C. Generative-discriminative basis learning for medical imaging. IEEE Trans Med Imag. 2012; 31(1):51–69.

Christensen GE, Johnson HJ. Consistent image registration. IEEE Trans Med Imag. 2001; 20(7):568–582.

Chung MK, et al. A unified statistical approach to deformation-based morphometry. Neuroimage. 2001; 14(3):595–606. [PubMed: 11506533]

Davatzikos C, et al. Voxel-based morphometry using the RAVENS maps: methods and validation using simulated longitudinal atrophy. Neuroimage. 2001; 14(6):1361–1369. [PubMed: 11707092]

Davatzikos C, Prince JL, Bryan N. Image registration based boundary mapping. IEEE Trans Med Imag. 1996; 15(1):112–115.

Dong A, Honnorat N, Gaonkar B, Davatzikos C. CHIMERA: Clustering of Heterogeneous Disease Effects via Distribution Matching of Imaging Patterns. IEEE Trans Med Imag. 2016; 35(2):612–621.

Doshi J, et al. MUSE: multi-atlas region segmentation utilizing ensembles of registration algorithms and parameters, and locally optimal atlas selection. NeuroImage. 2016; 127:186–195. [PubMed: 26679328]

Eavani H, et al. Identifying sparse connectivity patterns in the brain using rest*ing-state fMRI*. NeuroImage. 2015; 105:286–299. [PubMed: 25284301]

Fan Y, Shen D, Gur RC, Gur RE, Davatzikos C. COMPARE: classification of morphological patterns using adaptive regional elements. IEEE Trans Med Imag. 2007; 26(1):93–105.

Gaonkar B, Davatzikos C. Analytic estimation of statistical significance maps for support vector machine based multi-variate image analysis and classification. Neuroimage. 2013; 78:270–283. [PubMed: 23583748]

Honnorat N, et al. GraSP: Geodesic graph-based segmentation with shape Priors for the functional parcellation of the cortex. Neuroimage. 2015; 106:207–221. [PubMed: 25462796]

Joshi S, et al. Unbiased diffeomorphic atlas construction for computational anatomy. Neuroimage. 2004; 23(Suppl. 1):S151–S160. [PubMed: 15501084]

Joshi SC, Miller MI. Landmark matching via large deformation diffeomorphisms. IEEE Trans Image Process. 2000; 9(8):1357–1370. [PubMed: 18262973]

Lao Z, Shen D, Xue Z, Karacali B, Resnick SM, Davatzikos C. Morphological classification of brains via high-dimensional shape transformations and machine learning methods. Neuroimage. 2004; 21(1):46–57. [PubMed: 14741641]

Miller MI, et al. Mathematical textbook of deformable neuroanatomies. Proc Natl Acad Sci USA. 1993; 90:11944–11948. [PubMed: 8265653]

Ou Y, Sotiras A, Paragios N, Davatzikos C. DRAMMS: deformable registration via attribute matching and mutual-saliency weighting. Med Image Anal. 2011; 15(4):622–639. [PubMed: 20688559]

Shen D, Davatzikos C. HAMMER: hierarchical attribute matching mechanism for elastic registration. IEEE Trans Med Imag. 2002; 21(11):1421–1439.

Sotiras A, Resnick SM, Davatzikos C. Finding imaging patterns of structural covariance via non-negative matrix factorization. NeuroImage. 2015; 108:1–16. [PubMed: 25497684]

Thompson P, Toga AW. A surface-based technique for warping three-dimensional images of the brain. IEEE Trans Med Imag. 1996; 15(4):402–417.

Thompson PM, et al. Growth patterns in the developing brain detected by using continuum mechanical tensor maps. Nature. 2000; 404:190–193. [PubMed: 10724172]

Varol E, Sotiras A, Davatzikos C, Alzheimer's Disease Neuroimaging Initiative. HYDRA: Revealing heterogeneity of imaging and genetic patterns through a multiple max-margin discriminative analysis framework. NeuroImage. 2016

Vercauteren T, et al. Diffeomorphic demons: efficient non-parametric image registration. Neuroimage. 2009; 45(1 Suppl):S61–S72. [PubMed: 19041946]

Wang H, et al. Multi-atlas segmentation with joint label fusion. IEEE Trans Pattern Anal Mach Intell. 2013; 35(3):611–623. [PubMed: 22732662]

Zhang TH, Davatzikos C. ODVBA: optimally-discriminative voxel-based analysis. IEEE Trans Med Imag. 2011; 30(8):1441–1454.

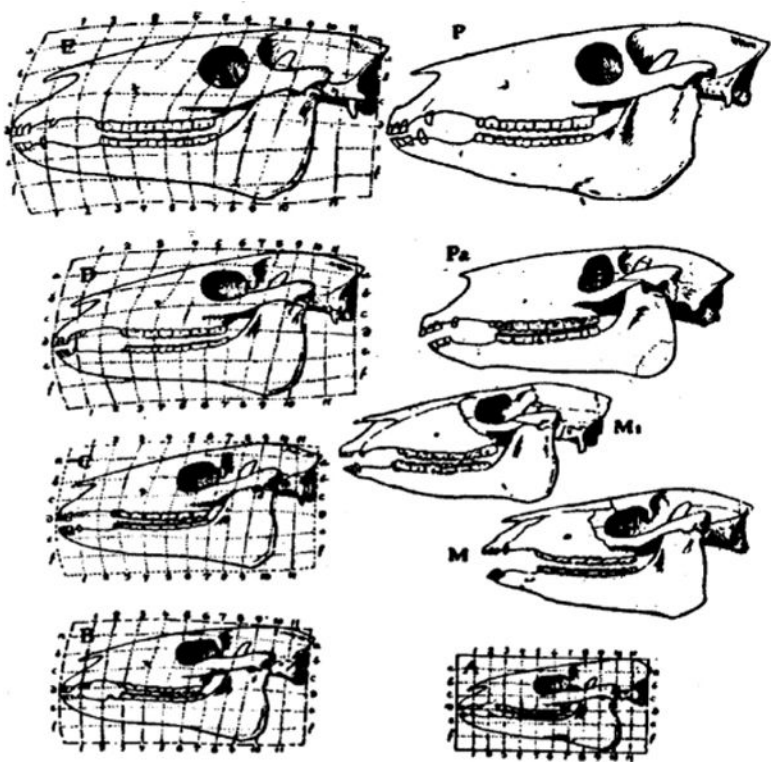

Fig. 1.

First attempts by D'Arcy Thompson in 1917 to measure shape via 2D transformations of a grid, which can be considered as the root of modern computational anatomy.

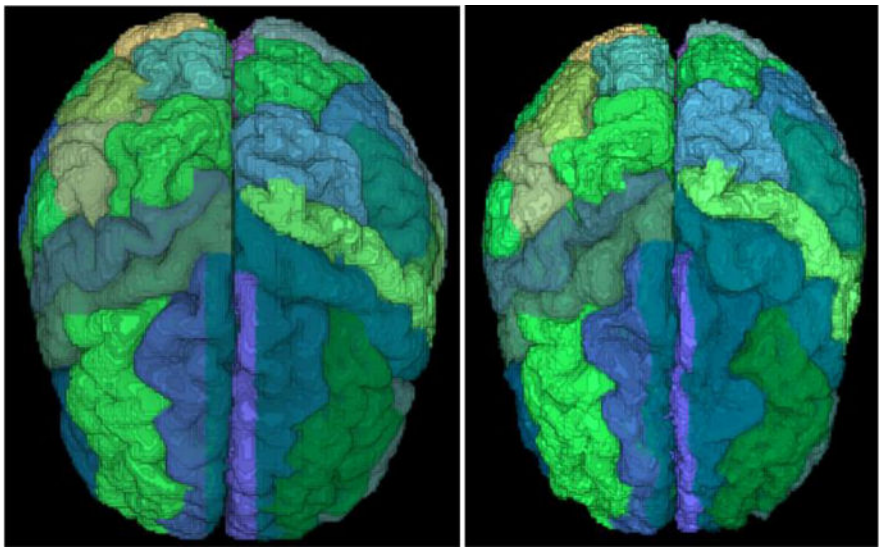

Fig. 2.
Example of an anatomical template (left), and its warped version that matches an individual's MRI scan and transfers the labels to it. Reprinted from Lao et al. (2004).

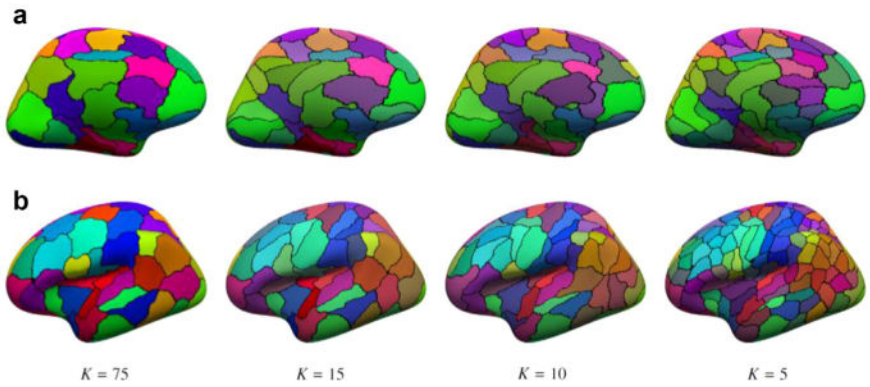

Fig. 3.

Example of multi-scale extraction of rsfMRI features from the brain, by partitioning the cortical surface into regionally homogeneous sets of functionally-coherent points. Parameter *K* controls the scale of the partitioning. Figure reprinted from Honnorat et al. (2015).

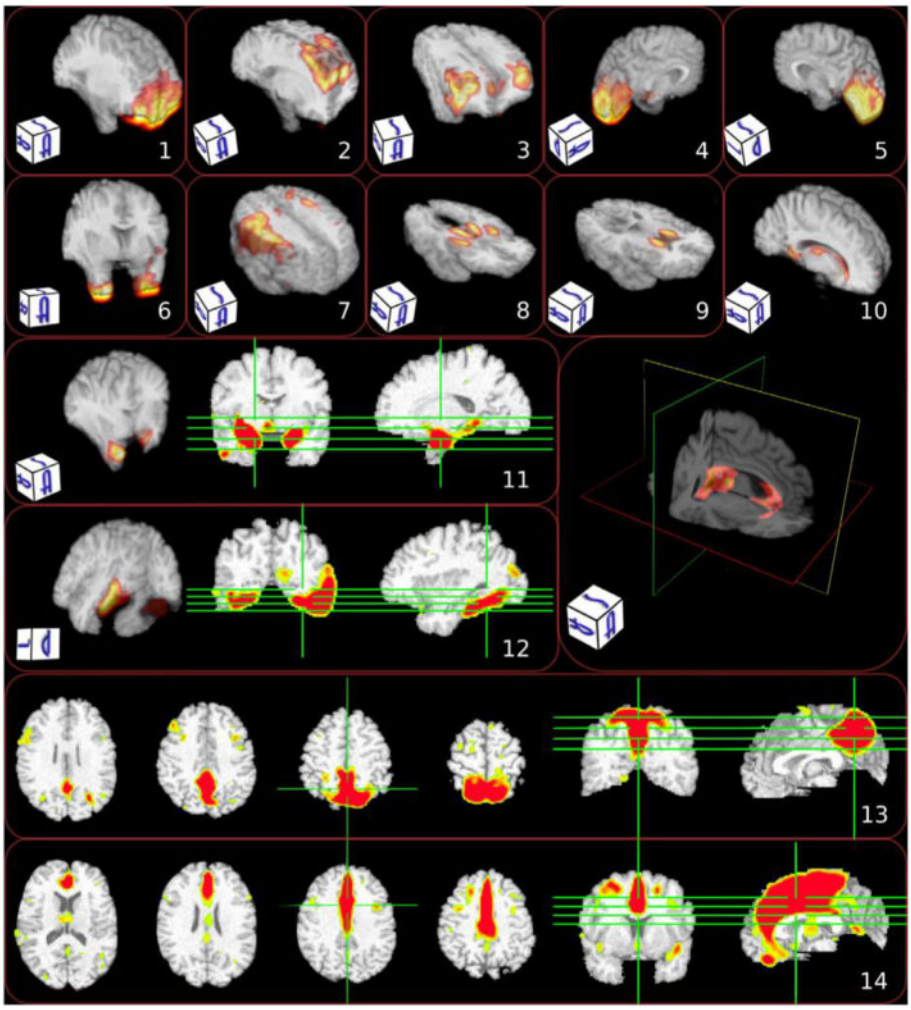

Fig. 4.

Components of structural co-variance identified via non-negative matrix factorization of structural MRI images reflect that a small number of anatomical units might be able to summarize high-dimensional imaging data in effective and highly interpretable ways. Figure reprinted from Sotiras et al. (2015).