What's in a functional brain parcellation?

Anonymous Author(s)
Affiliation
Address
email

Abstract

To communicate, to ground hypotheses and models, to analyse data, neuroscientists often refer to divisions of the brain. Here we consider atlases used to parcellate the brain when studying brain function. We discuss the meaning and the validity of these parcellations, from a conceptual point of view as well as by running various analytical tasks on popular functional brain parcellations.

Breaking up the human brain in territories is a longstanding tradition, dating back to Brodmann areas [5]. While these subdivisions of the brain originated from a desire to reveal homogeneous neural populations, they have become a major communication tool in human neurosciences. As such, a brain parcellation shapes how we think about units of the brain. When studying brain function, *eg* the neural implementations of mental processes, the macroscopic brain structures that we investigate implicitly shape our decompositions of mental function and the models we fit to data. Even the seemingly-simple study of brain responses to stimuli is tied to a choice of functional units. For instance, in the ventral visual stream, the Fusiform Face Area (FFA) is generally associated to face recognition [15], yet it is also strongly associated with visual-expertise in object recognition [10], and neurons supporting face recognition overlap in a distributed way with place recognition [14]. Beyond stimuli response, brain-wide models call for macroscopic units, for instance to study small-word properties of brain functional connectivity [1], to build coupled-oscillator dynamical models of neural activity, [8], or fitting spiking-neuron network models to brain-activity data [17].

Brain imaging has brought many different divisions of the brain into areas, available as brain atlases or parcellations. However, their extraction from brain-imaging data all entail different modeling choices. Even for atlases based on anatomical structures, there is weak concordance across different atlases used [4]. As functional subdivisions are not always simply visible in the microstructure or the brain anatomy, the problem is even more pronounced for the choice of functional units. Yet, the choice of functional parcellation impacts models of brain functions that are learned from these [26].

Here, we discuss principles to guide these choices. We first consider the meaning and construct validity of functional brain units. We then conduct an empirical study using analytic questions that probe different aspects of functional parcels, on 6 popular brain-imaging functional parcellations.

1 How to think about functional brain parcellations?

What are functional units? A neurocognitive model entails brain units associated with specific functions. For instance, a model of spatial navigation could position long-term spatial-memory representations in the hippocampus and short-term visuospatial representation in the intra-parietal sulcus. Both structures are fairly big, and encompass many millions of neurons. Should they be subdivided? What should be their specific boundaries? The problem is particularly pronounced in the cortex which is a continuous sheet of neurons. Models of vision provide a paradigmatic example of successful cognitive decomposition of a functional system. We understand vision well-enough to break it down in elementary constituents that can be mapped precisely to neural supports [13]. Studying local descriptors of the visual field, edge detection, color, or orientation, leads to

Submitted to the NeuroAI workshop at NeurIPS 2019, Vancouver, Canada.

retinotopic maps that reveal neural populations with local gradients and yet large-scale organizations in functional modules: from low-level functional units V1, V2... all the way to mid and high-level representations as in IT or the FFA. From a cognitive modeling perspective, these visual areas can be seen as implementing an elementary operation, analogous to layers of an artificial neural network [9]. In an ideal world, a functional brain parcellation would capture such units.

Can a large-scale division of the brain true, or merely useful? There is very active research on full-brain functional parcellations, used for instance to model functional connectivity. Multimodal imaging data has been used to seek a division of the brain in units with homogeneous neurobiological properties [11]. One goal is to reflect intrinsic brain structure with parcels that characterize better brain locations than stereotactic coordinates. Yet, representing brain function on a few hundred parcels is a vast simplification compared to the 100 billions neurons in the human brain, or the 100 thousands voxels in an fMRI volume. Even prototypical functional areas well known for sharp intrinsic functional properties, such as V1, have finer topological functional organization, as with retonotopic maps or ocular-dominance columns. Sharp boundaries are not present in other, higher-level, paradigmatic functional areas, such as the FFA, characterized by functional properties –face recognition– that partly overlap with neighboring areas [14]. The picture of functionally-uniform units is a convenient simplification with no intrinsic truth. And yet, it is very useful. A functional parcellation is a crucial component to build a rich picture of the neural basis of mental function: At the level of a study, it provides a necessary data reduction to fit full-brain models to the data. At the level of the field, the brain structures that it delineates define common objects of study.

Better MR-based functional parcellations Some brain parcellations are however more useful than others. Nodes adapted to the functional signal give better models of functional connectivity [22, 7]. Given the weak concordance between anatomical atlases [4], grounding a functional analysis on anatomical labels brings little benefit. Rather, the functional subdivisions can be learned from large-scale fMRI data [20].

2 An empirical investigation of some functional atlases for MR imaging

Measuring the analytic utility of an atlas To serve as a common object in the field, a good functional brain parcellation should be well suited for a variety of analytic tasks that model brain activity and its relationship to behavior.

a. Mapping brain responses Standard analysis in fMRI strives to detect difference in brain responses. Performing it on parcels mitigates multiple comparisons and inter-subject spatial variability [23]. Good parcels lead to detecting brain structure that match at the voxel level with well-powered analysis on the original volumetric data.

b. Decoding brain function Assign a functional label to a brain structure calls for decoding: predicting mental processes from observed brain activity [19]. Good functional units would have clear-cut functional labels and thus help decoding performance.

c. Fidelity to the original signal Summarizing brain activity on large parcels necessarily leads to signal loss. A good parcellation should minimize this distortion for a given number of regions.

d. Functional-connectivity biomarkers Brain parcellations are often used to define the nodes of functional-connectivity models. An independent validation of such a model is whether it can be well associated with variations of behavioral or clinical traits across subjects [7].

Popular functional atlases We investigate 6 popular atlases derived from fMRI, detailed in Table 1. The atlases differ in their number of regions, whether these regions are defined with continuous or binary maps, and which method was used to extract them.

Experimental procedures

a. Consistency in detection of neural responses We run an fMRI standard analysis at the parcel level and compare the overlap of detected brain territories to detections at the voxel level. We calibrate the noise level with the consistency of single subjects with regards to group-level results.

b. Decoding brain function We compare the performance of an SVM decoders trained on the data extracted on various parcellations and at the voxel level.

Name	# regions	Fuzzy	Extraction method	Reference
BASC	64, 122, 197, 325, 444	No	Hierarchical clustering	[3]
Craddock	200, 400	No	Spectral clustering	[6]
FIND	90, 499	Yes	ICA; Ward clustering	[21, 2]
Gordon	333	No	Local-gradient approach	[12]
UKBB ICA	21, 55	Yes	Selected ICA components	[16]
Schaefer	100, 200, 300, 400, 500, 600, 800, 1000	No	Gradient-weighted Markov Random Field (gwMRF)	[20]

Table 1: Functional atlases that we investigate

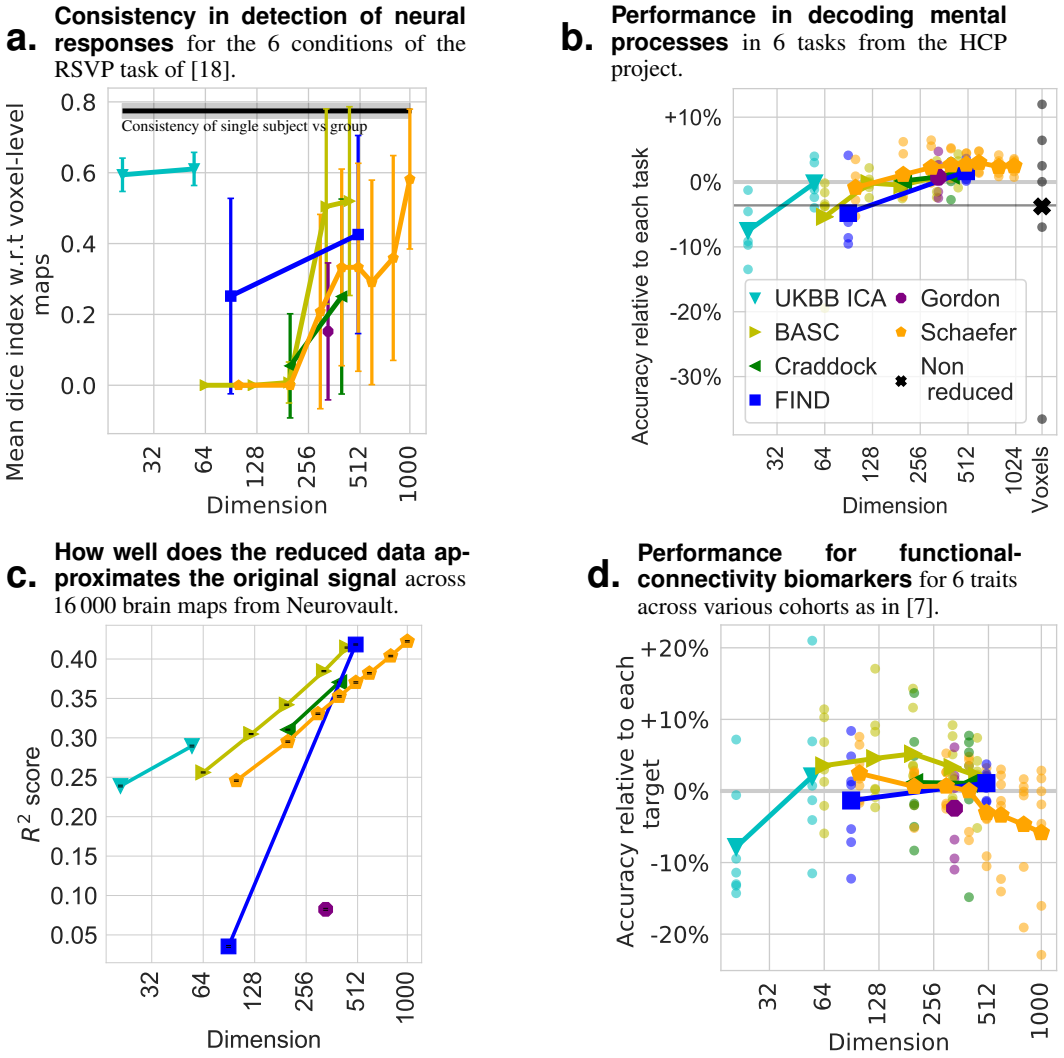

Figure 1: **Usefulness of different atlases for various analytic questions** across different datasets

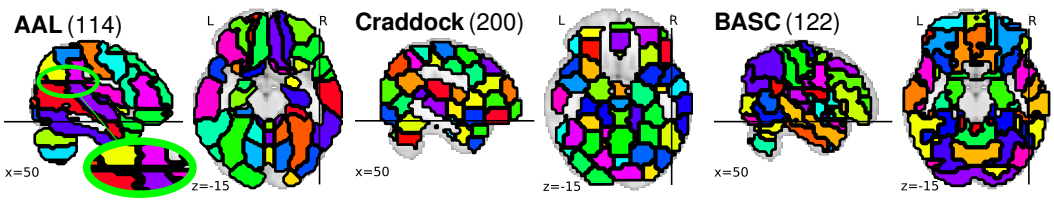

Figure 2: **Some brain parcellations used in functional-connectivity models** – Several boundaries in the AAL, an anatomical parcelation hand-drawn from one subject, are straight lines which is anatomical improbable. The Craddock parcellation does not capture the shape of local brain structures.

c. Fidelity to the original signal We compute the fraction of variance of the original signal explained by the summarizing brain images on parcels.

d. Functional biomarkers We measure the prediction performance using a standard functional-connectivity prediction pipeline [7] on the different parcellations.

Results: which atlases lead to clear analysis For standard analysis (Figure 1a) and data approximation (Figure 1c), where analysis on the voxel-level data defines the gold standard, the larger the number of regions, the better the performance. On the other hand, for predictive tasks –decoding brain responses (Figure 1b) or biomarkers from functional connectivity (Figure 1d)– a reduced number of regions acts as a regularizer and using a few hundred regions outperforms voxel-level analysis. At a low dimensionality, ICA-derived functional atlases performed comparatively well, confirming the usefulness of continuously defined node reported in [7]. At higher dimensions, BASC [3] gave the best overall compromise, aside from very high-dimensional settings (> 500) only covered by the Schaefer parcellation [20]. Overall, for most analytic tasks, very high-dimensional atlases –with a number of regions neighboring the thousand– are beneficial. Only to build connectomes is it useful to limit the number of nodes to a few hundreds. This can be explained because the number of edges grows quadratically, and quickly encounters the curse of dimensionality.

3 Conclusion

Atlases defining functional regions can lead to better constructs and better data analysis. Validating these functional regions is challenging, yet a guiding principle is that they should improve statistical modeling. More research is needed in high-dimensional atlases, which prove very useful. Consensus and adoption of easily-accessible functional atlases is important. Indeed, in the mean time, many computational modeling studies [1, 8] default to the AAL, an atlas that is neither functionnal, nor captures well anatomical boundaries (Figure 2).

But the evidence to ground the choice of a functional atlas is subtle and there is not simple story. From a pure signal-processing point of view (Figure 1 a and c), the best option at low dimensionality are ICA-derived modes, and at higher dimensionality the BASC atlas, built by clustering fMRI with weak spatial constraints. Decoding performance (Figure 1b) gives a useful measure of the functional specificity of the units defined [24]. In this respect, for higher dimensionalities the Schaefer atlas [20] –built from clustering with more spatial constraints– appears to give the most useful functional units. Finally, the best option to build brain-connectivity models (Figure 1d) is the BASC atlas. Continuously-defined modes, as opposed to hard parcellations, give excellent expressive power for low dimensionality, however no such atlas is currently available at high dimensionality.

It is unclear how close any of these atlases get to actual functional units, when these exist. In some regions for which functional organization is well known, [25] confirmed the face validity of parcels extracted from clustering fMRI. When an atlas is used to define the units of models of brain dynamics and function, these units should ideally capture coherent neural populations. Yet, when the data modeled are fMRI, these come with many measurement imperfections, including intersubject variability. Multimodal approaches promise to define parcellations that capture this variability [11]. Yet, as they entail a significant increase in complexity, most studies prefer to use a predefined atlas. Such an atlas is a simplified view on brain architecture at the population level. Its choice should be guided by its suitability to an analytic task, as studied here: all atlases are wrong, some are useful.

References

[1] S. Achard, R. Salvador, B. Whitcher, J. Suckling, and E. Bullmore. A resilient, low-frequency, small-world human brain functional network with highly connected association cortical hubs. *J Neurosci*, 26:63, 2006.

[2] A. Altmann, B. Ng, S. M. Landau, W. J. Jagust, and M. D. Greicius. Regional brain hypometabolism is unrelated to regional amyloid plaque burden. *Brain*, 138:3734, 2015.

[3] P. Bellec, P. Rosa-Neto, O. Lyttelton, H. Benali, and A. Evans. Multi-level bootstrap analysis of stable clusters in resting-state fMRI. *NeuroImage*, 51:1126, 2010.

[4] J. W. Bohland, H. Bokil, C. B. Allen, and P. P. Mitra. The brain atlas concordance problem: quantitative comparison of anatomical parcellations. *PloS one*, 4(9):e7200, 2009.

[5] K. Brodmann. *Vergleichende Lokalisationslehre der Grosshirnrinde in ihren Prinzipien dargestellt auf Grund des Zellenbaues*. Barth, 1909.

[6] R. C. Craddock, G. A. James, P. E. Holtzheimer, X. P. Hu, and H. S. Mayberg. A whole brain fMRI atlas generated via spatially constrained spectral clustering. *Human brain mapping*, 33:1914, 2012.

[7] K. Dadi, M. Rahim, A. Abraham, D. Chyzhyk, M. Milham, B. Thirion, and G. Varoquaux. Benchmarking functional connectome-based predictive models for resting-state fMRI. *NeuroImage*, 192:115 – 134, 2019.

[8] G. Deco, V. Jirsa, A. R. McIntosh, O. Sporns, and R. Kötter. Key role of coupling, delay, and noise in resting brain fluctuations. *Proceedings of the National Academy of Sciences*, 106(25):10302–10307, 2009.

[9] M. Eickenberg, A. Gramfort, G. Varoquaux, and B. Thirion. Seeing it all: Convolutional network layers map the function of the human visual system. *NeuroImage*, 152:184–194, 2017.

[10] I. Gauthier, M. J. Tarr, A. W. Anderson, P. Skudlarski, and J. C. Gore. Activation of the middle fusiform'face area'increases with expertise in recognizing novel objects. *Nature neuroscience*, 2(6):568, 1999.

[11] M. F. Glasser, T. S. Coalson, et al. A multi-modal parcellation of human cerebral cortex. *Nature*, 536(7615):171, 2016.

[12] E. M. Gordon, T. O. Laumann, et al. Generation and Evaluation of a Cortical Area Parcellation from Resting-State Correlations. *Cereb Cortex*, 26:288, 2014.

[13] K. Grill-Spector and R. Malach. The human visual cortex. *Annu. Rev. Neurosci.*, 27:649–677, 2004.

[14] J. V. Haxby, I. M. Gobbini, M. L. Furey, et al. Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science*, 293:2425, 2001.

[15] N. Kanwisher, J. McDermott, and M. M. Chun. The fusiform face area: a module in human extrastriate cortex specialized for face perception. *Journal of neuroscience*, 17:4302, 1997.

[16] K. L. Miller, F. Alfaro-Almagro, et al. Multimodal population brain imaging in the UK biobank prospective epidemiological study. *Nature Neuroscience*, sep 2016.

[17] T. T. Nakagawa, M. Woolrich, et al. How delays matter in an oscillatory whole-brain spiking-neuron network model for meg alpha-rhythms at rest. *Neuroimage*, 87:383–394, 2014.

[18] A. L. Pinho, A. Amadon, T. Ruest, et al. Individual Brain Charting, a high-resolution fMRI dataset for cognitive mapping. *Scientific Data* , 5:180105, 2018.

[19] R. A. Poldrack. Inferring mental states from neuroimaging data: from reverse inference to large-scale decoding. *Neuron*, 72:692, 2011.

[20] A. Schaefer, R. Kong, et al. Local-Global Parcellation of the Human Cerebral Cortex from Intrinsic Functional Connectivity MRI. *Cerebral Cortex*, 28:3095, 2018.

[21] W. Shirer, S. Ryali, E. Rykhlevskaia, V. Menon, and M. Greicius. Decoding subject-driven cognitive states with whole-brain connectivity patterns. *Cerebral Cortex*, 22:158, 2012.

[22] S. Smith, K. Miller, G. Salimi-Khorshidi, M. Webster, C. Beckmann, T. Nichols, J. Ramsey, and M. Woolrich. Network modelling methods for fMRI. *Neuroimage*, 54:875, 2011.

[23] B. Thirion, G. Flandin, P. Pinel, A. Roche, P. Ciuciu, and J. Poline. Dealing with the shortcomings of spatial normalization: Multi-subject parcellation of fMRI datasets. *Hum brain map*, 27:678, 2006.

[24] G. Varoquaux, Y. Schwartz, R. A. Poldrack, B. Gauthier, D. Bzdok, J.-B. Poline, and B. Thirion. Atlases of cognition with large-scale human brain mapping. *PLoS computational biology*, 14(11):e1006565, 2018.

[25] B. Yeo, F. Krienen, J. Sepulcre, M. Sabuncu, et al. The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *J Neurophysio*, 106:1125, 2011.

[26] A. Zalesky, A. Fornito, I. H. Harding, L. Cocchi, M. Yücel, C. Pantelis, and E. T. Bullmore. Whole-brain anatomical networks: does the choice of nodes matter? *Neuroimage*, 50:970, 2010.