

# **Multi-task learning of cognitive concepts**

## **Application to brain image meta-analysis**

### **Summary**

An important part of scientific research is now evolving toward a discovery science paradigm, where knowledge is acquired by using inductive reasoning, given a sufficient amount of data. In spite of its premises, this is a hard endeavor, as it rests on the availability of (public) data, shared descriptions/ontologies that can be associated with these data and on learning procedures that are used to draw inference based on the data and ontologies. As these conditions are being realized for the first time in the rapidly evolving field of brain imaging, we wish to design the analysis framework that will make it possible to create a cognitive atlas from publicly available functional Magnetic Resonance Imaging datasets. With this framework, we hope to create a strong incentive toward more data sharing in the brain imaging community.

Based on our current experience, we propose to improve the inference procedures used to learn the atlas: first, we will increase the flexibility of linear models used to associate cognitive terms of interest with brain activation maps. Our second aim is to design a more powerful decoder that will rest on multi-task learning to better take advantage of the information available across heterogeneous datasets. In order to deal with data scarcity, we will set up a semi-supervised learning framework. The code produced and resulting brain cognitive atlas will be made public to the community. Whenever this is possible, algorithmic improvements (e.g. stochastic gradient descent approaches to perform prediction on large data sets) will also be shared more widely based on open-source Python software.

# Project

## Context

**Cognitive ontologies and brain mapping** Neuroimaging produces huge amounts of complex data that are used to better understand the relations between brain structure and various cognitive functions. In the long term, the results of this field will be used to better understand, characterize, predict and diagnose brain diseases. Although some aspects of the acquisition and analysis of these data are gradually being standardized, the neuroimaging community is still largely missing appropriate tools to store and organize the knowledge related to the data. As a consequence, there is little data reuse [1], which implies that the neuroimaging community wastes important resources, and only the corpus of publications gives an account of the accumulated results. An important reason for this is that **there is no proper inference framework that can link existing data with a proper ontology of cognitive processes**. Finding a mechanism to gather and update the information on brain structure and function carried by these data is an important challenge, as the images carry much more information than the reports given in the neuroimaging literature (typically activation peaks coordinates in a reference space).

**A learning framework for cognitive mapping** Current techniques for meta-analysis based on published coordinates [2] are not satisfactory, as the abstraction of brain maps into a few peak activation coordinates discards essential information; using instead the original statistical parametric maps is more sensitive and more consistent with the current practice of researchers in this field [3]. What we propose is thus to automate the learning of brain functional organization based on image data. This involves three main building blocks:

- (Public) image data provided with sufficient annotations to compare them with other data.
- An ontology that relates the image annotations and makes it possible to draw inferences on those.
- A learning framework that associates concepts with the image data. The main difficulty here is to take advantage of the shared information between the experimental data and the concepts involved.

**Opportunity: development of large-scale learning tools and public repositories** An opportunity to address this problem is provided by the current development of public data repositories, such as OpenfMRI<sup>1</sup>, large-scale datasets (Human Connectome Project HCP<sup>2</sup>) and the NeuroVault<sup>3</sup> site, in which neuroscientists can upload activation maps with annotations with a minimal amount of efforts. Altogether, these initiatives provide an opportunity to develop a multilevel framework for learning the brain activation patterns associated with a large number of experimental conditions and help the neuroimaging community to organize the knowledge acquired along twenty years of brain imaging.

Regarding the ontology, two models are currently available [4, 5]. As we are mostly limited by the lack of diversity of the concepts used frequently enough across current public datasets, we will first pragmatically blend these two ontologies as needed. In the long term, more ambitious descriptions will be used, possibly based on those of the BrainSpell or Neurosynth framework<sup>4</sup>.

What is needed are thus learning tools that can use these resources. Developing those is the aim of CogLearn.

## Proposed work

### Mapping cognitive terms to brain maps

The first analytic step consists in mapping cognitive terms to brain activation data. Cognitive terms are related to the experimental conditions used in experiments and their combinations. They are supposed to model cognitive components involved in the performance of a given task. Such occurrences are hard to define, because the ontologies that support this description are not very accurate, and the segmentation of basic cognitive components,

---

<sup>1</sup><https://openfmri.org>

<sup>2</sup><http://www.neuroscienceblueprint.nih.gov/connectome>

<sup>3</sup><http://neurovault.org>

<sup>4</sup><http://brainspell.org>, <http://neurosynth.org/>

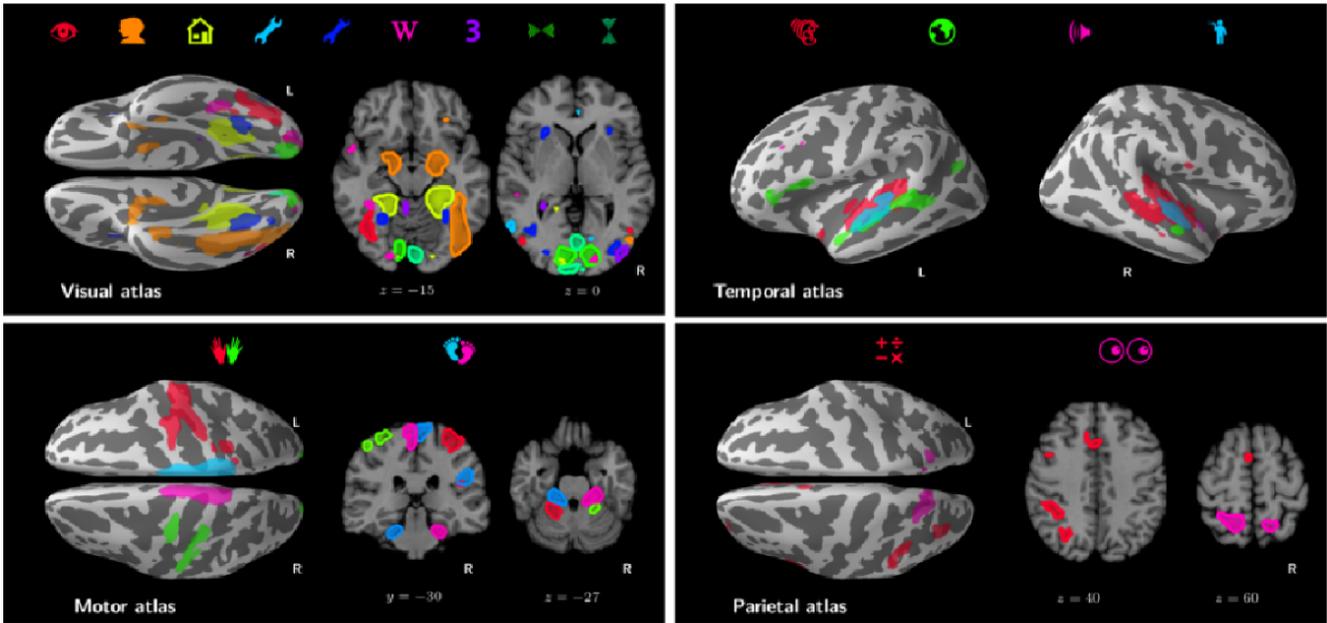


Figure 1: Based on 34 datasets, one can obtain a mapping of a large collection of functional responses, whereby the displayed maps represent regions that are statistically associated with cognitive terms across all these studies. Besides the richness of these data, cross-studies results present a consensus across labs, hence are more reliable.

a.k.a. the *cognitive atoms* involved in a given task, is a topic of active research [6]. The second difficulty is that the occurrence information alone is not sufficient to quantitatively map cognitive atoms to brain function.

Current achievements have shown that a simple binary occurrence matrix could be used to map some cognitive components to brain activation patterns through a classical linear model (the so-called *forward model* in [7, 8]), that is fit in a voxel-by-voxel fashion, by pooling the information across studies, see Fig. 1 for an illustration. Yet this model has clearly not a sufficient level of detail to deal with labeling/modeling issues.

The first way to improve it is to rely on a diagnosis approach, by considering current misfits, the structure that remains in the model residuals, the respective importance of different variability terms. Once the main origin of signal misfit will be identified, we will improve the encoding model. A simple yet effective way of improving the model quality is to *learn* a weighting that quantifies the importance of the different cognitive terms in the experimental conditions used for mapping (see e.g. [9]). The post-doc candidate will devise some strategies to learn these weights and assess the gain. As this problem entails a classical bias/variance trade-off, it is important to choose a parametrization that generalizes well to unseen data and even possibly unseen datasets.

An underlying technical aspect is that linear model fitting across very large datasets should be made **efficient**. An obvious solution in that respect is to rely on summary statistics approaches, where intermediate results (mean effects and variances) are obtained at the study level and then combined across studies. Mixed-effects models are a natural solution to that problem. By decreasing the memory demand of the estimator, the summary statistics approach can yield impressive time savings. It is actually a pre-requisite to further model improvements. Further gains will be obtained using online estimators and will be implemented if time permits.

### Multivariate learning framework

The second view on the problem is a *decoding* view, in which the inference consists in predicting terms from the ontology based on their occurrence in association with activation images.

**Multi-task learning** A major novelty of the proposed framework with respect to standard decoding approaches in functional brain imaging [10] is that the problem is now *multi-label*, since all the labels in the ontology can

potentially be used to tag a target image. However, our choice here is also handle it as a *multi-task problem*, in which the statistical association between tasks can be leveraged to improve the decoding. More specifically, we have recently shown that sharing information across a large set of tasks benefits to decoding accuracy and to the interpretability of the model [11]. The large-scale decoding experiments described here is thus an excellent opportunity to assess the gain brought by multi-task learning. More generally, relying on more powerful classification architectures makes little sense if the problem is limited to simple binary classification problems for which convex solvers perform well, but it brings significant improvements for complex tasks [12]. Our strategy is thus to increase the complexity of the decoding model together with that of the problem.

**Semi-supervised learning** Another observation is that the use of unlabeled data, which introduces a semi-supervised learning framework, may be beneficial to both performance and interpretability, as this provides additional information on the structures underlying the data, which remain hard to characterize in a purely supervised setting [11]. In the case of functional neuroimaging, unlabeled observations are ubiquitous: in particular, a wealth of resting-state data are available and are known to provide meaningful decomposition of brain images into networks [13]. We will thus investigate the impact of augmenting the learner with an unsupervised layer that tries to build sensible approximations of unlabeled data.

**Interpretability of the results** Finally, an important concern is the interpretability of the final weight maps (assuming that linear models are used). The usefulness of multivariate learners is not only characterized by the prediction performance, but also by the reliability and stability of the estimates of the predictive patterns (the so-called *recovery* or *identification* problem). For instance, elastic net regularization is known to yield very good accuracy, but is a poor approach for recovering truly predictive features [14].

We have developed recently two frameworks to overcome this difficulty: one based on randomized parcellations of the brain volume that are combined to yield an accurate ensemble model [15, 16]. Second, we have introduced spatially regularized models that can be used in a convex setting [14, 17]. We plan to benchmark these two families of models regarding recovery and computation time on the large-scale experiment performed within CogLearn.

**Code development** The code developed in the framework of this project will rely on the existing open source scientific Python tool stack and in particular Scikit-learn<sup>5</sup> [18]; the post-doc will also use and develop scientific Python tools that have been designed specifically for neuroimaging, such as Nilearn<sup>6</sup>, PySurfer etc.

## Outcome of the project

The main outcome of the project CogLearn is a learner that will be available as open-source code and that will rely on open data to infer classification rules that predict cognitive terms based on test data. We will develop the project in tight collaboration with NeuroVault and OpenfMRI, so that the authors and users of these tools can also probe the learning procedure and reproduce the same kind of inference on other data.

We are addressing important bottlenecks of such models that are key to the development of community-level meta-analyses. The expected impact is that this will showcase the interest of public data repositories and incite researchers to share the data, enhancing the reproducibility of scientific results in neuroimaging [19].

## Skilled required from the post-doc candidate

- Understanding of cognitive functional imaging (with prior experience in that field, if possible).
- Expertise on modern machine learning tools.
- Good knowledge of statistical inference techniques.
- Knowledge of Python language.
- Interest in Open Software development.

---

<sup>5</sup><http://scikit-learn.org>

<sup>6</sup><http://nilearn.github.io>

## References

- [1] J.-B. Poline, J. L. Breeze, S. Ghosh, K. Gorgolewski, Y. O. Halchenko, M. Hanke, C. Haselgrove, K. G. Helmer, D. B. Keator, D. S. Marcus, R. A. Poldrack, Y. Schwartz, J. Ashburner, D. N. Kennedy, and Others, “Data sharing in neuroimaging research,” *Frontiers in Neuroinformatics*, vol. 6, no. April, p. 1–13, 2012.
- [2] P. T. Fox and J. L. Lancaster, “Opinion: Mapping context and content: the brainmap model.,” *Nat Rev Neurosci*, vol. 3, pp. 319–321, Apr 2002.
- [3] G. Salimi-Khorshidi, S. M. Smith, J. R. Keltner, T. D. Wager, and T. E. Nichols, “Meta-analysis of neuroimaging data: a comparison of image-based and coordinate-based pooling of studies.,” *Neuroimage*, vol. 45, pp. 810–823, Apr 2009.
- [4] J. Turner and A. Laird, “The cognitive paradigm ontology: design and application,” *Neuroinformatics*, vol. 10, p. 57, 2012.
- [5] R. A. Poldrack, A. Kittur, D. Kalar, E. Miller, C. Seppa, Y. Gil, D. S. Parker, F. W. Sabb, and R. M. Bilder, “The cognitive atlas: toward a knowledge foundation for cognitive neuroscience,” *Frontiers in neuroinformatics*, vol. 5, 2011.
- [6] R. Poldrack, A. Kittur, D. Kalar, E. Miller, C. Seppa, Y. Gil, P. D.S., F. Sabb, and R. Bilder, “The cognitive atlas: toward a knowledge foundation for cognitive neuroscience.,” *Front Neuroinform*, vol. 5, p. 17, 2011.
- [7] Y. Schwartz, B. Thirion, and G. Varoquaux, “Mapping cognitive ontologies to and from the brain,” in *Advances in Neural Information Processing Systems*, (Lake Tahoe, Nevada), John Lafferty, Dec. 2013.
- [8] Y. Schwartz, B. Thirion, R. Poldrack, and G. Varoquaux, “Isolating brain functions using ontology-based decoding,” *Submitted to PNAS*, 2015.
- [9] B. T. T. Yeo, F. M. Krienen, S. B. Eickhoff, S. N. Yaakub, P. T. Fox, R. L. Buckner, C. L. Asplund, and M. W. L. Chee, “Functional specialization and flexibility in human association cortex.,” *Cereb Cortex*, Sep 2014.
- [10] J.-D. Haynes and G. Rees, “Decoding mental states from brain activity in humans.,” *Nat Rev Neurosci*, vol. 7, pp. 523–534, Jul 2006.
- [11] D. Bzdok, M. Eickenberg, O. Grisel, B. Thirion, and G. Varoquaux, “Semi-Supervised Factored Logistic Regression for High-Dimensional Neuroimaging Data,” *Neural Information Processing Systems*, Dec. 2015.
- [12] Y. Bengio, “Learning deep architectures for ai,” *Found. Trends Mach. Learn.*, vol. 2, pp. 1–127, Jan. 2009.
- [13] S. M. Smith, P. T. Fox, K. L. Miller, D. C. Glahn, P. M. Fox, C. E. Mackay, N. Filippini, K. E. Watkins, R. Toro, A. R. Laird, and C. F. Beckmann, “Correspondence of the brain’s functional architecture during activation and rest.,” *Proc Natl Acad Sci U S A*, vol. 106, pp. 13040–13045, Aug 2009.
- [14] A. Gramfort, B. Thirion, and G. Varoquaux, “Identifying predictive regions from fMRI with TV-L1 prior,” in *Pattern Recognition in Neuroimaging (PRNI)*, (Philadelphia, États-Unis), IEEE, June 2013. ANR grant BrainPedia, ANR-10-JCJC 1408-01, FMJH Program Gaspard Monge in optimization and operation research with support from EDF.
- [15] G. Varoquaux, A. Gramfort, and B. Thirion, “Small-sample brain mapping: sparse recovery on spatially correlated designs with randomization and clustering,” in *International Conference on Machine Learning* (L. John and P. Joelle, eds.), (Edimbourg, Royaume-Uni), Andrew McCallum, June 2012.
- [16] A. Hoyos-Idrobo, Y. Schwartz, G. Varoquaux, and B. Thirion, “Improving sparse recovery on structured images with bagged clustering,” in *International Workshop On Pattern Recognition In Neuroimaging (PRNI), 2015*, (Palo alto, United States), June 2015.
- [17] E. Dohmatob, A. Gramfort, B. Thirion, and G. Varoquaux, “Benchmarking solvers for TV-l1 least-squares and logistic regression in brain imaging,” in *Pattern Recognition in Neuroimaging (PRNI)*, (Tübingen, Germany), IEEE, June 2014.
- [18] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, “Scikit-learn: Machine learning in Python,” *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.
- [19] K. S. Button, J. P. A. Ioannidis, C. Mokrysz, B. A. Nosek, J. Flint, E. S. J. Robinson, and M. R. Munafò, “Power failure: why small sample size undermines the reliability of neuroscience.,” *Nat Rev Neurosci*, vol. 14, pp. 365–376, May 2013.