Inria Saclay, Équipe Parietal http://team.inria.fr/parietal NeuroSpin, bat 145, CEA Saclay

Knowledge and Representation Integration on the Brain

Research theme: machine learning, life sciences
Keywords: machine learning, deep learning, NLP, brain imaging, heterogeneous representations, inference.
Duration & salary: 3 years months, about 2600 €net monthly, depending on experience.
Research team: Parietal (INRIA Saclay and CEA)
Adviser: Bertrand Thirion
Contact: bertrand.thirion@inria.fr
Application: Interested candidate should send CV along with the list of publications, a statement of interests in the position, and the names and contact information of at least three professional referees.

Summary

Cognitive science describes mental operations, and functional brain imaging provides a unique window into the brain systems that support these operations. A growing body of neuroimaging research has provided significant insight into the relations between psychological functions and brain activity. However, the aggregation of cognitive neuroscience results to obtain a systematic mapping between structure and function faces the roadblock that cognitive concepts are ill-defined and may not map cleanly onto the computational architecture of the brain.
To tackle this challenge, we propose to leverage rapidly increasing data sources: text and brain locations described in neuroscientific publications, brain images and their annotations taken from public data repositories (NeuroVault, Cognitive atlas), and several reference datasets. Our aim here is to develop multi-modal machine learning techniques to bridge these data sources. This project develops representation techniques for noisy data to couple brain data with descriptions of behavior or diseases, in order to extract semantic structure.

Objective

Existing multi-modal machine learning techniques have been developed for relatively abundant data, with overall high *signal-to-noise ratio* (SNR): text, natural images, videos, sound. These data are most often non-ambiguous, while brain data typically are, due to the low SNR per image and, more crucially, *poor annotation quality*. We propose to tackle this by adapting machine learning solutions to this low-SNR regime: introduction of priors, aggressive dimension reduction, aggregation approaches and data augmentation to reduce overfit.
While data sources contain lots of implicit information that could be used as targets in supervised learning, there is most often no obvious way to extract it. We propose to tackle this by using additional, ill- or not-annotated data, relying on *self-supervision methods*.
This project is part of the Karaib chair. It develops representation techniques for low-SNR data to couple brain data with descriptions of behavior or diseases in order to extract semantic structure.

Eventually, one should be able to reason about the information extracted within this project. For this, we will develop dedicated statistical, causal and the formal (ontology-based), namely the Neurolang system.

Proposed work

We here outline two possible directions of research that may be addressed:

Leveraging Modern language models of psychological literature NeuroSynth [6, 3] and Neuroquery use text analysis to extract brain activation locations from the neuroimaging literature: The strength of these data is the amount of different paradigms covered, as well as the full-text publications associated with each sample that provide a rich yet complex description of behavior. In this context, the data are not images, but term occurrences in publications and corresponding brain locations. The setting is close to natural language processing (NLP) where many semantic models rely on co-occurrence matrices. We will adapt these models to incorporate brain locations, as in [5].

NLP has recently been revolutionized by the advent of powerful architectures, such as BERT [2], based on an encoder named *Transformer*, that is obtained from attentional mechanisms. Importantly, we will leverage these transformers, simply specializing them to the vocabulary of cognitive neuroscience and psychiatry. To inject information from brain mapping in these models, we will generate co-occurrence matrices by considering pairs of publications with their terms, and assigning a weight corresponding to the overlap of their activation peaks (generated from the convolution of reported locations with a Gaussian kernel) [5, 4, 3].

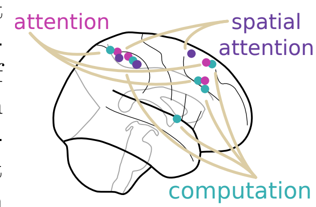

Deep learning for coordinated representation of brain topographic information and psychological concepts Deep learning is well suited to our goals of learning intermediate representations between brain signals and text-based description of mental processes — following [1], we call them *coordinated representations*. For this purpose we will develop a framework that unites in one deep learning formulation the task of estimating brain structures, cognitive concepts, and their links. Brain structures and cognitive concepts will appear as intermediate representations responsible for linking brain activity to observed behavior.

Qualification and skills

- PhD in computer science, physics, mathematics, artificial intelligence, biomedical engineering, computational neuroscience or related quantitative fields
- Exceptionally strong computational skills in biomedical signal processing, machine learning, statistical modeling.
- Solid experience in Python
- Independent, self-motivated with a proven track record of productive research

- Excellent verbal and written communication skills
- Ability to work effectively both independently and in collaboration with multiple investigators
- Strong publication record and excellent academic credentials
- Interest in cognitive neurosciences

[1] T. Baltrušaitis, C. Ahuja, and L. Morency. Multimodal machine learning: A survey and taxonomy. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(2):423–443, Feb 2019.

[2] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: pre-training of deep bidirectional transformers for language understanding. *CoRR*, abs/1810.04805, 2018.

[3] Jérome Dockès, Russell Poldrack, Romain Primet, Hande Gözükan, Tal Yarkoni, Fabian M Suchanek, Bertrand Thirion, and Gaël Varoquaux. NeuroQuery, comprehensive meta-analysis of human brain mapping. *eLife*, March 2020.

[4] Jérôme Dockès, Demian Wassermann, Russell Poldrack, Fabian M. Suchanek, Bertrand Thirion, and Gaël Varoquaux. Text to brain: predicting the spatial distribution of neuroimaging observations from text reports. In *MICCAI 2018*, pages 1–18, Granada, Spain, September 2018.

[5] Timothy Rubin, Oluwasanmi O Koyejo, Michael N Jones, and Tal Yarkoni. Generalized correspondence-LDA models (GC-LDA) for identifying functional regions in the brain. In *NIPS*, pages 1118–1126. Curran Associates, Inc., 2016.

[6] Tal Yarkoni, Russell A Poldrack, Thomas E Nichols, David C Van Essen, and Tor D Wager. Large-scale automated synthesis of human functional neuroimaging data. *Nature methods*, 8:665, 2011.