

Machine Learning for NeuroImaging, learn the basics before going deeper.

Janaina Mourao-Miranda Co Organizer

University College London
Computer Science
London, London
United Kingdom

Christophe Phillips Organizer

University of Liege
GIGA Institute
Liège
Belgium

The application of machine learning techniques to neuroimaging data has increased substantially in the last years, from classic “support vector machines” to more recent “deep learning”, leading to a large body of publications.

Over the years the “machine learning” field (at the border of statistics and engineering) has delivered families of tools, which have been adapted to investigate cognitive and clinical neuroscience questions. Depending on the research question asked, experimental design and data modality, it is important that the experimenter knows which tools to use and how to draw reliable conclusions. The course will focus on subject and/or patient predictions (for cognitive and clinical applications), including classification and regression models. The usual functional and structural MRI modalities will be covered but the presented approaches could also be applied on other types of data.

Model validation and statistical inference are particularly crucial as these notions somewhat differ from the standard univariate statistics usually applied to analyse neuroimaging data (e.g. General Linear Model) and should thus be specifically addressed. After introducing the basic foundations of machine learning in neuroimaging, the remaining talks will tackle more advanced methodological points, such as classic machine learning, unsupervised learning and deep learning applications pertaining to neuroimaging.

At the end of the course, the neuroscientist should have a global understanding of machine learning approaches, how to apply these tools to his/her own data to address new questions in a rigorous way, and how to interpret the outcomes of these analyses as well as how to draw reliable conclusions.

Objective

he course is organized such that the participant acquire knowledge about

- the various machine learning tools available and how to make rigorous statistical inferences with these;
- the “Do’s and Don’ts” of the technique in neuroimaging, their limitations, and how to interpret their results depending on their question of interest;
- the application of these methods on different modalities such as fMRI and sMRI, as well as various applications in cognitive and clinical neuroscience.

Target Audience

The target audience would consist of neuroscientists with intermediate to advanced knowledge of standard neuroimaging analysis techniques but with no or intermediate knowledge of pattern recognition/machine learning techniques. Researchers with a methodological or ‘machine learning’ background could also benefit from the course by being exposed to specificities of neuroimaging data and the questions typically asked in neuroimaging applications.

Presentations

Machine learning in neuroimaging, what are we talking about?

This introductory lecture describes the basic principles of machine learning (ML) and how these tools can be applied to neuroimaging data. These will focus on label or parameter prediction, i.e. classification or regression models, of subjects’ whole brain images. First these approaches will be compared to standard univariate GLM approaches. Then data representation, as samples and sets of features, and the principles of supervised & unsupervised machine learning are introduced. The talk will also introduce standard ML approaches, e.g. kernel methods, and how their performance is assessed. The issues of generalizability and inference are central: how well can we predict unseen data? and are we performing above chance level? Most of these points will be further developed in the following talks.

Presenter

Christophe Phillips, University of Liege
GIGA Institute
Liège
Belgium

Cross-validation, how do we assess predictive performance?

Decoding calls for measuring the predictive power of a decoder. It may be used to tune parameters of the decoder to improve it. Cross-validation is the central tool for these purposes. This tutorial will go over cross-validation for decoding and address a few fundamental questions: – What exactly is cross-validation? Why do we need it in our machine learning studies? – What are the different types of CV, and how do they compare to each other? – How to measure and compare predictive performance, and how not to? – What are the key terminologies one needs to be aware of (overfitting, different performance metrics, sources of bias, etc.)? The tutorial will also be interactive and

participative wherever possible.

Presenter

Pradeep Reddy Raamana, PhD, Baycrest Health Sciences
Rotman Research Institute
Toronto, ON
Canada

Confounding variables, how do we account for them?

The presence of confounds is a common problem in neuroimaging studies. Confounds, also known as covariates, are variables that affect the imaging data but whose effect we are not interested in studying (e.g. scanner, gender, age). There is a lot of evidence that confounds can strongly affect the performance of neuroimaging-based machine learning models and a number of approaches have been proposed to deal with the problem. This talk will describe the problem of confounds focusing on clinical neuroimaging studies, present different approaches to deal with the problem and discuss their advantages and limitations.

Presenter

Janaina Mourao-Miranda, University College London
Computer Science
London, London
United Kingdom

How do we test our hypothesis? A permutation approach

With the adoption of novel analysis methods beyond the classical GLM-based approach including multivariate pattern analysis, there is an increasing need for hypothesis tests that make only weak distributional assumptions, in particular permutation tests. In this talk I will give an introduction to permutation tests from the ground up, giving practitioners clear guidelines on how to design and implement them. The talk will cover the following topics: 1/ A basic example and the formal procedure to perform a permutation test; 2/ Exchangeability as the assumption underlying permutation tests; 3/ Examples: univariate and multivariate two-sample and paired tests; 4/ Limited exchangeability in fMRI time series, randomization test and exchangeability of run-wise estimates; 5/ Further topics: multiple comparisons, extensions based on mirror symmetry and sphericity. To improve audience engagement, after the first, second, and fourth section listeners will be given the opportunity to ask questions (in addition to questions after the whole talk).

Presenter

Carsten Allefeld, City, University of London

How do we test our hypotheses? a Bayesian approach.

The decoding of neural correlates is a task of primary interest in many neuroscientific investigations where typically, the subject is presented with a set of stimuli and a classification algorithm is used to extract stimulus-related information from the neuroimaging data. Frequently, in those investigations, the final outcome of the inferential process is a test of hypotheses to quantify the evidence extracted by the classifier to support the neuroscientific claims. In this tutorial, we present the Bayesian way to test hypotheses on the results of a classification-based pipeline of analysis. Specifically, we describe the comparison of Bayesian hypothesis testing with the classical test of hypotheses, discussing the pros and cons. Moreover, we show how to conduct the Bayesian test in case of multiple categories of stimulus and multiple subjects.

Presenter

Emanuele Olivetti, Ph.D., Bruno Kessler Foundation
NeuroInformatics Laboratory (NILab)
Trento, Trento
Italy

What makes a good multivariate model for fMRI-based decoding?

Standard multivariate analyses yield two pieces of information: what experimental manipulations are indeed reflected in brain activity, and with due caution, a characterization of the brain activity patterns associated with these manipulations. Performing some inference with pattern analysis model raises some caveats, in particular regarding the reliability of the multivariate patterns obtained in this framework. In this talk, we will discuss in depth the differences between classical univariate and multivariate analyses, and analyze the pros and cons of either approach. We will describe the criteria that can be used to assess the classifiers in terms of accuracy, interpretability and stability. We will then describe cheap, yet effective ways to improve the quality of the models using fast clustering schemes for spatial regularization and model averaging. We will illustrate these concepts on datasets and accessible to the audience as Jupyter notebooks.

Presenter

Undefined labels? Try unsupervised approaches!

Unsupervised methods such as Partial Least Squares (PLS) and Canonical Correlation Analysis (CCA) can allow to overcome the limitations that arise with classification when classes are not well defined. In the field of psychiatry, neurobiological abnormalities largely overlap between disorders, which make the boundaries between diagnostic categories unreliable and challenge the current nosology. PLS/CCA can be used to identify neuroimaging patterns and clinical symptoms that are shared across disorders. These methods aim to maximize the covariance (PLS) / correlation (CCA) between two data matrices by deriving latent factors that are optimally weighted linear combinations of the original variables. They are particularly powerful in identifying multivariate relationships between two high dimensional data matrices, as they do not require dimensionality reduction. Furthermore, unsupervised approaches can also be deployed in concert with supervised methods, by extracting meaningful/interpretable features which can then be used to predict clinical scores (e.g., disease severity), or to classify between two groups (e.g., patients vs. controls).

Presenter

Valeria Kebets, National University of Singapore Singapore, Singapore
Singapore

Deep Learning for Neuroimaging: What are we talking about?

Linear models have been the gold standard for both encoding and decoding models in functional neuroimaging for many years, and in certain cases may outperform more complex classifiers. However, even linear models are difficult to interpret and their performance may depend significantly on feature selection. In contrast, deep learning approaches allow multiple layers of processing to discover the underlying representations in the data across levels of abstraction, and have recently brought about numerous breakthroughs in computer vision and related fields. Here, we introduce the key theoretical principles related to deep learning, and review how their evolution over time has been inspired by neural computation in biology.

Presenter

Pamela Douglas, UCLA/UCF
Psychiatry, Modeling & Simulation
Los Angeles, CA
United States

Undefined labels? Try unsupervised approaches!

When labels are not clearly defined (e.g., in psychiatry, where boundaries between diagnostic categories can be unreliable), unsupervised methods can be useful for uncovering latent relationships between different data. We will first show how unsupervised methods can overcome certain limitations of supervised methods. Partial Least Squares (PLS) and Canonical Correlation Analysis (CCA) will be used to illustrate how to find multivariate relationships between neuroimaging and behavioral data in the form of latent components. These algorithms aim to maximize the covariance (PLS) / correlation (CCA) between two data matrices by deriving latent components that are optimally weighted linear combinations of the original variables. We will then show how to test the significance and generalizability of these latent components, using permutation testing and cross-validation. We will also demonstrate how to determine the robustness of the neuroimaging and behavioral variables that drive the latent relationships, using bootstrapping.

Presenter

Ninon Burgos, CNRS – Brain and Spine Institute (ICM) Paris
France

Why is this educational course not fully devoted to deep learning?

Despite the resounding success of deep learning in many fields, recent studies have suggested that for certain applications, classical machine learning algorithms might achieve comparable performance at significantly lower computational cost. In this educational course, we will start by discussing such examples, including earthquake prediction and our own work on functional connectivity prediction of behavior. Along the way, we will also discuss various issues that might arise from comparing different algorithms and I will also speculate about various classes of neuroimaging problems, in which DNNs might excel in relative to classical approaches.

Presenter

B.T. Thomas Yeo, National University of Singapore
Electrical & Computer Engineering
Singapore, South West
Singapore
