A practical guide on how to incorporate machine learning and/or deep learning in your data analysis

Date: 29th of September, 2018
Location: MNI, 3801 University St., Montreal, QC, Canada

Organizers:
Amir Shmuel, McConnell Brain Imaging Centre, MNI, McGill University, QC, Canada
Andrew Doyle, McGill Centre for Integrative Neuroscience, Montreal, QC, Canada

The aim of the workshop is to introduce the basic theory and software packages that can be used for machine learning and deep learning of neuroimaging data, with either no background in programing or just basic understanding of programing. The attendees will be able to (but do not have to) download data and scripts before the workshop, and apply the pipelines using their laptops. We will record all lectures and tutorials, making it possible for all attendees of the workshop, including those that do not pursue the analysis during the workshop itself, to follow up step by step once they are back in their laboratories.

Registration fee: $90 CAD.
This fee will cover the course material, including recordings of the lectures, coffee breaks and lunch. The registration for this workshop can be done separately from the registration to the educational workshop or the main conference.

Schedule

7:45-8:20
Registration, coffee

8:20-10:20
Pradeep Reddy Raamana
Baycrest Health Sciences, Toronto, ON, Canada

Title: Practical introduction to machine learning for neuroimaging:
classifiers, dimensionality reduction, cross-validation and *neuropredict*

An alternative title: How to apply machine learning to your data, even if you do *not* know how to program

Objectives:
1. Learn **what** machine learning is, and get a high-level overview of popular types of classification and dimensionality reduction methods. Learn (without any math) how support vector machines work.

2. Learn **how** to plan a predictive analysis study on your own data. What are the key steps of the workflow? What are the best practices, and which cross-validation scheme to choose? How to evaluate and report classification accuracy?
3. Learn **which** toolboxes to use when, with a practical categorization of toolboxes. This is followed by a detailed demo of and *neuropredict*, for automatic estimation of predictive power of different features or classifiers without needing to code at all.

Bio
Dr. Pradeep Reddy Raamana is a postdoctoral fellow in Prof. Stephen Strother's lab. His research interests include the development of 1) robust imaging biomarkers and algorithms for early detection and differential diagnosis of brain disorders, and 2) easy-to-use software to lower and remove the barriers for predictive-modelling and quality control for neuroimagers. He is also interested in characterizing the impact of different methodological choices at different stages of medical image processing (preprocessing and prediction). He blogs at crossinvalidation.com and tweets at @raamana_.

10:20-10:40
Coffee break

10:40-1:00 PM
Gael Varoquaux
Inria, Paris, France

Jake Vogel
MNI, Montreal, Canada

Title: *Nilearn*: An accessible tool for statistical learning with neuroimages

An alternative title: You do *not* need to be a hacker in order to incorporate machine learning in your data analysis

The tutorial will showcase how to use nilearn to run full analysis pipelines for resting state functional images, including supervised and unsupervised machine learning approaches.

A strong background in programming is *not* necessary.
Familiarity with neuroimaging is strongly recommended.

1:00 PM-1:45 PM
Lunch break

1:45-4:45 (including a 20 minute coffee break)

Robert Allan Brown and Robert Fratila
MNI, McGill University, Montreal, QC, Canada

SickKids, Toronto, OM, Canada

Title: Introduction to- and application of - deep learning in neuroimaging

This mini-workshop aims to provide an intuitive understanding of the mechanisms underlying deep learning, and back this up with a practical exercise in which you will create and train a system that can segment the corpus callosum from a central sagittal slice of an MRI. The workshop will challenge some of the preconceptions about deep learning, such as the need for thousands of training examples, powerful GPUs for processing, and a long time to train.

The theory portion of the workshop is aimed at a general audience. While some knowledge of Python is an asset for the practical part, if you haven't ever programmed a computer, join a group and soak up some coding experience along with your deep learning.

Notes and demos, as well as setup instructions if you'd like to follow along on your own computer, will be available for download for all attendees of the workshop. If you do want to follow along, please follow the installation and testing instructions well before the workshop.