

Prediction of Individual Differences from Neuroimaging Data

Guest Editors: Vince D. Calhoun, Stephan M. Lawrie, Janaina Mourao-Miranda, Klaas E. Stephan

Brain imaging has been extensively applied to many different areas of health and disease, with many remarkable successes that, collectively, have profoundly shaped our understanding of brain function. However, most previous work has focused on identifying differences between groups or striving for aggregate results across multiple subjects. In most cases mass-univariate analyses, such as the general linear model (GLM), are employed to make local inferences about the effect of interest, typically at the group level. This approach is limited though as it does not capture linked relationships across voxels and is also not directly able to classify individuals into diagnostic groups based upon neuroimaging findings. As a consequence, neuroimaging has so far contributed only relatively few tools that have had practical impact on real-world problems, such as clinical diagnosis and prognosis. This poor translational track record has triggered a recent interest in machine learning approaches that can be applied to predict clinically relevant outcomes, such as diagnosis or treatment response, from individual subject neuroimaging measures.

This special issue on individual subject prediction from neuroimaging data highlights recent work on individual subject predictions, with the goal of identifying the limitations so far and the challenges going forward. In addition to original papers, we include several review articles which discuss methodological issues related to individual subject prediction. In brief, more than 20 papers in this special issue cover a wide range of themes as described in the following.

The first six articles provide some exciting reviews of the field. To start us off, Arbabshirani et al. provide an extensive overview of over 200 studies which have attempted to perform prediction at the level of the single subject. They also discuss the overall approach and illuminate through examples where group difference studies and individual subject studies can diverge. Importantly, they also discuss possible biases (many of which exist in published work) as well as some of the promises and pitfalls of such work. Varoquaux et al. focus on 'brain decoders' and include guidelines and caveats especially with respect to cross-validation strategies. They also highlight conditions under which the widely used leave-one-out approach can lead to biased results. A complementary approach to "classical" machine learning is the use of generative models as "computational assays" that infer, from neuroimaging and behavioral data, disease mechanisms within individual subjects. Stephan et al. review the rationale of this approach, illustrating how computational assays of pathophysiology and cognition can accommodate heterogeneity in mental illness and leverage the strengths of machine learning by providing mechanistically interpretable features for single-subject predictions. One clinical area where heterogeneity poses difficult challenges involves individuals with lesions due to stroke. Price et al. provide a detailed look at ten challenges in this area and lay out some proposed solutions. They highlight the need for multivariate lesion analyses to provide a more comprehensive characterization of how the impact of brain damage depends on interactions between multiple factors – such as the location, extent and age of lesions – and how these interactions explain variability across individuals. Beyond statistical predictions of individual outcome, they highlight the need for explanatory models of stroke-inflicted neural network dysfunction that are constrained by individual anatomical and physiological data. Gifford et al. tackle the important question of how well neuroimaging measures can predict the onset of psychosis. They review previous key findings and lay out the challenges that lie ahead in translating such work to the clinic. The combination of multiple imaging modalities has shown great promise in further increasing the sensitivity of predicting clinically important variables. Sui et al. present a generalized framework for prediction of cognitive and symptoms scores from the joint information provided by multiple modalities (functional MRI, structural MRI, diffusion imaging). They also show how such work can provide more information about localized brain deficits than conventional approaches.

The next five articles provide a range of new methodologies that are applied to study patients with psychiatric disorders. Frangou et al. discuss the clinically important issue of distinguishing bipolar disorder (BD) and major depression (MD), which sometimes cannot be fully disambiguated using clinical observations alone. They applied Gaussian process classifiers to whole-brain fMRI data from a working memory task and showed that individual patients with BD and MS, respectively, can be differentiated with a relatively high accuracy. Brambilla et al. introduce an approach based on multiple kernel learning to account for potentially confounding variables directly within a classification framework. Using cortical thickness data acquired at multiple field strengths to dissociate first episode schizophrenia patients from controls, their results provide localized classification information, highlighting frontal and temporal areas in particular. Nieuwenhuys et al. targets the important question whether multi-center structural brain imaging data can help predicting illness course in first episode psychosis patients. While predicting sex from brain structure was possible with high accuracy, illustrating the potential of the approach, predicting illness course proved to be more challenging. The article discusses several strategies for improving predictive accuracy. Mwangi et al. ask whether distinct clinical phenotypes of individuals with bipolar disorder are also biologically distinct. Using neurocognitive measures and fractional anisotropy data from diffusion weighted imaging, the phenotypes could be differentiated from one another as well as from healthy controls. The analysis also identified which neurocognitive measurements and white matter regions played the largest role. This strategy highlights the potential for imaging data to help inform us about subgroups in conventional clinical classifications and their neurobiological implications. Finally, Steele et al. show that it is possible to discriminate different levels of a personality disorder, psychopathy, using structural MRI data. Applying support vector machine classification to voxel-based morphometry (VBM) data, they report between 70% and 80% accuracy for predicting different levels of psychopathy. This lays the groundwork for future improvements of diagnosis and prognosis in personality disorders.

The next three articles focus on some key clinical application questions that pose particular challenges, including the study of pain prediction, and minimally conscious or behaviorally non-responsive patients. Lindquist et al. introduce an interesting multi-voxel pattern analysis (MVPA) approach which utilizes group information from a previous study as a prior in order to improve the ability to classify single subjects in new data. They show that the model has the ability to predict pain on a trial-by-trial basis, and that the results can be used to provide additional benchmarks for population-level maps. The use of noninvasive brain imaging to provide a personalized window into the brain function in (minimally) conscious patients is considered by Noirhomme et al. They discuss factors which determine the success of training and validating models for predicting clinical outcomes disorders in consciousness. Naci et al. describe additional methodological challenges in using fMRI to predict conscious experiences in behaviorally non-responsive patients. They propose the use of naturalistic paradigms which engage attention naturally and do not require compliance. They present an interesting clinical case in which they decode conscious experiences from fMRI measurements in an individual who has been non-responsive for several years.

Going beyond clinical applications we have three articles focused on prediction of task, movie scenes, or personality traits with a variety of machine learning methods. The use of deep learning approaches to analyze brain imaging data is an area of growing interest. Lee et al. propose one of the first approaches using deep neural networks on whole brain 4D fMRI data to predict multiple tasks

performed by an individual. Using the information contained within the network's 'hidden layer', they demonstrate the potential of such flexible models for enhancing single-subject applications of fMRI. The use of naturalistic paradigms is explored in Güçlü et al., also in the context of deep learning algorithms. They show that when a deep network is trained for action recognition in movies, it can predict fMRI activity in dorsal stream areas when an individual watches natural movies. Moreover, they report a remarkable correspondence between representations in the layers of the model and in dorsal stream areas. Fernandes et al. show that prediction of individual negative affect in response to threat stimuli is possible with fMRI data. The authors used a multiple kernel learning (MKL) approach to combine information from different brain regions hierarchically in a whole brain predictive model. The MKL approach allows investigating the contribution of different brain regions to the predictive model. The results demonstrate the potential of machine learning to predict personality traits from fMRI data in the context of emotional stimuli.

The final four articles each focus on a unique challenge and proposed solution including a) disease heterogeneity, b) the impact of pipeline choices, c) the need for realistic network models, and d) the benefits of an extremely large data set. A key challenge for individual subject predictions is the biological heterogeneity of conventionally defined brain diseases. Varol et al. present a framework that can account for both group common as well as heterogeneous effects, enabling both classification and subtype identification. The utility of this framework is demonstrated using structural MRI data from the Alzheimer's Disease Neuroimaging Initiative (ADNI). The impact of preprocessing pipelines has received increasing attention in the context of the impact of motion and other confounds. Using both a simulation framework and real data, Vergara et al. present an approach in which multiple pipeline choices are optimized in the context of two different prediction examples, mild traumatic brain injury and substance abuse. They show that preprocessing choices should be considered an important element of a prediction study. The work by Jirsa et al. illustrates the potential of biophysical network models for simulating whole-brain activity and understanding its central determinants. They use an individualized virtual brain model, informed by a patient's anatomical data, to examine some key aspects of epilepsy. They show that such an approach enables the exploration of parameters, validation of the simulated model against real data, and opens up the possibility of the development of novel personalized approaches to both intervention and therapy. Finally, one of the major limitations in developing good predictive models is access to large data sets. The ENIGMA consortium links together numerous sites from around the world and has already performed meta-analyses on over 30,000 MRI scans. The article by Thompson et al. describes ENIGMA and discusses how the results can be used to inform individual subject analyses.

We hope that this collection of articles will be helpful for accelerating the further development of methods and applications for single-subject predictions based on neuroimaging.