NeuroImage 145 (2017) 274–287

Contents lists available at ScienceDirect

NeuroImage

journal homepage: www.elsevier.com/locate/ynimg

Full Length Article

Group-regularized individual prediction: theory and application to pain☆,☆☆

Martin A. Lindquist [a], Anjali Krishnan [b,f], Marina López-Solà [b], Marieke Jepma [b], Choong-Wan Woo [b], Leonie Koban [b], Mathieu Roy [c], Lauren Y. Atlas [d], Liane Schmidt [g,h], Luke J. Chang [b], Elizabeth A. Reynolds Losin [b,e], Hedwig Eisenbarth [b], Yoni K. Ashar [b], Elizabeth Delk [b], Tor D. Wager [b,*]

[a] Johns Hopkins University, USA
[b] University of Colorado Boulder, USA
[c] Concordia University, USA
[d] National Center for Complementary and Integrative Health, National Institutes of Health, USA
[e] University of Miami, USA
[f] Brooklyn College of the City University of New York, USA
[g] INSEAD, France
[h] Cognitive Neuroscience Laboratory, INSERM U960, Department of Cognitive Sciences, Ecole Normale Supérieure, Paris, France

ARTICLE INFO

Available online 17 November 2015

Keywords:
Machine learning
Statistical learning
Pain
MVPA
Empirical Bayes
Prediction
fMRI
Mega-analysis
Meta-analysis
Shrinkage

ABSTRACT

Multivariate pattern analysis (MVPA) has become an important tool for identifying brain representations of psychological processes and clinical outcomes using fMRI and related methods. Such methods can be used to predict or 'decode' psychological states in individual subjects. Single-subject MVPA approaches, however, are limited by the amount and quality of individual-subject data. In spite of higher spatial resolution, predictive accuracy from single-subject data often does not exceed what can be accomplished using coarser, group-level maps, because single-subject patterns are trained on limited amounts of often-noisy data. Here, we present a method that combines population-level priors, in the form of biomarker patterns developed on prior samples, with single-subject MVPA maps to improve single-subject prediction. Theoretical results and simulations motivate a weighting based on the relative variances of biomarker-based prediction—based on population-level predictive maps from prior groups—and individual-subject, cross-validated prediction. Empirical results predicting pain using brain activity on a trial-by-trial basis (single-trial prediction) across 6 studies (N = 180 participants) confirm the theoretical predictions. Regularization based on a population-level biomarker—in this case, the Neurologic Pain Signature (NPS)—improved single-subject prediction accuracy compared with idiographic maps based on the individuals' data alone. The regularization scheme that we propose, which we term group-regularized individual prediction (GRIP), can be applied broadly to within-person MVPA-based prediction. We also show how GRIP can be used to evaluate data quality and provide benchmarks for the appropriateness of population-level maps like the NPS for a given individual or study.

Introduction

Tremendous progress has been made in fMRI research over the past 10 years. Much of the benefit has resulted from multivariate pattern analysis (MVPA) based studies of mental representations, which have enhanced the ability to identify brain patterns that are predictive of behavioral and psychological outcomes (Chang et al., in press; Davis and Poldrack, 2013; Haxby et al., 2001, 2014; Kay et al., 2008; Poldrack et al., 2009; Wager et al., 2013; Woo et al., 2015). In standard brain mapping analyses, many regions of the brain might respond to a given task. However, for a pattern of brain activity to be considered useful as a *representation* of a psychological or behavioral state, it must be predictive of (i.e., be sensitive and specific to) that state.

Recent studies have identified provisional representations for many kinds of psychological states, including perception of low level visual features (Kamitani and Tong, 2005) and higher-order object properties (Haxby et al., 2001), knowledge of semantic categories (Huth et al., 2012; Mitchell et al., 2008), memory (Kuhl et al., 2011; Rissman et al., 2010; Xue et al., 2010), affective states such as pain (Brodersen et al.,

☆ Author note: We are grateful for the funding support of NIH, which supported this work under grants R01DA035484 (TDW), 2R01MH076136 (TDW), R01DA027794 (TDW), R01 EB016061 (MAL) and P41 EB015909 (MAL).

☆☆ Matlab code for all analyses is available at: http://www.columbia.edu/cu/psychology/tor/

* Corresponding author at: Department of Psychology and Neuroscience, University of Colorado Boulder, 345 UCB, Boulder, CO 80309, USA.
E-mail address: tor.wager@colorado.edu (T.D. Wager).

http://dx.doi.org/10.1016/j.neuroimage.2015.10.074
1053-8119/

2012; Cecchi et al., 2012; Marquand et al., 2010; Wager et al., 2013), and emotion (Baucom et al., 2012; Chang et al., in press; Kassam et al., 2013), and identification of individuals with clinical disorders (Arbabshirani et al., 2013; Craddock et al., 2009; Doehrmann et al., 2013; Fu et al., 2008; Siegle et al., 2006; Whelan et al., 2014). Once representations of specific percepts (e.g., objects) or experiences (e.g., emotion) are identified, studies can examine how these representations are shaped by contextual, psychological, and neurobiological processes—e.g., how object representations are maintained in working memory during a delay (Harrison and Tong, 2009), how items are recollected during memory recall and compete with other memories (Kuhl et al., 2011), or how pain representations are modified by cognitive reappraisal (Woo et al., 2015). Identifying patterns of fMRI activity that can serve as proxies for representations requires multivariate analyses that are predictive of outcomes in individual subjects. In this paper, we develop a method for improving such single subject, MVPA-based predictions.

Most single subject predictive analyses utilize only data from one participant in developing the predictive model (e.g., Horikawa et al., 2013). The theory behind this approach is that brain representations are idiographic (i.e., different individual subjects have different multivariate brain patterns that predict outcomes). For example, the pattern of fMRI activity within V1 that predicts line orientation, may be different for different individuals (Freeman et al., 2011; Kamitani and Tong, 2005), and only patterns at broader spatial scales may be conserved across individuals (Heeger and Ress, 2002; Norman et al., 2006). If brain topography is truly idiographic and varies dramatically across individuals, individualized training to derive the best predictive multivariate brain pattern is likely the optimal strategy. However, often, there is information at multiple spatial scales, including much information conserved across individuals (Chang et al., in press; Kassam et al., 2013; Poldrack et al., 2009; Rissman et al., 2010; Shinkareva et al., 2008; Wager et al., 2013; Woo et al., 2014). In addition, the quantity and quality of fMRI data are limited in single subject datasets, and often high-quality single subject prediction requires hours of scanning for each individual over multiple days (Gonzalez-Castillo et al., 2012; Nishimoto et al., 2011). Often, perhaps surprisingly, models that are trained to predict out-of-sample individuals perform as well or better than models trained on individual subject data (Chang et al., in press; Poldrack et al., 2009; Shinkareva et al., 2008; Wager et al., 2013) when the spatial topography of predictive information is shared across individuals. In such cases, using normative group maps based on other individuals may help to regularize single subject predictive patterns using information conserved across subjects, constraining the single subject solution in ways that improve prediction accuracy and prevent overfitting.

This paper develops a principled scheme for combining normative group maps based on previously defined predictive patterns (i.e., signatures) and single subject idiographic maps. In addition to improving prediction accuracy, this procedure regularizes individual subject maps towards prior expectations, therefore improving the quality of single subject predictive maps, and allowing for a principled updating of normative population-based maps as more data are accumulated. This weighting can be expressed in both frequentist and Bayesian frameworks, which are shown to be mathematically equivalent. Marquand et al (2014) addressed a similar problem, by recasting the decoding problem in a multi-task learning framework, allowing them to extract information from the data by sharing information between subjects. This was found to be extremely beneficial when only a small number of trials were available for each subject.

The method we develop here, which we term group-regularized individual prediction (GRIP), combines group and idiographic maps in proportion to their respective variances, in accordance with theory on empirical Bayes estimation. It can be applied prospectively to individual subjects' data to improve prediction accuracy and stabilize individual-subject predictive maps. Thus, one main use is in improving single subject MVPA-based prediction accuracy. In addition, it can be used to provide quality control estimates and benchmarks for a given individual or study paradigm. The quality of idiographic predictions can be used to benchmark data quality for individual persons, and the accuracy of prediction using a population-level map can provide benchmarks on the appropriateness of the map for a given population, sample, or study paradigm. Such cross-study metrics are valuable as fMRI data are increasingly used in multi-site and translational settings.

The GRIP method can be applied to any domain and is agnostic with respect to the training algorithm used. However, in this paper, we evaluate its utility in predicting pain intensity ratings. Pain is an interesting application domain for three reasons. First, it is associated with enormous cognitive, social, and economic costs (IOM, 2011), but its neurological bases are not yet well understood (Tracey, 2011). Developing brain models capable of predicting pain intensity and dissociating different types of neurological contributions to pain is a high-priority. Second, pain is currently assessed primarily by means of self-report, a behavioral measure of subjective experience that is compromised in many vulnerable populations (e.g., the very old or very young, persons with cognitive impairment, and those who are minimally conscious) and influenced by a number of complex sociocultural factors. Brain-based predictive models could complement self-report by providing measures of neurophysiological systems that contribute to pain, and ultimately identify sub-types of pain and sub-types of patients based directly on brain information. And third, population-level maps predictive of pain intensity are available (Wager et al., 2013), providing priors to use in regularizing individual-subject predictions. Several groups have published innovative work on single subject prediction (Brodersen et al., 2012; Cecchi et al., 2012; Marquand et al., 2010). Complementing these approaches, we have developed a normative population-based pattern that classifies stimuli differing moderately in pain intensity with over 90% accuracy, across multiple sites and scanners and in new, out-of-sample individuals (Wager et al., 2013). Here, we combine information from this population-normed signature pattern—called the Neurologic Pain Signature (NPS)—with idiographic MVPA maps to improve the accuracy of predicting pain intensity from brain activity.

We begin by developing the statistical theory underlying empirical Bayes regularization and the GRIP model. We then present brief theoretical simulations that characterize the conditions under which weighting towards individuals versus group maps is optimal. Then we apply the method to combine data from six studies of experimental thermal pain (N = 180), comparing the accuracy of (a) cross-validated idiographic predictive maps, (b) a population-level map, the NPS, and (c) the GRIP combination of the NPS prior and idiographic maps (see Fig. 1 for an overview). Predictions are made about single trials, i.e., individual periods of thermal stimulation lasting 1.85 to 15 s, using time series-appropriate cross-validation methods. The results show that the GRIP estimator outperforms both the population-level NPS map and the idiographic, single-subject prediction map.

Method

Theory

Suppose we have a set of observations from m trials of a certain stimulus applied to a single subject, which we denote $(\boldsymbol{x}_j, y_j)$ for $j = 1, \ldots m$. Here, $\boldsymbol{x}_j$ is a vector of features of length V, and y_j is a scalar outcome variable. In our example, we assume that each trial consists of a thermal stimulus. Thus, $\boldsymbol{x}_j$ is a summary of the brain response, and y_j is the reported pain corresponding to that trial.

Now, suppose we seek to use these observations to create a predictive model from which we can estimate pain report from brain activation for the subject in question. Using standard machine learning techniques (the approach is agnostic to the specific type of technique, though we assume that it is linear in the continuation) we can find a set of idiographic brain weights $\hat{\boldsymbol{w}}_I$ that can be used to predict the