

# Nonlinear functional mapping of the human brain

Allgaier, N.A.<sup>a,b,c,d,\*</sup>, Banaschewski, T.<sup>e</sup>, Barker, G.J.<sup>f</sup>, Bokde, A.L.W.<sup>g</sup>, Bongard, J.C.<sup>b,c</sup>, Bromberg, U.<sup>h</sup>, Büchel, C.<sup>h</sup>, Cattrell, A.<sup>i</sup>, Conrod, P.<sup>j,k</sup>, Danforth, C.M.<sup>a,b</sup>, Desrivieres, S.<sup>i</sup>, Dodds, P.S.<sup>a,b</sup>, Flor, H.<sup>l</sup>, Frouin, V.<sup>m</sup>, Gallinat, J.<sup>n</sup>, Gowland, P.<sup>o</sup>, Heinz, A.<sup>p</sup>, Ittermann, B.<sup>q</sup>, Mackey, S.<sup>d</sup>, Martinot, J-L.<sup>s</sup>, Murphy, K.<sup>r</sup>, Nees, F.<sup>l</sup>, Papadopoulos-Orfanos, D.<sup>m</sup>, Poustka, L.<sup>e,t</sup>, Smolka, M.N.<sup>u</sup>, Walter, H.<sup>p</sup>, Whelan, R.<sup>v</sup>, Schumann, G.<sup>i</sup>, Garavan, H.<sup>d</sup>, IMAGEN Consortium

<sup>a</sup>Department of Mathematics and Statistics, University of Vermont, Burlington, Vermont, USA

<sup>b</sup>Vermont Complex Systems Center, Vermont Advanced Computing Core, Burlington, Vermont, USA

<sup>c</sup>Department of Computer Science, University of Vermont, Burlington, Vermont, USA

<sup>d</sup>Departments of Psychiatry and Psychology, University of Vermont, 05405 Burlington, Vermont, USA

<sup>e</sup>Department of Child and Adolescent Psychiatry and Psychotherapy, Central Institute of Mental Health, Medical Faculty Mannheim, Heidelberg University, Square J5, 68159 Mannheim, Germany

<sup>f</sup>Centre for Neuroimaging Sciences, Institute of Psychiatry, Psychology & Neuroscience, Kings College London, United Kingdom

<sup>g</sup>Discipline of Psychiatry, School of Medicine and Trinity College Institute of Neurosciences, Trinity College Dublin, Ireland

<sup>h</sup>University Medical Centre Hamburg-Eppendorf, House W34, 3.OG, Martinistr. 52, 20246, Hamburg, Germany

<sup>i</sup>Medical Research Council - Social, Genetic and Developmental Psychiatry Centre, Institute of Psychiatry, Psychology & Neuroscience, Kings College London, United Kingdom

<sup>j</sup>Department of Psychiatry, Université de Montreal, CHU Ste Justine Hospital, Canada

<sup>k</sup>Department of Psychological Medicine and Psychiatry, Institute of Psychiatry, Psychology & Neuroscience, King's College London

<sup>l</sup>Department of Cognitive and Clinical Neuroscience, Central Institute of Mental Health, Medical Faculty Mannheim, Heidelberg University, Square J5, Mannheim, Germany

<sup>m</sup>Neurospin, Commissariat à l'Energie Atomique, CEA-Saclay Center, Paris, France

<sup>n</sup>Department of Psychiatry and Psychotherapy, University Medical Center Hamburg-Eppendorf (UKE), Martinistrasse 52, 20246 Hamburg

<sup>o</sup>Sir Peter Mansfield Imaging Centre School of Physics & Astronomy, University of Nottingham, University Park, United Kingdom

<sup>p</sup>Department of Psychiatry and Psychotherapy, Campus Charité Mitte, Charité, Universitätsmedizin Berlin, Germany

<sup>q</sup>Physikalisch-Technische Bundesanstalt (PTB), Braunschweig and Berlin, Germany

<sup>r</sup>Cardiff University Brain Research Imaging Centre (CUBRIC), School of Psychology, Cardiff University, Cardiff, United Kingdom

<sup>s</sup>Institut National de la Santé et de la Recherche Médicale, INSERM Unit 1000 Neuroimaging & Psychiatry, University Paris Sud, University Paris Descartes - Sorbonne Paris Cité; and Maison de Solenn, Paris, France

<sup>t</sup>Department of Child and Adolescent Psychiatry and Psychotherapy, Medical University of Vienna, Austria

<sup>u</sup>Department of Psychiatry and Neuroimaging Center, Technische Universität Dresden, Dresden, Germany

<sup>v</sup>Department of Psychology, University College Dublin

## Abstract

The field of neuroimaging has truly become data rich, and novel analytical methods capable of gleaning meaningful information from large stores of imaging data are in high demand. Those methods that might also be applicable on the level of individual subjects, and thus potentially useful clinically, are of special interest. In the present study, we introduce just such a method, called *nonlinear functional mapping* (NFM), and demonstrate its application in the analysis of resting state fMRI (functional Magnetic Resonance Imaging) from a 242-subject subset of the IMAGEN project, a European study of adolescents that includes longitudinal phenotypic, behavioral, genetic, and neuroimaging data. NFM employs a computational technique inspired by biological evolution to discover and mathematically characterize interactions among ROI (regions of interest), without making linear or univariate assumptions. We show that statistics of the resulting interaction relationships comport with recent independent work, constituting a preliminary cross-validation. Furthermore, nonlinear terms are ubiquitous in the models generated by NFM, suggesting that some of the interactions characterized here are not discoverable by standard linear methods of analysis. We discuss one such nonlinear interaction in the context of a direct comparison with a procedure involving pairwise correlation, designed to be an analogous linear version of functional mapping. We find another such interaction that suggests a novel distinction in brain function between drinking and non-drinking adolescents: a tighter coupling of ROI associated with emotion, reward, and interoceptive processes such as thirst, among drinkers. Finally, we outline many improvements and extensions of the methodology to reduce computational expense, complement other analytical tools like graph-theoretic analysis, and allow for voxel level NFM to eliminate the necessity of ROI selection.

## Keywords:

resting state fMRI, modeling, nonlinear, machine learning, genetic programming, symbolic regression

## 1. Introduction

Many advances in our understanding of brain function have been achieved through analysis of fMRI data. Though the BOLD (blood oxygen level dependent) signal obtained from fMRI is a proxy, physiological confounds such as breathing and heart rate are separable from neuronal-induced signal, as demonstrated in Birn et al. (2009). Inter-subject differences in vascular reactivity can be modeled as shown in Murphy et al. (2011), and BOLD has been directly shown to provide a reliable measure of neuronal activity in specific circumstances, as in Mukamel et al. (2005). The many years of successful research before and since support that assessment. Accomplishments include localization of regions responsible for particular tasks, such as episodic memory in Nolde et al. (1998) and human face recognition in Kanwisher et al. (1999), assessment of the risk of postoperative motor defect in patients with tumors in Mueller et al. (1996), analysis of the effects of acupuncture in Hui et al. (2000), and recently, identification of neural markers for both current *and future* alcohol use among adolescents in Whelan et al. (2012) and Whelan et al. (2014).

These examples, and indeed the majority of fMRI studies, make use of the GLM (general linear model) to determine neural correlates for various tasks and stimulus responses. Though typical analyses have been performed at the group level with a univariate approach, other recent work reported in Rio et al. (2013) has extended the capabilities of the GLM to analyze multivariate signal in the Fourier domain to reduce confounds from time-correlated noise, thus improving the suitability of the GLM for subject level analysis. Despite these advances, however, the GLM can only *confirm* hypothesized nonlinear models of function, not *discover* them.

Group-level inferences from fMRI have also been performed using linear ICA (independent component analysis), as described in Calhoun et al. (2001). Though ICA and the GLM can be used in conjunction, for example in Liu et al. (2010) to investigate the neural effects of stimulation of a particular acupoint, ICA is particularly useful in circumstances that preclude the use of the GLM, such as the analysis of resting-state data, for which there is no task or stimulus regressor. Covarying networks have been suggested by ICA of resting-state fMRI in Smith et al. (2009), and functional, hierarchical classification of these networks has been automated through HCA (hierarchical cluster analysis) of aggregated experimental metadata in Laird et al. (2011). However, it was determined early on, for example in McKeown and Sejnowski (1998), that nonlinear interactions within the brain need to be addressed in order to properly determine functional architecture.

Although ICA algorithms that employ nonlinear mixing functions exist, severe restrictions on those functions

are required to avoid non-uniqueness of solutions, as explained in Hyvärinen and Pajunen (1999). Due to this failing, other methodologies have been employed in the attempt to account for nonlinearity. Examples include various forms of nonlinear regression, as in Kruggel et al. (2000), and dynamic causal modelling, as described in Friston et al. (2003). In each of these, a particular nonlinear form must be posited *a priori*, and thus the capability to *discover* previously unknown nonlinear interactions within the brain is diminished. As a result, a fuller picture of the nature of intra- and inter-network functional connectivity within the brain is missing from the literature.

Here we introduce a methodology designed to accomplish such a mathematical characterization, provide insight at the group, subject, and ROI levels, and to avoid linear and univariate assumptions. With some modification, analysis of higher dimensional data is likely attainable, allowing for eventual application at the voxel scale and eliminating the necessity of ROI selection. After standard preprocessing (slice-timing and motion correction, normalization, smoothing, etc.), our procedure consists of ROI selection, inter-ROI *symbolic regression* (a model-free form of nonlinear regression), accomplished by an evolutionary algorithm called *genetic programming* (GP; a form of stochastic optimization), and statistical analysis of the resulting models. We demonstrate our technique on a 242-subject collection of resting-state data from the IMAGEN project, though analysis of task or stimulus experiments can be accomplished with little or no modification. The IMAGEN project is described in detail in Schumann et al. (2010).

We organize the paper as follows. In Section 2, we discuss the data and selection of ROI, provide some background on GP, and describe the procedural details of NFM by symbolic regression. In Section 3, we report results of applying the technique to the IMAGEN data, including statistical and hierarchical visualizations, comparison with previous results for cross-validation, effects of nonlinearity, and an example of group-level variation. We discuss the results and potential applications of the technique in Section 4, and conclude the paper in Section 5.

## 2. Materials and methods

In this section, we first briefly describe the source of the data for our study, and then provide the details of ROI selection that allow for comparison with recent work. We then provide some background on the GP algorithm in general and the specific implementation employed here, along with the method by which it is applied to BOLD signal time series extracted from the selected ROI. Finally, we describe the statistical technique used to interpret the roughly quarter of a million mathematical models that result from the application of GP to all 52 ROI time series extracted from each of the 242 subjects.

---

\*Corresponding author.

Email address: [nicholas.allgaier@uvm.edu](mailto:nicholas.allgaier@uvm.edu) (Allgaier, N.A.)

URL: [www.imagen-europe.com](http://www.imagen-europe.com) (IMAGEN Consortium)

## 2.1. Data

The data investigated here are a subset of the fMRI scans from the IMAGEN study, a European research project with the goal of better understanding teenage psychological and neurobiological development. The project is longitudinal, and utilizes several forms of high and low-tech experimental protocols including self-report questionnaires, behavioral assessment, interviews, neuroimaging, and blood sampling for genetic analyses. Each of the 2000 participating adolescents was 14 when entering the study, which itself commenced in late 2007, and data collection continues today.

More specifically, the data for the present study are 6-minute resting-state fMRI time series of 242 of the adolescent subjects who were asked to keep their eyes open while in the scanner, but were presented with no other task or stimulus. To allow for comparison with previous work, locations of the ROI were chosen based on results from Laird et al. (2011), in which statistical analysis across thousands of previous imaging studies (both stimulus/task-based and resting-state) was used to identify networks of brain regions that tend to activate together, termed ICN (intrinsic connectivity networks). The ICN were determined by ICA, from which  $z$ -statistic maps were derived. To select ROI for this study, a  $z$ -statistic threshold was set for each ICN to determine the number of regions in the network, and ROI were defined as rough spheres with radii of 3 voxels (9mm) and centered at the location of peak  $z$ -statistic in each region.




Figure 1: ROI Selection. (Red) ROI from within the default mode network, with radii of 3 voxels and centers corresponding to the highest  $z$ -statistics (green) in each region as determined in Laird et al. (2011).

We provide a cut-out illustrating ROI selection for the default mode network (ICN 13) in Figure 1, and Figure 2 contains axial cross sections showing many of the ROI derived from the 18 non-artifactual ICN in Laird et al.

(2011). In Appendix A, Table A.1 we list all 52 ROI by number, give their anatomical names, indicate the ICN from within which they were defined, and provide visual representations of their locations within the brain.

Subsequent to ROI definition, a gray matter mask was applied to assure that only appropriate voxels were contained within each ROI. In some cases this resulted in a considerable reduction of ROI voxels, but the majority maintained the full complement of about 100 voxels. For each of the 242 subjects, time series were extracted from each of the 52 resulting ROI by averaging the BOLD signal over all voxels within the ROI. These time series then form the input to the GP algorithm.




Figure 2: Visualization of ROI. Axial cross sections showing many of the ROI derived from the ICN in Laird et al. (2011).

## 2.2. Genetic programming

GP is a biologically inspired, population-based machine learning algorithm. It is most commonly employed for symbolic regression: the algorithm searches for models explaining some quantity of interest (e.g., average BOLD signal from an ROI in the brain) as a function of some other possibly related observable quantities, statistics, or summary data (e.g., BOLD signals from other ROI). The algorithm proceeds by evolving the functional forms of a population of potential models, which are initially constructed at random from user-specified mathematical building blocks (available variables, arithmetic functions, parameter constants, etc.). In brief, the models that better explain the data produce more offspring, leading to a gradual reduction of error within the population. We show a representative set of models produced by this approach in Figure 3(a). A key advantage of the technique is that no assumption (e.g., linearity) is imposed on the form of solutions, other than the choice of building blocks from which

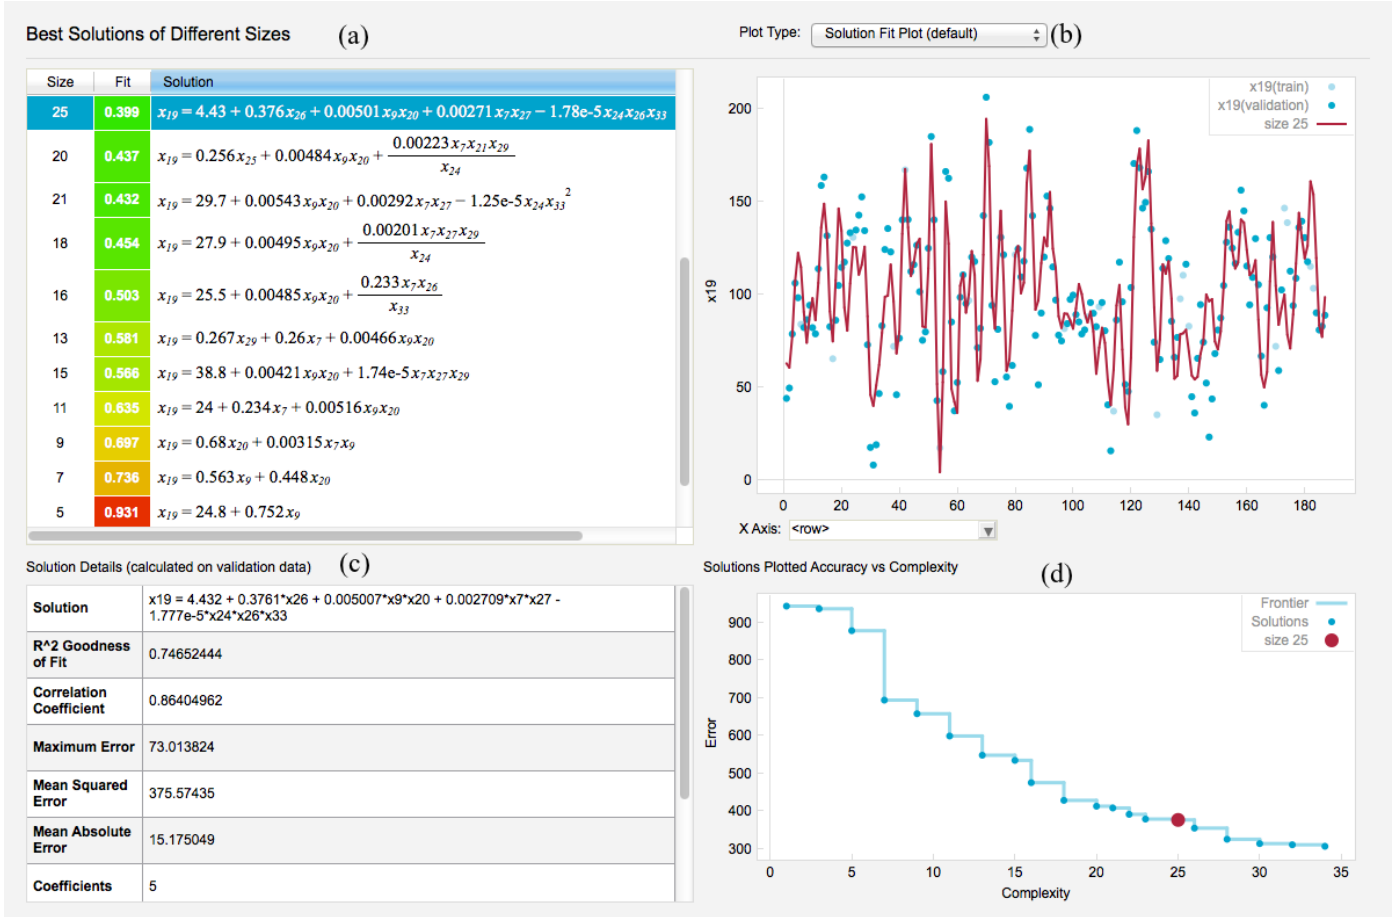


Figure 3: Screenshot of the GP package Eureqa during a search for models of the activity in ROI 19 in a single subject, as a function of activity in the other 51 regions. (a) The current set of models along the Pareto front of accuracy vs. parsimony, shown in (d) where each point represents a model and the red point represents the highlighted model. (b) Data from ROI 19 (points) over the 6-minute time series for this subject ( $x$ -axis in scans, 2 seconds each). The highlighted model is shown in red, and statistics for this model’s fit appear in (c).

they can be made (we use arithmetic operations in the present work).

Typically, some measure of error (e.g., root mean square error) constitutes a model’s explanatory fitness, and some measure of its size (e.g., number of operators, constants, and variables in the equation) represents its parsimony. The next generation of potential models is obtained by mutation (e.g., a single change of variable or operator) and recombination (i.e., swapping of function components between models) of the current set of *non-dominated* solutions: those models for which no simpler model in the population has less error. This set of non-dominated models is said to approach the ideal *Pareto front* of fitness versus parsimony as the population evolves. An important aspect of GP is that the result of a single search is this entire set of potential models, providing a trove of information for statistical analysis. Figure 3 is a screenshot of the off-the-shelf GP package Eureqa from Schmidt and Lipson (2009) performing a search (Eureqa version 0.97 Beta was used to generate the results reported in this study).

To apply GP to the fMRI data, for each of the 242 subjects we extract a single BOLD signal time series from each

of the 52 selected ROI by averaging over the voxels within that ROI. Then the GP algorithm is run 52 times, one for each ROI, using all other ROI as potential explanatory variables. Note that the algorithm has no knowledge of the hypothesized networks from which these regions were chosen.

We describe the computational expense of the algorithm in terms of core-hours, i.e., the number of hours required for a single processor core to perform the necessary computation. Specifically, twelve core-hours of search were performed for each region, amounting to 624 core-hours per subject, and over 17 total core-years of computation were required for the population of 242 subjects. This yielded roughly 12 thousand Pareto fronts comprised of a quarter million models for statistical analysis.

The results of this analysis characterize the entire population of 242 subjects. Alternatively, results can be aggregated over phenotypic groups to produce group-level characterizations, or many GP searches can be run for a single individual to produce a subject-level characterization. We report results of population- and group-level analyses in Section 3, and discuss an example subject-level






Figure 4: Functional interaction map. (a) Interaction map across all 242 subjects, and (b) map of RSD (relative standard deviation) of the interaction rates over 100 subsamples with 100 randomly selected subjects each. Solid outlines indicate ICN and dashed outlines indicate functional groupings of ICN from Laird et al. (2011).

analysis in Appendix B.

### 2.3. Analysis

The output of the GP algorithm poses a challenge for interpretation. Here we present a coarse statistical analysis of this rich mathematical characterization. For each ROI, we count the number of models for that ROI, across the Pareto fronts for all 242 subjects, that have a particular (other) region on the right-hand side of the equation. We compute this count for each of the other 51 regions.

For example, consider the GP search for models of ROI 19 within a single subject illustrated in Figure 3. Upon completion, all 20 of the models along the Pareto front for this subject had at least one term containing ROI 9, and 17 models had terms containing ROI 20. In the subject pool as a whole, the total counts are 2990 and 1984, respectively. Specifically, of the roughly 5000 models for ROI 19 across all subjects, about 60% have terms containing ROI 9, and about 40% have terms containing ROI 20. Note that these frequencies are not properly normalized, because most models contain several ROI. Thus we normalize by the sum of the counts for all ROI. In the case of ROI 19, this sum is 22016.

The result is a vector for each ROI that describes, in a statistical sense, its relative dependence on each of the other regions. We interpret this vector as a distribution of likely interaction, and define the computed values to be relative *interaction rates* (IR). Note that both linear and nonlinear interactions, as well as weakly and strongly weighted basis functions, are counted equally. We form an interaction map by stacking these IR row vectors

to visualize interaction across all 52 ROI, shown in Figure 4(a). The value in row 19 column 9, for example, is  $2990/22016 \approx 0.136$ , depicted as a yellow square.

Note that the IR map is not symmetric by construction (though it appears nearly so), and indeed the value in row 9, column 19 is  $0.148 \neq 0.136$ . We interpret a row of the IR map as a distribution of relative *dependence* of the corresponding ROI on each of the other regions. We interpret a column, on the other hand, as a measure of the *influence* of the corresponding ROI on each of the other regions. By averaging the IR map with its transpose, we produce a symmetric, *overall* IR map (not shown) that can be used in hierarchical analysis. We examine the interaction map, and provide results of hierarchical analysis, in the next section.

## 3. Results

Figure 4(a) shows the interaction map generated by the normalized frequency analysis of the NFM procedure, summarizing ROI interaction across all 242 subjects. To test the robustness of the computed interaction map, we form 100 random subsamples (with replacement) from the pool of 242 subjects, each with 100 subjects. For each sample, we perform the same counting procedure to produce the interaction map corresponding to that sample. A heat map of relative standard deviation (RSD) of IR over the 100 subsamples is shown in Figure 4(b).

The strong block-diagonal structure of the interaction map corresponds directly to the grouping of ROI into ICN. For example, regions 39-42, which form a partial block in the figure, are the four ROI that make up the default

mode network (ICN 13) in Laird et al. (2011). Robustness (across subjects) of intra-network interaction is supported by the matching block-diagonal structure of low subsampling RSD (mean intra-network RSD < 20%), for all but ICN 2 (ROI 3-4) and ICN 5 (ROI 10-18). In addition to the strong primary block-diagonal structure, there is a secondary structure of lighter blocks that group ICN together. For example, regions 32-38 are composed of the strong blocks 32-33, 34-35, and 36-38 (corresponding to ICN 10, 11 and 12 respectively). There is a lighter block structure that suggests interaction among these three ICN which are, in fact, *together* responsible for visual processing. The secondary structure visible for regions 19-31 is comprised of ICN 6-8, which perform motor and visuospatial tasks. Each of these examples shows a matching secondary structure of moderate subsampling RSD (mean inter-network RSD < 30%), indicating fairly robust inter-network interaction as well.

### 3.1. Hierarchical analysis

To further illustrate and clarify the hierarchical organization suggested by the interaction map, we generate the dendrogram in the top of Figure 5 by HCA (hierarchical cluster analysis, implemented in MATLAB with the nearest distance algorithm), using the reciprocal of the overall IR between each pair of ROI as the distance between them. For example, ROI 1 and 2 have an approximate overall IR of 0.2, and thus the distance between them is 5. We emphasize that the organization of ROI into networks, and clustering of those networks into functional groups described in Laird et al. (2011), are both captured by NFM. Some examples:

- The red group forms the visual cluster. ROI 32 and 33, the lateral occipital cortices, form one network (ICN 10), while ROI 34-35, the occipital poles, and ROI 36-38, the lingual gyrus, right cuneus and right fusiform gyrus, respectively, form two other networks (ICN 11 and 12) from within the visual cluster.
- Regions 39-42 (the orange group) form ICN 13, the default mode network, and interact with ROI 4, the ventromedial prefrontal cortex, from ICN 2.
- The green group to the far left includes all but one of the ROI from the motor and visuospatial complex. Interaction of this complex with the middle cingulate cortex (mCC, ROI 9) and the network composed of ROI 46 and 47, thought to be responsible for multiple cognitive processes such as attention and inhibition, is indicated as well, suggesting that this interaction was common among many of the subjects.
- ICN 1 (ROI 1,2), 3 (ROI 5,6), the first two regions from ICN 4 (ROI 7-9), ICN 14 (ROI 43-45), 16 (ROI 48,49), and 17 (ROI 50,51) are also indicated.

- Many of the regions from ICN 5 (ROI 10-18) interact with ICN 1 (ROI 1-2), and also form a loose interaction group with ICN 14 (ROI 43-45), the cerebellum, the most robust connection of which appears to be between ROI 16 and 43.

The robustness of each of the interactions discussed in this list is supported by low interaction rate subsampling RSD, shown in Figure 4(b).

### 3.2. Impact of nonlinearity

In this section we demonstrate that the NFM procedure both captures the hierarchical structure of ROI interaction indicated by linear analyses, *and* reveals nonlinear interactions not discoverable by such methods. To accomplish this, we compare the population-level hierarchy generated by NFM with the results of an analogous linear procedure involving pairwise correlation analysis. Furthermore, we validate nonlinear relationships suggested by NFM in a stepwise multiple regression, and an elastic net regularized regression, the results of which we describe at the end of this section.

#### 3.2.1. Comparison with correlation analysis

For each of the 242 subjects, we compute the correlation matrix for the 52 ROI time series. Squaring the elements of the correlation matrix and normalizing each row (after setting the diagonal to zero) provides the relative explained variance (relative  $R^2$ ) of the ROI corresponding to that row by each of the other 51 ROI. The average of the 242 normalized subject matrices is interpreted as the linear version of the population-level IR map generated by NFM. As with IR, the reciprocal of relative explained variance can be considered a distance between ROI (higher relative  $R^2$  means closer). The resulting hierarchy generated by HCA is shown in the bottom of Figure 5.

As expected, much of the large-scale structure revealed by NFM is also indicated by the linear correlation analysis. The similarity of the generated hierarchies supports the validity of the models discovered by GP (i.e., the algorithm is not excessively overfitting the data), and the subtle differences between them suggest potentially interesting interactions that are missed if linearity is assumed. In the following we investigate one of these differences.

Interaction of the mCC (ROI 9) with the motor visuospatial complex is evident in both hierarchies. However, in the linear analysis it appears more closely connected with its own ICN (ROI 7,8, the bilateral anterior insula), and only with the posterior dorsomedial prefrontal cortex (dmPFC, ROI 19) from ICN 6. In contrast, NFM reveals that activity in the mCC is related to more components of the motor system. The nonlinear models generated by GP show a strong connection between the mCC, posterior dmPFC, and the paracentral lobule (PL) of the primary motor cortex (ROI 29 from ICN 8) shown in red, green, and blue, respectively, in Figure 6. Specifically, about 20%





Figure 5: Hierarchical cluster analysis (HCA) of interaction among ROI, generated with NFM (top) and correlation analysis (bottom).

of all of the models generated for the activity in the posterior dmPFC, across all subjects and levels of complexity, contain *both* the mCC and the PL as explanatory variables.

For many of these models, the mCC and PL only show up as linear terms, so it is reasonable to wonder why the correlation analysis did not pick up this interaction. The vast majority of models containing mCC and PL in only linear terms *also* contain nonlinear terms in other ROI. It is the posterior dmPFC *along with* these nonlinear terms that is correlated with the mCC and PL. Thus the interaction is hidden from linear analyses. Furthermore, many of the models *do* contain nonlinear terms involving the mCC and PL. In fact, the product of the activity in these two regions shows up in 78 models for the posterior dmPFC across 21 different subjects, and it is always additive. This term is involved in models across the spectrum of complexity, including instances where it is the *only* term.




Figure 6: NFM reveals a nonlinear interaction among these three ROI: mCC (red), posterior dmPFC (green), and PL in the primary motor cortex (blue).

### 3.2.2. Validation of nonlinear terms

To validate first order nonlinearity (pairwise product and quotient terms, as well as reciprocals) suggested by NFM, we first randomly assign 100 subjects to a training group, and 100 different subjects to a testing group. NFM results are aggregated over the training group to produce an IR map and hierarchy (not shown) summarizing ROI interaction within the training group as a whole. The roughly 2000 specific models generated by NFM for each region (approximately 20 models per training subject) are then used to inform the modeling of ROI activity in that region within the testing group, by stepwise regression.

For each ROI and each testing subject, we first perform a standard stepwise linear regression using the other

51 ROI as regressors. We then perform a stepwise regression including all first order nonlinear terms suggested by NFM over the training group in addition to the 51 linear regressors. Statistics of the linear and nonlinear models are compared to determine the effect of including these first order terms. To illustrate, we describe results of the validation procedure for the posterior dmPFC (ROI 19) here.

We show a histogram of increase in the percentage of explained variance for the nonlinear versus linear regression models for the posterior dmPFC in Figure 7. The inclusion of first order nonlinear terms suggested by NFM over the training group increases the percentage of explained variance for every test subject, with a mean increase of 12.5% and maximum increase of 42%. The nonlinear models contain more terms (mean 46, compared with mean 19 for linear models), so a potential concern is that the increase in  $R^2$  might simply be a result of the additional degrees of freedom. However, for each test subject the nonlinear model  $F$ -statistic is also greater than that of the linear model (mean increase of 83.5, maximum increase of 1000), and comparisons of adjusted  $R^2$ , which account for differences in degrees of freedom, show only slightly smaller increases for all test subjects. This suggests that the increase in explained variance is due to explanatory power of the nonlinear terms, and not simply the additional degrees of freedom in the nonlinear models.




Figure 7: Histogram of increase in explained variance. The inclusion of first order nonlinear terms suggested by NFM over the training group, in a stepwise regression analysis for the posterior dmPFC in the testing group, increases the percentage of explained variance for every test subject, with a mean increase of 12.5% and maximum increase of 42%.

To further support the validity of the nonlinear terms suggested by NFM, we apply a similar testing approach using a machine learning algorithm called elastic net regularized regression. In contrast to stepwise regression, regularization allows for the inclusion of highly correlated explanatory variables, while simultaneously discounting regressors with very small coefficients. Regularized models can have more explanatory power or fewer terms (or both), with respect to those from stepwise regression. We

see each of these scenarios in the present case. Using the same training and testing groups, and modeling the same ROI (the posterior dmPFC), elastic net regularization produces linear models with an average of 45 terms that explain roughly the same amount of variance (on average) as the nonlinear models generated with stepwise regression, improving upon the explanatory power of their stepwise counterparts. However, the regularized *nonlinear* models provide that same explanatory power with a mean of only 26 terms. Furthermore, the regularized nonlinear model is preferable to the regularized linear model, as determined by the Akaike information criterion, *for every single test subject*.

### 3.3. Group-level variation

Variation among individuals (illustrated in Appendix B) suggests that statistics of interaction rates among ROI may differ between phenotypic groups. The hierarchical organization of ROI induced by IR might illuminate, in such cases, variation in functional dynamics associated with demographic, behavioral, or genetic characteristics. An example illustrating this potential is provided by the contrast between drinking (D) and non-drinking (ND) adolescents from the IMAGEN dataset. In Figure 8, we show hierarchies for the top and bottom 100 subjects in terms of lifetime drinking score, determined by self-report questionnaire, corresponding to those who have had 2 or more lifetime drinks, and those who have had 1 or fewer, respectively. The two hierarchies are similar to one another (and comparable to the population level hierarchy), but subtle differences between them suggest group-differentiating factors.

- The ROI pair 3,18, the subgenual anterior cingulate cortex (ACC) and fornix body, respectively, are coupled in both the D and ND groups. However, their arrangement in the hierarchies is different, as we'll describe in a moment, resulting from the following two distinguishing interaction rates.
- For the ND group, there is a 22% lower IR between ROI 6, the left globus pallidus, and the fornix body, ROI 18. We note that this reduced interaction is completely missed by pairwise correlation analysis, (which indicates a slightly reduced interaction among *drinkers*, see Appendix C), and thus appears to be an entirely nonlinear effect.
- In contrast, there is a 33% higher intra-network IR within ICN 2, comprised of ROI 3-4, the subgenual ACC and the ventromedial prefrontal cortex (vmPFC), respectively, among non-drinkers. Though this difference is also indicated by correlation analysis, only about half of the effect is captured (a 16% elevation).
- These two differences in interaction cooperate to shuffle the hierarchical arrangement of ROI in the D versus ND group. The subgenual ACC and fornix

body are most closely associated with the default mode network in non-drinkers, through the vmPFC. Among drinkers, in contrast, they are grouped directly with the bilateral globus pallidus of ICN 3 (ROI 5-6). In other words, in drinkers there is a tighter coupling among ROI most strongly linked to reward and thirst tasks as reported in Laird et al. (2011). The relevant ROI are shown in Figure 9.




Figure 9: Interaction between the subgenual ACC (top red) and vmPFC (bottom red) is lower among drinkers, who also show elevated interaction between the left globus pallidus (green) and fornix body (blue), an apparently nonlinear effect.

- The largest single difference between the D and ND groups is a 74% elevated IR between the right angular gyrus (ROI 41 in the default mode network) and ROI 11, the posterior cingulate cortex, among drinkers. These ROI are shown in red and green, respectively, in Figure 10. About half of this effect is captured by correlation analysis.

## 4. Discussion

The large extent of ICN reproduction, and their hierarchical organization into functional groups using an entirely different approach than that described in Laird et al. (2011), provides strong evidence for the analytical potential of NFM. Furthermore, the technique reveals nonlinear interactions that are not discoverable with standard linear techniques, or without prior hypotheses. Such relationships could provide a new window into brain function, and this highlights the potential of the methodology as a *hypothesis generator*. Of course proper care must be taken (with regard to independence of observations, etc.) in the ensuing investigations of such data-driven hypotheses. Nonetheless, hypothesis generation is a powerful tool







Figure 8: Hierarchies for groups with high (top) and low (bottom) alcohol consumption rates, defined by two or more lifetime drinks and one or fewer lifetime drinks, respectively.




Figure 10: Interaction between the right angular gyrus from the default mode network (red), and posterior cingulate cortex (green) is 74% higher among drinkers.

for scientific exploration, and has been used recently to inform biomedical research, such as in Abedi et al. (2012) and Spangler et al. (2014).

In addition to providing insight on its own, the NFM procedure complements other modes of analysis. A potentially promising extension, especially for a hybrid version capable of voxel level analysis (discussed in Appendix D), would be to use it in conjunction with graph-theoretic analyses such as those described in Bassett and Bullmore (2006), Stam and Reijneveld (2007), and van den Heuvel et al. (2008). The general technique, as detailed in Bullmore and Sporns (2009) and Rubinov and Sporns (2010), is to compute pairwise correlations among all voxels, set a threshold above which two voxels are considered connected, and calculate various network summary measures (e.g., degree distribution, assortativity, diameter, etc.). By simply replacing correlations in these networks with interaction rates determined by NFM, the assumption of linearity is left behind.

Finally, it is important to note that the specific forms of the models in the output of the GP algorithm have been analyzed simplistically in the present work. A major potential benefit of NFM is the insight that might be gained from precisely analyzing these mathematical descriptions of the relationships among ROI in the brain. Of course, ascribing meaning to any particular one of these models would have to be done cautiously. However, given the results we describe here, obtained by a coarse treatment, the *collection* of models determined by GP may offer a number of as yet undiscovered insights. This seems a potentially

fruitful avenue for future theoretical research.

## 5. Conclusions

Results produced in our study suggest that there is potential analytical power in the use of NFM, or some modification thereof, in the neuroimaging domain. The procedure we investigated here utilizes commercially available, out-of-the-box GP software, and preliminary statistical analysis of its output. Many improvements and extensions are possible, only some of which we have suggested in this work. Reproduction of recent results constitutes a measure of cross-validation, and the preliminary results presented demonstrate the unique capability of NFM to discover non-linear relationships among regions of the brain that hold promise for illuminating differences in brain function between subject groups. Further, the mathematical characterizations we have achieved, which are not limited by linear or univariate assumptions, are ripe for future investigation.

## Acknowledgements

This work was supported by DARPA grant FA8650-11-1-7155, and the Vermont Advanced Computing Core, NASA (NNX-08AO96G) at the University of Vermont, which provided High Performance Computing resources. The IMAGEN study receives research funding from the European Communitys Sixth Framework Programme (LSHM-CT-2007-037286). The authors also wish to acknowledge NASA for funding in support of NAA through the Vermont Space Grant Fellowship; Mike Schmidt, creator of the GP package Eureqa used in this study; and Ilknur Icke, Trevor Andrews and Richard Watts for crucial conceptual and technical support.

## Appendix A. Table of ROI

In Table A.1 we list all 52 ROI investigated in this study by number, give their anatomical names, indicate the ICN from within which they were defined, and provide visual representations of their locations within the brain. Due to its length, the table appears after the References.

## Appendix B. Subject-level variation

Figure B.11 contains individual subject interaction hierarchies for two different (randomly selected) subjects, generated by 100 random restarts of the GP algorithm and subsequent normalized frequency analysis. Though the two hierarchies are quite different from one another, they do show some network organization similar to that illustrated in the population level hierarchy (top of Figure 5).

- Portions of the visual cluster (ROI 32-38) are intact in each case.

- Many of the two-region networks remain together, e.g. ICN 1 (ROI 1,2), ICN 16 (ROI 48,49), and though associated with other ROI, also ICN 3 (ROI 5,6) and ICN 17 (ROI 50,51)
- The default mode network (ROI 39-42) is mostly intact in each subject.

The interaction profile of the default mode network illustrates an interesting distinction between the two subjects. For the top subject, the network is fully intact, interacting with ROI 4 (consistent with the population level hierarchy), and also interacting with ROI 26 from the motor visuospatial complex. For the bottom subject, three of the four ROI in the network remain together, but interact instead with several other ROI from the emotional interoceptive class instead of ROI 4, specifically ROI 10, 12, and 16, and a different ROI from the motor visuospatial complex as well (ROI 29 instead of 26). By themselves, these dendrogram comparisons offer no conclusive evidence regarding connections between cognitive processes. However, an experiment could be designed to test if any inferences can be made from such distinctions. For example, the administration of post-scan surveys might grant some interpretability to the specifics of these single-subject interaction hierarchies.

### Appendix C. Linear HCA of alcohol consumption

Here we demonstrate that the shuffling of the interaction hierarchy in drinking (D) versus non-drinking (ND) adolescents discovered by NFM is not uncovered by linear correlation analysis. To perform group-level correlation analysis, the normalized relative  $R^2$  matrices for each subject (described in Section 3.2) are averaged over the 100 subjects in each group. Recall that these matrices are generated for each subject by computing the correlation matrix for the 52 ROI time series, squaring the elements, and normalizing each row (after setting the diagonal to zero). The reciprocal of relative explained variance can be considered a distance between ROI (higher relative  $R^2$  means closer), and the resulting D and ND hierarchies generated by HCA are shown in the top and bottom, respectively, of Figure C.12.

Comparison with Figure 8 suggests that this linear analysis *partially* uncovers a distinguishing difference in interaction between drinking and non-drinking adolescents. Specifically, among non-drinkers, a higher intra-network interaction within ICN 2, comprised of the subgenual ACC and the vmPFC (ROI 3 and 4, respectively) is detected here. The result is an indirect coupling, within non-drinkers, of the default mode network (ROI 39-42) and the complex comprised of ROI 3,18,5,6, through the vmPFC.

The results of NFM provide further insight in two important ways. First, the elevated interaction within ICN 2 among non-drinkers is detected at twice the strength. Second, the main interaction responsible for grouping the

complex of ROI 3,18,5,6, specifically the interaction between the left globus pallidus (ROI 6) and fornix body (ROI 18), is lower among non-drinkers. This second effect is entirely missed by correlation analysis, suggesting that it is nonlinear in nature. The result of capturing these effects *together*, as shown in Figure 8, is a breakup of the complex in non-drinkers, for whom ROI 3,18 are separated from the bilateral globus pallidus of ICN 3 (ROI 5-6). This breakup is suggestive, as each of these ROI is associated with emotion, reward, and interoceptive processes such as thirst, and experiments reporting activity in ICN 5, including the fornix body (ROI 18), predominantly involved interoceptive stimulation, as reported in Laird et al. (2011).

### Appendix D. Improvements and modifications

The GP implementation we used for this study is the commercially available package Eureqa from Nutonian, as described in Schmidt and Lipson (2009). Though much of its behavior can be controlled through the interface or command line, it is proprietary code and thus somewhat of a black box. There are many reasons why a dedicated, open source implementation of GP would be more desirable.

A major challenge for this method of analysis is the computational expense of running a large number of GP searches. Generating the IR map for a single subject requires a large number of random restarts for each ROI. For example, running 100 restarts for each of the 52 ROI in this study, allowing 1 core-hour for each search, requires over 10 hours with access to 500 dedicated processors. The procedure as described here is likely computationally prohibitive for running analyses on large numbers of subjects, or for larger collections of ROI. Intelligent stopping criteria, and many other approaches to the mitigation of computational expense, have been reported at length in the GP literature, an example of which is the use of graphics processors reported in Harding and Banzhaf (2007). It may also be possible to determine an ideal (and smaller) number of restarts that balances computation time with the statistical power of the resulting IR map.

It should also be noted that for collections of ROI much larger than that considered here, in addition to the computational expense resulting from more required searches, each search will take much longer to produce meaningful models due to the larger number of possible explanatory variables. A hybrid method of symbolic regression employing a machine learning algorithm called FFX (Fast Function Extraction) described in McConaghy (2011) as a first pass, and then GP, has great potential for the treatment of higher dimensional data, e.g., large numbers of ROI. A prototype of this method was reported in Icke et al. (2014). FFX is a deterministic algorithm that builds up models with nonlinear terms (e.g., products of ROI signal) in a prescribed fashion and evaluates explanatory power at





Figure B.11: Example hierarchies for two different individual subjects. Note the large degree of variation between the two.





Figure C.12: Linear hierarchies for groups with high (top) and low (bottom) alcohol consumption rates, defined by two or more lifetime drinks and one or fewer lifetime drinks, respectively.



each stage. By ruling out ROI that are likely not explanatory at each stage, the algorithm reduces the dimensionality of the search. In other words, at the cost of reduced breadth in the search space, the algorithm provides huge reductions in computation time in addition to reducing the number of variables that will eventually be injected into the GP algorithm. Implemented effectively, this hybrid algorithm could eliminate the necessity of ROI selection completely by allowing direct regression over voxel signals.

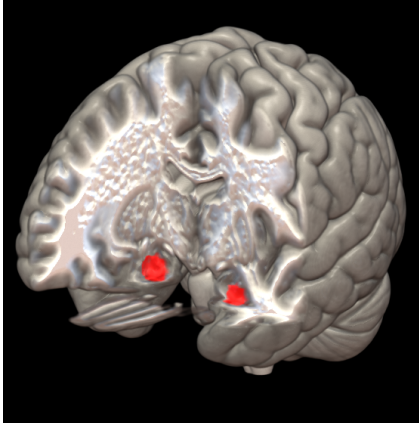
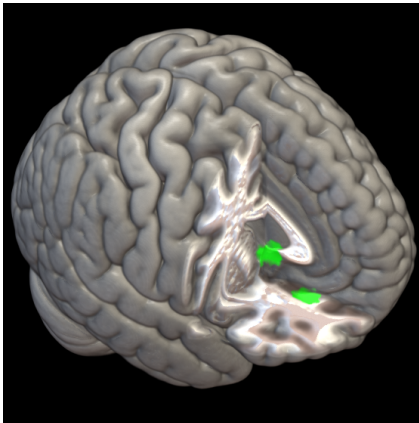
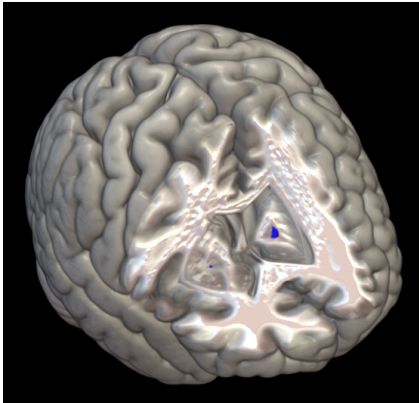
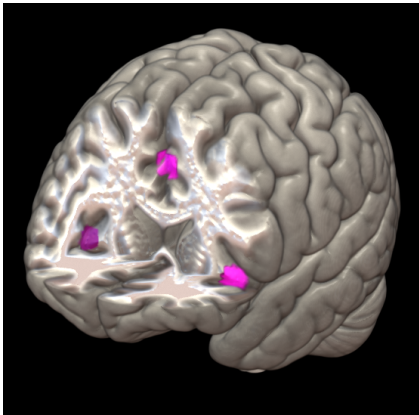
An ever-present concern in the analysis of fMRI is the level of noise in the data. Particularly in the case of regressing over voxel signals, low signal-to-noise ratio is a major challenge, and indeed GP efficacy is diminished in such circumstances. However, there has been some work on modifying the GP algorithm to better manage noisy data, an example of which is the inclusion of noise generators called *stochastic elements* with user-defined distributions (e.g., Gaussian or uniform) as potential explanatory “variables”. These generators can themselves end up inside complex functions within the models, providing those models the capability of reproducing realistic noise distributions more likely to be at play than the typical Gaussian. There is no guarantee that this modification will prove beneficial in the case of fMRI, but it has been shown, in Schmidt and Lipson (2007), to effectively identify exact underlying analytical models in the presence of nonlinear, non-Gaussian and nonuniform noise.

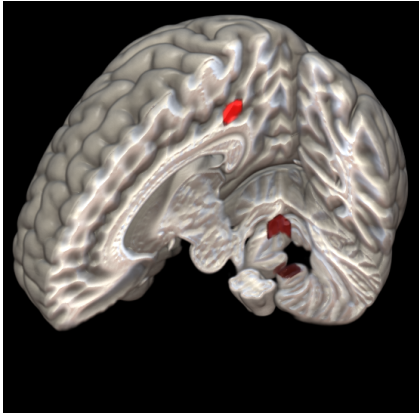

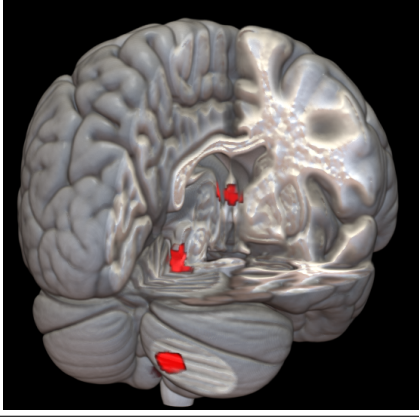
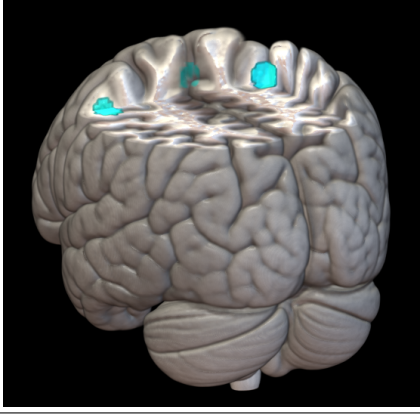
## References

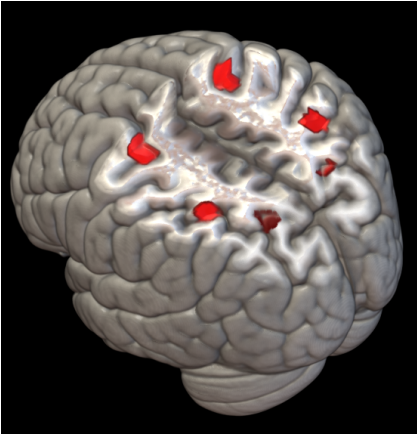
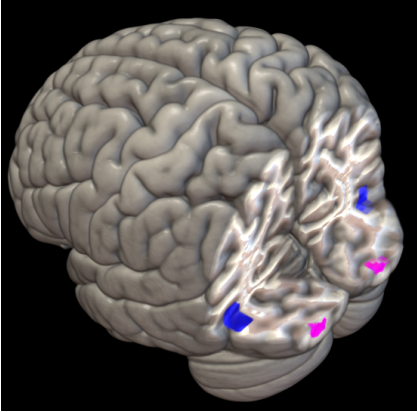
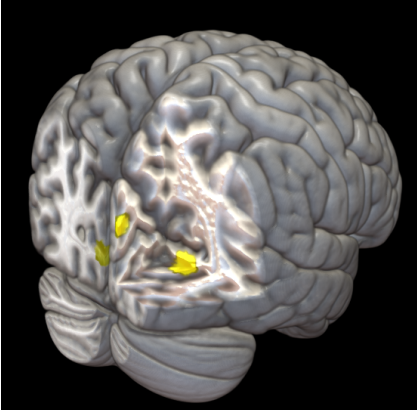
- Abedi, V., Zand, R., Yeasin, M., Faisal, F. E., 2012. An automated framework for hypotheses generation using literature. *BioData Min* 5 (1), 13.
- Bassett, D. S., Bullmore, E., Dec 2006. Small-world brain networks. *Neuroscientist* 12 (6), 512–523.
- Birn, R. M., Murphy, K., Handwerker, D. A., Bandettini, P. A., 2009. fmri in the presence of task-correlated breathing variations. *NeuroImage* 47 (3), 1092 – 1104, *brain Body Medicine*.
- Bullmore, E., Sporns, O., 03 2009. Complex brain networks: graph theoretical analysis of structural and functional systems. *Nat Rev Neurosci* 10 (3), 186–198.  
URL <http://dx.doi.org/10.1038/nrn2575>
- Calhoun, V., Adali, T., Pearlson, G., Pekar, J., 2001. A method for making group inferences from functional mri data using independent component analysis. *Human Brain Mapping* 14 (3), 140–151.  
URL <http://dx.doi.org/10.1002/hbm.1048>
- Friston, K., Harrison, L., Penny, W., 2003. Dynamic causal modelling. *NeuroImage* 19 (4), 1273 – 1302.
- Harding, S., Banzhaf, W., 2007. Fast genetic programming on gpus. In: Ebner, M., O'Neill, M., Ekrt, A., Vanneschi, L., Esparcia-Alczar, A. (Eds.), *Genetic Programming*. Vol. 4445 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg, pp. 90–101.
- Hui, K. K., Liu, J., Makris, N., Gollub, R. L., Chen, A. J., I. Moore, C., Kennedy, D. N., Rosen, B. R., Kwong, K. K., 2000. Acupuncture modulates the limbic system and subcortical gray structures of the human brain: Evidence from fmri studies in normal subjects. *Human Brain Mapping* 9 (1), 13–25.
- Hyvärinen, A., Pajunen, P., 1999. Nonlinear independent component analysis: Existence and uniqueness results. *Neural Networks* 12 (3), 429 – 439.
- Icke, I., Allgaier, N., Danforth, C. M., Whelan, R. A., Garavan, H. P., Bongard, J. C., 2014. A deterministic and symbolic regression hybrid applied to resting-state fmri data. In: Riolo, R., Moore, J., Kotanchek, M. (Eds.), *Genetic Programming Theory and Practice XI*. Springer, Ch. 9, pp. 155–173.
- Kanwisher, N., Stanley, D., Harris, A., 1999. The fusiform face area is selective for faces not animals. *NeuroReport* 10 (1).
- Kruggel, F., Zysset, S., von Cramon, D., 2000. Nonlinear regression of functional {MRI} data: An item recognition task study. *NeuroImage* 12 (2), 173 – 183.
- Laird, A. R., Fox, P. M., Eickhoff, S. B., Turner, J. A., Ray, K. L., McKay, D. R., Glahn, D. C., Beckmann, C. F., Smith, S. M., Fox, P. T., 2011. Behavioral interpretations of intrinsic connectivity networks. *J. Cognitive Neuroscience*, 4022–4037.
- Liu, P., Zhou, G., Zhang, Y., Dong, M., Qin, W., Yuan, K., Sun, J., Liu, J., Liang, J., von Deneen, K. M., Liu, Y., Tian, J., 2010. The hybrid glmica investigation on the neural mechanism of acupoint st36: An fmri study. *Neuroscience Letters* 479 (3), 267 – 271.
- McConaghy, T., 2011. Ffx: fast, scalable, deterministic symbolic regression technology. In: Riolo, R., Vladislavleva, E., Moore, J. (Eds.), *Genetic Programming Theory and Practice IX*. Springer, Ch. 13.
- McKeown, M. J., Sejnowski, T. J., 1998. Independent component analysis of fmri data: Examining the assumptions. *Human Brain Mapping* 6 (5-6), 368–372.
- Mueller, W. M., Yetkin, F. Z., Hammeke, T. A., Morris, G. L. I., Swanson, S. J., Reichert, K., Cox, R., Haughton, V. M., 1996. Functional magnetic resonance imaging mapping of the motor cortex in patients with cerebral tumors. *Neurosurgery* 39 (3).
- Mukamel, R., Gelbard, H., Arieli, A., Hasson, U., Fried, I., Malach, R., 2005. Coupling between neuronal firing, field potentials, and fmri in human auditory cortex. *Science* 309 (5736), 951–954.  
URL <http://www.sciencemag.org/content/309/5736/951.abstract>
- Murphy, K., Harris, A. D., Wise, R. G., 2011. Robustly measuring vascular reactivity differences with breath-hold: Normalising stimulus-evoked and resting state {BOLD} fmri data. *NeuroImage* 54 (1), 369 – 379.
- Nolde, S. F., Johnson, M. K., D’Esposito, M., 1998. Left prefrontal activation during episodic remembering: an event-related fmri study. *NeuroReport* 9 (15).
- Rio, D. E., Rawlings, R. R., Woltz, L. A., Gilman, J., Hommer, D. W., 2013. Development of the complex general linear model in the fourier domain: Application to fmri multiple input-output evoked responses for single subjects. *Computational and Mathematical Methods in Medicine* 2013, 16.  
URL <http://dx.doi.org/10.1155/2013/645043>
- Rubinov, M., Sporns, O., 2010. Complex network measures of brain connectivity: Uses and interpretations. *NeuroImage* 52 (3), 1059 – 1069, *computational Models of the Brain*.
- Schmidt, M., Lipson, H., 2009. Distilling free-form natural laws from experimental data. *Science* 324 (5923), 81–85.  
URL <http://www.sciencemag.org/content/324/5923/81.abstract>
- Schmidt, M. D., Lipson, H., 7-11 Jul. 2007. Learning noise. In: Thierens, D., Beyer, H.-G., Bongard, J., Branke, J., Clark, J. A., Cliff, D., Congdon, C. B., Deb, K., Doerr, B., Kovacs, T., Kumar, S., Miller, J. F., Moore, J., Neumann, F., Pelikan, M., Poli, R., Sastry, K., Stanley, K. O., Stutzle, T., Watson, R. A., Wegener, I. (Eds.), *GECCO '07: Proceedings of the 9th annual conference on Genetic and evolutionary computation*. Vol. 2. ACM Press, London, pp. 1680–1685.
- Schumann, G., Loth, E., Banaschewski, T., Barbot, A., Barker, G., Buechel, C., Conrod, P., Dalley, J., Flor, H., Gallinat, J., Garavan, H., Heinz, A., Itterman, B., Lathrop, M., Mallik, C., Mann, K., Martinot, J.-L., Paus, T., Poline, J.-B., Robbins, T., Rietschel, M., Reed, L., Smolka, M., Spanagel, R., Speiser, C., Stephens, D., Stroehle, A., Struve, M., Williams, S., Desrivieres, S., 2010. The imagen study: reinforcement-related behaviour in normal brain function and psychopathology. *Molecular Psychiatry* 15 (12), 1128 – 1139.
- Smith, S. M., Fox, P. T., Miller, K. L., Glahn, D. C., Fox, P. M., Mackay, C. E., Filippini, N., Watkins, K. E., Toro, R., Laird, A. R., Beckmann, C. F., 2009. Correspondence of the brain’s func-

- tional architecture during activation and rest. *Proceedings of the National Academy of Sciences* 106 (31), 13040–13045.  
 URL <http://www.pnas.org/content/106/31/13040.abstract>
- Spangler, S., Wilkins, A. D., Bachman, B. J., Nagarajan, M., Daryaram, T., Haas, P., Regenbogen, S., Pickering, C. R., Comer, A., Myers, J. N., Stanoi, I., Kato, L., Lelescu, A., Labrie, J. J., Parikh, N., Lisewski, A. M., Donehower, L., Chen, Y., Lichtarge, O., 2014. Automated hypothesis generation based on mining scientific literature. In: *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. KDD '14.* ACM, New York, NY, USA, pp. 1877–1886.  
 URL <http://doi.acm.org/10.1145/2623330.2623667>
- Stam, C. J., Reijneveld, J. C., 2007. Graph theoretical analysis of complex networks in the brain. *Nonlinear Biomed Phys* 1 (1), 3.
- van den Heuvel, M. P., Stam, C. J., Boersma, M., Hulshoff Pol, H. E., Nov 2008. Small-world and scale-free organization of voxel-based resting-state functional connectivity in the human brain. *Neuroimage* 43 (3), 528–539.
- Whelan, R., Conrod, P. J., Poline, J.-B., Lourdasamy, A., Banaschewski, T., Barker, G. J., Bellgrove, M. A., Buchel, C., Byrne, M., Cummins, T. D. R., Fauth-Buhler, M., Flor, H., Gallinat, J., Heinz, A., Ittermann, B., Mann, K., Martinot, J.-L., Lalor, E. C., Lathrop, M., Loth, E., Nees, F., Paus, T., Rietschel, M., Smolka, M. N., Spanagel, R., Stephens, D. N., Struve, M., Thyreau, B., Vollstaedt-Klein, S., Robbins, T. W., Schumann, G., Garavan, H., 06 2012. Adolescent impulsivity phenotypes characterized by distinct brain networks. *Nat Neurosci* 15 (6), 920–925.  
 URL <http://dx.doi.org/10.1038/nn.3092>
- Whelan, R., Watts, R., Orr, C. A., Althoff, R. R., Artiges, E., Banaschewski, T., Barker, G. J., Bokde, A. L. W., Buchel, C., Carvalho, F. M., Conrod, P. J., Flor, H., Fauth-Buhler, M., Frouin, V., Gallinat, J., Gan, G., Gowland, P., Heinz, A., Ittermann, B., Lawrence, C., Mann, K., Martinot, J.-L., Nees, F., Ortiz, N., Paillere-Martinot, M.-L., Paus, T., Pausova, Z., Rietschel, M., Robbins, T. W., Smolka, M. N., Strohle, A., Schumann, G., Garavan, H., the IMAGEN Consortium, 08 2014. Neuropsychosocial profiles of current and future adolescent alcohol misusers. *Nature* 512 (7513), 185–189.  
 URL <http://dx.doi.org/10.1038/nature13402>

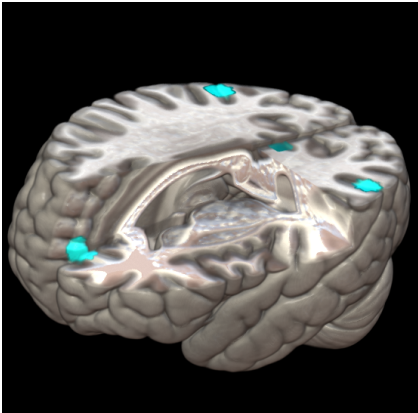
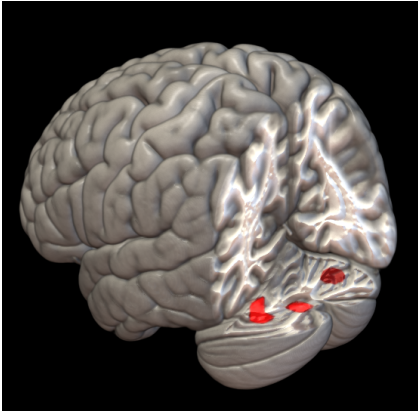
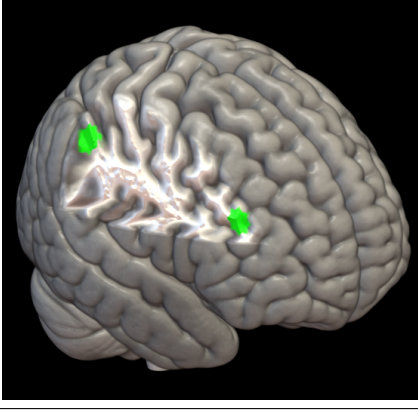
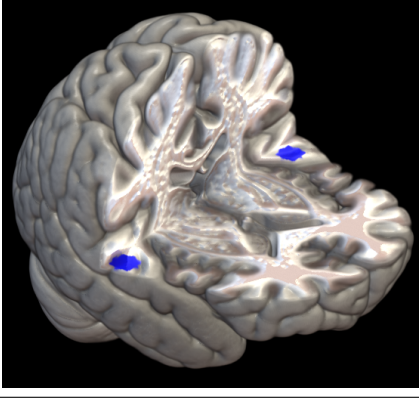
Table A.1: Table of ROI

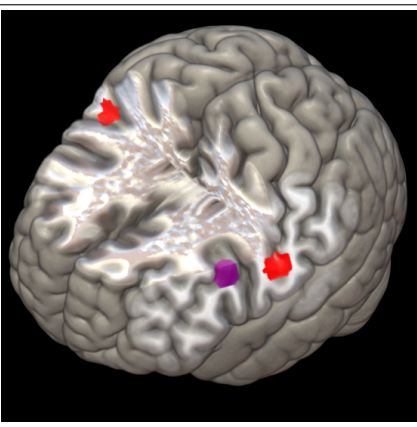
ROI	ICN	ROI Description	Visualization
1	1	left anterior hippocampus	
2	1	right anterior hippocampus	
3	2	subgenual anterior cingulate cortex, anterior caudate	
4	2	ventromedial prefrontal cortex, medial frontal gyrus	
5	3	right globus pallidus	
6	3	left globus pallidus	
7	4	right anterior insula	
8	4	left anterior insula	
9	4	middle cingulate cortex, dorsomedial prefrontal cortex	

ROI	ICN	ROI Description	Visualization
10	5	inferior cerebellum	
11	5	posterior cingulate cortex	
12	5	inferior vermis	
13	5	inferior vermis	
14	5	anterolateral cerebellum	
15	5	anterolateral cerebellum	
16	5	posterior cerebellum	
17	5	inferior colliculus, anterior vermis	
18	5	fornix (body)	
19	6	posterior dorsomedial prefrontal cortex	
20	6	left superior precentral gyrus	
21	6	right posterior superior parietal cortex	

ROI	ICN	ROI Description	Visualization
22	7	left superior parietal cortex	
23	7	right precuneus	
24	7	left superior parietal cortex	
25	7	right superior parietal cortex	
26	7	left posterior dorsolateral prefrontal cortex	
27	7	right posterior dorsolateral prefrontal cortex	
28	8	left postcentral gyrus	
29	8	paracentral lobule	
30	8	anterior inferior parietal cortex	
31	8	right postcentral gyrus	
32	10	right lateral occipital cortex	
33	10	left lateral occipital cortex	
34	11	left occipital pole	
35	11	right occipital pole	
36	12	lingual gyrus	
37	12	right cuneus	
38	12	right fusiform gyrus	



ROI	ICN	ROI Description	Visualization
39	13	posterior cingulate cortex	
40	13	left angular gyrus	
41	13	right angular gyrus	
42	13	anterior dorsomedial prefrontal cortex	
43	14	right superior cerebellum	
44	14	vermis	
45	14	left superior cerebellum	
46	15	right middle frontal gyrus	
47	15	right supramarginal gyrus	
48	16	left superior temporal gyrus	
49	16	right superior temporal gyrus	

ROI	ICN	ROI Description	Visualization
50	17	right inferior pre and post central gyrus	
51	17	left inferior pre and post central gyrus	
52	18	left inferior frontal gyrus	