Reverse Engineering the Brain with Eureqa

Nicholas Allgaier
University of Vermont
33 Colchester Ave
Burlington, VT 05405
nicholas.allgaier@uvm.edu

Ryan McDevitt
University of Vermont
33 Colchester Ave
Burlington, VT 05405
ryan.mcdevitt@uvm.edu

ABSTRACT

A detailed characterization of the human brain, its structural and functional underpinnings, remains on the frontier of modern science. Neurological research is important not only for its intrinsic interest, but for the purpose of better understanding (diagnosing and treating) neurological disorder as well. Happily, along with many other fields, neuroscience is entering an era of "Big Data" in which a new approach is possible: start from the data, *then* get to a theory (rather than testing a theory by collecting data).

The introduction and exploration of a methodology designed to implement this approach is the subject of this study. Evolutionary symbolic regression is performed on brain imaging data, using software called Eureqa, in order to discover and characterize interaction between regions of the brain. The technique is applied to two data sets: (1) a single subject performing two different, but related, tasks; and (2) a different subject performing no task, whose brain is said to be in a "resting state". Results suggest that the methodology provides meaningful information, in that active regions are correctly predicted and tasks classified in (1), and many of the known resting-state interactions are uncovered in (2). Furthermore, the technique characterizes interactions as either linear or nonlinear, providing more information than current methods (which make a linear assumption).

As each data set comes from only one subject, and some details of the methodology require principled refinement, we emphasize the exploratory nature of this study. Results presented suggest the potential of the methodology, but should be taken as preliminary and unvalidated.

Categories and Subject Descriptors

I.2.1 [**Artificial Intelligence**]: Applications and Expert Systems—*Medicine and Science*

General Terms

Experimentation

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
BURLINGTON '11 Burlington, VT USA
Copyright 20XX ACM X-XXXXX-XX-X/XX/XX ...$10.00.

Keywords

Brain, fMRI, neuroimaging, symbolic regression, Eureqa

1. INTRODUCTION

The detailed workings of the human brain remain a mystery to modern science. Though some things are understood in a fuzzy sense, for example the primary regions responsible for motor, visual and auditory response, the connectivity and interaction within and between such regions has not yet been fully described. This study is a preliminary investigation of an evolutionary technique for (1) teasing out such interactions; and (2) describing the *nature* of those interactions.

A Genetic Programming (GP) package called Eureqa, developed by Schmidt et. al. [[8, 9]], is used to regress functional relationships between 17 selected regions of the brain of a subject performing two slightly different tasks in response to a visual stimulus. The goal is to determine, for each region, which other regions it "depends" on, and also to determine the qualitative nature of that dependence, i.e., linear vs. nonlinear. Also, the technique is applied to uncover interactions between regions of the brain when a subject is given no task, and the brain is in a so-called "resting state", which is somewhat of a misnomer.

The General Linear Model (GLM) is the current method of analysis for determining correlations of neuronal activity among regions of the brain [1]. In contrast with the GLM, the method proposed here avoids the assumption of linearity. Furthermore, application of the GLM requires a priori knowledge of a stimulus administered to the subject, whereas no stimulus is required at all for the evolutionary approach. This means that the GLM can't be used (without modification) to study the resting state of the brain. Details of the GLM will be presented in the Background section.

Measurement of activity in the brain is accomplished by functional Magnetic Resonance Imaging (fMRI). 3-D images of a subject's brain are captured real-time, and show contrast in blood-oxygen-level dependence (BOLD). The BOLD signal (resulting from a time series of fMRI images) is taken as a proxy for neuronal activity [11]. The approach proposed in this study consists of using GP to evolve models that express the BOLD signal from a particular region of the brain *as a function of* the concurrent BOLD signals from other regions. The interpretation, however, is not of dependence necessarily, but of interaction. It is correlation, and the nature of that correlation, that is sought, not causation.

The rest of the paper is structured as follows: Section 2 contains background information describing (1) fMRI in

more detail, including a discussion of the risks in taking the BOLD signal as a proxy for neuronal activity; (2) the application of the GLM, and its limitations; and (3) the GP package Eureqa. Section 3 details the experiments performed to test the evolutionary approach to uncovering brain interactions, and results appear in Section 4 along with a discussion. Finally, we conclude the paper in Section 5.

2. BACKGROUND

Before describing the evolutionary approach to discovering interactions in the brain, we present some background information about neuroimaging with fMRI in Section 2.1, detail the current GLM method for studying brain activation in Section 2.2, and give some specifics about the GP package Eureqa used to implement the evolutionary approach in Section 2.3.

2.1 Neuroimaging with fMRI

Direct measurement of electrical neuron activity is not easily accomplished. Electroencephalograms (EEGs) can be recorded by placing electrodes on the surface of subject's head, and are adequate for measuring an averaged whole-brain electrical signal. However, targeting specific regions of the brain, especially at the resolution of a neuron, is not possible with current technology. Highly invasive procedures involving the placement of electrodes directly on the surface (or even inside) of the brain do exist, but this is obviously not practical for large scale studies, nor for measuring activity at many locations in a single brain. This difficulty in direct measurement is the motivation for neuroimaging.

One of the most recent methods of neuroimaging, developed about 20 years ago, is functional MRI. 3-D images of the brain captured by magnetic field indicate changes in blood flow, called the *hemodynamic response*, that correlate with neuronal activity. The reason for this correlation is that active neurons require more glucose and oxygen than inactive ones, and thus neuronal activity is said to be blood-oxygen-level dependent, giving rise to the name BOLD for the fMRI signal [11]. It is important to note that there is some danger in using the BOLD signal as a proxy for neuronal activity. Many other factors contribute to increased blood flow in the brain, including contraction/dilation of vasculature, the structure of that vasculature (which can vary considerably with demographic), and even subject-specific general health and behavioral habits. It is basically impossible to separate these factors. There are some techniques for reducing the effect of conflation [7], and also modifications to the standard fMRI procedure that effectively use other proxies for activity [4], but these methods are outside the scope of this paper.

Very high image resolutions are attainable with fMRI. For this work, the data collected fall on a grid with voxels (3D pixels) approximately 3mm on a side, corresponding to about 50,000 voxels in a single whole-brain image. Typically, a region of the brain is represented by an averaged BOLD signal over some number of contiguous voxels, in order to smooth the noise present in a single voxel. The temporal resolution at which data can be collected depends on how much of the brain is being imaged (and the equipment). Unfortunately, imaging the whole-brain takes about 2 seconds on the equipment used for this study, and even state-of-the-art equipment doesn't perform much better. This has two important implications: first, many processes in the brain occur at a timescale not resolvable by fMRI; and second, a single 3D image is taken a slice at-a-time over the 2 seconds, so that the interpretation of concurrency is loose, to say the least. A technique called time-slice correction, which we won't describe here, is used to give the best estimate of brain-state in the 2-second window.

Other challenges in the collection of fMRI data include the necessity of motion correction resulting from unavoidable head movement during the scan [5], and the transformation of the image into "standard brain-space" for the purposes of comparison between subjects [2]. Neither of these are trivial, but the techniques used to address these challenges are also outside the scope of this paper. Despite proxy data issues, correction for motion and temporal misalignment, required transformations, and probably other issues we neglect here, fMRI has proven to reveal meaningful and reproducible results for 20 years.

2.2 The General Linear Model

The GLM is the current method used to determine if a region of the brain is activated in response to a stimulus signal. A multi-variable linear regression is performed, whereby the mean BOLD signal B from the region is fitted to a linear combination of the expected hemodynamic response to the stimulus, which we call s, and any nuisance variables n_i that are not of interest, for example fluctuations due to the heart beating. Thus,

$$B = \beta s + \alpha_1 n_1 + \alpha_2 n_2 + ... + \epsilon \quad (1)$$

where ϵ is an error term. If, over many trials, β is bounded away from zero in a statistically significant sense, then the region is determined to be activated by the stimulus [1]. An example of a BOLD signal, response to stimulus, and possible nuisance variables appears in Figure 1. To perform the linear regression, note that the form of the stimulus signal must be known a priori.

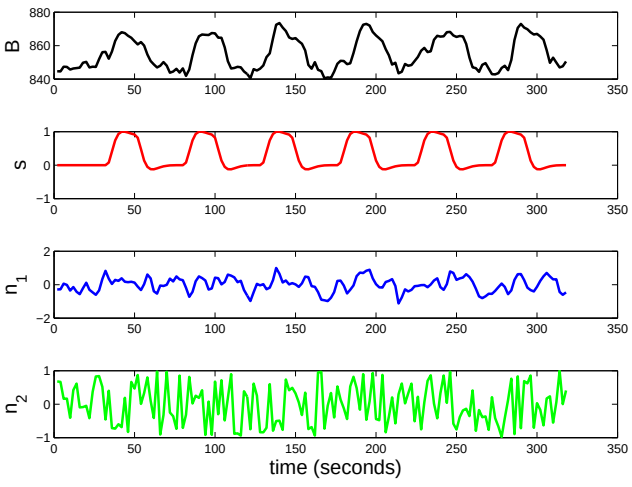

Figure 1: Examples of a BOLD signal B, expected hemodynamic response s to a visual stimulus, and illustrative examples of possible nuisance variables $n_{1,2}$, top to bottom, respectively. The units of the BOLD signal are unimportant here, and the bottom 3 signals into which B is to be linearly decomposed are scaled to have maximum amplitude of 1.

Regions of the brain that co-activate in response to a particular stimulus, or type of stimulus, are often interpreted as being related, or interacting. However, the nature of their interaction or relationship can't be determined using the GLM, as it is assumed to be both linear and a direct result of an applied stimulus. By performing symbolic regression, the assumption of linearity can be avoided, and regressing brain regions with respect to one another, instead of a stimulus, allows the discovery of interactions in response to unknown stimuli, for example the self-stimuli in intrinsic ("resting state") brain activity. This is the motivation for the evolutionary technique presented in this work, which is carried out with Eureqa.

2.3 Eureqa

Eureqa uses symbolic regression to find an analytical solution to explain experimental data. Symbolic regression, as implemented in Eureqa, creates an initial population of functional forms from user-specified building blocks, which are stored as the operator and terminal sets. These building blocks can consist of a range of operators, including arithmetic, trigonometric, exponential, etc. Using a graph-based representation (parse tree) the genotype is arranged such that the top and middle of the tree is created from members of the operator set, and the leaves consist of members of the terminal set. The phenotype is then the function created from this parse tree.

Once the initial population has been created, standard EA techniques (mutation/crossover) are used to generate offspring. Eureqa uses single point crossover, where a random branch is swapped between two individuals to form two offspring. Mutation, which occurs exclusive of crossover, picks a random location in the tree and swaps an operator for another operator of the same arity, or swaps a terminal node with another member of the terminal set. Typically mutation will occur with a probability $\approx 1\%$ and crossover with a probability $\geq 50\%$. These offspring then need to be evaluated for fitness. Since the goal of symbolic regression is to find a functional form that explains the data, a natural fitness evaluation is the error between the functional output, and the provided data.

Eureqa is in continuous development, and as such it implements new genetic programming techniques when possible. For this paper, beta version .84 was used, which incorporates the island model for parallelization and age-fitness Pareto for selection. As these are improvements to the basic GP technique, a brief discussion of each techniques follows.

The island genetic algorithm model [12],[6] was originally developed as a method for parallelizing genetic algorithms, but it has the added benefit of helping to prevent premature convergence. In the island model, as implemented in Eureqa [3], the initial population is divided into separate "islands" that are allowed to evolve independently. At fixed intervals, a percentage of the population on each island "migrates" to another random island. By keeping the populations separate, each island is able to search a different part of the search space, with the migration serving to increase new genetic material to each island which helps to prevent premature convergence. Since these populations evolve independently and only need to interact during the migrations, each island can be run on its own processor.

Within each island, Eureqa uses age-fitness Pareto [10] as a second method for preventing premature convergence. In this method, a single-objective problem (e.g. error minimization) is turned into a multi-objective problem by adding the age of the genetic material as a second objective to be optimized. This is done by by giving all genetic material an initial age of 1, and during crossover, the children inherit the age of the oldest parent. The best solution in this problem will then be the highest fitness solution with the minimum age.

Using these methods, Eureqa generates a number of solutions, each offering a potential function that explains the target data by input data. The concept behind this study is to use symbolic regression to determine the functional forms relating regions of interest (ROI's) in an fMRI scan of a subject performing two different tasks, and also while no task is being performed. Ideally this method would be able to find the connections that are already known to exist, but would also show non-linear connections that the GLM cannot account for.

3. EXPERIMENTS

The main outcome for this project was the creation of a methodology for using Eureqa to analyze the fMRI data. This section focuses on procedure used to analyze the two different data sets. In the first, fMRI data from a subject performing two different tasks is analyzed using Eureqa to determine if known connections within the brain can be determined using symbolic regression. In the second, the connections between known networks in resting state data are analyzed.

In the first task of the first data set, the subject is asked to watch a blinking checkerboard (called a visual metronome) and tap their finger continuously. In the second task, the subject again watches the visual metronome and taps their fingers, but this time they are asked to count to five and then stop tapping, take a short break, and then repeat. This counting leads to different regions of the brain being active. The fMRI data for the two different tasks, labeled test 1 and test 2 (where the authors do not know which includes counting), has been parsed to 17 regions of interest, labeled x_1 through x_{17}. These regions are known to contain some areas that are expected to show increased activity during the tasks, as well as some regions that are just noise. For each of these regions there is a time series of 159 data points, collected in 2 second intervals. To find the relationship between each region, each time series was modeled as a function of the other 16 time series (i.e $x_1 = f(x_2, x_3 ... x_{17})$). The relevant parameters for each of these runs is summarized in Table 1. The operator set includes the basic arithmetic operators, basic trigonometric functions and real valued constants. Each time series was run for 8 core hours using RMSE as the target for minimization.

The resting state data contains 52 regions that are known to be part of 17 different networks. Figure 2 shows one of these networks that contains four regions, which appear as the red regions highlighted on the brain image at the bottom of the figure. Using a methodology similar to that presented previously, one region from each of the 17 networks is run as a function of the other 51 regions and the resulting functional forms are saved. For networks that contain multiple regions, a second run is done with the regions found in the first run eliminated. By eliminating the regions that are highest fitness, Eureqa is able to find other regions that are less correlated.

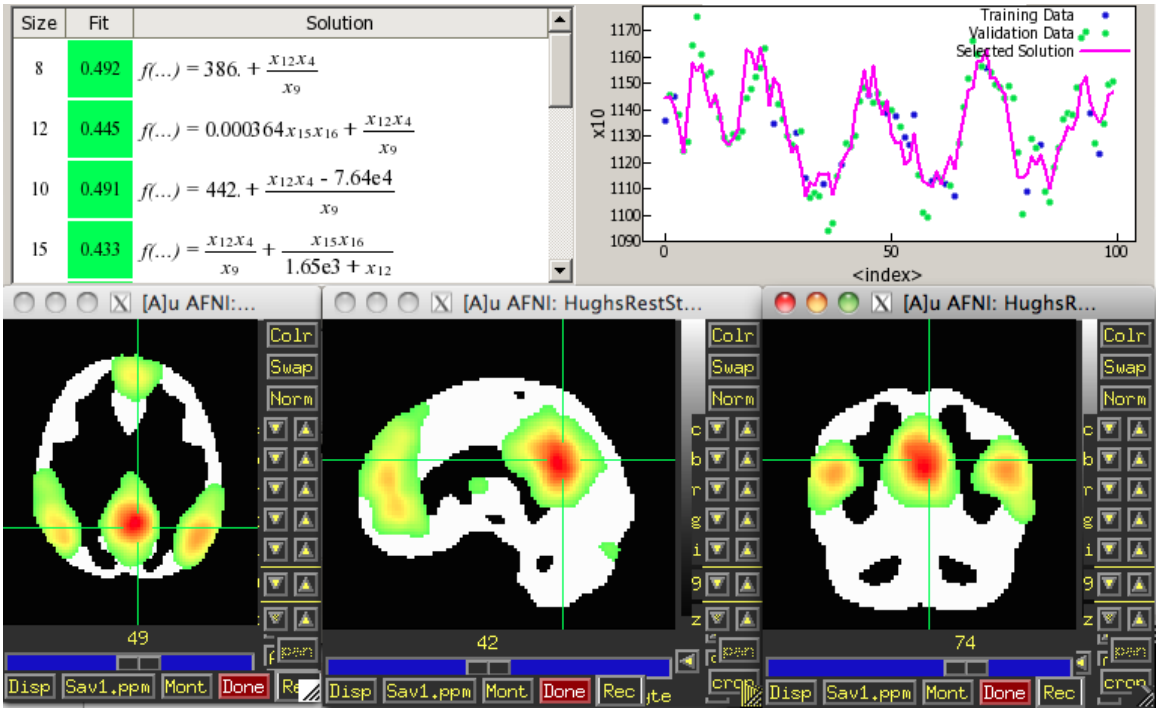

Figure 2: Visualization of the resting state data analyzed in this study. The image at the top of the figure is a sample of Eureqa output as a network of correlated resting state regions is run.

Table 1: Summary of Experimental Setup

Solution Population Size	512
Selection Method	Age-Fitness Pareto
Parallelization Method	Island Model
P(mutation)	.015625
P(crossover)	.5
Solution Encoding	Operation List (graph)
Operator Set	$+, -, *, /, \sin, \cos$, const
Terminal Set	x_1-x_{17}
Crossover	single point
Mutation	single point
Run Time	8 core hours
Training Data	100 data points
Validation Data	59 data points
Error Minimization	RMSE

For each run, Eureqa generates an accuracy-parsimony Pareto front of potential solutions. To find the most likely relationships between regions, the top 5 solutions on this Pareto front were selected as candidate functions. Within these candidate functions, the most frequently occurring regions were determined. Finally, the candidate function with the lowest (most optimal) fitness that contained all of the most frequent regions, and only those regions, was selected as the functional form for that data point.

As an example, the candidate functions for the x_1 region of test 1 are shown in Figure 3. In these candidate functions, the regions x_2 and x_4 are the most frequently occurring regions. The candidate function with the highest fitness that contains both of these regions, and only these regions, is the second function on the list. In this example, the functional form is a linear combination of the two points, but this is not the case for all the functional forms, a point which will be explored more in the results section.

4. RESULTS AND DISCUSSION

The evolutionary technique for analysis of fMRI BOLD signals was applied to two data sets. The first comes from the single-subject experiment in which the participant was asked to perform finger-tapping tasks described earlier in the paper, and the second comes from a subject who was given no task or stimulus. We present and discuss the results of the analyses in the Finger Tapping and Resting State sections, respectively.

4.1 Finger tapping

The functional form selected for each region of interest (ROI), as described in the previous section, indicates three things: (1) what other regions are depended upon by the ROI, (2) the nature of the dependence (linear or nonlinear) for each region on which the ROI depends, and (3) a measure of our confidence in the dependence, as represented by the fitness of the selected functional form. These results are displayed graphically in Figure 4 for test 1, (top panel), and test 2 (bottom panel).

As an informal proof of concept (or suggestion of meaningfulness), the results pictured in Figure 4 were used in an attempt to determine (1) which regions were activated during each task; and (2) which of test 1 or test 2 involved counting in addition to finger-tapping. Inferences were made blind, (without knowledge of the physical locations of the regions), and then verified afterward by the experimenter who

Complexity(Size)	Accuracy(Fitness)	Formula
7	0.361477	f(x2,x3,x4,x5,x6,x7,x8,x9,x10,x11,x12,x13,x14,x15,x16,x17,conv) = 0.309468*x2 + 0.76333*x4
9	0.349936	f(x2,x3,x4,x5,x6,x7,x8,x9,x10,x11,x12,x13,x14,x15,x16,x17,conv) = 102.196 + 0.323408*x2 + 0.636109*x4
11	0.336521	f(x2,x3,x4,x5,x6,x7,x8,x9,x10,x11,x12,x13,x14,x15,x16,x17,conv) = 0.37481*x2 + 0.0953432*x16 + 0.562044*x4
5	0.384217	f(x2,x3,x4,x5,x6,x7,x8,x9,x10,x11,x12,x13,x14,x15,x16,x17,conv) = x4 + 0.110148*x3
19	0.292892	f(x2,x3,x4,x5,x6,x7,x8,x9,x10,x11,x12,x13,x14,x15,x16,x17,conv) = 0.32793*x2 + 0.13052*x16 + 0.574124*x4 + 3.17257*sin(0.186004*x14)

Figure 3: Sample of top 5 candidate functions for the x_1 region for Test 1

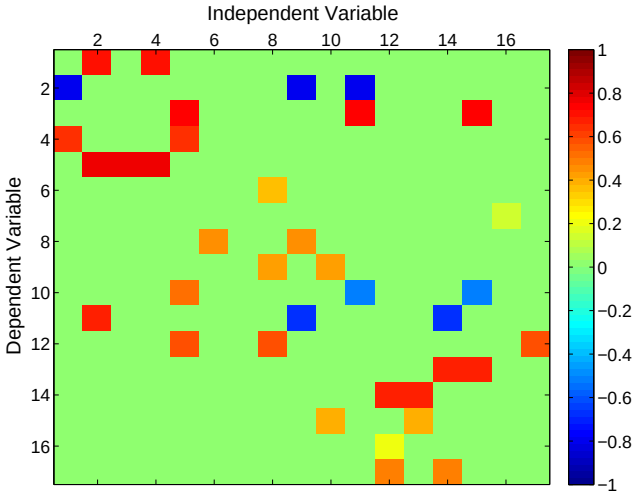

(a) Test 1

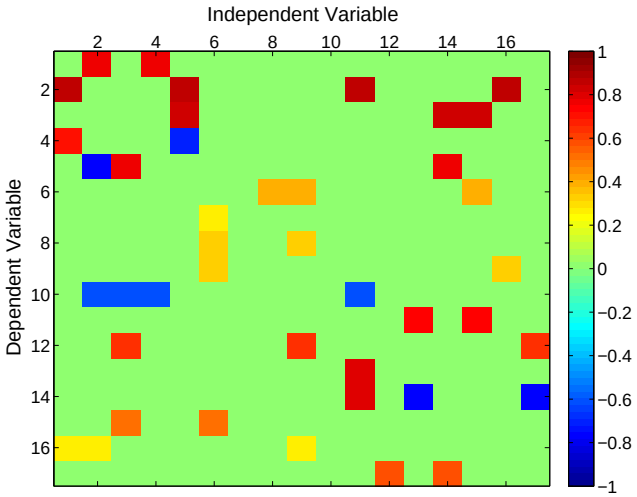

(b) Test 2

Figure 4: Results for (a) test 1 and (b) test 2. Regions for which functional forms were determined are on the vertical axis, and nonzero elements indicate dependence on the corresponding region on the horizontal axis. Fitness is color magnitude, where fitness has been scaled to [0,1] and 1 is best. The sign of the color indicates the nature of the dependence, positive for linear and negative for nonlinear. For example, the first row for test 1 can be read as x_1 is a linear combination of x_2 and x_4, with confidence approximately 0.7.

collected the data. The following paragraphs summarize the inferences and reasoning behind them.

Based on the idea that regions activating in response to a stimulus should have codependent BOLD signals, the symmetric block of high fitness codependence among regions $1-5$ present in both test 1 and test 2 suggests that they are activated in both tasks. This leads to the conclusion that they represent visual and/or motor response centers, which are expected to be active in both tests. Likewise, the poor fitness for regions $6-9$ in both tests suggests that these regions are unrelated to either task. Regions $15-17$ are likely inactive in test 1, but the fact that there is some high fitness dependence on these regions in test 2 suggests that they may be active for test 2.

The generally higher activity and fitness in test 2, indicated by more colored points and darker colors, respectively, suggest that this is the task that includes counting. Further supporting this conclusion, regions 10, 12, and 15 also show dependence on the regions we have concluded are visual or motor areas, suggesting that in test 2 some other task (like counting) has become linked with the visual or motor response. The fitness of functional forms for these regions is somewhat low, however, so we call this inference "uncertain".

It is hard to say anything about regions 11, 13 or 14. These regions may be active in both tasks, as their functional forms have reasonable fitness in each case. Again, there is some high fitness dependence on 11, and 14, particularly for test 2. For all 3 regions, we note that the dependence changes markedly from test 1 to test 2, but don't infer anything based on this observation.

Table 2 gives the names and and related functions of the 17 regions, along with their expected activity during the two tasks, as provided by the experimenter. The last column is a summary of the inferences made blindly by the authors, based on Figure 4, as described in the previous paragraphs. The inferences line up fairly well with the expected activity of the regions. Furthermore, the two regions that should best distinguish between tasks, (regions 10 and 15), were both regions that became dependent on the visual and motor control centers in test 2, and thus the inference that this was the counting task was not only correct, but also suggests that the proposed methodology does in fact uncover interactions between regions of the brain. Also, since region 10 interacts nonlinearly with regions 2-4, (and 11), it is unclear whether or not this relationship would be picked up by the GLM.

4.2 Resting state

As shown in the finger tapping experiments, Eureqa is able to successfully identify some regions that are active in response to a known stimulus. As a second check of the ability of this method to provide meaningful results, resting state data is analyzed by seeking functional forms for

Table 2: 17 regions by name, showing related function, expected activity, and inferred activity.

#	Region name	Related function(s)	Expected activity	Inferred activity
1	Occipital Lobe	visual response	both tasks	both
2	Primary and Supplementary Motor	motor control	both tasks	both
3	Right Insula	motor control	both tasks	both
4	Left Cerebellum	motor control	both tasks	both
5	Left Insula	motor control	both tasks	both
6	Left Frontal Lobe	none related	neither task	neither
7	Right Hippocampus	none related	neither task	neither
8	Anterior Cingulate Cortex	none related	neither task	neither
9	Left Temporal Lobe	none related	neither task	neither
10	Right Dorsolateral Prefrontal Cortex	motor planning	more with counting	test 2 (uncertain)
11	Left Temporoparietal Junction	self-other distinction	minor in both tasks	both (uncertain)
12	Right Putamen	motor learning, prep., sequences	both tasks	test 2 (uncertain)
13	Left Temporoparietal Junction	self-other distinction	minor in both tasks	both (uncertain)
14	Right Temporoparietal Junction	self-other distinction	minor in both tasks	both (uncertain)
15	Right Inferior Frontal Gyrus	go/no go tasks	more with counting	test 2 (uncertain)
16	Right Occipital Lobe	visual response	both tasks	test 2 (uncertain)
17	Left Putamen	motor learning, prep., sequences	both tasks	test 2 (uncertain)

one region from each of 17 networks (small groups of regions known to co-activate). Since the subject is not being presented with a stimulus, it can be challenging to regress the regions that are active using the GLM. With Eureqa, the regions that are in the same network should show up as part of the same functional form. Regions that are in the same network are numbered sequentially, and a functional form is sought for the first region of each network. In Figure 5 there is a clear trend of highly fit regions appearing just above the diagonal. In addition, there exist a number of out-of-network connections. As in the finger tapping experiments, the functional forms that Eureqa produced include linear and non-linear terms.

Eureqa was able to find at least one other region within the network for many of the 17 networks analyzed. For each of the larger networks, Eureqa is able to find multiple points within the network, either during the first or second run. Further refinement of each region beyond the initial point in the network, as well as re-running each point to find all of the connections would fill out the image. Generating these plots for a number of patients should demonstrate if the secondary networks that Eureqa is identifying are meaningful.

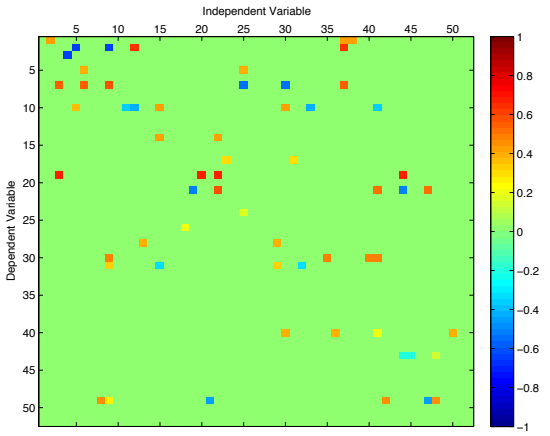

Figure 5: Results for Resting State data. Regions in the same network are numbered consecutively, so the highly fit correlations just above the diagonal show Eureqa's ability to detect these connections.

5. CONCLUSIONS

The main deliverable for this work is a methodology for using an evolutionary technique to (1) discover interactions between regions of the brain; and (2) describe the nature of those interactions qualitatively, as either linear or nonlinear. This study was exploratory in nature, and thus the results presented here are meant to *suggest the potential* of the methodology. No claim can yet be made about the validity of the results, or even the methodology itself. The next step of this work is, naturally, to pursue a rigorous validation by applying the technique to data from many subjects.

Even given the exploratory nature of the work accomplished so far, however, the authors are very encouraged by the suggested potential of the method! In the finger tapping analysis, with very little domain knowledge and no details about which locations of the brain the data came from, we were able to (1) give a reasonably accurate account of which regions were active in each task; (2) classify which of the tasks was being performed; and (3) describe the interactions between regions that led us from (1) to (2), including nonlinear interactions that may be difficult to detect using current methods. Furthermore, in the case of the resting state analysis, the method discovered expected interactions among networked regions of the brain (giving further credence to the validity of the technique), and uncovered unexpected inter-network relationships that may provide new information to domain specialists, provided they can be validated across subjects.

Tackling the seemingly infinite complexity of the human brain requires techniques well-suited to such complexity. The ever-growing store of neuroimaging data begs for large-scale analysis, and the developing tools of evolutionary computation were born, one might say, to perform just this type of analysis. Upon validation of the method proposed here, we hope that it may provide new useful information to domain specialists, not only for a better general understanding of the brain, but also potentially for the diagnosis and treatment

of neurological disorder.

6. ACKNOWLEDGMENTS

We'd like to thank Hugh Garavan and Richard Watts for collecting and explaining the fMRI data, Kevin Murphy and Clare Kelly for their help in processing the data and explaining the details of neuroimaging, and Chris Danforth and Josh Bonguard for strategic advice.

7. REFERENCES

[1] V. Calhoun, T. Adali, V. McGinty, J. Pekar, T. Watson, and G. Pearlson. fmri activation in a visual-perception task: Network of areas detected using the general linear model and independent components analysis. *NeuroImage*, 14(5):1080 – 1088, 2001.

[2] D. L. Collins, P. Neelin, T. M. Peters, and A. C. Evans. Automatic 3d intersubject registration of mr volumetric data in standardized talairach space. *J. Comput. Assist. Tom.*, 18:192–205, 1994.

[3] F. Francisco, S. Giandomenico, T. Marco, and V. Leonardo. Parallel genetic programming. In A. Enrique, editor, *Parallel Metaheuristics*, pages 127–153. Ed. Hoboken, 2005.

[4] T. J. Huppert, R. D. Hoge, S. G. Diamond, M. A. Franceschini, and D. A. Boas. A temporal comparison of bold, asl, and nirs hemodynamic responses to motor stimuli in adult humans. *Neuroimage*, 29(2):368–382, January 2006.

[5] M. Jenkinson, P. Bannister, M. Brady, and S. Smith. Improved optimization for the robust and accurate linear registration and motion correction of brain images. *NeuroImage*, 17(2):825 – 841, 2002.

[6] S. Mahfoud. *Niching methods for genetic algorithms.* PhD thesis, University of Illinois at Urbana-Champaign, 1995.

[7] K. Murphy, A. D. Harris, and R. G. Wise. Robustly measuring vascular reactivity differences with breath-hold: normalising stimulus-evoked and resting state bold fmri data. *Neuroimage*, 54(1):369–379, January 2011.

[8] M. Schmidt and H. Lipson. Coevolution of Fitness Predictors. *IEEE Transactions on Evolutionary Computation*, 12:736–749, December 2008.

[9] M. Schmidt and H. Lipson. Distilling free-form natural laws from experimental data. *Science*, 324(5923):81–85, June 2009.

[10] M. Schmidt and H. Lipson. Age-fitness Pareto optimization. *Genetic Programming Theory and Practice VIII*, pages 129–146, 2011.

[11] M. Singh, S. Kim, and T.-S. Kim. Correlation between bold-fmri and eeg signal changes in response to visual stimulus frequency in humans. *Magnetic Resonance in Medicine*, 49:108–114, 2003.

[12] D. Whitley, S. Rana, and R. Heckendorn. The island model genetic algorithm: On separability, population size and convergence. *Journal of Computing and Information Technology*, 7:33–48, 1999.