

Network, anatomical, and non-imaging measures for the prediction of ADHD diagnosis in individual subjects

Jason W Bohland*¹, Sara Saperstein², Francisco Pereira³, Jérémy Rapin³ and Leo Grady³

¹Department of Health Sciences, Boston University, Boston, MA 02215

²Graduate Program for Neuroscience, Boston University, Boston, MA 02215

³Siemens Corporation, Corporate Research and Technology, Princeton, NJ 08540

Email: Jason W Bohland* - jbohland@bu.edu; Sara Saperstein - saperste@bu.edu; Francisco Pereira - francisco-pereira@siemens.com; Jeremy Rapin - jeremy.rapin@cea.fr; Leo Grady - leo.grady@siemens.com;

*Corresponding author

Abstract

Brain imaging methods have long held promise as diagnostic aids for neuropsychiatric conditions with complex behavioral phenotypes such as Attention-Deficit/Hyperactivity Disorder. This promise has largely been unrealized, at least partly due to the heterogeneity of clinical populations and the small sample size of many studies. A large, multi-center dataset provided by the ADHD-200 Consortium affords new opportunities to test methods for individual diagnosis based on MRI-observable structural brain attributes and functional interactions observable from resting state fMRI. In this study, we systematically calculated a large set of standard and new quantitative markers from individual subject datasets. These features (**>12,000 per subject**) consisted of local anatomical attributes such as cortical thickness and structure volumes, and both local and global resting state network measures. Three methods were used to compute graphs representing interdependencies between activations in different brain areas, and a full set of network features was derived from each. Of these, features derived from the inverse of the time series covariance matrix, under an L1-norm regularization penalty, proved most powerful. Anatomical and network feature sets were used individually, and combined with non-imaging phenotypic features from each subject. Machine learning algorithms were used to rank attributes, and performance was assessed under cross-validation and on a separate test set of 168 subjects for a variety of feature set combinations. While non-imaging features gave highest performance in cross-validation, the addition of imaging features in sufficient numbers led to improved generalization to new data. Stratification by gender also proved to be a fruitful strategy to improve classifier performance. We describe the overall approach used, compare the predictive power of different classes of features, and describe the most impactful features in relation to the current literature.

Keywords: ADHD, fMRI, network analysis, functional connectivity, resting state, machine learning

1 INTRODUCTION

2 Attention-Deficit/Hyperactivity Disorder (ADHD) is a complex developmental neuropsychiatric disorder
3 characterized by abnormal inattentiveness, impulsivity, and hyperactivity. Recent estimates based on meta-

1 analyses from the literature suggest a worldwide prevalence rate of $\sim 5.29 \pm 0.28\%$ in children 18 years of
2 age or younger (Polanczyk et al., 2007), making it among the most common childhood disorders. Many
3 children diagnosed with ADHD continue to exhibit symptoms throughout adulthood. The Diagnostic and
4 Statistical Manual-IV Text Revision (DSM-IV-TR) describes three different types of ADHD: a *predominantly*
5 *inattentive* type, a fairly uncommon *predominantly hyperactive-impulsive* type, and a most common *combined*
6 *type* that includes features from each of the other two types (American Psychiatric Association, 2000). The
7 biology of ADHD, including its genetics (Faraone et al., 2005; Banaschewski et al., 2010) and neurobiology
8 (Tripp and Wickens, 2009), has received considerable attention but remains relatively poorly understood
9 (see e.g., Casey et al., 2007; Bush, 2010).

10 **Diagnosis.** There is no single, standard test for ADHD in children, and thus diagnosis requires the
11 extended involvement of mental health professionals to accurately assess the existence and range of behavioral
12 evidence and to differentiate ADHD from other disorders with overlapping symptomatology or from typically
13 occurring behaviors. This process is costly and time-consuming. The use of non-invasive brain imaging
14 methods coupled with advanced image analytics techniques holds the promise of great benefit for expediting
15 or adding certainty to this diagnostic process. While this hope exists for essentially all neuropsychiatric
16 disorders which rely on behavioral evidence for diagnosis, algorithms for objective classification of patients
17 may hold special value in ADHD due to its heterogeneity, high prevalence, and particularly controversial
18 diagnosis (Wolraich, 1999).

19 **Gender and IQ differences.** A number of demographic factors appear to be related to positive diag-
20 nosis of ADHD and may be useful in informing diagnostic algorithms. ADHD is diagnosed at a significantly
21 higher rate in boys than in girls (Polanczyk et al., 2007). In 2007 in the United States, based on parent
22 reports of any ADHD diagnosis in children ages 4-17, ADHD had been diagnosed in 13.2% of boys compared
23 with 5.6% of girls¹. Further, multiple studies have reported gender differences in the symptom profiles of
24 children with ADHD (Gaub and Carlson, 1997; Newcorn et al., 2001; Gershon and Gershon, 2002), sug-
25 gesting possible sex-specific mechanisms or manifestations of the pathophysiology of the disorder. Cognitive
26 measures including Full Scale IQ as well as Verbal and Performance IQ are also reliably different between
27 individuals with ADHD and typically developing controls (TDCs) (Frazier et al., 2004).

28 **Neuroimaging correlates of ADHD.** Family and twin studies of ADHD have established high degrees
29 of heritability (Faraone et al., 2005; Burt, 2009), supporting the existence of a biological and genetic basis for
30 the disorder. Brain imaging may then be viewed as a method for providing quantitative or semi-quantitative
31 *endophenotypes* (Doyle et al., 2005), measures which are theoretically more closely related to the underlying
32 biological etiology than are the behavioral signs and symptoms. To this end, a wide range of anatomical and
33 functional brain imaging studies have been conducted comparing children with ADHD to typically developing
34 children, and have described a number of relatively consistent results (Giedd et al., 2001; Durston, 2003;
35 Bush, 2010). These range from gross findings that total cerebral volume may be reduced by $\sim 3 - 4\%$ (Valera
36 et al., 2007) and that global cerebral glucose metabolism is substantially reduced (Zametkin et al., 1990),
37 to results demonstrating reduced cortical thickness in the right superior frontal gyrus across the lifespan
38 (Almeida et al., 2010), and numerous reports of altered anatomical or functional connectivity (Konrad and
39 Eickhoff, 2010; Liston et al., 2011) in individuals with ADHD relative to controls. The breadth of the
40 available functional and structural imaging studies, which are too numerous to review here, have generally
41 implicated prefrontal cortex (including dorsolateral and ventrolateral prefrontal areas), anterior cingulate
42 cortex, parietal cortex, striatum, and cerebellum.

43 Despite the promise of brain imaging for aiding clinical diagnosis, currently no imaging techniques are
44 recommended for this purpose (Bush, 2010). One possible explanation for this unfulfilled promise could be
45 that the measures necessary for accurate diagnosis may be high-dimensional and not readily observable from

¹<http://www.cdc.gov/ncbddd/adhd/data.html>

1 classical univariate image analysis methods. Further, studies conducted in small samples may not sufficiently
2 generalize to larger populations.

3 **Complex brain networks.** It is long established that alterations in inter-regional neuronal connectivity,
4 as in the case of so-called *disconnection syndromes*, can underlie complex brain disorders (Geschwind, 1965).
5 Recent theories of the basis for neuropsychiatric disorders have reinvigorated these conceptualizations (e.g.,
6 Mega and Cummings, 1994; Tekin and Cummings, 2002; Geschwind and Levitt, 2007). Concomitant with
7 such theories, advances in brain imaging and data analytic methods (as well as the rise of the more general
8 domain of systems biology), have enabled the generation and quantitative analysis of *complex brain networks*
9 built from structural and/or functional imaging data (Bullmore and Sporns, 2009). Such networks, including
10 those constructed based on *resting state fMRI (rs-fMRI)*, are now commonly used in the study of normal
11 and abnormal cognitive function.

12 Resting-state fMRI, based on low-frequency BOLD signal fluctuations that occur while the subject is
13 resting and performing no explicit task, has garnered significant recent interest as a tool for finding clinically
14 relevant biomarkers and/or for measuring responses to treatment (Greicius, 2008). Networks built from
15 correlations or related measures calculated across instances of functional imaging time series obtained during
16 rest may be interpreted, at least partly, to reflect the intrinsic *functional connectivity* between different
17 brain areas, and their properties may be relevant for understanding typical and atypical variability across
18 the population. A number of studies have provided sparse evidence for altered connectivity in ADHD, but
19 much further work is necessary to fully characterize network phenotypes as well as intersubject variability
20 (Castellanos et al., 2009).

21 Currently a large number of methods exist for the construction of functional connectivity networks.
22 Simulation studies performed by Smith et al. (2011) have demonstrated direct evidence that not all methods
23 are equivalent in their ability to estimate the existence of underlying inter-regional connections. Thus the
24 potential for discovering network-based measures that correlate with the presence or absence of ADHD may
25 hinge on the methods used to define each individual network's elements, including nodes, presence / absence
26 of edges between node pairs, and any weights assigned to those edges.

27 **Network analytics.** The *structure* of a system (physical or biological) is an abstract concept that can
28 be difficult to quantify in a manner that can be used to predict its characteristics or to distinguish between
29 different types of systems. However, by representing the structure of a system with a network model it
30 becomes possible to quantify various measurements of the network that may be used to characterize the
31 system. For example, a series of network measures has been used to determine whether the configuration of
32 a network of cells derived from histological section images can predict the presence of cancer (Gunduz et al.,
33 2004; Demir et al., 2005; Khurd et al., 2011; Chekkoury et al., 2012). Historically, these network measures
34 have been widely used in chemical graph theory for a very long time (e.g., Wiener, 1947) to predict various
35 structural and behavioral properties of molecules (for reviews of this type of usage see Hansen and Jurs,
36 1988; Mihalić and Trinajstić, 1992; Bonchev and Rouvray, 1991). More recently, a wide variety of measures
37 have been explored and used for the structural quantification of systems spanning scientific disciplines (see
38 Costa et al., 2007). Furthermore, network-level properties such as graph efficiency (Latora and Marchiori,
39 2001) have already been demonstrated to be useful markers of ADHD (Wang et al., 2009).

40 ADHD CHALLENGE

41 In this paper we describe our efforts to use attributes derived from MR images as well as non-imaging
42 phenotypic measures to predict the presence or absence of an ADHD diagnosis in child and adolescent
43 subjects. This opportunity was made possible by the availability of a large dataset comprising structural

1 and resting-state functional MRI scans and associated non-imaging phenotypic data (e.g., gender, age,
2 cognitive testing measures) from 776 children and young adults. These data were provided by the ADHD-
3 200 Consortium, a “self-organized, grassroots initiative, dedicated to accelerating the scientific community’s
4 understanding of the neural basis of ADHD through the implementation of discovery-based science” as part of
5 the *ADHD-200 Global Competition*². **This competition invited researchers from all disciplines to participate**
6 **in an effort to produce the highest performance imaging-based diagnostic classification algorithm for ADHD.**
7 **Scoring was based on a pre-specified point system that involved both base diagnosis and diagnosis of ADHD**
8 **subtype. A separate award was provided for the most innovative neuroscientific examination of ADHD.** Our
9 group finished fifth overall in the classification competition, and the present paper describes our approach
10 and continued efforts to improve and characterize classification methods and results.

11 The open availability of large N datasets with compatible, commonly coded primary data and metadata is
12 critical to successfully fulfilling the promise of exploratory and machine learning approaches for the discovery
13 of principles of normal and disordered brain function (Biswal et al., 2010; Milham, 2012). The ADHD-
14 200 sample represents a starting point for this approach in ADHD research and presents a test bed for
15 utilizing large sets of anatomical, network, and non-imaging measures for objective diagnosis of complex
16 neurobehavioral disorders that currently require extensive, continued behavioral testing for diagnosis and
17 lack clear biomarkers.

18 APPROACH

19 We approached the diagnosis problem by examining the predictive power of three sets of features or attributes:
20 (i) non-imaging phenotypic features, (ii) anatomical features derived from structural brain images, and (iii)
21 network features derived from graphs depicting functional connectivity during rs-fMRI. A set of over 12,000
22 features was computed for each individual subject; these features were provided, in the groups described
23 above, to train classifiers. These classifiers were evaluated using a cross-validation approach, and then used
24 to predict the presence or absence of ADHD in a separate group of test subjects.

25 Prediction using the non-imaging phenotypic features (e.g., age, gender, IQ) provides a baseline perfor-
26 mance level which, in the current dataset, is well above chance. We anticipated that the addition of certain
27 anatomical and network features would add predictive power, increasing overall performance of the classi-
28 fiers. Based on previously observed gender differences, we also anticipated that *separate classifiers* may be
29 needed to maximize predictive power for boys vs. girls. Below we describe the overall results of classification
30 using this approach and discuss the power of the different classes of features.

31 METHODS

32 ADHD-200 DATASET

33 Anatomical and resting state functional MRI scans were performed at 8 different facilities on children and
34 young adults ages 7 to 21 years (mean: ~ 12 years), approximately half ($\sim 53\%$) male. Participants were
35 diagnosed as either typically developing or ADHD-Hyperactive, ADHD-Inattentive, or ADHD-combined
36 type. These data and various metadata describing subject phenotypic traits (including diagnosis) were made
37 available to download through the ADHD-200 Consortium. **Too few cases of ADHD-Hyperactive type were**
38 **available for any practical analysis and, in general, the other two subtypes are not distinguished in our analysis**
39 **below except to assess subtype-specific performance. In other words, we describe classifiers built to determine**

²http://fcon_1000.projects.nitrc.org/indi/adhd200/

1 the presence or absence of ADHD irrespective of subtype, but we were able to analyze post hoc whether
2 their performance differed depending on the subtype. While the ADHD Global Competition included scoring
3 based on subtype diagnosis, and while being able to differentiate subgroups within a diagnostic category is
4 of high potential clinical value, we chose to focus efforts here on the problem of *primary* diagnosis, applying
5 a large and diverse feature set, combined with gender-stratified training that limited the number of training
6 examples that would be available for each individual subtype.

7 Imaging parameters (i.e., repetition time, number of volumes acquired, other MR acquisition parameters)
8 differed somewhat across sites³, and some subjects were imaged more than once. For subjects with multiple
9 rs-fMRI sessions, time series were concatenated after pre-processing.

10 Non-imaging phenotypic attributes that were included in the dataset and used in the classifiers included:

- 11 1. Age
- 12 2. Gender
- 13 3. Handedness
- 14 4. Verbal IQ
- 15 5. Performance IQ

16 OVERALL MACHINE LEARNING FRAMEWORK

17 Our approach to diagnosing Attention-Deficit/Hyperactivity Disorder combined the use of anatomical mark-
18 ers, non-imaging phenotypic data (above), and network analytics computed from graphs constructed from
19 each individual’s resting state fMRI data. We calculated a variety of standard and new quantitative markers
20 and applied machine learning algorithms to perform the ADHD classification.

21 The anatomical and network features (described below) were normalized to have zero mean and unit
22 standard deviation across all subjects in the dataset. Any features with constant values (across subjects)
23 were excluded at this stage. The non-imaging phenotypic features were used without any normalization,
24 and missing values of Verbal or Performance IQ were replaced by the respective population average. We
25 augmented these with several binary features: *NoIQ* (1 if the subject was missing IQ scores, 0 otherwise)
26 and *Site1* through *Site8* (1 if the subject was imaged at that site, 0 otherwise).

27 We performed a two-fold cross-validation procedure over the released training data. Examples (subjects)
28 were sorted by site, classification label (i.e., diagnosis), gender, and age, and were assigned in round-robin
29 fashion to folds 1 and 2; this was done to ensure that each fold contained equal proportions of examples with
30 similar values for those attributes. Next the dataset was separated by gender (474 boys and 280 girls were
31 available in the final training data), and all the steps described below were performed separately for the two
32 resulting datasets, with results combined at the end. **The assessments of overall diagnostic performance were
33 made after pooling all individual classification results across the gender groups (i.e., treating the outputs as
34 if they had come from a single classifier).**

35 We used three methods to rank features (Guyon, 2003), and used the resulting rankings to select between
36 5 and 6000 (or *all* in the cases where more than 6000 features were available) features using each method,
37 with the selection procedures performed inside the training data of each cross-validation fold. **Two filter**

³The individual parameters for each site are available at the competition web site:
<http://fcon.1000.projects.nitrc.org/indi/adhd200/>.

1 methods were used to score each feature individually: (i) a 2-sample t-test (comparing the feature values
2 for ADHD participants vs. controls) and (ii) the accuracy of a classifier trained and tested in nested cross-
3 validation over that single feature. We also applied a wrapper method, recursive feature elimination, which
4 also made use of nested cross-validation. This method consisted of training a linear support vector machine
5 classifier on all available features, then scoring features by the magnitude of the weights assigned to them.
6 The bottom 50% features were then eliminated from consideration and the procedure repeated until there
7 were 10 or fewer features. Final feature scores were a combination of the last round in which each feature
8 survived and, within that, the magnitude of the weight assigned by the classifier, such that the last surviving
9 features had the higher scores.

10 The three feature selection methods were parallel equivalents of one another (i.e., one replaces the other,
11 each resulting in a *ranking* of the overall feature set). Because features were ranked using each of the above
12 methods throughout a cross-validation process, we computed the average rank of each feature (across folds)
13 for each method. These average rankings were then used for selecting the features used to train classifiers
14 using examples from the *entire* training dataset. Specifically, we chose the top K ranked features for a
15 variety of values of K . Performance was then assessed using the *separate test set* released by the ADHD-
16 Consortium, which was not used for either training or feature selection. A linear SVM classifier (LIBSVM
17 Chang and Lin, 2011) with regularization parameter $\lambda = 1$ was used as the classification algorithm with all
18 sets of selected features; other classification algorithms were tested in preliminary studies, but provided
19 similar or inferior overall performance.

20 **Evaluating classification performance.** Comparing the value of different classifiers requires a mea-
21 sure capable of representing the utility of one classifier over another. One natural measure is the *accuracy*
22 which quantifies the probability that the classifier will make a correct prediction of ADHD vs. TDC. How-
23 ever, under differing practical scenarios, it may be more important to be confident that the classifier provides
24 a correct diagnosis of ADHD positive (high true positive rate) or that the classifier provides confidence in
25 ruling out an ADHD positive diagnosis (low false positive rate). A mechanism for characterizing the value of
26 a classifier under such different scenarios is the *Receiver Operating Characteristic* (ROC) curve, which plots
27 the probability of predicting a true ADHD positive given a tolerance for a certain percentage of false ADHD
28 positive results. Consequently the *area under the ROC curve* (AUC) can be used to measure the value of
29 one classifier compared to another, regardless of the practical scenario. An AUC of 1.0 indicates a perfect
30 classifier (i.e., a true positive is always obtained without sacrificing any false positives) while an AUC of 0.5
31 indicates that the classifier does no better than chance in predicting the presence / absence of ADHD. In
32 the present study, all results are reported using this measure and are provided for the cross-validation stage
33 (on the non-training folds) as well as for the separate test stage.

34 In order to compute the ROC curve for a classifier, we ranked the examples by the magnitude of the
35 LIBSVM decision value output for each. This was obtained for each example by multiplying the weight
36 assigned to each feature by the value it took in that example and adding over features. We then computed
37 the true positive and false positive rates obtained when setting the classification threshold at each point in
38 the ranking. For the set of results using all feature types, we also provide accuracy scores (which represent
39 one point on the ROC curve that maximizes the overall percent correct in binary diagnosis).

40 ANATOMICAL FEATURES

41 All structural MRI scans (T1-weighted MPRAGE volumes, anonymized using a “defacing” algorithm to
42 protect patient confidentiality) were processed through the FreeSurfer software package⁴, version 5.0.0 using
43 the typical “recon-all” procedures. Specifically, this software was used to perform intensity normalization,

⁴<http://surfer.nmr.mgh.harvard.edu>

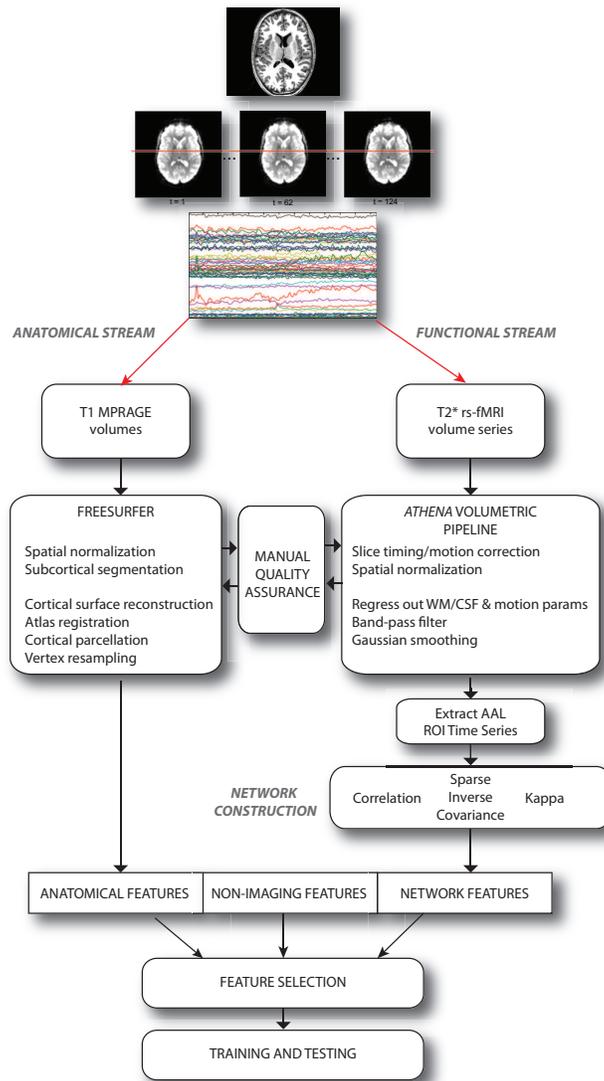


Figure 1: Schematic describing the overall pipeline through which data were processed and submitted to machine learning algorithms for classification. Note that anatomical and functional / network streams were largely independent.

1 skull stripping, white matter segmentation, and tessellation and reconstruction of the cortical surface in
2 each hemisphere. In addition, individual surfaces were registered to a spherical atlas space, and the cortex
3 was parcellated into macroanatomical regions. Furthermore, from the T1 volumes, a set of subcortical
4 structures were segmented, and a variety of morphometric measures were estimated. Technical details of
5 these procedures are described elsewhere (Dale et al., 1999; Fischl et al., 1999a,b). The quality of MR
6 images, Talairach registration, pial surface demarcation, and surface inflation were assessed using a manual
7 inspection protocol. Approximately two percent (14 of 776) of the images failed this stage of quality assurance
8 and were removed from the subsequent analyses. Cortical surface-based features (thickness and curvature)
9 were computed for each subject and resampled onto an icosahedral surface model defined in the atlas space.
10 This surface consists of 2,562 locations (vertices) in each hemisphere, equally spaced around the inflated
11 sphere. Based on initial experiments, we discarded average curvature features and focused on local thickness
12 features as possible ADHD diagnostic aids. Thus, for each subject we calculated a total of 5,124 local cortical
13 thickness features.

14 Additionally, automated surface-based cortical parcellations and volume-based subcortical structure seg-
15 mentations were computed for each subject (Fischl et al., 2002, 2004), and a series of statistics were calculated
16 for the individual structures (average cortical thickness, surface area, volume, mean curvature, and standard
17 deviations of these measures for each cortical region of interest). The volumes of various subcortical gray
18 and white matter structures were also estimated, and normalized by each individual’s total intracranial
19 volume (ICV) to help control for age effects. FreeSurfer also calculated the volumes of subcortical areas
20 with hypointensities in gray or white matter; these were also normalized by ICV and included in the overall
21 feature set.

22 NETWORK FEATURES

23 **Preprocessing rs-fMRI data.** Individual subject resting state functional connectivity networks were gener-
24 ated (using three different network construction methods, see below) from preprocessed rs-fMRI time series
25 data, and a large set of network measures were calculated from these networks. Functional MRI preprocess-
26 ing relied on scripts provided publicly by the NeuroBureau⁵, specifically using the so-called *Athena Pipeline*.
27 All raw rs-fMRI data were reprocessed using these scripts, adapted to our local computing environment,
28 which used methods from the publicly available AFNI (Cox and Hyde, 1997) and FSL (Smith et al., 2004)
29 software packages. This pipeline⁶ included steps for normalization of anatomical volumes to an age-specific
30 (4.5-18.5 years) template brain volume in MNI space (Fonov et al., 2011) (contrast with surface based reg-
31 istration in our FreeSurfer-based anatomical pipeline) using a low-dimensional non-linear deformation, and
32 realignment and co-registration of functional images to this space. The first 4 EPI volumes in each rs-fMRI
33 scan were discarded due to T1-equilibration effects. Slice timing correction was performed, and the mean
34 activation time courses from white matter (WM) and cerebrospinal fluid (CSF) as well as estimated motion
35 parameters and a set of low-order polynomials were used as nuisance regressors. Resulting voxel-wise time
36 courses were band pass filtered ($0.009 \text{ Hz} < f < 0.08 \text{ Hz}$) according to common practice in rs-fMRI analysis
37 (Cordes et al., 2001). Region-specific average time courses were extracted from each subject’s data using
38 the Automated Anatomical Labeling (AAL) template atlas (Tzourio-Mazoyer et al., 2002), which consists
39 of 116 brain regions-of-interest (ROIs) demarcated in MNI-space based on sulcal and gyral landmarks in the
40 MNI single-subject template atlas.

41 **Network construction.** Several different methods have been proposed for inferring functional connec-
42 tivity from a resting state time series acquisition. Smith et al. (2011) used simulation studies to test a wide
43 variety of methods for inferring connections from fMRI time series; following these results, we deployed three

⁵<http://neurobureau.projects.nitrc.org/ADHD200/Data.html>

⁶For details of the Athena pipeline see <http://www.nitrc.org/plugins/mwiki/index.php/neurobureau:AthenaPipeline>

1 of the best-performing methods for estimating weighted networks from the AAL time course data above:

- 2 1. Correlation, with correction for temporal autocorrelation (*Corr*)
- 3 2. Sparse regularized Inverse Covariance (*SIC*)
- 4 3. Patel’s Kappa (*Kappa*)

5 *Corr* networks were based on calculations of Pearson correlation coefficients between the average time
6 series for pairs of AAL regions. These were then converted into P-values under the null hypothesis of
7 no correlation using a Fisher transformation, and taking into account temporal autocorrelation. False-
8 discovery rate (Benjamini and Hochberg, 1995) was used to correct for multiple comparisons at a rate of
9 0.01. Edges representing significant correlation between nodes (AAL regions) were assigned weights equal
10 to the corresponding pairwise correlation coefficient; edges for which the corrected correlations were not
11 significant were set to zero.

12 The *SIC* networks were created using methods from the Sparse Learning with Efficient Projections
13 (SLEP) toolbox (Liu et al., 2009). In particular, the inverse of the AAL time series covariance matrix was
14 computed under an L1-norm regularization penalty (see also Friedman et al., 2008; Huang et al., 2010),
15 yielding a measure of partial correlation. Based on Smith et al. (2011) and exploratory testing, we chose a
16 regularization parameter of $\lambda = 0.1$ in all cases. **L1-regularization enforces sparsity in the inverse estimate,
17 and thus these networks contained many edges with weight values that were close to zero. Nonzero edge
18 weights were real valued, between 0 and 1.**

19 *Kappa* networks were computed based on the κ measure described in Patel et al. (2006), extended to
20 continuous (non-binary) measurements as described in Smith et al. (2011). This is a measure of connection
21 strength based on conditional states of pairs of normalized time series. *Kappa* networks had continuous-
22 valued edge weights and were not subjected to an edge threshold.

23 In definitions below we refer to a graph $G = \{V, E\}$ consisting of a set of vertices (or nodes) V , and edges
24 E . We denote an individual i^{th} vertex as $v_i \in V$ and an edge spanning vertex v_i and v_j as $e_{ij} \in E$. We
25 denote the weight assigned to edge e_{ij} as w_{ij} . In all cases, inferred edges were weighted using real-values
26 (i.e., they were not binarized, as is common in the literature). The *absolute values* of edge weights were
27 used to calculate network measures. These edge weights computed in our network construction methods are
28 *affinity weights*, which are larger if two nodes are more strongly connected. Therefore, in order to compute
29 meaningful measures based on paths, it was important to convert the edge weights to *distance weights*, which
30 are small if nodes are similar. The appropriate relationship between affinity and distance weights was given
31 in Grady and Polimeni (2010) as:

$$w_{\text{distance}} = \frac{1}{w_{\text{affinity}}}, \quad (1)$$

32 which is the same relationship between resistance (distance) and conductance (affinity) in an electrical circuit;
33 see Grady and Polimeni (2010) for more discussion of this point. In the following sections, we will specify
34 whether the affinity or distance weights were used to compute each measure. All networks were *undirected*
35 ($w_{ij} = w_{ji}$) and contained 116 nodes (corresponding to the regions of the AAL atlas) and 6,670 possible
36 edge weights.

37 We subjected each network to a wide range of feature analysis in order to capture specific markers which
38 might aid in predicting the presence / absence of ADHD. We examined a large number of standard measures

1 from the resting state network literature (see e.g., Sporns, 2010). However, since it was unclear which network
2 characterizations might provide insight into ADHD, we broadened the range of network features considered
3 to include features derived from the literature outside of neuroscience, such as those reviewed in Costa et al.
4 (2007) and Grady and Polimeni (2010). Specifically, in addition to examining the standard “scale-free” and
5 “small-world” properties of the networks, we also looked for measures that quantify network connectivity
6 (both single path and parallel path), network separability, network cycle structure and the recurrence of
7 certain network motifs.

8 *Standard network measures.*

9 We calculated measures of graph connectivity that are derived from the literature on “scale-free” and
10 “small-world” networks. To characterize small-world properties, we computed the average path length,
11 diameter, radius, and mean/max/min clustering coefficients for each network (Grady and Polimeni, 2010).
12 These measures would help us distinguish if the functional connectivity network of ADHD subjects more
13 strongly resembled a small-world network (as first described in Watts and Strogatz, 1998). **It has been**
14 **suggested in the literature to also examine the related measures of global efficiency and the harmonic mean**
15 **of the path lengths (Latora and Marchiori, 2001), which we have also included in our study.**

16 Similarly, to examine the scale-free properties of the network, we computed the entropy of the degree
17 structure and the assortativity (Newman, 2002). **These computations could reveal if the functional con-**
18 **nectivity networks of ADHD patients were more or less scale-free than typically developing controls. Since**
19 **scale-free networks (in which the node degree distribution approximately follows a power law) are modeled**
20 **through a process of preferential node attachment, a difference in the scale-free properties could suggest an**
21 **underlying difference in the process of network / circuit formation for ADHD patients.**

22 We also examined several conventional measures that were computed for each node. These features
23 would reveal if the role of a particular node within the overall network were different for ADHD subjects.
24 Specifically, for each node we computed its degree, betweenness, eccentricity and central point dominance,
25 providing a set of measures of how the node is situated in the overall network (Grady and Polimeni, 2010).
26 The weighted degree of a node (region) can be interpreted as a measure of its overall functional connectivity
27 with the rest of the brain, while betweenness has a possible interpretation as indicating the “importance” of
28 a region in the overall flow of information. These per-node features were also condensed into a small set of
29 measures characterizing the overall network such as mean/max/min degree, mean/max/min betweenness,
30 and the entropy of the betweenness values. **For the novice reader, a brief introduction to a set of network**
31 **measures is provided in the Supplementary Materials.**

32 All of these measures depend on shortest paths and were therefore computed using distance edge weights.
33 Since the clustering coefficient measures are typically defined for unweighted graphs, we computed these
34 coefficients by treating a connection between nodes of any weight as a connection (effectively setting all edge
35 weights to unity).

36 ***Parallel connectivity measures.*** One criticism of the conventional measures described above is that
37 they generally rely on measuring shortest paths between pairs of nodes in the network. However, information
38 which spreads through multiple channels (such as a diffusion process) is more sensitive to the collection of
39 strong *parallel* paths connecting a pair of nodes rather than the presence or absence of a strong *single* path
40 connecting the pair. One way of measuring parallel paths is through the concept of *effective resistance*,
41 which treats the network as a linear resistive circuit (Klein and Randić, 1993). Note that it is important to
42 treat all edge weights as affinity weights in an electrical circuit interpretation (Grady and Polimeni, 2010).

43 To quantify the strength of parallel connections between each node and the rest of the network, we
44 calculated the eccentricities of each node with respect to the resistance distance instead of the conventional

1 shortest-path distance. These per-node eccentricities can also be used to quantify the overall parallel con-
 2 nectivity of the network via the radius and diameter of the effective resistance as well as the Kirchhoff Index
 3 used in chemical graph theory (Bonchev et al., 1994; Diudea and Gutman, 1998).

4 **Network separability.** Some networks have one strongly interconnected component while others have
 5 multiple different functional clusters. The separability (or modularity) of a network is not easily characterized
 6 by the previous measures which depend on examining strong paths between node pairs. To quantify the
 7 separability of a network, we computed two classical measures. Specifically, we calculated the Fiedler values
 8 of the unnormalized and normalized Laplacian matrix. The unnormalized Laplacian matrix is defined as
 9 $L = D - W$ where W is the weighted adjacency matrix (with affinity edge weights) and D is a diagonal
 10 matrix of weighted node degrees such that $D_{ii} = \sum_i W_i$. The normalized Laplacian matrix is defined as
 11 $\bar{L} = D^{-1}L$. The Fiedler values are the second-smallest eigenvalues of these matrices, which are known
 12 to reflect the separability of the graph into two pieces (Fiedler originally termed this value the *algebraic*
 13 *connectivity* to reflect this property; Fiedler, 1973). We also computed the third-smallest eigenvalue of these
 14 matrices as a feature, since this value reflects the separability of the network into three components.

15 A feature that describes separability of the graph is its *isoperimetric number* (Mohar, 1989), which
 16 describes the smallest ratio of

$$\iota(S) = \frac{\text{cut}(S, \bar{S})}{|S|}, \quad (2)$$

17 for any node subset $S \subset V$ such that $0 < |S| \leq \frac{1}{2}|V|$, where

$$\text{cut}(A, B) = \sum_{e_{ij}, \text{ s.t. } v_i \in A, v_j \in B} w_{ij}. \quad (3)$$

18 Similarly, a normalized isoperimetric number of a network is also defined as the minimum ratio

$$\overline{\iota(S)} = \frac{\text{cut}(S, \bar{S})}{\text{Vol}(S)}, \quad (4)$$

19 for any node subset $S \subset V$ such that $0 < \text{Vol}(S) \leq \frac{1}{2}\text{Vol}(V)$ where $\text{Vol}(S) = \sum_{v_i \in S} d_i$. Unfortunately,
 20 calculation of either isoperimetric number for a graph is NP-Hard (Mohar, 1989). Therefore, we performed
 21 an estimation of the isoperimetric number using different clustering methods to find good candidate sets S
 22 and taking as features the values of $\iota(S)$ and $\overline{\iota(S)}$ that are smallest over all clustering methods. Specifically,
 23 we applied the spectral clustering and isoperimetric clustering algorithms (Grady and Schwartz, 2006) to
 24 estimate the isoperimetric number. Note that if the network is disconnected, the Fiedler values and isoperi-
 25 metric numbers will all be zero. However, in our experiments with these data and set of network construction
 26 methods it was uncommon to find disconnected networks.

27 **Other measures of modularity** have been explored in the literature (e.g., the Q measure; Rubinov and
 28 Sporns, 2011) that characterize the goodness of a particular node partitioning in a network. However, these
 29 measures are dependent on the partitioning algorithm and, since these networks were generally quite dense,
 30 greedy algorithms (e.g., Newman, 2004) are unlikely to provide meaningful results. Due to the dependence
 31 on partitioning algorithm, these measures were not included as features in our study.

32 **Cycle measures.** The most common way to quantify network structure is from the standpoint of
 33 connectivity between various node pairs. Although these connectivity measures give an indication of the

1 ability to pass signals between nodes, they fail to characterize the structure of feedback loops in the network.
 2 In fact, the cycles of complex networks have recently been shown to contain substantial information about
 3 certain types of networks (Khurd et al., 2011). Consequently, we followed Khurd et al. (2011) to produce
 4 a set of features characterizing the cyclic structure of these networks based on computing a minimum cycle
 5 basis for each network (Horton, 1987; Kavitha et al., 2008). All cycle lengths were computed using distance
 6 edge weights.

7 Most cycles in the minimal cycle basis are triangles (i.e. only three nodes and edges). Our quantification
 8 of the cycle structure was done by calculating the percentage of non-triangle cycles, their mean/max length
 9 and the sum of all cycle lengths.

10 **Sparsity measures.** Network sparsity gives some indication of the overall synchronization of the network
 11 and the overall energy expended by the network. A natural measure of sparsity is to count the number of
 12 nonzero edges in a network, but this measure *assumes that all nodes are valid*. Any connection between two
 13 ROIs, however, might potentially be subdivided into a string of smaller nodes. Similarly, any ROI might
 14 potentially be subdivided into a number of small, tightly coupled nodes. If this subdivision was accidentally
 15 made for one network but not for another, then by taking the number of nonzero edges as the sparsity
 16 measure, the subdivided network would appear more sparse (if it were a connection that was subdivided) or
 17 less sparse (if an ROI were subdivided) than the unsubdivided network.

18 To build a measure that is robust to the subdivision problem, we examined a sparsity measurement
 19 derived from linear algebra. Specifically, let P be a permutation matrix representing a node ordering in
 20 the network. Then, the sparsity of the Cholesky factor for the matrix $\tilde{L} = PLP^T$ may be compared to
 21 the sparsity of the original Laplacian matrix, L , to determine the amount of “fill-in” (new edges) created
 22 by the ordering. It has been shown (e.g., Grady and Polimeni, 2010) that Gaussian elimination of a node
 23 (row/column) in the Laplacian matrix creates a new reduced Laplacian matrix, representing a graph in
 24 which the eliminated node is removed, and a connection is created between all neighbors of the eliminated
 25 node. Consequently, if the removed node is part of a path then the two neighbors of the eliminated node are
 26 connected by a single edge in the reduced graph, causing no fill-in. Similarly, if the eliminated node is part
 27 of a fully connected clique, then no new edges are created in the reduced graph, since all of its neighbors
 28 were already connected. Therefore, we believe that by comparing the sparsity of the Laplacian matrix to
 29 the amount of fill-in created by the Cholesky decomposition of the reordered matrix, a measure of *intrinsic*
 30 sparsity may be obtained which is robust to subdivisions of connections into paths or ROIs into tightly
 31 connected clusters of many nodes.

32 The intrinsic sparsity that we have defined is substantially dependent on the node ordering that defines
 33 the permutation matrix P . Unfortunately, finding an ordering that produces a *minimum* fill-in is known to be
 34 an NP-Hard problem (Papadimitriou, 1976). Consequently, the field of numerical linear algebra has produced
 35 several ordering strategies that are known to provide low fill-in for different types of matrices/networks, such
 36 as the Approximate Minimum Degree (AMD), Cuthill-McKee, and Dulmage-Mendelsohn orderings. To
 37 quantify intrinsic sparsity of the networks, we used as network features the sparsity of the original Laplacian
 38 matrix and the fill-in obtained by the above orderings, as well as the fill-in produced by the lexicographic
 39 ordering as a reference. Since the sparsity measurements look only at the structure of the network, rather
 40 than the edge weights, we computed these sparsity features for each network when the affinity weights were
 41 thresholded at the levels of $\{0, 0.1, 0.3, 0.5, 0.7, 0.9\}$.

42 **Network motifs.** The previous features all examined structural properties of the networks in terms of
 43 paths between node pairs, node centrality, degree distribution, cycles, separability and sparsity. However,
 44 it may be possible that what distinguishes ADHD subjects from TDC subjects is the presence/absence of a
 45 particular circuit of connections in the brain. Unfortunately, measuring the presence of all possible circuits

1 in the network is combinatorically prohibitive, even for our networks of 116 nodes. For example, even the
2 number of possible subgraphs of three nodes is $\sim 250,000$, which is far too many to meaningfully explore
3 without an enormous test dataset. Consequently, we examined the more tractable set of the 6670 possible
4 2-cliques (edges), using each affinity edge weight as a feature.

5 FEATURE IMPACT

6 For features of all types, *feature impact weights* were calculated after training the classifier on the complete
7 training set (all folds). The linear SVM learns weights associated with each feature, and it is common to use
8 the absolute value or squared value of these weights as measures of importance or impact in the classification
9 problem. However, since features were normalized to have zero mean, it was important to consider the sign
10 of feature values in assessing impact. Therefore we calculated feature impact for each class (ADHD vs.
11 TDC) by multiplying the learned feature weight by the mean value of that feature *within that class*. The
12 sign of the feature impact indicates whether those features were, on average, driving the classifier toward
13 that class diagnosis (for positive impacts) or away from that diagnosis (negative impacts).

14 RESULTS

15 The results detailed below describe the predictive power of individual classes of features (non-imaging phe-
16 notypic, anatomical, and network features) as well as the potential for improving predictions by combining
17 different feature classes. By grouping features in this way, we are able to examine the promise for different
18 avenues of diagnostic aids (e.g., structural features vs. functional network features, which necessitate ad-
19 ditional MR scan time). We describe a series of classifiers for each grouping of features, which were built
20 using the multiple feature selection methods described above, with different numbers of selected features.
21 All results describe both performance in cross-validation (testing on the fold not used for training) and on
22 the separate test set (training on all folds of the cross-validation / training set). In nearly all cases, classi-
23 fiers were able to predict the presence / absence of ADHD in individual subjects from the test set at above
24 chance levels as measured by the area under the ROC curve. **Unless otherwise specified, results are based**
25 **on a dataset consisting of 755 training examples, which is 21 fewer than the 776 provided in the original**
26 **data. These datasets were excluded due to processing or quality assurance problems in either our anatomical**
27 **stream (14 subjects) or our network stream (7 additional subjects).**

28 GENDER DIFFERENCES

29 Figure 2 shows a summary scatter plot of most of the non-imaging phenotypic information available for
30 subjects, plotted separately for males and females across all performance sites. These plots summarize age,
31 IQ, handedness, and diagnosis. By inspection it was clear that the ratio of control subjects to ADHD subjects
32 was different between genders, and that within each gender that ratio was non-uniform across ages. Based
33 on these determinations, we decided to stratify the dataset, treating each gender as a separate classification
34 problem. The results reported below, unless otherwise noted, used this approach, and results represent the
35 total performance levels across the gender-specific classifiers.

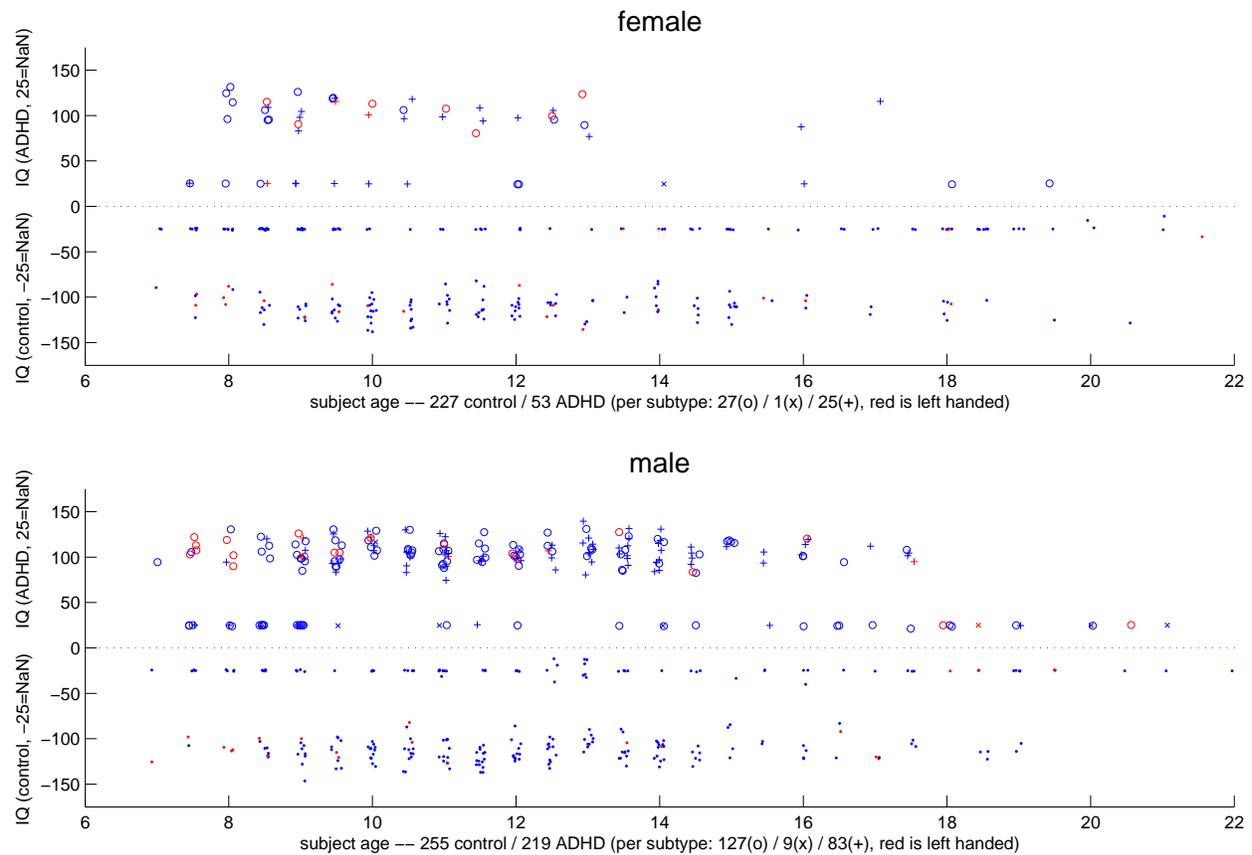


Figure 2: **Top:** Scatter plot representing female subjects from all sites, their DX status, age, combined IQ (average of the two available measures) and handedness. The x-axis corresponds to the subject age. Along the y-axis, typically developing control subjects are plotted below the $y = 0$ line, and ADHD subjects above (i.e., $|IQ| = IQ$, with negative values simply denoting TDC subjects). The position along the y-axis corresponds to the average IQ measure, with subjects whose records did not include IQ plotted around $y = -25$ and $y = 25$. The marker for each subject corresponds to diagnosis status (dot is control, o is ADHD-combined type, x is ADHD-hyperactive/impulsive, + is ADHD-inattentive), and its color to handedness (red is left-handed). Coordinates are jittered slightly to improve visualization. **Bottom:** Same as the plot above, but for male subjects.

1 PREDICTING ADHD FROM NON-IMAGING PHENOTYPIC FEATURES

2 The non-imaging phenotypic feature set (age, gender, handedness, verbal and performance IQ, and binary
 3 site variables) provided substantial predictive power. Only 14 phenotypic features were available, and the
 4 results for predicting ADHD diagnosis using only these features (or subsets of these features) are described in
 5 Table 1. Performance according to the AUC measure is well above chance on the folds not used for training
 6 during cross-validation, even for 5 features, and reaches a level of $AUC \approx 0.81$ (maximum possible value of
 7 1.0 for a perfect binary classifier) using all available phenotype features. Performance is similar for the three
 8 different feature selection methods. The AUC drops considerably, however, on the separate test set (when
 9 trained on the complete training set), with maximum values of ~ 0.72 for each feature selection method.

10 PREDICTING ADHD FROM ANATOMICAL FEATURES

11 The anatomical feature set included cortical thicknesses at uniformly sampled locations in the spherical
 12 atlas space, average overall cortical thickness, and volumes of individual cortical and subcortical structures.

Table 1: **Non-imaging phenotypic features only.** Results summarizing ADHD prediction using non-imaging phenotypic features only. Entries indicate the area under the ROC curve (AUC) for classifiers built using different feature selection methods (rows) and different numbers of features (columns). *Top*: Results on leave-out folds during cross-validation. *Bottom*: Results on separate test set based on training across all examples in the training / cross-validation set.

Cross-Validation	5	10	All (14)
2 sample t-test	0.70	0.78	0.81
Nested CV	0.72	0.76	0.81
Recursive FE	0.73	0.80	0.81
Test Set			
2 sample t-test	0.71	0.71	0.72
Nested CV	0.70	0.70	0.72
Recursive FE	0.69	0.66	0.72

1 Using these features alone, classifier performance is again substantially above chance and comparable to
2 the baseline performance established by using the non-imaging phenotypic features only. These results are
3 summarized in Table 2. Maximum performance on the cross-validation folds is achieved using the full feature
4 set ($N = 5081$ anatomical features), with $AUC \approx 0.77$, slightly lower than the maximum of ~ 0.81 achieved in
5 cross-validation for the non-imaging phenotypic features. Performance again drops somewhat to $AUC \approx 0.74$
6 on the separate test set. However, it should be noted that this value indicates slightly *better* generalization
7 of these classifiers to new test subjects than is observed for classifiers based on the non-imaging phenotypic
8 features only (see Table 1).

Table 2: **Anatomical features only.** Results summarizing ADHD prediction using anatomical features only. Entries indicate the area under the ROC curve (AUC) for classifiers built using different feature selection methods (rows) and different numbers of features (columns). *Top*: Results on leave-out folds during cross-validation. *Bottom*: Results on separate test set based on training across all examples in the training / cross-validation set.

Cross-Validation	10	20	50	100	200	400	800	1200	2000	3000	4000	All (5081)
2 sample t-test	0.64	0.67	0.63	0.58	0.62	0.69	0.71	0.72	0.74	0.76	0.77	0.77
Nested CV	0.62	0.64	0.65	0.63	0.62	0.64	0.68	0.70	0.75	0.77	0.76	0.77
Recursive FE	0.64	0.63	0.63	0.67	0.69	0.72	0.74	0.76	0.77	0.77	0.77	0.77
Test Set												
2 sample t-test	0.74	0.76	0.73	0.69	0.61	0.61	0.68	0.70	0.70	0.72	0.74	0.74
Nested CV	0.67	0.70	0.74	0.67	0.59	0.70	0.71	0.70	0.71	0.71	0.73	0.74
Recursive FE	0.68	0.63	0.52	0.54	0.63	0.71	0.74	0.75	0.75	0.74	0.74	0.74

9 PREDICTING ADHD FROM NETWORK FEATURES

10 As described above, three types of functional connectivity networks were constructed based on the filtered
11 rs-fMRI time course data extracted from 116 brain regions in each subject. For each network (*Corr*, *SIC*, and
12 *Kappa*; see *Methods*), the complete set of network features were computed and provided to feature selection
13 and classification methods.

14 Table 3 summarizes the predictive power of the extracted network features alone for each of the three
15 network construction methods. By comparison with Tables 1 and 2 it is clear that network features do not,
16 on their own, achieve the same predictive power that either non-imaging phenotypic features or anatomical

1 features achieve.

Table 3: **Network features only.** Results summarizing ADHD prediction using network features only, as calculated from 3 different network construction methods. Entries indicate the area under the ROC curve (AUC) for classifiers built using different network construction methods (major groupings), feature selection methods (rows within each grouping), and different numbers of features (columns). Within each network construction method, *Top*: Results on leave-out folds during cross-validation; *Bottom*: Results on separate test set based on training across all examples in the training / cross-validation set.

<i>CORR NETWORKS</i>													
Cross-Validation	10	20	50	100	200	400	800	1200	2000	3000	4000	6000	All (7150)
2 sample t-test	0.64	0.66	0.64	0.58	0.57	0.60	0.63	0.65	0.65	0.66	0.66	0.67	0.67
Nested CV	0.64	0.63	0.60	0.58	0.56	0.56	0.62	0.64	0.63	0.65	0.66	0.67	0.67
Recursive FE	0.55	0.57	0.54	0.59	0.63	0.66	0.65	0.65	0.67	0.67	0.67	0.67	0.67
Test Set													
2 sample t-test	0.72	0.68	0.67	0.60	0.53	0.53	0.57	0.62	0.66	0.69	0.70	0.67	0.67
Nested CV	0.64	0.64	0.68	0.65	0.59	0.50	0.60	0.62	0.65	0.70	0.71	0.68	0.67
Recursive FE	0.58	0.59	0.53	0.56	0.60	0.64	0.65	0.63	0.65	0.65	0.66	0.67	0.67
<i>SIC NETWORKS</i>													
Cross-Validation	10	20	50	100	200	400	800	1200	2000	3000	4000	6000	All (7345)
2 sample t-test	0.63	0.64	0.62	0.58	0.64	0.67	0.69	0.70	0.71	0.72	0.72	0.74	0.74
Nested CV	0.65	0.67	0.67	0.62	0.62	0.65	0.67	0.68	0.71	0.73	0.73	0.73	0.74
Recursive FE	0.65	0.64	0.64	0.70	0.71	0.72	0.74	0.74	0.73	0.73	0.74	0.74	0.74
Test Set													
2 sample t-test	0.64	0.72	0.66	0.65	0.60	0.58	0.62	0.63	0.65	0.67	0.70	0.71	0.71
Nested CV	0.49	0.58	0.62	0.64	0.57	0.58	0.66	0.68	0.70	0.69	0.70	0.70	0.71
Recursive FE	0.66	0.67	0.61	0.61	0.66	0.71	0.71	0.72	0.71	0.71	0.72	0.71	0.71
<i>KAPPA NETWORKS</i>													
Cross-Validation	10	20	50	100	200	400	800	1200	2000	3000	4000	6000	All (7344)
2 sample t-test	0.66	0.64	0.61	0.61	0.60	0.61	0.61	0.62	0.62	0.63	0.64	0.66	0.66
Nested CV	0.59	0.60	0.62	0.61	0.60	0.61	0.60	0.62	0.62	0.65	0.64	0.64	0.66
Recursive FE	0.62	0.61	0.60	0.62	0.64	0.64	0.65	0.65	0.65	0.66	0.66	0.66	0.66
Test Set													
2 sample t-test	0.63	0.63	0.63	0.60	0.57	0.56	0.56	0.54	0.56	0.59	0.59	0.61	0.61
Nested CV	0.60	0.60	0.56	0.56	0.59	0.59	0.55	0.57	0.59	0.60	0.60	0.61	0.61
Recursive FE	0.56	0.59	0.51	0.57	0.55	0.59	0.58	0.59	0.60	0.61	0.62	0.61	0.61

2 Features extracted from Sparse regularized Inverse Covariance (SIC) networks performed best overall,
 3 yielding a maximal AUC in cross-validation of ~ 0.74 when using all network features, dropping to ~ 0.71
 4 when testing on the separate test set. Neither Correlation-based network features nor features derived from
 5 networks built from the Kappa statistic provided as much predictive power as the *SIC* networks, either
 6 in cross-validation or on the test set. Based on this result, only features derived from *SIC* networks were
 7 considered when combining feature sets as described below.

8 PREDICTING ADHD FROM COMBINED FEATURE SETS

9 The classification results obtained through using *all* of the available features from different combinations
 10 of feature sets are summarized in Figure 3, in which the x-axis is sorted by AUC on the *test set*. The

1 best performing classifier (all of which used stratification by gender, see *Methods*) used all features from
 2 all classes; in general combining feature types improved performance in cross-validation and particularly so
 3 on the test set. We note, however, that adding network features to the non-imaging phenotype features
 4 markedly reduced performance during cross-validation (0.81 vs. 0.76), but resulted in improved performance
 5 (0.74 vs. 0.72) on the separate test set when using all features from both classes. Indeed the difference
 6 between cross-validation performance and performance on the test set is largest when using the non-imaging
 7 phenotype features only. Tables 4 through 7 detail the results for the classifiers that used more than one
 8 feature set, for different numbers of features used, and for each of the three feature selection methods.

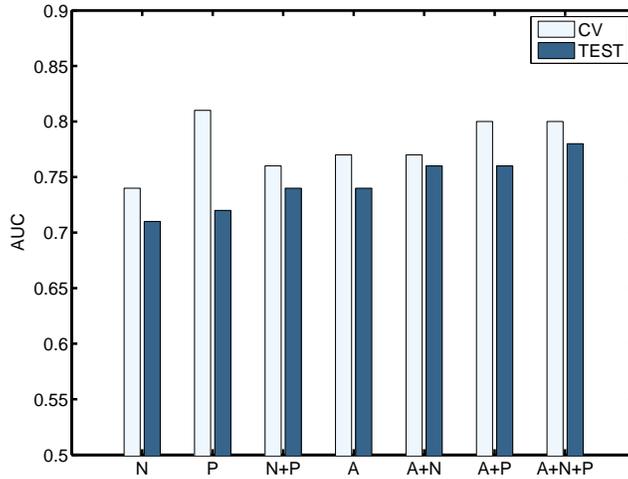


Figure 3: Summary of best performing classifiers using all features in the set or combination of sets below, sorted by performance on the test set (dark bars). Light bars indicate performance on non-training folds during cross-validation. A = Anatomical features, P = Non-imaging Phenotype features, and N = Network features.

Table 4: **Combined Non-imaging phenotype and anatomical features.** Results summarizing ADHD prediction using non-imaging phenotypic features combined with anatomical features (but not network features). Entries indicate the area under the ROC curve (AUC) for classifiers built using different feature selection methods (rows) and different numbers of features (columns). *Top*: Results on leave-out folds during cross-validation. *Bottom*: Results on separate test set based on training across all examples in the training / cross-validation set.

Cross-Validation	10	20	50	100	200	400	800	1200	2000	3000	4000	All (5095)
2 sample t-test	0.69	0.70	0.67	0.63	0.66	0.74	0.74	0.76	0.78	0.79	0.80	0.80
Nested CV	0.76	0.79	0.75	0.68	0.64	0.70	0.74	0.76	0.79	0.80	0.80	0.80
Recursive FE	0.60	0.61	0.56	0.63	0.73	0.77	0.78	0.79	0.80	0.80	0.80	0.80
Test Set												
2 sample t-test	0.69	0.66	0.65	0.63	0.66	0.69	0.71	0.70	0.73	0.74	0.75	0.76
Nested CV	0.74	0.74	0.70	0.66	0.65	0.69	0.72	0.71	0.73	0.76	0.75	0.76
Recursive FE	0.60	0.61	0.62	0.63	0.68	0.71	0.73	0.75	0.75	0.76	0.76	0.76

9 Combining all three feature sets yields the overall best prediction performance with an AUC ≈ 0.78 on
 10 the test set (~ 0.80 on leave-out folds in cross-validation). For these results, which are depicted in Table 7,
 11 we have also included classifier *accuracy* results. As noted above, these values are based on one point on the
 12 ROC curve, which optimizes total percent correct for the binary diagnosis. We see that overall accuracy is
 13 above chance, but – particularly on the test set – not at a level that would allow, at this stage, a confident
 14 binary diagnosis for all example subjects.

15 Figures 4 and 5 provide visualizations of the feature impact weights for *cortical thickness* features, which

Table 5: **Combined non-imaging phenotypic and network features.** Results summarizing ADHD prediction using non-imaging phenotypic features combined with SIC network features (but not anatomical features). Entries indicate the area under the ROC curve (AUC) for classifiers built using different feature selection methods (rows) and different numbers of features (columns). *Top*: Results on leave-out folds during cross-validation. *Bottom*: Results on separate test set based on training across all examples in the training / cross-validation set.

Cross-Validation	10	20	50	100	200	400	800	1200	2000	3000	4000	6000	All (7359)
2 sample t-test	0.68	0.69	0.66	0.60	0.66	0.70	0.71	0.72	0.73	0.75	0.75	0.76	0.76
Nested CV	0.77	0.80	0.72	0.67	0.68	0.69	0.69	0.70	0.73	0.75	0.75	0.75	0.76
Recursive FE	0.63	0.62	0.59	0.64	0.72	0.74	0.75	0.75	0.75	0.76	0.76	0.76	0.76
Test Set													
2 sample t-test	0.69	0.72	0.69	0.67	0.60	0.56	0.64	0.65	0.68	0.70	0.73	0.74	0.74
Nested CV	0.60	0.61	0.64	0.65	0.58	0.59	0.65	0.70	0.71	0.71	0.72	0.72	0.74
Recursive FE	0.67	0.67	0.61	0.63	0.72	0.73	0.72	0.74	0.76	0.75	0.75	0.74	0.74

Table 6: **Combined anatomical and network features.** Results summarizing ADHD prediction using anatomical features combined with SIC network features (but not non-imaging phenotypic features). Entries indicate the area under the ROC curve (AUC) for classifiers built using different feature selection methods (rows) and different numbers of features (columns). *Top*: Results on leave-out folds during cross-validation. *Bottom*: Results on separate test set based on training across all examples in the training / cross-validation set.

Cross-Validation	10	20	50	100	200	400	800	1200	2000	3000	4000	6000	All (12426)
2 sample t-test	0.63	0.64	0.61	0.60	0.62	0.66	0.71	0.72	0.75	0.75	0.76	0.77	0.77
Nested CV	0.62	0.64	0.65	0.63	0.62	0.64	0.68	0.70	0.75	0.77	0.76	0.77	0.77
Recursive FE	0.65	0.65	0.65	0.69	0.72	0.73	0.76	0.76	0.76	0.77	0.77	0.77	0.77
Test Set													
2 sample t-test	0.72	0.74	0.72	0.66	0.66	0.70	0.67	0.71	0.70	0.70	0.70	0.73	0.76
Nested CV	0.49	0.63	0.62	0.64	0.57	0.69	0.68	0.69	0.70	0.72	0.74	0.73	0.76
Recursive FE	0.74	0.73	0.69	0.67	0.69	0.70	0.72	0.72	0.76	0.75	0.77	0.77	0.76

made up the majority of the anatomical feature set. These figures show the feature impacts for the ADHD-positive class; equivalent visualizations for the control class are available as Supplementary Materials. Despite the strong predictive power observed in non-imaging phenotype features alone, we observed that the combination of anatomical and network features *without* inclusion of the non-imaging phenotype features yielded the second best performing classifier (tied with using anatomical and non-imaging phenotype features together) on new data, with an AUC ≈ 0.76 on the test set.

Table 7: Combination of all feature classes. Results summarizing ADHD prediction using combination of all feature types (non-imaging, anatomical, and network features extracted from SIC networks). Entries indicate the area under the ROC curve (AUC) for classifiers built using different feature selection methods (rows) and different numbers of features (columns). *Top*: Results on leave-out folds during cross-validation. *Bottom*: Results on separate test set based on training across all examples in the training / cross-validation set. The corresponding accuracy results are shown below.

AUC results													
Cross-Validation	10	20	50	100	200	400	800	1200	2000	3000	4000	6000	All (12440)
2 sample t-test	0.65	0.66	0.64	0.62	0.65	0.69	0.75	0.75	0.77	0.78	0.79	0.79	0.80
Nested CV	0.77	0.79	0.75	0.68	0.64	0.70	0.74	0.76	0.79	0.80	0.80	0.80	0.80
Recursive FE	0.66	0.64	0.63	0.66	0.75	0.77	0.78	0.78	0.79	0.80	0.80	0.80	0.80
Test Set													
2 sample t-test	0.72	0.74	0.76	0.70	0.66	0.73	0.70	0.73	0.71	0.72	0.72	0.76	0.78
Nested CV	0.60	0.66	0.64	0.63	0.62	0.69	0.67	0.69	0.71	0.74	0.75	0.76	0.78
Recursive FE	0.63	0.63	0.61	0.60	0.69	0.72	0.73	0.75	0.77	0.79	0.78	0.78	0.78
Accuracy results													
Cross-Validation	10	20	50	100	200	400	800	1200	2000	3000	4000	6000	All (12440)
2 sample t-test	0.68	0.65	0.63	0.63	0.64	0.67	0.69	0.69	0.74	0.72	0.72	0.73	0.74
Nested CV	0.74	0.74	0.70	0.66	0.65	0.69	0.72	0.71	0.73	0.76	0.75	0.74	0.74
Recursive FE	0.66	0.63	0.65	0.66	0.70	0.73	0.73	0.73	0.73	0.73	0.73	0.73	0.74
Test Set													
2 sample t-test	0.69	0.67	0.69	0.65	0.63	0.66	0.63	0.66	0.62	0.61	0.60	0.63	0.67
Nested CV	0.61	0.62	0.60	0.60	0.58	0.65	0.63	0.64	0.64	0.67	0.70	0.65	0.67
Recursive FE	0.62	0.62	0.62	0.64	0.64	0.64	0.67	0.65	0.66	0.68	0.68	0.68	0.67

7 STRATIFICATION

In order to evaluate the importance of stratifying by gender (which was used throughout our approach), classification using features from all feature sets (non-imaging, anatomical, and network) was repeated *without any stratification*. These results are shown in Table 8. Here we observed a dropoff in performance in cross-validation, but particularly a relatively large dropoff in AUC on the test set (AUC ~ 0.70 using all features) in comparison with results that used separate gender-specific classifiers. Figure 6 demonstrates that the improvement due to stratification is systematic, plotting AUC on the test set for a range of numbers of selected features, and showing that stratification improves performance for the large majority of cases.

In addition Figure 6 plots the results obtained from *randomly* ranking features, then training classifiers (separately for boys and girls) using the top N features from the random list (in which rankings were held constant across folds). This plot is an average across 10 randomly generated rankings. The results show that our data-driven methods for selecting N features improve performance over simply choosing features at random (for all N), but interestingly, using gender-specific classifiers with *randomly selected features* outperforms “intelligent” feature selection without gender stratification when the number of features used

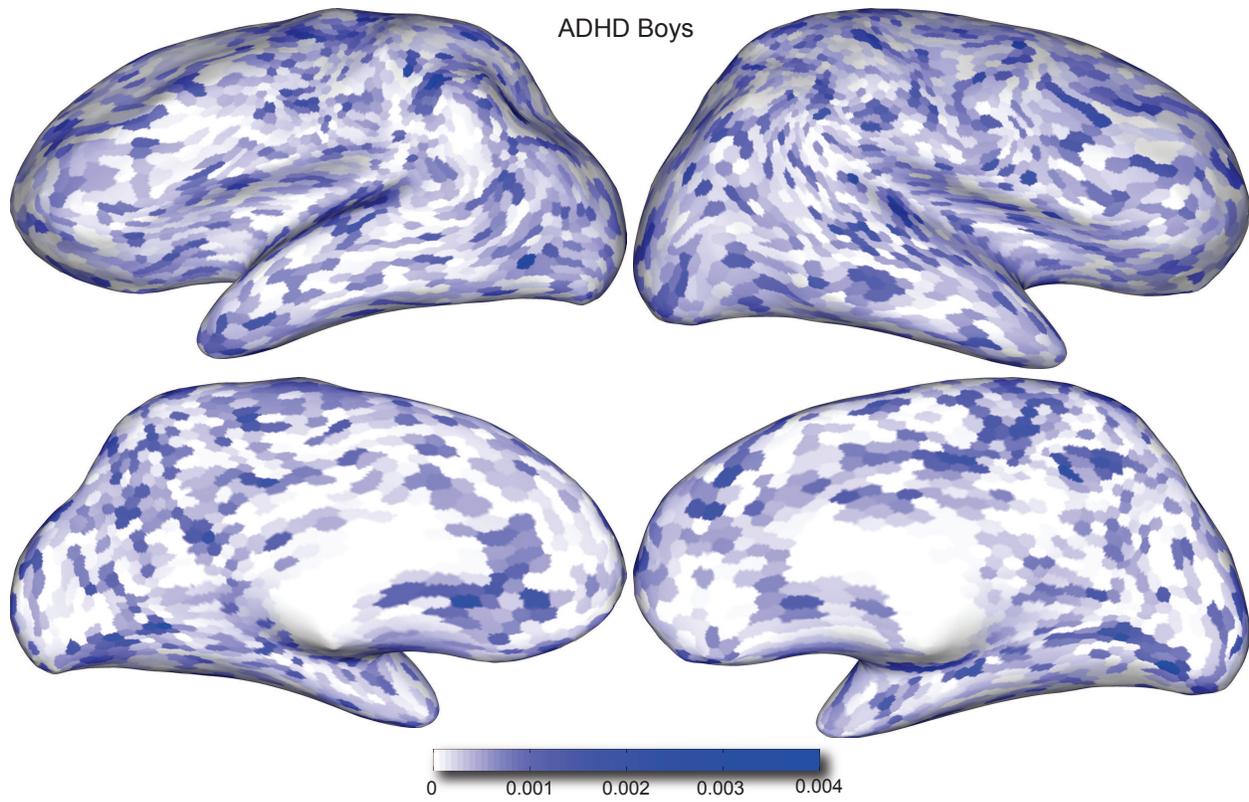


Figure 4: Classification feature impact weights for cortical thickness features in the classification of boys with positive ADHD diagnosis. Darker blue values indicate higher impact weights.

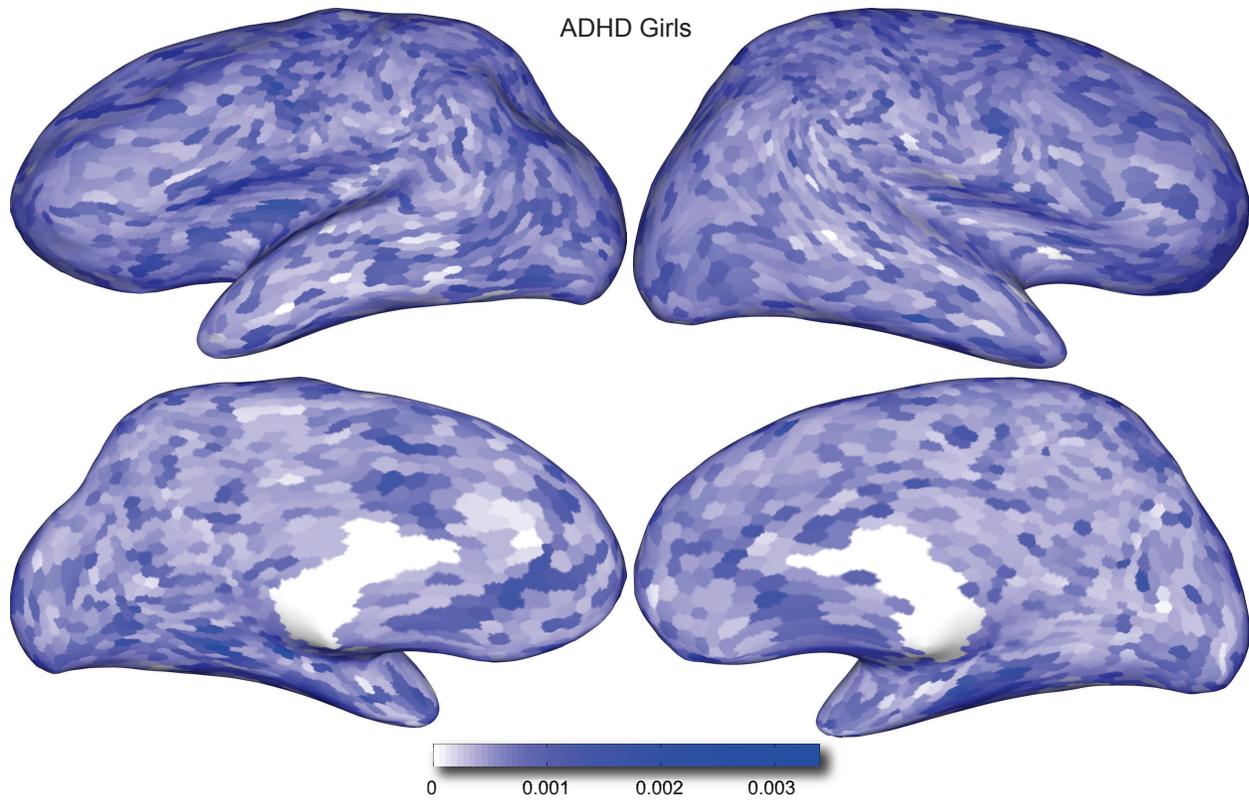


Figure 5: Classification feature impact weights for cortical thickness features in the classification of girls with positive ADHD diagnosis. Darker blue values indicate higher impact weights.

1 is large.

Table 8: **Combination of all feature classes without gender stratification.** Results summarizing ADHD prediction using combination of all feature types (non-imaging, anatomical, and network features extracted from SIC networks) *without stratification by gender*. Entries indicate the area under the ROC curve (AUC) for classifiers built using different feature selection methods (rows) and different numbers of features (columns). *Top*: Results on leave-out folds during cross-validation. *Bottom*: Results on separate test set based on training across all examples in the training / cross-validation set.

Cross-Validation	10	20	50	100	200	400	800	1200	2000	3000	4000	6000	All (12440)
2 sample t-test	0.67	0.68	0.64	0.63	0.66	0.68	0.69	0.71	0.73	0.74	0.74	0.75	0.77
Nested CV	0.77	0.80	0.77	0.73	0.62	0.64	0.67	0.67	0.71	0.74	0.74	0.77	0.77
Recursive FE	0.63	0.64	0.65	0.67	0.68	0.71	0.73	0.73	0.75	0.76	0.76	0.77	0.77
Test Set													
2 sample t-test	0.71	0.73	0.65	0.69	0.66	0.66	0.65	0.63	0.64	0.68	0.69	0.70	0.70
Nested CV	0.60	0.65	0.66	0.68	0.65	0.50	0.60	0.61	0.65	0.63	0.62	0.66	0.70
Recursive FE	0.64	0.63	0.61	0.60	0.60	0.66	0.68	0.69	0.71	0.69	0.70	0.70	0.70

2 Because stratifying by gender was quite powerful, we tested adding an additional level of stratification,
 3 by age. In this case, classifiers were constructed for three different age groups (the 0th-25th percentile, 25th-
 4 75th percentile, and 75th-100th percentile of age) within each gender, thus resulting in 6 total classifiers.
 5 The results of applying this method to the full set of features are given in Table 9. **Cross-validation per-**
 6 **formance was qualitatively similar compared to stratification by gender alone, but performance on the test**
 7 **set was slightly improved, reaching AUC \approx 0.80 when using all features (which was slightly better than the**
 8 **performance predicted in cross-validation – AUC \approx 0.79).**

Table 9: **Combination of all feature classes with stratification by gender and age.** Results summarizing ADHD prediction using combination of all feature types (non-imaging, anatomical, and network features extracted from SIC networks) *with stratification by gender and age*. Entries indicate the area under the ROC curve (AUC) for classifiers built using different feature selection methods (rows) and different numbers of features (columns). *Top*: Results on leave-out folds during cross-validation. *Bottom*: Results on separate test set based on training across all examples in the training / cross-validation set.

Cross-Validation	10	20	50	100	200	400	800	1200	2000	3000	4000	6000	All (12440)
2 sample t-test	0.63	0.64	0.61	0.66	0.68	0.72	0.75	0.75	0.77	0.77	0.78	0.78	0.79
Nested CV	0.75	0.77	0.68	0.63	0.67	0.72	0.75	0.77	0.79	0.80	0.80	0.79	0.79
Recursive FE	0.62	0.62	0.67	0.73	0.75	0.76	0.77	0.78	0.78	0.78	0.78	0.79	0.79
Test Set													
2 sample t-test	0.72	0.71	0.67	0.68	0.67	0.66	0.71	0.73	0.76	0.77	0.79	0.79	0.80
Nested CV	0.67	0.65	0.67	0.62	0.67	0.69	0.70	0.72	0.75	0.76	0.77	0.78	0.80
Recursive FE	0.59	0.55	0.64	0.64	0.72	0.76	0.78	0.80	0.80	0.80	0.80	0.80	0.80

9 PREDICTION RESULTS BY SUBTYPE

10 In the present study we did not focus on predicting ADHD subtypes. However, most patients in the provided
 11 dataset were categorized as either ADHD-Combined Type (DX 1) or ADHD-Inattentive Type (DX 3), and we
 12 were interested in the question of whether one of these subtypes was easier to distinguish from controls than
 13 the other. We thus calculated performance of the classifier using all available features (e.g., as presented in
 14 Table 7), over two different datasets, each containing all the controls and patients from one of the subtypes.

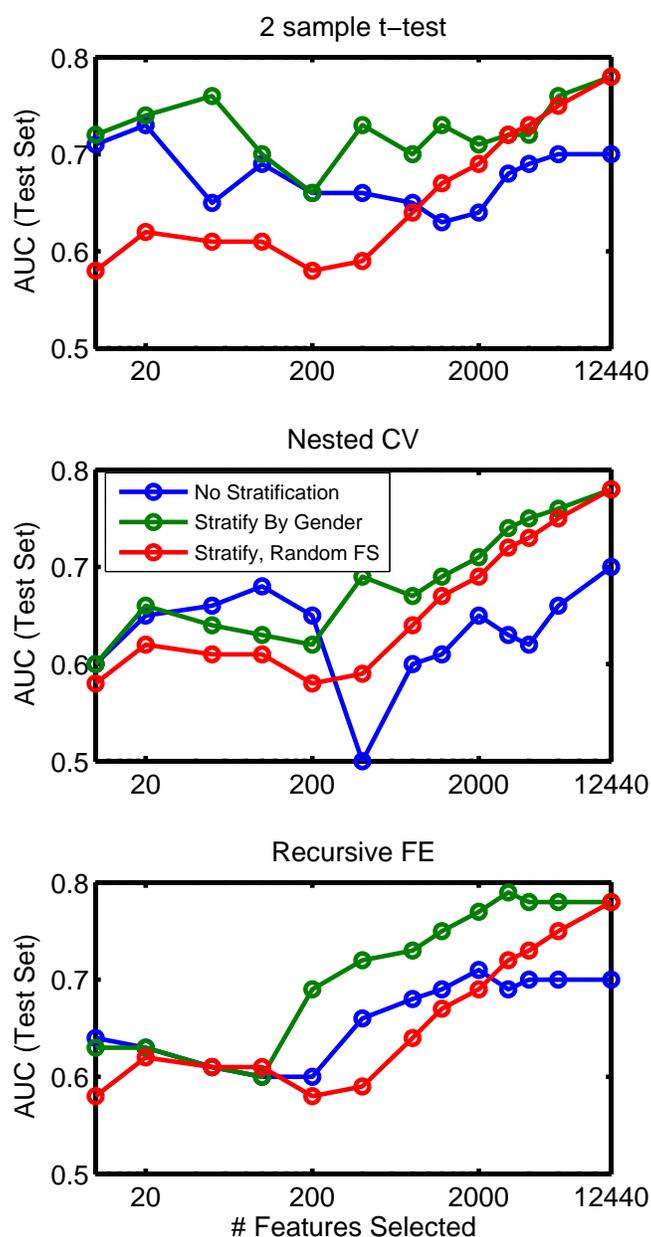


Figure 6: Comparison of AUC performance on the test set as a function of number of features used for classifiers built from a combination of all feature classes, with (green line) and without (blue line) stratification by gender. Each subplot represents results based on a different feature selection method. Also shown (red line, same in all plots) is the result obtained when features are ranked randomly (separately for each stratum); this result is averaged over classifiers built using 10 random feature rankings.

1 Figure 7 shows the results on the test set for each subtype, using each feature selection method, as a
 2 function of the number of features selected; each color is associated with one of the methods, with solid and
 3 dashed lines corresponding to ADHD-Combined and ADHD-Inattentive results, respectively. Regardless of
 4 the feature selection method used, AUC performance is almost always higher in the task of distinguishing
 5 ADHD-Combined from controls than it is for distinguishing ADHD-Inattentive from the same controls.

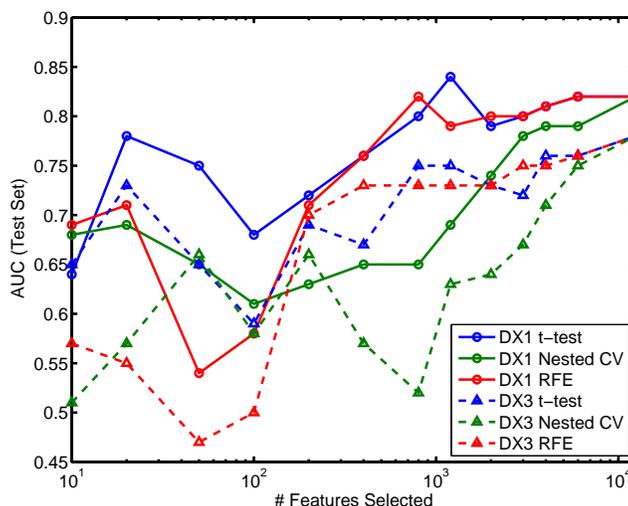


Figure 7: Comparison of AUC performance on test set as a function of number of features used for the two predominant ADHD subtypes. DX1 is the ADHD-Combined subtype; DX3 is the ADHD-Inattentive subtype.

1 SIGNIFICANCE OF CLASSIFICATION RESULTS

2 The various sets of results are reported in the above tables without an explicit measure of statistical signifi-
 3 cance. There is no simple analytical model for the distribution of the AUC under the null hypothesis that the
 4 classifier is performing at chance, as there is for accuracy. One of the main goals of this paper is to contrast
 5 the results obtained using different feature types, and also the extent to which information is present across
 6 many features and, thus, we believe it is useful to report these results comprehensively. This poses a multiple
 7 comparisons problem, however, which is further complicated by the fact that the results obtained with each
 8 feature ranking method – a table row – are potentially correlated, since the top 20 features contain the top
 9 10, the top 50 contain the top 20, and so on. Given that we do want to report results comprehensively, the
 10 canonical solution of using cross-validation within the training set to determine a single number of features
 11 to use is not desirable.

12 Perhaps the most elegant approach to test significance for an entire table of results would be to use
 13 permutation testing (Golland and Fischl, 2003). Unfortunately, given that each experimental run, using
 14 nested cross-validation for feature selection, is computationally expensive, it is impractical to run a sufficiently
 15 large number of permutations for each of the experimental conditions. Here we have opted to perform a
 16 smaller number ($N = 100$) of permutations of the results in Table 7 (which describes the best performing
 17 classifiers based on using all feature types), and we report the mean AUC and standard deviation of the
 18 estimate of the mean AUC in Table 10, as well as the analogous numbers for accuracy in Table 11. In each
 19 training set, the category labels used to select features and train a classifier were permuted, separately for
 20 each fold in the cross-validation case; in this manner, the balance of examples in each class was maintained,
 21 as was the stratification approach. It is important to note that we should expect permutation test accuracy
 22 results to be somewhat above 0.5 because there are unequal numbers of examples in each class.

23 These results, in comparison with Table 7, demonstrate that our best performing classifiers are performing
 24 well above empirically defined chance levels. For example, the mean AUC and accuracy (ACC) scores for
 25 classifiers using all 12440 features under class label permutations is $AUC = 0.602 \pm 0.006$ and $ACC =$
 26 0.626 ± 0.002 under cross-validation, compared to an observed values of $AUC = 0.80$ and $ACC = 0.74$ for
 27 the classifier trained with correct class labels. This large difference holds up in the test set as well, where the
 28 empirical chance values are $AUC = 0.60 \pm 0.01$ and $ACC = 0.582 \pm 0.003$ and the correctly trained classifier

Table 10: **Permutation tests of Table 7 results (AUC)**. Results summarizing ADHD prediction using combination of all feature types (non-imaging, anatomical, and network features extracted from SIC networks) *with stratification by gender*. Entries indicate the area under the ROC curve (AUC) for classifiers built using different feature selection methods (rows) and different numbers of features (columns). *Top*: Results on leave-out folds during cross-validation. *Bottom*: Results on separate test set based on training across all examples in the training / cross-validation set.

Cross-Validation	10	20	50	100	200	400	800	1200	2000	3000	4000	6000	All (12440)
mean across 100 permutations													
2 sample t-test	0.581	0.575	0.557	0.536	0.525	0.542	0.560	0.572	0.582	0.589	0.594	0.598	0.602
Nested CV	0.593	0.583	0.561	0.542	0.525	0.543	0.558	0.564	0.574	0.583	0.585	0.589	0.602
Recursive FE	0.565	0.552	0.530	0.563	0.583	0.593	0.598	0.598	0.600	0.601	0.601	0.602	0.602
standard deviation of mean across 100 permutations													
2 sample t-test	0.005	0.004	0.003	0.003	0.003	0.003	0.004	0.004	0.005	0.005	0.005	0.005	0.006
Nested CV	0.006	0.005	0.004	0.003	0.003	0.003	0.004	0.004	0.004	0.005	0.005	0.005	0.006
Recursive FE	0.004	0.004	0.003	0.004	0.005	0.005	0.005	0.005	0.005	0.006	0.006	0.006	0.006
Test Set	10	20	50	100	200	400	800	1200	2000	3000	4000	6000	All (12440)
mean across 100 permutations													
2 sample t-test	0.566	0.580	0.565	0.549	0.527	0.529	0.544	0.554	0.562	0.569	0.579	0.586	0.600
Nested CV	0.563	0.553	0.549	0.534	0.522	0.529	0.543	0.554	0.564	0.577	0.587	0.592	0.600
Recursive FE	0.562	0.561	0.543	0.527	0.554	0.570	0.584	0.589	0.594	0.596	0.599	0.599	0.600
standard deviation of mean across 100 permutations													
2 sample t-test	0.010	0.009	0.008	0.006	0.005	0.005	0.006	0.006	0.007	0.008	0.008	0.009	0.010
Nested CV	0.007	0.006	0.006	0.006	0.005	0.005	0.006	0.006	0.007	0.008	0.009	0.009	0.010
Recursive FE	0.008	0.008	0.007	0.006	0.007	0.008	0.009	0.009	0.009	0.010	0.010	0.010	0.010

Table 11: **Permutation tests of Table 7 results (accuracy)**. Results summarizing ADHD prediction using combination of all feature types (non-imaging, anatomical, and network features extracted from SIC networks) *with stratification by gender*. Entries indicate the accuracy for classifiers built using different feature selection methods (rows) and different numbers of features (columns). *Top*: Results on leave-out folds during cross-validation. *Bottom*: Results on separate test set based on training across all examples in the training / cross-validation set.

Cross-Validation	10	20	50	100	200	400	800	1200	2000	3000	4000	6000	All (12440)
mean across 100 permutations													
2 sample t-test	0.630	0.620	0.573	0.564	0.575	0.591	0.606	0.613	0.617	0.620	0.621	0.622	0.626
Nested CV	0.626	0.623	0.583	0.565	0.577	0.595	0.605	0.610	0.616	0.621	0.621	0.621	0.626
Recursive FE	0.592	0.593	0.609	0.615	0.618	0.620	0.625	0.625	0.625	0.624	0.626	0.626	0.626
standard deviation of mean across 100 permutations													
2 sample t-test	0.002	0.001	0.002	0.002	0.002	0.002	0.002	0.002	0.002	0.002	0.002	0.002	0.002
Nested CV	0.002	0.003	0.002	0.002	0.002	0.002	0.002	0.002	0.001	0.002	0.002	0.002	0.002
Recursive FE	0.002	0.002	0.002	0.002	0.002	0.002	0.002	0.002	0.002	0.002	0.002	0.002	0.002
Test Set	10	20	50	100	200	400	800	1200	2000	3000	4000	6000	All (12440)
mean across 100 permutations													
2 sample t-test	0.557	0.569	0.570	0.541	0.544	0.558	0.565	0.570	0.575	0.580	0.582	0.586	0.582
Nested CV	0.558	0.553	0.540	0.529	0.532	0.536	0.556	0.561	0.565	0.575	0.577	0.581	0.582
Recursive FE	0.557	0.555	0.562	0.568	0.575	0.584	0.583	0.582	0.579	0.580	0.581	0.582	0.582
standard deviation of mean across 100 permutations													
2 sample t-test	0.002	0.003	0.004	0.004	0.004	0.004	0.003	0.003	0.003	0.003	0.003	0.003	0.003
Nested CV	0.003	0.003	0.003	0.004	0.004	0.003	0.004	0.004	0.003	0.004	0.003	0.003	0.003
Recursive FE	0.004	0.004	0.004	0.003	0.003	0.003	0.003	0.003	0.004	0.003	0.003	0.003	0.003

1 values are $AUC = 0.74$ and $ACC = 0.67$.

2 PREDICTION RESULTS BY SITE

3 We examined the performance of the *best classifier* (with stratification by gender only, based on using all
4 features from all feature classes) as a function of the *site of performance* of the imaging session. Here we
5 used accuracy as the outcome measure, and separated the corresponding errors into false positives and false
6 negatives. These results represent a single point from an ROC curve. Our goal was to show the variability
7 in performance across sites and the types of errors made within each site. These results are depicted in
8 Table 12 for both cross-validation and on the separate test set.

Table 12: Results of best overall accuracy predictor by site for predictions based on the full set of combined features.

Cross-Validation		Site	FP error	FN error	accuracy	# Subjects
1		Peking University	0.07	0.18	0.75	189
3		Kennedy Krieger Institute	0.10	0.22	0.69	83
4		NeuroIMAGE Sample	0.10	0.27	0.62	48
5	New York University Child Study Center		0.14	0.18	0.68	213
6	Oregon Health & Science University		0.07	0.37	0.56	73
7	University of Pittsburgh		0.03	0.00	0.97	89
8	Washington University in St. Louis		0.02	0.00	0.98	59
Test Set		Site	FP error	FN error	accuracy	# Subjects
1		Peking University	0.06	0.32	0.62	50
3		Kennedy Krieger Institute	0.00	0.18	0.82	11
4		NeuroIMAGE Sample	0.00	0.35	0.65	23
5	New York University Child Study Center		0.10	0.22	0.68	41
6	Oregon Health & Science University		0.09	0.18	0.74	34
7	University of Pittsburgh		0.11	0.44	0.44	9
8	Washington University in St. Louis		N/A	N/A	N/A	0

9 DISCUSSION

10 The above results demonstrate that it is possible to predict the diagnosis of Attention-Deficit/Hyperactivity
11 Disorder within the set of available subjects to a level of certainty that is well above what would be expected
12 by chance (naively $AUC \approx 0.5$, but empirically evaluated under certain conditions as shown in Table 10)
13 using many combinations of non-imaging phenotype features and/or MRI-based anatomical or resting-state
14 functional network features. We took a systematic approach to testing the predictive power of each of the
15 three main feature sets on their own (Tables 1-3) and then in combination (Tables 4-7). The results using
16 all features from those sets are summarized in Figure 3. In general, the small set of non-imaging phenotypic
17 features were able to provide a large fraction of the available predictive power. However, imaging features,
18 when used in large numbers, were able to boost performance by improving generalization to a separate
19 test set of subjects. Such anatomical and rs-fMRI network features were also able to provide very strong
20 predictions on their own, while using only gender as a non-imaging basis for stratifying the classification
21 problem.

22 In this paper we measured prediction performance using the area under the ROC curve (AUC) rather than
23 the simpler measure of *accuracy* (i.e., rate of correct classification) because accuracy can be a misleading
24 measure under certain circumstances. In particular, when classes are unbalanced, as is often the case in
25 clinical diagnostics (i.e., there will be many more control cases than disorder cases), high accuracy can be

1 achieved simply by biasing the classifier heavily toward the control class. ROC curves thus are an important
2 approach for evaluating clinical tests (Zweig and Campbell, 1993) because they plot the true positive rate
3 against the false positive rate, establishing the area under the ROC curve as a measure of discrimination
4 between classes across the full operating range of the classifier. By comparison, the use of accuracy as the
5 measure of performance in this two-class problem forces evaluation of the classifier at a particular threshold,
6 above which one call is made, and below which the other is made. In a real-world clinical setting it may be
7 useful to *tune* the classifier to highlight either *sensitivity* to detect true cases of ADHD at the expense of a
8 higher false positive rate, or to increase *specificity* to reduce false positive rate at the expense of a higher
9 false negative rate. Thus AUC provides a more general measure that is applicable in the current context.
10 **Still, because accuracy is a well-recognized and commonly reported performance measure, we have included**
11 **such results in the case of classifiers built from all feature types (Table 7).**

12 STRATIFICATION BY GENDER

13 Based on current knowledge of ADHD, there are compelling reasons to believe that diagnostic aids might
14 perform better if tailored to boys separately from girls, and this is reflected in our results. First, ADHD is
15 diagnosed at a significantly higher rate in boys than in girls (Polanczyk et al., 2007), and thus the overall
16 probability of an ADHD-positive diagnosis – averaged across other factors – should be different between
17 genders. Still, it is possible that the same *features* (and weightings of those features) are important in
18 diagnosing ADHD in both genders, and that those features are simply present in different proportions
19 between genders. If this were the case, then a single classifier trained using those features should perform
20 equally well compared with gender-specific classifiers, but this was not the case. Tables 7 and 8 show
21 results using all feature sets with and without stratification by gender. The performance is notably reduced,
22 particularly in the test set, using a single classifier. Figure 6 demonstrates that this stratification effect is
23 quite consistent for different numbers of selected features.

24 These results are consistent with the literature suggesting neurobehavioral gender differences in individ-
25 uals with ADHD that go beyond mere prevalence (Gaub and Carlson, 1997; Newcorn et al., 2001; Gershon
26 and Gershon, 2002; Castellanos et al., 2001; Mahone and Wodka, 2008). In the present study, the gender-
27 specific classifiers that combined all non-imaging phenotypic, anatomical, and network features weighted
28 the available imaging features substantially differently. **For the classifiers using the entire feature set, the**
29 **Pearson correlation coefficient between the feature weight vector for boys and the feature weight vector for**
30 **girls was $r = 0.0757$.** Although small, this represented a significant correlation in this high dimensional space
31 ($p < 1 \times 10^{-13}$). Qualitative comparison of the feature impact weights points to many differences (see Sup-
32 plementary Materials for a full list of features and their impact weights for each class, with local attributes
33 linked to anatomy through the AAL nomenclature (Tzourio-Mazoyer et al., 2002) for network features or
34 the nomenclature of the Desikan et al. (2006) FreeSurfer parcellation system). To provide one such example,
35 highly impactful features for boys with ADHD include a series of network edge weights connecting cortical
36 and striatal structures (5 of the top 50 ranked features) whereas none of these features appear in the top 50
37 for girls with ADHD.

38 An additional level of stratification was tested for the classifiers that used all feature sets, in this case
39 including sub-classifiers for age groups. These groups were defined by percentiles (0-25, 25-75, 75-100) within
40 the age range spanned by the training set within each gender group. While this set of stratified classifiers
41 yielded quantitatively the *best* overall performance (AUC ~ 0.79 in both cross-validation and on the separate
42 test set), we focused more here on gender-only stratification because further splitting of the dataset results in
43 relatively small sample-size partitions, and we lack overall confidence that such results will continue to hold
44 over new, larger datasets. Still, stratification by age is a sensible approach given sufficient data given the

1 considerable developmental changes that are occurring over the window of ages considered here, including
2 the maturation of resting state functional connectivity networks in typically developing children (Fair et al.,
3 2008).

4 PHENOTYPIC FEATURES

5 Non-imaging phenotype features included gender, age, handedness (which was coded as an integer – left /
6 right / ambidextrous – for most sites, but as a decimal based on the Edinburgh Handedness Inventory for
7 patients scanned at NYU) and performance and verbal IQ (which were missing for some patients and were
8 thus replaced with the population mean value when training). In addition, binary features were created to
9 indicate the absence of an IQ score as well as eight features to indicate at which of the eight possible sites
10 the scan data were obtained.

11 These features had substantial predictive power and provided a high baseline for image feature-based
12 classifiers to surpass. The maximal performance using these non-imaging attributes only was achieved using
13 all 14 available phenotypic features (both in cross-validation and on the separate test set), but performance
14 dropped off substantially from the 0.81 AUC obtained in cross-validation to 0.72 on the separate test set.
15 This reduction in performance may indicate that the distributions of these phenotypic features in the test set
16 did not sufficiently mirror those found in the training set. Overall, the population of patients and typically
17 developing controls for whom data were made available were not necessarily representative of the population
18 at large. For example, ~ 36% of training samples had an ADHD-positive diagnosis, and the male-to-female
19 ratio within that sample was not reflective of overall diagnosis rates. While the dataset used here was
20 *large* relative to many single studies of ADHD, the numbers are still potentially too small given the overall
21 population variance in measures like IQ to enable completely effective prediction based on this small set of
22 features.

23 IMAGING FEATURES

24 A large number of features were calculated based on each patient’s imaging data. These measures were
25 subdivided into *anatomical features*, which were processed in one stream, and *network features*, which were
26 processed in another (see Figure 1).

27 ***Anatomical features.*** The anatomical feature set was dominated by estimates of cortical thickness
28 at 2,562 locations per hemisphere, a small fraction of which were excluded because they represented non-
29 cortical areas and were set to be constant across subjects. The cortical thickness locations used corresponded
30 to the vertices comprising an icosahedral approximation to a sphere in the FreeSurfer spherical atlas space
31 that was used for anatomical inter-subject registration. This yielded a relatively low-resolution resampling
32 of each subject’s thickness data, which was performed in order to reduce dimensionality to accommodate
33 machine learning techniques, to improve estimates of local thickness by effectively averaging over small
34 neighborhoods, and to allow for the expected small registration errors without destroying subject-to-subject
35 feature correspondences.

36 Cortical thickness has previously been shown to be a relevant biomarker in ADHD (e.g., Shaw et al.,
37 2006; Makris et al., 2007). Many cortical thickness values were assigned relatively large feature weights (see
38 Supplementary Materials) in the classifiers built using all features from all feature sets; the *feature impact*
39 *weights* (see *Methods*) for ADHD-positive diagnosis for boys (Figure 4) and girls (Figure 5) were rendered
40 on inflated cortical surface models for purposes of visualization. While some of the particular nodes assigned
41 high impact are in areas previously shown to have differences in cortical thickness in subjects with ADHD,

1 in general the thickness impacts are distributed across cortex and not focused in one or a few brain areas,
2 and thus are difficult to interpret succinctly in a biological context. Within the group of cortical thickness
3 features with impacts at least two standard deviations above the median (i.e., those that were particularly
4 important in diagnosis), $\sim 53\%$ were in the right hemisphere in boys whereas only $\sim 42\%$ were in the right
5 hemisphere in girls.

6 The caudate nucleus has been considered a structure of interest in ADHD due to reports of volume
7 differences in patients (Castellanos et al., 1994). In our results, left and right caudate volume (normalized
8 by intracranial volume) each were assigned high impact for ADHD diagnosis in girls (237th and 345th
9 most highly impactful features, respectively, each of which represented an impact that was more than two
10 standard deviations above the median feature impact weight for that class), but not in boys. We did not
11 directly include measures of *asymmetry* (though in principle this could be computed implicitly by the existing
12 classifier), which may be particularly useful, as it has been claimed that the degree of caudate asymmetry
13 may reflect severity of ADHD symptoms in children (Schrimsher et al., 2002). The only other structure
14 volume feature to receive similarly high impact was the volume of the right inferior lateral ventricle in boys,
15 which ranked 450th in impact for the control class and 506th in impact for the ADHD-positive class. This
16 feature did not rank highly for girls.

17 **Network features.** Our corpus of network-based features included a variety of standard features found
18 commonly in functional connectivity studies as well as a series of features drawn from other scientific domains
19 or developed for this study. These included measures of small-worldness and scale-free properties, node-level
20 connectivity or other measures of “importance” (using single path and parallel path measures), network
21 separability, network sparsity, and network motifs. This broad base of network feature types allowed us
22 to systematically examine the power of features by exploring which were selected and/or had large impact
23 on the classification problem. Plots showing the mean and standard deviation feature impacts for each
24 gender, calculated across *categories* of features (including both network and anatomical features) based on
25 the classifier built using all features of all types is available in the Supplementary Materials.

26 We calculated the set of network features after building rs-fMRI functional connectivity networks using
27 three different methods: correlation (*Corr*), the inverse of the covariance matrix under L1-norm regulariza-
28 tion to promote sparsity (*SIC*), and Patel’s Kappa measure (*Kappa*). These network construction methods
29 were chosen largely based on the results of Smith et al. (2011), as they were among the top performing
30 methods in simulation studies for correctly inferring underlying connectivity (without direction). In the
31 present results (see Table 3) we observed a rather clear advantage for using network measures derived from
32 the SIC networks relative to the other two methods; this advantage was predicted by better overall results
33 throughout cross-validation (particularly for large numbers of features), as well as better results on the sep-
34 arate test set ($AUC \approx 0.71$ vs. 0.67 for *Corr* networks and 0.61 for *Kappa* networks, when using all network
35 features).

36 Many network features were among those that had high impact weights for classifiers that used all features
37 (with stratification by gender). Because each edge weight (of 6670 possible edges in the AAL network) was
38 used as a feature, these tended to dominate lists of highly ranked features, but other derived network
39 properties were also impactful. For convenience, we again examined impact weights that were greater than
40 two standard deviations from the mean for that class. While individual features are difficult to interpret, a
41 broad picture emerges pointing to a role for functional connectivity attributes involving multiple cerebellar
42 areas and primarily frontal cortical areas, especially in the right hemisphere.

43 Beyond individual edge weights, weighted node degrees (which are a measure of the overall estimated
44 functional connectivity of that brain area with all other regions), and node betweenness (which can be
45 thought of as a node’s *importance* if the network represents flow of information or activity) of frontal cortical

1 and cerebellar nodes often had high impact. The second most impactful feature for diagnosis of girls with
2 ADHD was the degree of the left median cingulate cortex, while the impact of betweenness of that same
3 node ranked fifth overall. Highly impactful node degree features were found within the cerebellum in both
4 girls (right lobules III, VI, VIIB and Crus II; left lobule III; and vermal lobules I, II, and X) and boys (right
5 cerebellum lobules VI and IX, vermal lobule III), as were node betweenness features (left lobule III; right
6 lobule VI and Crus II, vermal lobules I, II, and X in girls; right lobule VI, vermal lobule III in boys). Yu-Feng
7 et al. (2007) found that children with ADHD had reduced amplitude of low frequency fluctuations (in the
8 approximate frequency band analyzed here) bilaterally in the cerebellar cortex and vermis, which might
9 drive reductions in weights of the functional connections with cerebellar regions in our networks, resulting
10 in differences in subtle group-level features like node degree or betweenness.

11 Additional non-standard network features that had relatively high impact included central point dom-
12 inance measures (e.g., in the left middle temporal pole, left middle and superior frontal cortex, and left
13 superior occipital cortex in boys with ADHD), eccentricity measured by effective resistance (of the left pre-
14 central gyrus in girls with ADHD) and a variety of network sparsity features. In particular, sparsity measures
15 based on lexicographic, Dulmage-Mendelsohn, and approximate minimum degree orderings (without apply-
16 ing an edge threshold) in girls with ADHD (and to a lesser extent boys with and without ADHD) were
17 associated with relatively large feature impacts.

18 Overall, performance for classifiers based on network features alone was somewhat lower than performance
19 using non-imaging phenotypic features or anatomical features. This seems to indicate that the individually
20 derived features may be somewhat noisy, and perhaps an even larger sample size will be required to develop
21 more robust network feature based classifiers. It is possible that other derived features could hold substan-
22 tially more predictive power than those used here, **including, for example, additional measures of the degree
23 of modularity or goodness or partitioning (e.g., Rubinov and Sporns, 2011)**, but the set of attributes we
24 deployed was quite large and heterogeneous, for example including many features at both the local (node)
25 level and the global (network) level.

26 Another important consideration for network-level features, beyond the network construction method we
27 explored here, is the definition of the brain regions that comprise nodes. In the current project we defined
28 each node as an anatomical region-of-interest using the AAL volumetric brain atlas (Tzourio-Mazoyer et al.,
29 2002), which is a macro-anatomical parcellation of the MNI-space single subject brain. Each ROI is then
30 represented by a mean time course (after rs-fMRI preprocessing), which is assumed to be representative of
31 regional activity. Smith et al. (2011) demonstrated, however, that improperly defined regions of interest can
32 substantially degrade inference of connections (in any network construction method). Thus one avenue of
33 further study to improve predictive power is to attempt to define regions of interest (and therefore rs-fMRI
34 networks) in a more data-driven manner, perhaps by first clustering voxels into ROIs based on similarity of
35 response.

36 COMBINING FEATURE SETS

37 Figure 3 demonstrates classifier performance (in both cross-validation and on the separate test set) for
38 classifiers that were built using *all available features* for different combinations of feature sets (defined as
39 non-imaging phenotype, anatomical, and network sets). The best overall performance in cross-validation,
40 which should generally be predictive of performance on the separate test set, occurred for the classifier built
41 from non-imaging phenotype features only. This result was surprising (and somewhat discouraging for efforts
42 to base diagnosis on biological measurables from imaging data), but did not robustly generalize when results
43 were assessed on the separate test set. In particular, generalization of the non-imaging features seemed
44 relatively weak, whereas adding additional features from the imaging sets (in sufficient numbers) appeared

1 to improve the classifiers' ability to generalize to new patient data. Adding only a small number of imaging
2 features to the phenotype set actually *degraded* performance in almost all cases, but as features continued
3 to be added, there was a slow, systematic increase in performance, generally reaching maximum when using
4 *all* available features.

5 As seen in Figure 3, the best overall performance on the *test set* came using an SVM that weighted
6 all 12,000+ features across the three feature sets. It is somewhat unusual that such a large feature set
7 (relative to the number of training examples / subjects) would provide optimal performance. We see this as
8 an indication of the diffuse nature of ADHD. Our results make it clear that there is substantial predictive
9 power in the set of features derived from either anatomical scans, resting state fMRI networks built from
10 sparse inverse covariance estimates, or combinations of the two, but it also appears that, in general, no small
11 set of these features can provide better predictions than the non-imaging phenotypic features on their own.
12 Instead, predictions are honed and made more robust to new data by a combination of many features, each
13 with small overall contributions to the final class prediction.

14 Upon initial review, these results may bring into question the relative merits - particularly from a cost /
15 benefit point of view - of using neuroimaging-derived characteristics for the clinical diagnosis of ADHD, and
16 perhaps more generally for diagnosis of neuropsychiatric disorders. While our results indicate an improve-
17 ment in AUC of less than ~10% in new test subjects when incorporating imaging features (both anatomical
18 and network features) above and beyond the non-imaging feature set, we feel that this is an important
19 improvement. Because ADHD diagnosis is controversial (Wolraich, 1999), any added ability to point to
20 objective, biological measures is of high potential significance. One major objective for algorithmic diagnosis
21 of behaviorally-defined disorders is to find biomarkers that enable diagnosis in a manner that is indistin-
22 guishable from a group of physicians. In this dataset we only have one diagnosis, which was presumably
23 given by one physician or a small group of physicians, and thus we cannot relate our results to the variability
24 that might exist in a set of blind diagnoses from a group of physicians. Thus while we do not feel that the
25 current measures are ready for the clinic, the appearance of modest but highly significant improvements of
26 diagnosis prediction based on imaging features gives hope for such methods serving as a diagnostic aid in
27 the future.

28 IMPORTANCE OF LARGE DATASETS

29 The ADHD-200 dataset has provided a new opportunity to apply exploratory data analysis, data mining, and
30 machine learning tools to the incredibly challenging problem of neuropsychiatric disorder diagnosis. While
31 there is great value in meta-analyses across sets of previously conducted studies, computational scientists
32 place great value in the ability to apply the same methods directly to commonly coded datasets made publicly
33 available in a common format. To achieve results that would directly impact ADHD diagnosis, it may be
34 important to collect data from even more subjects.

35 We observed substantial differences in classifier performance across scan sites (Table 12); however, data
36 were insufficient to properly train classifiers for each site, and we interpret these results cautiously due
37 to relatively small numbers of examples. The variability in classifier performance by site may arise from
38 at least two sources. First, it may reflect aspects of the patient populations that systematically differed
39 by site. Furthermore, different sites used somewhat different acquisition protocols; standardizing these
40 protocols and/or developing additional algorithmic methods to account for such differences may improve
41 overall performance and reduce classifier differences across sites. The types of data assembled here would be
42 well complemented with a variety of other measures, possibly including task-driven fMRI in, for example,
43 set shifting or cognitive / inhibitory control tasks that highlight difficult patterns of behavior for subjects
44 with ADHD. Further, ADHD has a significant genetic component (Faraone et al., 2005), and combining

1 genotype or gene expression information with possible endophenotypes from brain imaging (i.e., imaging
2 genetics Durston et al., 2009) is a promising avenue that will also require collection of both brain imaging
3 data using standardized protocols as well as genetic materials from a large number of patients and matched
4 controls.

5 CONCLUSION

6 The overall framework presented here, which combines different feature sets, each processed in distinct
7 software streams, provides a flexible and extensible means for studying diagnostic measures in large clinical
8 datasets. While it was somewhat surprising that classifiers, in general, achieved best performance using *all*
9 available features, this points to the distributed nature of pathology in complex neuropsychiatric disorders,
10 while establishing the need to combine many diverse attributes for best outcomes. The general approach
11 and set of features described here will be useful in future examinations of nervous system disorders and also
12 potentially for predictions of various brain states based on functional connectivity networks in task-active
13 fMRI.

14 ACKNOWLEDGEMENTS

15 The authors would like to thank Cameron Craddock and the Neuro Bureau for providing pre-processed
16 datasets and scripts to recreate these intermediate results. We would also like to thank Lauren Drake and
17 Natasha Desai for assistance in quality assurance during data processing. Support for this project was
18 provided by the Center of Excellence for Learning in Education, Science, and Technology (CELEST; NSF
19 SMA-0835976).

20 Author contributions

21 JWB and LG conceived the study and overall approach. SS, LG, and JWB conducted the initial data
22 analyses. LG provided algorithms for network measures. FP conducted additional analyses and developed
23 and greatly enhanced the machine learning framework. JR provided advice on network construction methods.
24 JWB, LG, FP, and SS wrote the manuscript.

References

- Almeida, L., Ricardo-Garcell, J., Prado, H., Barajas, L., Fernández-Bouzas, A., Ávila, D., and Martínez, R. (2010). Reduced right frontal cortical thickness in children, adolescents and adults with ADHD and its correlation to clinical variables: A cross-sectional study. *Journal of Psychiatric Research*, 44(16):1214–1223.
- American Psychiatric Association (2000). *Diagnostic and Statistical Manual of Mental Disorders: DSM-IV-TR*. American Psychiatric Publishing, Inc.
- Banaschewski, T., Becker, K., Scherag, S., Franke, B., and Coghill, D. (2010). Molecular genetics of attention-deficit/hyperactivity disorder: an overview. *European Child & Adolescent Psychiatry*, 19(3):237–257.
- Benjamini, Y. and Hochberg, Y. (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 289–300.
- Biswal, B., Mennes, M., Zuo, X., Gohel, S., Kelly, C., Smith, S., Beckmann, C., Adelstein, J., Buckner, R., Colcombe, S., et al. (2010). Toward discovery science of human brain function. *Proceedings of the National Academy of Sciences*, 107(10):4734–4739.
- Bonchev, D., Balaban, A., Liu, X., and Klein, D. (1994). Molecular cyclicity and centrality of polycyclic graphs. I. Cyclicity based on resistance distances or reciprocal distances. *International Journal of Quantum Chemistry*, 50(1).
- Bonchev, D. and Rouvray, D. H., editors (1991). *Chemical Graph Theory — Introduction and Fundamentals*. Gordon & Breach, New York.
- Bullmore, E. and Sporns, O. (2009). Complex brain networks: graph theoretical analysis of structural and functional systems. *Nat Rev Neurosci*, 10(3):186–198.
- Burt, S. (2009). Rethinking environmental contributions to child and adolescent psychopathology: A meta-analysis of shared environmental influences. *Psychological Bulletin*, 135(4):608.
- Bush, G. (2010). Attention-deficit/hyperactivity disorder and attention networks. *Neuropsychopharmacology*, 35(1):278–300.
- Casey, B., Nigg, J., and Durston, S. (2007). New potential leads in the biology and treatment of attention deficit-hyperactivity disorder. *Current Opinion in Neurology*, 20(2):119.
- Castellanos, F., Giedd, J., Berquin, P., Walter, J., Sharp, W., Tran, T., Vaituzis, A., Blumenthal, J., Nelson, J., Bastain, T., et al. (2001). Quantitative brain magnetic resonance imaging in girls with attention-deficit/hyperactivity disorder. *Archives of General Psychiatry*, 58(3):289.
- Castellanos, F., Giedd, J., Eckburg, P., Marsh, W., et al. (1994). Quantitative morphology of the caudate nucleus in attention deficit hyperactivity disorder. *The American Journal of Psychiatry*.
- Castellanos, F. X., Kelly, C., and Milham, M. P. (2009). The restless brain: attention-deficit hyperactivity disorder, resting-state functional connectivity, and intrasubject variability. *Can J Psychiatry*, 54(10):665–672.
- Chang, C.-C. and Lin, C.-J. (2011). LIBSVM: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology*, 2:27:1–27:27.
- Chekkoury, A., Khurd, P., Ni, J., Bahlmann, C., Kamen, A., Patel, A., Grady, L., Singh, M., Groher, M., Navab, N., Krupinski, E., Johnson, J., Graham, A., and Weinstein, R. (2012). Automated malignancy detection in breast histopathological images. In *Proc. of SPIE*, volume 8315.
- Cordes, D., Haughton, V. M., Arfanakis, K., Carew, J. D., Turski, P. A., Moritz, C. H., Quigley, M. A., and Meyerand, M. E. (2001). Frequencies contributing to functional connectivity in the cerebral cortex in “resting-state” data. *AJNR Am J Neuroradiol*, 22(7):1326–1333.
- Costa, L., Rodrigues, F., Travieso, G., and Boas, P. (2007). Characterization of complex networks: A survey of measurements. *Advances in Physics*, 56(1):167–242.
- Cox, R. and Hyde, J. (1997). Software tools for analysis and visualization of fMRI data. *NMR in Biomedicine*, 10:171–178.
- Dale, A. M., Fischl, B., and Sereno, M. I. (1999). Cortical surface-based analysis. I. Segmentation and surface reconstruction. *Neuroimage*, 9(2):179–194.

- Demir, C., Gultekin, S., and Yener, B. (2005). Learning the topological properties of brain tumors. In *Proc. of IEEE/ACM Transactions on Computational Biology and Bioinformatics (TCBB)*, volume 2, pages 262–270.
- Desikan, R., Ségonne, F., Fischl, B., Quinn, B., Dickerson, B., Blacker, D., Buckner, R., Dale, A., Maguire, R., Hyman, B., et al. (2006). An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest. *Neuroimage*, 31(3):968–980.
- Diudea, M. V. and Gutman, I. (1998). Wiener-type topological indices. *Croatica Chemica Acta*, 71(1):21–51.
- Doyle, A., Willcutt, E., Seidman, L., Biederman, J., Chouinard, V., Silva, J., and Faraone, S. (2005). Attention-deficit/hyperactivity disorder endophenotypes. *Biological Psychiatry*, 57(11):1324–1335.
- Durston, S. (2003). A review of the biological bases of ADHD: what have we learned from imaging studies? *Mental Retardation and Developmental Disabilities Research Reviews*, 9(3):184–195.
- Durston, S., de Zeeuw, P., and Staal, W. (2009). Imaging genetics in ADHD: a focus on cognitive control. *Neuroscience & Biobehavioral Reviews*, 33(5):674–689.
- Fair, D., Cohen, A., Dosenbach, N., Church, J., Miezin, F., Barch, D., Raichle, M., Petersen, S., and Schlaggar, B. (2008). The maturing architecture of the brain’s default network. *Proceedings of the National Academy of Sciences*, 105(10):4028.
- Faraone, S., Perlis, R., Doyle, A., Smoller, J., Goralnick, J., Holmgren, M., and Sklar, P. (2005). Molecular genetics of attention-deficit/hyperactivity disorder. *Biological Psychiatry*, 57(11):1313–1323.
- Fiedler, M. (1973). Algebraic connectivity of graphs. *Czechoslovak Mathematical Journal*, 23(98):298–305.
- Fischl, B., Salat, D. H., Busa, E., Albert, M., Dieterich, M., Haselgrove, C., van der Kouwe, A., Killiany, R., Kennedy, D., Klaveness, S., Montillo, A., Makris, N., Rosen, B., and Dale, A. M. (2002). Whole brain segmentation: automated labeling of neuroanatomical structures in the human brain. *Neuron*, 33:341–355.
- Fischl, B., Sereno, M. I., and Dale, A. (1999a). Cortical surface-based analysis: II: Inflation, flattening, and a surface-based coordinate system. *NeuroImage*, 9(2):195 – 207.
- Fischl, B., Sereno, M. I., Tootell, R. B., and Dale, A. M. (1999b). High-resolution intersubject averaging and a coordinate system for the cortical surface. *Human Brain Mapping*, 8(4):272–284.
- Fischl, B., van der Kouwe, A., Destrieux, C., Halgren, E., Segonne, F., Salat, D. H., Busa, E., Seidman, L. J., Goldstein, J., Kennedy, D., Caviness, V., Makris, N., Rosen, B., and Dale, A. M. (2004). Automatically parcellating the human cerebral cortex. *Cerebral Cortex*, 14(1):11–22.
- Fonov, V., Evans, A., Botteron, K., Almli, C., McKinstry, R., Collins, D., et al. (2011). Unbiased average age-appropriate atlases for pediatric studies. *NeuroImage*, 54(1):313–327.
- Frazier, T., Demaree, H., and Youngstrom, E. (2004). Meta-analysis of intellectual and neuropsychological test performance in attention-deficit/hyperactivity disorder. *Neuropsychology*, 18(3):543.
- Friedman, J., Hastie, T., and Tibshirani, R. (2008). Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*, 9(3):432–441.
- Gaub, M. and Carlson, C. (1997). Gender differences in ADHD: A meta-analysis and critical review. *Journal of the American Academy of Child & Adolescent Psychiatry*, 36(8):1036–1045.
- Gershon, J. and Gershon, J. (2002). A meta-analytic review of gender differences in ADHD. *Journal of Attention Disorders*, 5(3):143.
- Geschwind, D. and Levitt, P. (2007). Autism spectrum disorders: developmental disconnection syndromes. *Current Opinion in Neurobiology*, 17(1):103–111.
- Geschwind, N. (1965). Disconnexion syndromes in animals and man. *Brain*, 88(3):585–585.
- Giedd, J., Blumenthal, J., Molloy, E., and Castellanos, F. (2001). Brain imaging of attention deficit/hyperactivity disorder. *Annals of the New York Academy of Sciences*, 931(1):33–49.
- Golland, P. and Fischl, B. (2003). Permutation tests for classification: towards statistical significance in image-based studies. In *Information processing in medical imaging*, pages 330–341. Springer.
- Grady, L. and Polimeni, J. R. (2010). *Discrete Calculus: Applied Analysis on Graphs for Computational Science*. Springer.

- Grady, L. and Schwartz, E. L. (2006). Isoperimetric graph partitioning for image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(3):469–475.
- Greicius, M. (2008). Resting-state functional connectivity in neuropsychiatric disorders. *Current Opinion in Neurology*, 21(4):424.
- Gunduz, C., Yener, B., and Gultekin, S. (2004). The cell graphs of cancer. In *Proc. of International Symposium on Molecular Biology ISMB/ECCB*, volume 20, pages 145–151.
- Guyon, I. (2003). An introduction to variable and feature selection. *Journal of Machine Learning Research*, 3:1157–1182.
- Hansen, P. and Jurs, P. (1988). Chemical applications of graph theory: Part I. Fundamentals and topological indices. *Journal of Chemical Education*, 65(7):574–80.
- Horton, J. (1987). A polynomial-time algorithm to find the shortest cycle basis of a graph. *SIAM Journal on Computing*, 16:358.
- Huang, S., Li, J., Sun, L., Ye, J., Fleisher, A., Wu, T., Chen, K., and Reiman, E. (2010). Learning brain connectivity of Alzheimer’s disease by sparse inverse covariance estimation. *Neuroimage*, 50(3):935–949.
- Kavitha, T., Mehlhorn, K., Michail, D., and Paluch, K. (2008). An algorithm for minimum cycle basis of graphs. *Algorithmica*, 52(3):333–349.
- Khurd, P., Grady, L., Kamen, A., Gibbs-Strauss, S., Genega, E., and Frangioni, J. V. (2011). Network cycle features: Application to computer-aided gleason grading of prostate cancer histopathological images. In *Proc. of ISBI*, pages 1632–1636. IEEE.
- Klein, D. J. and Randić, M. (1993). Resistance distance. *Journal of Mathematical Chemistry*, 12(1):81–95.
- Konrad, K. and Eickhoff, S. (2010). Is the ADHD brain wired differently? A review on structural and functional connectivity in attention deficit hyperactivity disorder. *Human Brain Mapping*, 31(6):904–916.
- Latora, V. and Marchiori, M. (2001). Efficient behavior of small-world networks. *Physical Review Letters*, 87(19):198701.
- Liston, C., Malter Cohen, M., Teslovich, T., Levenson, D., and Casey, B. (2011). Atypical prefrontal connectivity in attention-deficit/hyperactivity disorder: Pathway to disease or pathological end point? *Biological psychiatry*.
- Liu, J., Ji, S., and Ye, J. (2009). *SLEP: Sparse Learning with Efficient Projections*. Arizona State University. <http://www.public.asu.edu/~jye02/Software/SLEP>.
- Mahone, E. and Wodka, E. (2008). The neurobiological profile of girls with ADHD. *Developmental Disabilities Research Reviews*, 14(4):276–284.
- Makris, N., Biederman, J., Valera, E., Bush, G., Kaiser, J., Kennedy, D., Caviness, V., Faraone, S., and Seidman, L. (2007). Cortical thinning of the attention and executive function networks in adults with attention-deficit/hyperactivity disorder. *Cerebral Cortex*, 17(6):1364–1375.
- Mega, M. and Cummings, J. (1994). Frontal-subcortical circuits and neuropsychiatric disorders. *The Journal of Neuropsychiatry and Clinical Neurosciences*.
- Mihalić, Z. and Trinajstić, N. (1992). A graph-theoretical approach to structure: Property relationships. *Journal of Chemical Education*, 69(9):701–712.
- Milham, M. (2012). Open neuroscience solutions for the connectome-wide association era. *Neuron*, 73(2):214–218.
- Mohar, B. (1989). Isoperimetric numbers of graphs. *Journal of Combinatorial Theory, Series B*, 47:274–291.
- Newcorn, J., Halperin, J., Jensen, P., Abikoff, H., Arnold, L., Cantwell, D., Conners, C., Elliott, G., Epstein, J., Greenhill, L., et al. (2001). Symptom profiles in children with ADHD: effects of comorbidity and gender. *Journal of the American Academy of Child & Adolescent Psychiatry*, 40(2):137–146.
- Newman, M. (2002). Assortative mixing in networks. *Physical Review Letters*, 89(20):208701.
- Newman, M. (2004). Fast algorithm for detecting community structure in networks. *Physical Review E*, 69(6):066133.
- Papadimitriou, C. (1976). The NP-completeness of the bandwidth minimization problem. *Computing*, 16(3):263–270.

- Patel, R., Bowman, F., and Rilling, J. (2006). A Bayesian approach to determining connectivity of the human brain. *Human Brain Mapping*, 27(3):267–276.
- Polanczyk, G., Silva de Lima, M., Lessa Horta, B., Biederman, J., and Augusto Rohde, L. (2007). The worldwide prevalence of ADHD: A systematic review and metaregression analysis. *American Journal of Psychiatry*, 164:942–948.
- Rubinov, M. and Sporns, O. (2011). Weight-conserving characterization of complex functional brain networks. *Neuroimage*.
- Schrimsher, G., Billingsley, R., Jackson, E., and Moore, B. (2002). Caudate nucleus volume asymmetry predicts attention-deficit hyperactivity disorder (ADHD) symptomatology in children. *Journal of Child Neurology*, 17(12):877–884.
- Shaw, P., Lerch, J., Greenstein, D., Sharp, W., Clasen, L., Evans, A., Giedd, J., Castellanos, F., and Rapoport, J. (2006). Longitudinal mapping of cortical thickness and clinical outcome in children and adolescents with attention-deficit/hyperactivity disorder. *Archives of General Psychiatry*, 63(5):540.
- Smith, S., Jenkinson, M., Woolrich, M., Beckmann, C., Behrens, T., Johansen-Berg, H., Bannister, P., De Luca, M., Drobnjak, I., Flitney, D., Niazy, R., Saunders, J., Vickers, J., Zhang, Y., De Stefano, N., and Brady, J.M. and Matthews, P. (2004). Advances in functional and structural MR image analysis and implementation as FSL. *NeuroImage*, 23:208–219.
- Smith, S., Miller, K., Salimi-Khorshidi, G., Webster, M., Beckmann, C., Nichols, T., Ramsey, J., and Woolrich, M. (2011). Network modelling methods for fMRI. *Neuroimage*, 54(2):875–891.
- Sporns, O. (2010). *Networks of the Brain*. The MIT Press.
- Tekin, S. and Cummings, J. (2002). Frontal-subcortical neuronal circuits and clinical neuropsychiatry: An update. *Journal of Psychosomatic Research*, 53(2):647–654.
- Tripp, G. and Wickens, J. (2009). Neurobiology of ADHD. *Neuropharmacology*, 57(7):579–589.
- Tzourio-Mazoyer, N., Landeau, B., Papathanassiou, D., Crivello, F., Etard, O., Delcroix, N., Mazoyer, B., and Joliot, M. (2002). Automated anatomical labeling of activations in SPM using a macroscopic anatomical parcellation of the mni mri single-subject brain. *Neuroimage*, 15(1):273–289.
- Valera, E., Faraone, S., Murray, K., and Seidman, L. (2007). Meta-analysis of structural imaging findings in attention-deficit/hyperactivity disorder. *Biological Psychiatry*, 61(12):1361–1369.
- Wang, L., Zhu, C., He, Y., Zang, Y., Cao, Q., Zhang, H., Zhong, Q., and Wang, Y. (2009). Altered small-world brain functional networks in children with attention-deficit/hyperactivity disorder. *Human Brain Mapping*, 30(2):638–649.
- Watts, D. and Strogatz, S. (1998). Collective dynamics of ‘small-world’ networks. *Nature*, 393(6684):440–442.
- Wiener, H. (1947). Structural determination of paraffin boiling points. *J. Am. Chem. Soc.*, 69:17–20.
- Wolraich, M. (1999). Attention deficit hyperactivity disorder: the most studied and yet the most controversial diagnosis. *Mental Retardation and Developmental Disabilities Research Reviews*, 5:163–168.
- Yu-Feng, Z., Yong, H., Chao-Zhe, Z., Qing-Jiu, C., Man-Qiu, S., Meng, L., Li-Xia, T., Tian-Zi, J., and Yu-Feng, W. (2007). Altered baseline brain activity in children with ADHD revealed by resting-state functional MRI. *Brain and Development*, 29(2):83–91.
- Zametkin, A., Nordahl, T., Gross, M., King, A., Semple, W., Rumsey, J., Hamburger, S., and Cohen, R. (1990). Cerebral glucose metabolism in adults with hyperactivity of childhood onset. *New England Journal of Medicine*, 323(20):1361–1366.
- Zweig, M. and Campbell, G. (1993). Receiver-operating characteristic (ROC) plots: a fundamental evaluation tool in clinical medicine. *Clinical Chemistry*, 39(4):561–577.

List of Figures

Figure 1	Schematic describing the overall pipeline through which data were processed and submitted to machine learning algorithms for classification. Note that anatomical and functional / network streams were largely independent.	7
Figure 2	Top: Scatter plot representing female subjects from all sites, their DX status, age, combined IQ (average of the two available measures) and handedness. The x-axis corresponds to the subject age. Along the y-axis, typically developing control subjects are plotted below the $y = 0$ line, and ADHD subjects above (i.e., $ IQ = IQ$, with negative values simply denoting TDC subjects. The position along the y-axis corresponds to the average IQ measure, with subjects whose records did not include IQ plotted around $y = -25$ and $y = 25$. The marker for each subject corresponds to diagnosis status (dot is control, o is ADHD-combined type, x is ADHD-hyperactive/impulsive, + is ADHD-inattentive), and its color to handedness (red is left-handed). Coordinates are jittered slightly to improve visualization. Bottom: Same as the plot above, but for male subjects.	14
Figure 3	Summary of best performing classifiers using all features in the set or combination of sets below, sorted by performance on the test set (dark bars). Light bars indicate performance on non-training folds during cross-validation. A = Anatomical features, P = Non-imaging Phenotype features, and N = Network features.	17
Figure 4	Classification feature impact weights for cortical thickness features in the classification of boys with positive ADHD diagnosis. Darker blue values indicate higher impact weights.	20
Figure 5	Classification feature impact weights for cortical thickness features in the classification of girls with positive ADHD diagnosis. Darker blue values indicate higher impact weights.	20
Figure 6	Comparison of AUC performance on the test set as a function of number of features used for classifiers built from a combination of all feature classes, with (green line) and without (blue line) stratification by gender. Each subplot represents results based on a different feature selection method. Also shown (red line, same in all plots) is the result obtained when features are ranked randomly (separately for each stratum); this result is averaged over classifiers built using 10 random feature rankings.	22
Figure 7	Comparison of AUC performance on test set as a function of number of features used for the two predominant ADHD subtypes. DX1 is the ADHD-Combined subtype; DX3 is the ADHD-Inattentive subtype.	23