Classifying Psychopathy Patients Using Machine Learning Methods on Magnetic Resonance Imaging (MRI) Data

Margaret Pearce, *Graduate Student,* and Isaac Sung, *Graduate Student*

***Abstract*—Machine learning methods have the potential to be an effective tool to help doctors diagnose patients with various mental disorders. Our specific case study looks at MRI data from prisoners in Wisconsin and attempts to classify patients as being psychopathic or not. We were provided with two datasets, one with a large number of features and one with a much smaller set of features. For each dataset we tested several machine learning methods, such as support vector machines (SVM), decision trees, and k-nearest neighbor. Then, we attempted to improve our accuracy by implementing ensemble methods. We found that SVM performed the best with around 73% accuracy. Overall, our results appear promising, and with additional work, these methods could help identify patients as being psychopathic.**

I. Introduction

RECIDIVISM is a major issue in the criminal justice system. It is defined by Merriam-Webster as a tendency to relapse into a previous condition or mode of behavior, especially relapse into criminal behavior. The cost of recidivism to governments and citizens is high. The United States spent over $52 billion in 2011 on corrections [14]. Studies have also revealed a close link between recidivism and psychopathy [6]. Those diagnosed with a high degree of psychopathy have been found to be more likely to re-offend. Although no psychiatric or psychological organization has defined a condition known as "psychopathy", Hervey M. Cleckley introduced the idea and the categorization has been widely used in the criminal justice system. Psychopathy is generally considered to be a condition in which the person acts uninhibitedly and brashly, exhibits less empathy and remorse than the average person, and is antisocial. Because of this link between psychopathy and recidivism, we wish to design a system that can help identify prisoners who are more likely to reoffend and put them on a track to succeed after being released.

A common practice of assessing psychopathy in an individual a procedure called the Psychopathy Checklist-Revised (PCL-R) [5]. The PCL-R is a rating scale that uses 60-90 minute semi-structured interviews and case history information to evaluate an individual on 20 traits (Table I). Each trait is scored on a three-point scale (0, 1, or 2) based on the extent to which the individual exhibits the trait. The total score ranges from 0 to 40 and is calculated by summing each of the 20 items. Total PCL-R scores are then used to classify inmates as non-psychopathic (score of 20 or less) or psychopathic (score above 20).

TABLE I
PCL-R Items [16]

1. Glibness/superficial charm
2. Grandiose sense of self worth
3. Need for stimulation/proneness to boredom
4. Pathological lying
5. Conning/manipulative
6. Lack of remorse or guilt
7. Shallow affect
8. Callous/lack of empathy
9. Parasitic lifestyle
10. Poor behavioral controls
11. Promiscuous sexual behavior
12. Early behavior problems
13. Lack of realistic, long-term goals
14. Impulsivity
15. Irresponsibility
16. Failure to accept responsibility for own actions
17. Many short-term marital relationships
18. Juvenile delinquency
19. Revocation of conditional release
20. Criminal versatility

There is growing evidence to suggest that neurological structural and functional abnormalities are associated with psychopathy. The prefrontal cortex, particularly the ventromedial and anterior cingulate sectors, is believed to be a strong factor [8]. Prior work on has correlated damage to the ventromedial prefrontal cortex with several emotional and decision-making defects, such as diminished guilt, shame, and empathy; irritability; irresponsibility; and failure to learn from punishment. Further studies have demonstrated a link between ventromedial prefrontal cortex damage and changes in moral judgment and economic decision-making. However, Koenigs makes a distinction that patients with ventromedial prefrontal cortex damage are not typically assessed as psychopathic; in particular, patients do not exhibit criminal and antisocial behaviors.

The anterior cingulate cortex (ACC) is a second area of interest in the prefrontal cortex with respect to psychopathy. Koenigs notes that functions such as reward, punishment, pain, negative affect, empathy, error detection, performance

M. Pearce and I. Sung are with the Department of Computer Sciences, University of Wisconsin-Madison, Madison, WI, USA

Mike Koenigs and Rick Wolf are with the Koenigs Lab at the University of Wisconsin-Madison.

Manuscript received May 10, 2015.

monitoring, and cognitive control are related to activity in the ACC. Studies of patients with damage to the ACC describe patients as exhibiting greater irritability, social disinhibition, and reduced motivation.

We aim to replicate the classification findings of this assessment by applying machine learning techniques to MRI data obtained from inmates at correctional facilities. We employ a variety of classification algorithms, including support vector machines, k-nearest neighbors, and decision trees. Next, we explore the impact of using ensemble methods. We discuss the accuracy of each approach as well as their sensitivities to false positives and false negatives. Finally, we consider the implications of using two different forms of the data with each algorithm: voxel-wise data and anatomical data. The original voxel-wise data set contains several thousand features for each MRI scan, while the second anatomical data set has been preprocessed to a smaller size by domain experts.

II. Related Work

Using machine learning techniques on MRI data is not a new practice. Previous work has yielded high diagnostic accuracy on schizophrenia and depression classification [13]. A study on Alzheimer's disease used support vector machines trained on MRI brain scans from Alzheimer's patients and elderly normal persons [7]. Work has also been done to distinguish between patients with schizophrenia and bipolar disorder using support vector machines on structural MRIs [18]. Prior work has also used pattern recognition techniques against MRI data to distinguish psychopaths and healthy controls [17]. However, the data set used in this study was measured using a different technique that focuses on whole-brain gray matter vowels, not fractional anisotropy. Furthermore, the sample size was limited to 30 individuals (15 non-psychopathic, 15 psychopathic).

III. Materials and Methods

A. Data Acquisition

The data used in this study was acquired from ongoing work at the Koenigs Lab at the University of Wisconsin-Madison [20]. Participants were adult male inmates aged 18-45 at a medium-security correctional facility in Wisconsin. Eligibility criteria for participants included: IQ >70, no psychotropic medication use, no history of loss of consciousness lasting longer than 30 minutes, and no history of psychosis or bipolar disorder. The final sample included 231 inmates (74 non-psychopathic, 96 mid-psychopathic, 61 psychopathic) scanned between 2010 and 2013. All participants provided oral and written informed consent.

MRI scans for each patient were acquired by a Siemens mobile MRI system that was brought to the correctional facility grounds. Mind Research Network's Siemens 1.5 T Avanto Mobile MRI System is equipped with a 12-element head coil with diffusion sensitizing gradients. For each participant, five interleaved volumes were collected (repetition time TR = 9200 ms, echo time TE = 84 ms, field of view FOV = 256 mm x 256 mm, matrix size = 128 x 128, slice thickness = 2 mm, voxel size = 2 mm x 2 mm x 2 mm, 70 slices). To improve the signal-to-noise ratio, the sequence was repeated twice with data from each sequence combined.

Psychopathy levels were assessed by trained research staff using the Psychopathy Checklist-Revised (PCL-R) in a 60-90 minute semi-structured interview. Inter-rater reliability ratings were completed for 13 participants with high intra-class correlation for total PCL-R scores (r=.99). The total PCL-R score was categorized into three labels: non-psychopathic (score of 20 or lower), mid-psychopathic (score between 20 and 30), and psychopathic (score of 30 or higher). Both the PCL-R scores and categorical labels were available, though we chose to focus on categorical classification for this project.

In order to obtain access to the data, we were required to complete Human Subjects Protection for Biomedical Researchers training through the University of Wisconsin-Madison Collaborative Institutional Training Initiative (CITI). Upon completion, PI Mike Koenigs added study team members to an existing IRB protocol. Next, we visited the Wisconsin Psychiatric Institute and Clinics (WisPIC) at Research Park (west of the UW-Madison campus) to obtain credentials within the psychiatry department and remote access to servers storing the MRI data.

Due to the sensitivity of the data and the IRB protocol, all work with the data was required to be done on WisPIC servers. Study team members were not allowed to store the data in any form on personal or university-owned machines. All code written for this project was copied and run remotely on WisPIC servers with MATLAB Release 2014a.

B. Data Preprocessing

MRI data for each patient is available in two formats: a larger voxelwise set with 124,607 features and a condensed anatomical set with 48 features (both including participant ID numbers). Feature values in the larger data set are considered voxelwise fractional anisotropy values from diffusion tensor images. Features are named "x#y#z#" where the x, y, and z values correspond to the location of the voxel in 3D space of the brain. This allows neuroscientists to determine which anatomical structures in the brain are significant based on classification results. Only voxels that correspond to white matter tracts are represented in the values for a total of 124,607 voxels.

The smaller data set has combined voxels that are concentrated in particular anatomical regions. Values in this data set represent the average of fractional anisotropy values by each region. Feature names are the name of the anatomical region as well as its hemisphere (left or right). An example image showing the external capsule anatomical region is shown in Figure 1. For each example in the data sets, the PCL-R score is given in addition to a classification as non-psychopathic (-0.5), mid-psychopathic (0), and psychopathic (0.5).

C. Classification Approach

We applied a variety of ensemble methods to encourage diversity in classifier errors and strengthen the overall accuracy. To begin, we applied several learning algorithms to provide a baseline on the algorithm's performance, including support

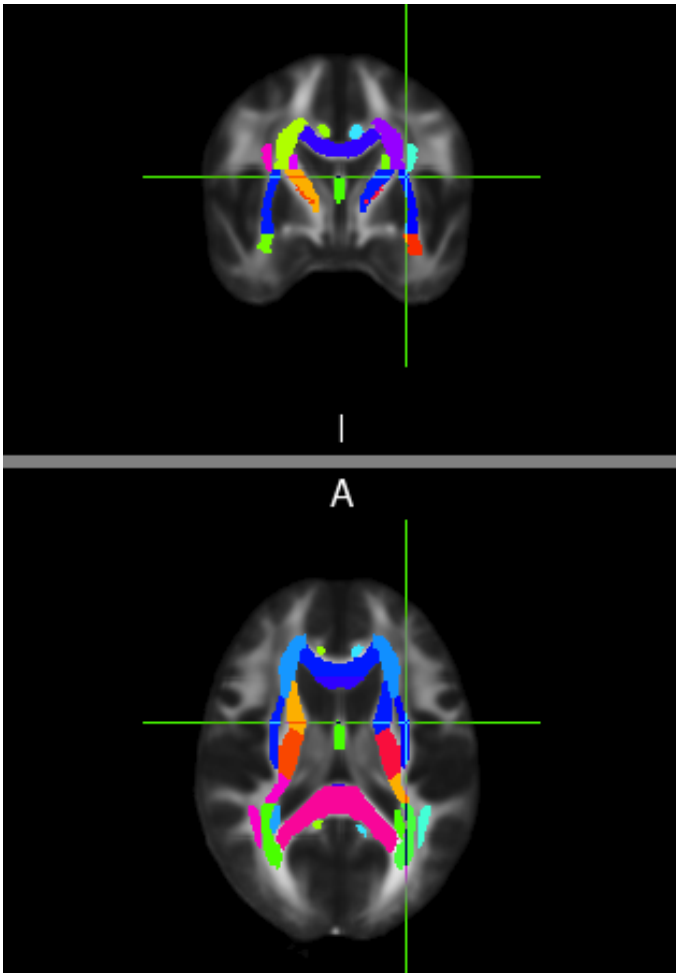

Fig. 1. External capsule region of the brain (marked by the cross-hairs). The top image shows the brain through the coronal plane, while the bottom image shows the brain through the axial plane.

vector machines, decision trees, k-nearest neighbor, and deep learning.

D. Base Learners

1) Support Vector Machines: Widely used in bioinformatics, support vector machines (SVM) are known for their high accuracy, ability to deal with high-dimensional data, and flexibility in modeling diverse sources of data [1]. SVMs employ kernel functions that compute a dot product on the data, possibly translating the data to a higher or lower dimensional space. In order for SVMs to be effective, it is critical to select the right parameters and the right kernel function. For this study, we utilized the libSVM library for MATLAB [3]. We evaluated several kernel functions to determine which one would result in the best fit, including linear, polynomial, radial basis, and sigmoid functions. Ultimately, we found the highest accuracy using a radial basis function with parameters C = 1 and gamma set as the reciprocal of the number of features.

2) Decision trees: Decision trees use a divide-and-conquer approach to learn patterns in data for classification [12]. They can be used to identify features and patterns that are significant for making predictions. Because they yield solutions that are intuitively interpretive, decision trees are an attractive option when trying to evaluate relationships between feature values and classification labels. We leveraged a built-in decision tree learner within MATLAB for this study [9].

3) K-Nearest Neighbor: Classification using k-nearest neighbor works by finding k training examples that are the most similar to the instance being classified. Similarity can be measured using a variety of distance functions depending on the format of the data (all nominal, all discrete, mix), such as Euclidean distance, Manhattan distance, Hamming distance, or a normalized mix of two approaches for mixed nominal/ discrete data[15]. Again, we utilized a built-in knn learner within MATLAB for this study [10]. In both data sets, all features are continuous, so we used the Euclidean distance to establish similarity between examples.

E. Ensembles

Next, we sought to enhance these base algorithms by ensemble techniques. For each of the base learners, we trained the model using different subsamples of the training set in a bagging algorithm. For the decision tree classifier, we applied an algorithm called AdaBoost to explore different probability distributions over the training instances. Finally, we selected a randomly drawn small set of feature splits to generate multiple decision trees as per the random forests technique.

1) Bagging: Bagging is a method that constructs multiple hypothesis by training a model using slightly different training data in each round [4]. If a training set with m examples is provided, a new training set will be derived by randomly picking m examples from the original training set with replacement. The result of this procedure is a resampled data set where some examples will appear more than once and others may not be present in each set. A classifier is independently trained on each of the new training sets. When new examples need to be classified, each classifier will predict a response and the class with the majority vote is chosen as the ensemble's predicted label. We were not certain how stable the base learning algorithms were (whether or not they were sensitive to small changes in the data). Under the hypothesis that the algorithms could be sensitive, we implemented a bagging algorithm and applied it to each of the base learners.

2) Boosting: Contrasting the independent nature of learning multiple hypotheses in bagging, boosting ensemble algorithms are considered additive. Additive models predict the class of a new example by taking a weighted sum from the set of models [4]. In this study, we used a built-in MATLAB extension to run the Adaboost algorithm. Adaboost generates hypotheses that aim to minimize the classification error on a weighted training set. Hypotheses are incrementally added one at a time with the goal of creating a weighted sum

$$H(x) = \sum_k w_k h_k(x)$$

matching the correct label. Adaboost performs well on data that doesn't include a high level of mislabeled training data points. In our data, each example's label is based on a

thorough evaluation by trained psychology research staff, so the probability of a mislabeled example is relatively low.

3) Random Forests: Another ensemble we tried was random forests. The random forest method improves upon decision trees by creating an ensemble of trees that uses a random selection of features to split at each decision node [2]. We used this method because the number of features in our dataset were very large, especially in our voxel-wise data. We hoped to improve the effectiveness of decision trees by randomly selecting a small sample set of the features. The built-in TreeBagger class in MATLAB was used to implement random forests. This algorithm randomly select square root of n features at each node, where n is the number of total features [11]. We found this method improved the overall accuracy for each dataset, as expected from previous works done on random forests.

F. Binary versus Multiclass

Finally, we compared results of training classifiers on both the anatomical and voxelwise data sets. As provided, the data sets use multi-class labels (-0.5: non-psychopathic, 0: mid-psychopathic, and 0.5: psychopathic). We evaluated two approaches to convert the multi-class labels to binary labels. In binary method A, we combined mid-psychopathic and psychopathic labels in the positive class (0 and 0.5 map to 1) and treated non-psychopathic as the negative class (-0.5 maps to 0). This method distinguishes between individuals with some level of psychopathy and those not considered psychopathic. As an alternate approach, binary method B merged non-psychopathic and mid-psychopathic labels in the negative class (-0.5 and 0 map to 0) with psychopathic being the positive label (0.5 maps to 1). Under this scheme, individuals who are in the highest category of psychopathy are separated. Depending on the application, one labeling technique may be more useful than the other.

G. Train and Test Data Sets

Out of the n = 231 examples, we used n = 162 examples for training (70%) and n = 69 examples for testing (30%). Training and test sets were partitioned via stratified sampling. For the binary labeling techniques, stratified sampling was not completed until after the new labels were generated.

IV. Results

For all of the learners, we noticed a significant increase in average accuracy when we changed the data labels from being multiclass to binary. There seemed to be no general pattern to whether binary method A was better than binary method B. In some cases method A produced higher accuracy, and in others it did not.

A. Base Learners

As can be seen in Table II and Table III, for base learners, SVMs performed notably better than decision trees or kNN for binary labels. For multiclass labels, all three performed equally poorly. Overall, decision trees seemed to be the worst performer. There was not a significant difference in accuracy between the anatomical data and voxelwise data.

We generated receiver operating characteristic (ROC) curves for a higher performing classifier, SVM with binary A labels

TABLE II
Anatomical Data Set Accuracy, Base Learners

	SVM	Decision Trees	kNN
Multiclass labels	0.4203	0.3913	0.4203
Binary labels (A)	0.6811	0.4203	0.5652
Binary labels (B)	0.7391	0.6087	0.6522

TABLE III
Voxelwise Data Set Accuracy, Base Learners

	SVM	Decision Trees	kNN
Multiclass labels	0.4203	0.3043	0.3623
Binary labels (A)	0.6739	0.5870	0.5425
Binary labels (B)	0.7391	0.5652	0.6667

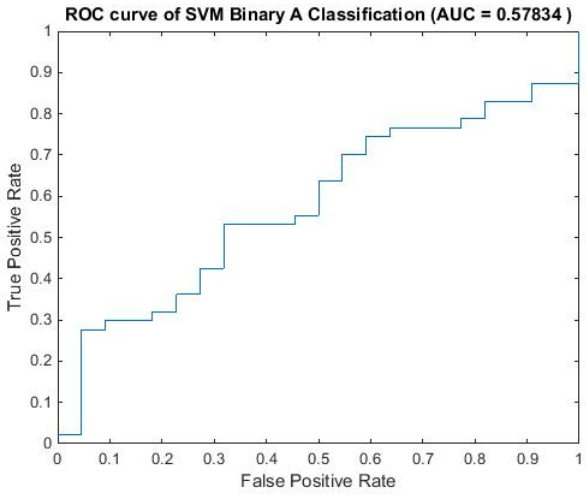

Fig. 2. ROC curve for SVM binary label A classification.

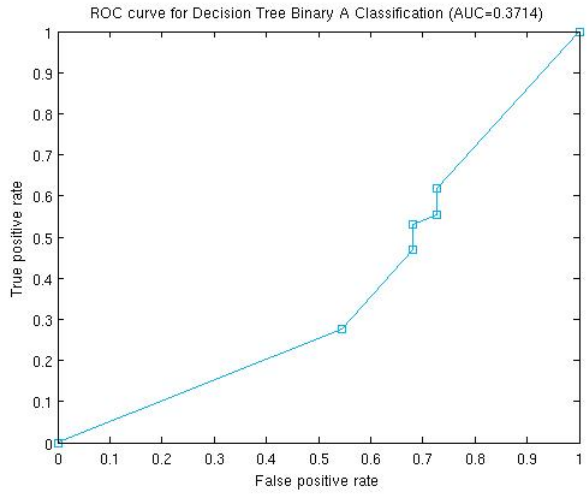

Fig. 3. ROC curve for decision tree binary label A classification.

TABLE IV
ANATOMICAL DATA SET ACCURACY, BAGGING ENSEMBLE

	SVM	Decision Trees	kNN
Multiclass labels	0.4203	0.2609	0.2754
Binary labels (A)	0.6812	0.6812	0.6957
Binary labels (B)	0.7391	0.4058	0.4348

TABLE V
VOXELWISE DATA SET ACCURACY, BAGGING ENSEMBLE

	SVM	Decision Trees	kNN
Multiclass labels	0.4203	0.3043	0.2609
Binary labels (A)	0.6812	0.6812	0.6812
Binary labels (B)	0.7391	0.3768	0.3768

TABLE VI
ANATOMICAL DATA SET ACCURACY, BOOSTING ENSEMBLE

	Decision Trees
Multiclass labels	0.4058
Binary labels (A)	0.5942
Binary labels (B)	0.6377

TABLE VII
VOXELWISE DATA SET ACCURACY, BOOSTING ENSEMBLE

	Decision Trees
Multiclass labels	0.3768
Binary labels (A)	0.5797
Binary labels (B)	0.6377

(see Fig 2), and an under-performing classifier, decision tree with binary A labels (see Fig 3). The lower performing classifier has an area under the curve (AUC) of <50%, while the higher performing classifier has an AUC of >50%. This confirms that the decision tree classifier with binary A labels is less accurate in correctly identifying the correct labels than the SVM classifier. One reason for this could be that the decision tree is overfitting the training data.

B. Bagging

SVMs once again achieved the highest average accuracy when using bagging, as shown in Tables IV and V. Though for both datasets, under binary label method A, both decision trees and kNN achieved similar results. However, SVMs were the most consistently accurate. For some reason, binary method B drastically reduced the performance of both decision trees and kNN, but increased the accuracy for SVM.

C. Boosting

Boosting generally helped increase the accuracy of decision trees, but it did not perform as well as the best performer of decision trees with bagging (Tables VI and VII).

TABLE VIII
RANDOM FORESTS ENSEMBLE

	Anatomical Data Set	Voxelwise Data Set
Multiclass labels	0.3043	0.3768
Binary labels (A)	0.6739	0.6304
Binary labels (B)	0.6522	0.7174

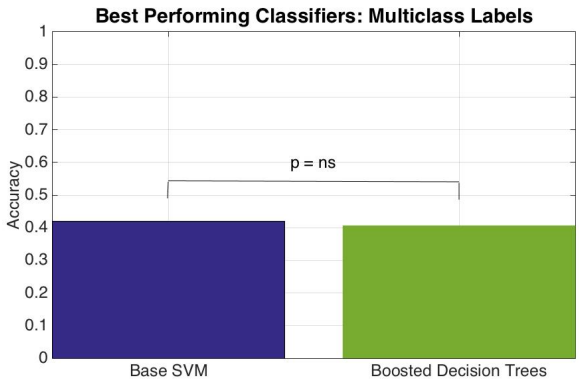

Fig. 4. Paired t-test for the best performing multiclass classifiers, base SVM and boosted decision trees.

D. Random Forests

Overall, random forests improved upon decision trees (Table VIII). The highest performing non-SVM classifier was a random forest ensemble with binary method B.

E. Comparing the Best Classifiers

It is clear from this analysis that not all algorithms and ensembles performed equally well. For each method of labeling the data, we sought to compare the top two classification methods to determine if one method has statistically significantly better performance than the other. We chose to make comparisons only across methods within the same labeling scheme (multiclass, binary method A, and binary method B) and data set (anatomical, voxelwise) to ensure that we are drawing a comparison purely on the algorithms. Furthermore, because the anatomical data set outperformed the voxelwise data set, we chose to focus our analysis on the anatomical data (with fewer features) for the sake of computational efficiency.

Anatomical data with multiclass labels (-0.5: non-psychopathic, 0: mid-psychopathic, 0.5: psychopathic) achieved the highest accuracy on the base SVM classifier (42.03%) and boosted decision trees (40.58%). Using 10-fold stratified cross validation on the anatomical data set, we computed the accuracy in each fold for the two methods. Next, we used these accuracies in a paired t-test to determine if one algorithm is better than the other. We chose a two-tailed test because we wanted to allow for the possibility that either method could be more accurate. We found no significant difference between the base SVM classifier and boosted decision tree classifier for multiclass labeled anatomical data (p >0.05), implying that the two methods are equally accurate (Fig 4).

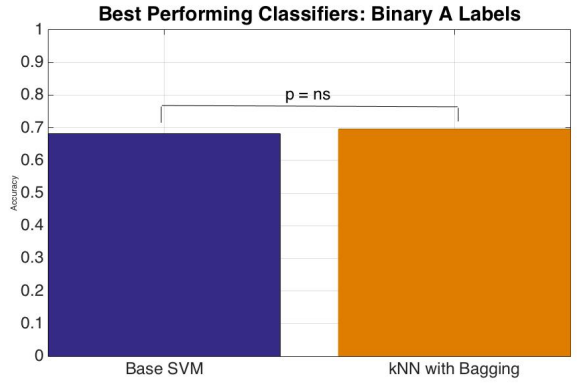

Fig. 5. Paired t-test for the best performing binary A classifiers, base SVM and kNN with bagging.

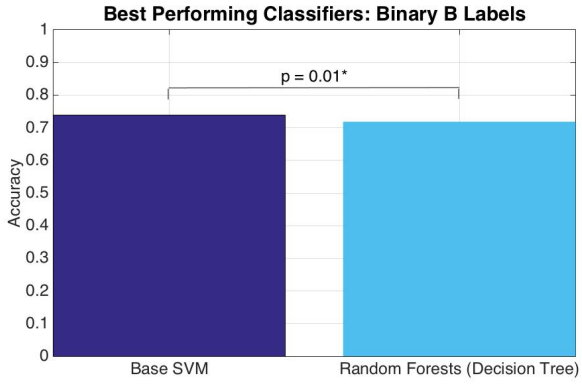

Fig. 6. Paired t-test for the best performing binary B classifiers, base SVM and random forests.

Next, we compared the top two algorithms for the Binary A method labels (0: non-psychopathic, 1: mid-psychopathic or psychopathic). The classifiers we compared were base SVM (68.11%) and kNN with Bagging (69.57%). Using the same method as in the multiclass data, we again found no significant difference between these two algorithms (Fig 5).

Finally, we compared the base SVM classifier (73.91% accuracy) with the random forests decision tree ensemble (65.22%) for Binary B method labels (0: non-psychopathic or mid-psychopathic, 1: psychopathic). In this case, the difference between the two algorithms was statistically significant (p = 0.01) (Fig 6). This tells us that the base SVM classifier has better performance than random forests with decision trees. However, this does not imply that there is a large-magnitude difference between the two algorithms.

V. Discussion

In this study, we explored the effectiveness of a variety of machine learning applications to classifying psychopathy severity. Overall, the most effective learner was support vector machines with binary labels separating non and mid-psychopathic instances from psychopathic instances. Although none of the methods performed as well as we had hoped with multiclass labels, we were able to achieve performance gains by mapping to binary labels. For weaker learners kNN and decision trees, using ensemble techniques greatly increase the accuracy.

One noteworthy takeaway from this study is that for the strongest learner (SVM), we did not see a considerable change in accuracy between the anatomical and voxelwise data set even though the anatomical data set had much fewer features. The anatomical data set that we were provided was derived by averaging the voxelwise values within anatomical regions. Because of its smaller number of features, every algorithm that we tested was much more efficient in training and/or classification compared to the voxelwise set.

Future work could explore different methods of consolidating the voxelwise data as well as incorporating more domain knowledge from neuroscientists. Prior work has hypothesized that particular areas of the brain - the ventromedial and anterior cingulate sectors in the prefrontal cortex - have a strong role in psychopathy. This could be factored into machine learning techniques through the use of prior probability distributions. Finally, it is worth noting that psychopathy is influenced by many components, including neurobiological, genetic, epidemiological and sociobiographical factors [19]. As suggested by Sato et al (2011), incorporating additional data that quantifies these factors may yield more accurate classifiers than considering each aspect individually.

Acknowledgments

The authors would like to thank Dr. Mike Koenigs and Rick Wolf for helping us understand their research and allowing us access to the MRI data. We would also like to thank Dr. David Page for his help in completing this project for his class, CS 760 Machine Learning.

References

[1] A. Ben-Hur and J. Weston. A user's guide to support vector machines. http://pyml.sourceforge.net/doc/howto.pdf, 2012.
[2] L. Breiman. Random forests. *Machine learning*, 45(1):5–32, 1999.
[3] C. C. Chang and C. J. Lin. Libsvm: A library for support vector machines. *AACM Transactions on Intelligent Systems and Technology*, 2:1–27, 2011.
[4] T. G. Dietterich. Ensemble learning. *The handbook of brain theory and neural networks*, 2:110–125, 2002.
[5] R. Hare, D. Clark, M. Grann, and D. Thornton. Psychopathy and the predictive validity of the pcl-r: an international perspective. *Behavioral Sciences & the Law*, 18:623–645, 2000.
[6] J. F. Hemphill, R. D. Hare, and S. Wong. Psychopathy and recidivism: A review. *Legal and Criminological Psychology*, 3:139–170, 1998.
[7] S. Klöppel, C. M. Stonnington, C. Chu, B. Draganski, R. I. Scahill, J. D. Rohrer, N. C. Fox, C. R. Jack, J. Ashburner, and R. S. Frackowiak. Automatic classification of mri scans in alzheimer's disease. *Brain*, 131:681–689, 2008.
[8] M. Koenigs. The role of prefrontal cortex in psychopathy. *Reviews in the neurosciences*, 23(3):253–262, 2012.
[9] Mathworks. Classification trees and regression trees (r2015a). http://www.mathworks.com/help/stats/classification-trees-and-regression-trees.html, 2015.
[10] Mathworks. Nearest neighbors (r2015a). http://www.mathworks.com/help/stats/nearest-neighbors.html, 2015.
[11] Mathworks. Treebagger. http://www.mathworks.com/help/stats/treebagger.html, 2015.
[12] A. J. Myles, R. N. Feudael, Y. Liu, N. A. Woody, and S. D. Brown. An introduction to decision tree modeling. *Journal of Chemometrics*, 18:275–285, 2004.
[13] I. Nouretdinov, S. Costafreda, A. Gammerman, A. Chervonenkis, V. Vovk, V. Vapnik, and C. Fu. Machine learning classification with confidence: Application of transductive conformal predictors to mri-based diagnostic and prognostic markers in depression. *NeuroImage*, 56:809–813, 2011.

[14] N. A. of State Budget Officers. State expenditure report 2010. http://www.nasbo.org/sites/default/files/2010%20State%20Expenditure%20Report_0.pdf, May 2015.
[15] D. Page. Instance-based learning. http://pages.cs.wisc.edu/~dpage/cs760/instance-based-learning.pdf, May 2015.
[16] C. Patrick. *Handbook of Psychopathy*.
[17] J. Sato, R. Oliveira-Souza, C. Thomaz, R. Basilio, I. Bramati, E. Amaro, F. Tovar-Moll, R. Hare, and J. Moll. Identification of psychopathic individuals using pattern classification of mri images. *Social Neuroscience*, 6:5–6, 2011.
[18] H. Schnack, M. Nieuwenhuis, N. van Haren, L. Abramovic, T. Scheewe, R. Brouwer, H. H. Pol, and R. Kahn. Can structural mri aid in clinical classification? a machine learning study in two independent samples of patients with schizophrenia, bipolar disorder and healthy subjects. *NeuroImage*, 84:299–306, 2014.
[19] S. Weber, U. Habel, K. Amunts, and F. Schneider. Structural brain abnormalities in psychopathsã -a review. *Behavioral Sciences & the Law*, 26:7–28, 2008.
[20] R. C. Wolf, M. S. Pujara, J. C. Motzkin, J. P. Newman, K. A. Kiehl, J. Decety, D. S. Kosson, and M. Koenigs. Interpersonal traits of psychopathy linked to reduced integrity of the uncinate fasciculus.