A Data-Driven Approach for Detecting Autism Spectrum Disorders

Manika Kapoor
Computer Engineering
San José State University
San José, CA, USA
manika.kapoor@sjsu.edu

David C. Anastasiu*
Computer Engineering
San José State University
San José, CA, USA
david.anastasiu@sjsu.edu

***Abstract*—Autism spectrum disorders (ASDs) are a group of conditions characterized by impairments in reciprocal social interaction and by the presence of restricted and repetitive behaviors. Current ASD detection mechanisms are either subjective (survey-based) or focus only on responses to a single stimulus. In this work, we develop machine learning methods for predicting ASD based on electrocardiogram (ECG) and skin conductance (SC) data collected during a sensory challenge protocol (SCP) in which the reactions to eight stimuli were observed from 25 children with ASD and 25 typically developing children between 5 and 12 years of age. The length of the time series makes it difficult to utilize traditional machine learning algorithms to analyze these types of data. Instead, we developed feature processing techniques which allow efficient analysis of the series without loss of effectiveness. The results of our analysis of the protocol time series confirmed our hypothesis that autistic children are greatly affected by certain sensory stimulation. Moreover, our ensemble ASD prediction model achieved 93.33% accuracy, which is 13.33% higher than the best of 8 different baseline models we tested.**

***Index Terms*—autism spectrum disorders, large time series, data-driven autism prediction, feature extraction from time series**

I. Introduction

Autism spectrum disorders (ASDs) are conditions which can lead to impairments in reciprocal social interaction and communication, and restricted and repetitive behaviors in subjects. These neurodevelopmental disorders do not have a cure, but their early detection increases the chances of patients being able to develop coping mechanisms that improve their ability to function in society. Current ASD detection mechanisms are focused on the observation of a subject's social interaction. The instruments used for such assessments are lengthy and require extensive training, which prevents them from being used on the overall population. Before referring the subjects for further evaluation, they are first identified as at-risk via a screening process which is sometimes not accurate [1]. The social responsiveness scale (SRS) test, the most popular of such screening instruments, was shown to only have 0.78 sensitivity and 0.67 specificity [2]. Recent work has identified autonomic and behavioral responses of children with autism to be different from those of typically developing (TD) children in response to auditory [3] or visual stimuli [4].

Our research project utilizes longitudinal physiological data collected from multiple sensors in response to a protocol involving eight stimuli sequentially administered to a mixed group of ASD and TD children. Each protocol took approximately one hour to execute and resulted in large amounts of time series data consisting of millions of correlated values across the length of the protocol. A subject may be affected by one stimulus and its residual effect may be present during the administration of the next stimulus, which suggests the sensor data should be analyzed as a time-dependent series. However, analyzing such large time series is a challenging task, both in terms of the time and the space requirements of the time series analysis methods. In our research, we develop several feature extraction techniques that transform the time series into a form which can be used for efficient analysis and prediction.

We hypothesized that autistic children would be greatly affected by certain sensory stimulation. While TD children can quickly recover to a normal state after the sensory trial, autistic children may be slower to return to normal. In this paper, we describe our experiments and the ASD prediction models we developed based on the ECG and SC response signals recorded during the sensory trials.

II. Literature Review

Current ASD detection mechanisms are based on the observation of a subject's social interaction by either close observers or behavioral therapists [5]. The ASD assessment instruments are often lengthy, require extensive training before they can be administered, and are in general not very accurate [2].

Some researchers have argued that PsNS activity can be used as an indicator for the presence of autism and machine learning-based approaches can be utilized to build predictive models for its detection. Laufer and Nemeth [6] used SC to predict user action, based on a neural network model, by collecting SC data while users were playing an arcade game. Changchun et al. [7] designed a therapist-like support vector machine (SVM)-based affective model as part of a computer-based ASD intervention tool for children using physiological responses that predicts autism with an accuracy of 82.9%.

Much of the existing research in the field of time series analysis is relevant for this study. Dynamic time warping (DTW) [8] is a popular technique that can be used to compare two time-dependent series with different time deformations and speeds. For example, Muda et al. [9] used DTW to create efficient voice recognition algorithms. Juang [10] used DTW based hidden markov models and linear predictive coding techniques to develop speech recognition models. To optimize DTW, Salvador and Chan introduced FastDTW [11], which

* Corresponding Author

is an approximation of DTW with linear time and space complexity and is thus comparatively fast. Mueen et al. have introduced several variants of DTW, including constrained DTW, multidimensional DTW and asynchronous DTW [12].

Piecewise linear approximation (PLA) is one of the most common ways to process time series. It works by approximating a time series of length l with n straight lines using different algorithms, such as the top-down, bottom-up and sliding window approaches. Keogh at el. [13] developed a sliding window and bottom-up algorithm as a means to derive PLA and perform segmentation of time series. Some methods represent time series using motifs, which are derived by identifying frequently occurring patterns in the time series and replacing each pattern with a symbol. Lonardi et al. introduced an algorithm, called enumeration of motifs (EoM) [14], that uses matrix approximation to locate repeated patterns in the time series. Lin et al. introduced the symbolic aggregate approximation (SAX) [15] method, which discretizes original time series data into strings and defines distance measures on the symbolic string representation. Looking for a way to characterize co-evolution patterns in time series, Anastasiu et al. [16] devised an optimal segmentation algorithm that segments users' individual series into varying length segments represented by one of k patterns shared by all the users.

III. Dataset

Our research is based on examining existing data from a study conducted by Dr. Megan C. Chang [3]. The data were collected from various sensors during a SCP [17] in which the reactions to multiple stimuli were observed from 25 children with ASD and 25 typically developing (TD) children between 5 and 12 years of age. Each protocol took 45–90 minutes including preparation, and had three phases: baseline, sensory challenge, and recovery. The baseline and recovery periods lasted 3 minutes each and did not include any stimulation. The sensory challenge consisted of six different sensory stimuli with a pseudorandom pause of 12–17 seconds between the stimuli. Each stimulus was administered for 3 seconds and was presented at least 8 times. Following are the six stimuli, listed in the order they were administered:

- auditory – continuous sound tone of 84 decibels
- visual – 20W strobe light at 10Hz
- auditory – interrupted sound siren at 78 decibels
- olfactory – wintergreen oil passed under the nose
- tactile – touch along the jaw bone with a feather
- vestibular – chair tilted back to a 30 degree angle

Physiological ECG and SC data were continuously collected from multiple sensors in response to the eight stimuli (including the baseline and recovery periods). To obtain an index of PsNS function, ECG activity was collected by placing sensors on the child's chest. To measure the SNS activity, galvanic skin response was measured by attaching sensors to the right hand of the child. The sweat glands secrete more sweat as the subject becomes excited or nervous, which in turn increases skin conductance. The ECG and SC data were collected at a frequency of 500Hz and 40Hz, respectively. This resulted in a very long multivariate time series consisting of approximately 3 million correlated values across the length of the series. Table I provides a description of the dataset that was collected from the 50 subjects.

TABLE I
Dataset Description

# Autistic samples	25
# TD samples	25
Average # data points per subject	2,981,476
Average # data points per stimulus	372,682

Fig. 1 shows an example of the ECG and SC data for a subject spanning 10 seconds. The left y-axis shows the ECG signal, measured in milli-Volts (mV), and the right y-axis shows SC intensities, measured in micro-Siemens (μSiemens).

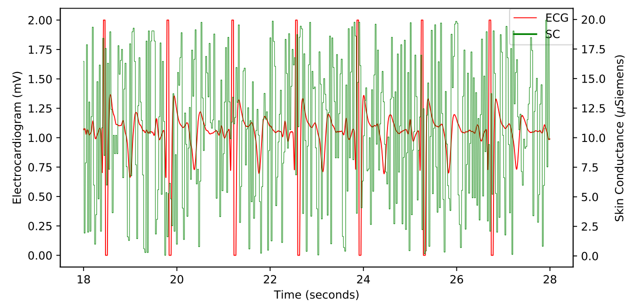

Fig. 1. Time series showing 10 seconds of ECG and SC signal for a subject. The figure is best viewed in color.

IV. Hypothesis and Supporting Evidence

We hypothesize that autistic children are greatly affected by certain sensory stimulation and thus may take longer to return to a normal state than TD children, who can quickly recover after the sensory trial. To test this, we compared the sensory data recorded during an initial baseline rest stage of the protocol, recorded prior to any stimulus being administered, with data recorded during the final recovery rest stage, 30 seconds after the final stimulus was administered. No stimulus was administered during either rest stage. For each subject, we compared the baseline and recovery rest stages by computing the Euclidean DTW distance of the ECG and SC time series recorded during the rest periods.

To analyze the differences between the baseline/recovery distances of autistic and TD children, we fit a Gaussian probability distribution function (PDF) over the distances between the baseline and recovery sensor time series data for autistic and TD children. Fig. 2 shows these functions for the ECG time series. Results show that autistic (solid green line) children exhibit substantially greater differences between their respective baseline and recovery phases than TD children (dashed red line). The PDF means for autistic and TD children were 1.25e+9 and 9.07e+8 and their standard deviations were 6.9e+8 and 4.03e+8, respectively. Results suggest that TD children recover faster, which would explain the shorter distances between the respective baseline and recovery phase time series.

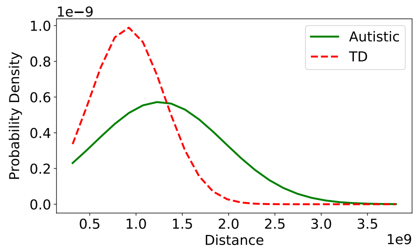

Fig. 2. ECG Gaussian probability density functions of DTW distances between the baseline and recovery stages for autistic and TD subjects.

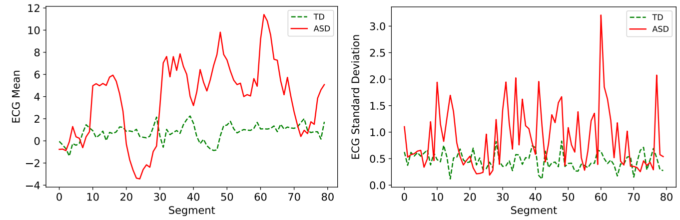

(a) (b)

Fig. 3. Plot showing ECG mean (a) and standard deviation (b) values for a TD subject (dashed) and an autistic subject (solid).

V. Methods

In the remainder of the paper we describe predictive models we developed for autism detection from the ECG and SC response signals recorded during the sensory trials.

A. Feature Extraction

As a means to improve analysis efficiency, we propose to transform the data in a form that is representative of the input signal but has much smaller uniform dimensionality. We devised three different methods to extract features that can be used to conduct specific experiments.

1) Equal Width Partitioning (EWP): During the SCP, a particular stimulus is administered in a specific number of contiguous trials at equal intervals. Thus, we can divide the data into sub-series and still capture the patterns or trends in the series. In this approach, for each subject, the ECG and SC data were first split into 8 equal parts representing the 8 stimuli. The data were then standardized using the mean and standard deviation of the baseline stage, i.e., the first of the 8 splits, which captures the normal ECG and SC signal for a subject prior to any stimulus. The data for each stimulus were then split into n equal parts, and two different approaches were used to encode the information in each split and create different machine learning models for ASD prediction in children using either ECG data, SC data, or both data types.

a) Mean and standard deviation (MSD) representation: In this approach, we represented the n splits for each stimulus using the mean and standard deviation of the data in that split. The final data vector consists of n ECG mean and standard deviation values followed by n SC mean and standard deviation values for each stimulus. Fig. 3 shows the ECG mean and standard deviation values for a TD subject (dashed green line) and for an autistic subject (solid red line) chosen at random. One can observe that the ECG mean and standard deviation values of the autistic subject are generally higher than those of the TD subject. The maximum mean value for the autistic subject is 9.52 and that for the TD subject is 5.08.

b) Slope and intercept (SI) representation: We assume that an autistic child gets more excited when a stimulus is administered as compared to a TD child. When a subject gets excited or nervous, his/her ECG values spike, showing higher maximum and minimum values, and his/her sweat glands secrete more sweat, which in turn increases skin conductance. Thus, we hypothesize that the trend and intensity of the signal contains sensitive information that can be used to predict ASD.

For each of the n splits and for each stimulus, we retrieved all peak (maximum) values and all valleys (minimum) values of ECG data in a cycle. A data point is considered a peak/valley value if its value is greater/smaller than the value of its neighboring data points. After retrieving all peak and valley values in a time series, we represented each split as the slope and intercept of the *best fit line* (BFL) for both peak and valley values. SC values fluctuate less than ECG values do, in general. Therefore, we represented the n splits for each stimulus with the slope and intercept of the BFL over the entire SC time series data in that split. The slope of the BFL captures the variation in trend and the intercept captures the intensity of the signal.

Fig. 4 shows the valley-based slope and intercept representation of the ECG time series and slope and intercept representation for the SC time series, for a TD subject (dashed green line) and for a subject with ASD (solid red line), chosen at random. Time series data represented in these figures were processed using $n = 10$. One can observe that the variation in slopes, especially for ECG valley points and SC data, is higher for the autistic subject as compared to the TD subject. SC data shows more discriminatory characteristics, with autistic subjects showing higher maximum and minimum slope values. One can also observe that the intensity of the signals (ECG and SC), as shown by the intercept graphs, is much higher for autistic subjects as compared to TD subjects.

2) Dynamic Time Warping (DTW): The approach we devised in Section V-A1 transforms the real time series data into a derived format, which may lead to some loss of information. DTW allows us to compare two time series in their raw format. As DTW automatically accounts for time deformations, it will identify similar patterns in two time series even if one of them is longer than the other. In this approach, we used FastDTW, which is an approximation of DTW that has linear time and space complexity [11] to compare the ECG or SC time series between two subjects. Due to the very large size of our time series, both the original DTW and the FastDTW methods failed to compute the distance between our time series for different stimuli on on our very large server with 24 GB of random access memory (RAM), both running out of available memory. We thus split the series into 8 sub-sequences with r% overlap, since each stimulus was repeated 8 times, computed distances between the ith sub-sequence of

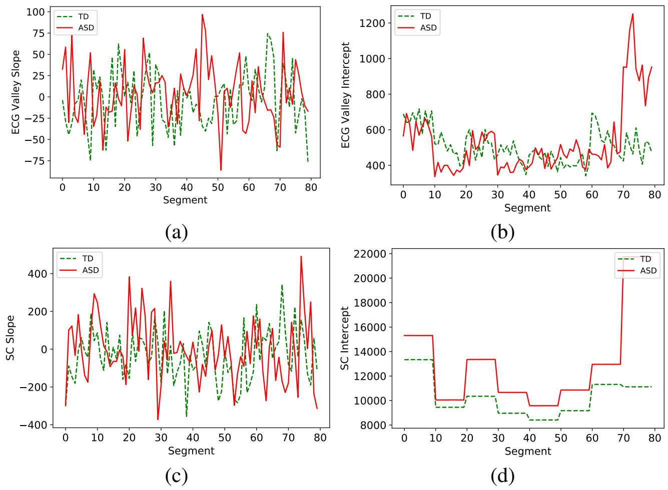

Fig. 4. Plot showing valley-based slope (a) and intercept (b) representation of the ECG time series and slope (c) and intercept (d) representation of the SC time series for a TD subject (dashed) and an autistic subject (solid).

the candidate sequences, and used the maximum of these subsequence distances as the distance between the two candidate sequences. We also tried the efficient DTW method introduced by Mueen et al. [12] and compared it with FastDTW. While it was marginally faster than FastDTW, it required more memory and most of our series could not be computed on our server due to lack of available memory.

3) Symbolic Representation of Time Series: In this approach, we used SAX [15] to represent each of the time series using a SAX vector with a given number of symbols and segments. To get the best representation, we tested with numbers of symbols in the range 2 to 10 and numbers of segments from 2 to 14, in increments of 1. After representing the time series using SAX, we computed pairwise Euclidean DTW distances. These distances were then used to create a KNN-based ASD prediction model.

B. Developing Prediction Models for Autism Detection

1) Base Models: In our experiments, we aim to classify the subject as either Autistic or TD. To perform this binary classification, we trained and tested models using the following methods:

- decision tree (DT)
- k-nearest neighbor (KNN)
- support vector machine (SVM)
- naïve Bayes (NB)
- random forest (RF)
- XGBoost (XGB)
- DTW-based KNN (DTWNN)
- SAX-based KNN (SAXNN)

The first six models consume the features generated using methods specified in Section V-A1. Separate models were created using the MSD and SI feature generation approaches. The DTWNN model is based on the method described in Section V-A2, which utilizes the raw time series for comparison and prediction. The SAXNN model is based on the method described in Section V-A3, which first transforms the raw time series data into its SAX representation before computing pairwise Euclidean DTW distances between the subjects. As we have both ECG and SC data, we wanted to understand how different physiological data help in predicting autism. Thus, we created different models using either only ECG data, SC data, or both ECG and SC data.

2) Ensemble Models: In Section V-B1, we executed experiments for each separate stimulus. After building the separate models for all stimuli, we combined them to build ensemble models and make additional predictions. We used three different approaches to create ensemble models.

a) Majority vote: In this approach, we combined the predictions from all the models for different stimuli and chose the majority predicted class as the final prediction. All the model outputs were given the same weight.

b) Weighted prediction: In this approach, instead of giving the same weight to all the model outputs, we weighed the classification output of each stimulus model with the prediction confidence of its associated model, which ranges between 0 and 1. Considering a vector $\mathbf{w}_c$ of weights associated with each stimulus and the vector $\mathbf{y}$ representing classification predictions of models associated with each stimulus, we compute the final prediction as the linear combination of vectors $\mathbf{w}_c$ and $\mathbf{y}$, $y_c = \mathbf{w}_c{}^T\mathbf{y}$. The vector $\mathbf{y}$ contains the predicted classes, +1 or -1, representing TD and autistic subjects, respectively. A negative y_c prediction value indicates that the models predicted the subject as autistic with higher confidence.

c) Stochastic gradient descent (SGD): In this approach, instead of using the prediction confidence scores from separate stimuli models as weights, as described in Section V-B2b, we learned the contribution of each stimulus towards predicting autism. Some stimuli may contribute positively towards correct prediction, while others may contribute negatively. This can be done by deriving a set of weights such that the linear combination of the weight vector and predictions from different stimulus models results in an accurate binary classification of autistic and TD children. The weight vector $\mathbf{w}_s$, is learned via the SGD algorithm applied to training set predictions. Then, the stimuli predictions in the test set are combined linearly with the weights to generate the final SGD predictions for test samples, computed as $y_s = \mathbf{w}_s{}^T\mathbf{y}_s$.

VI. Experiment Design

We used *accuracy* as the performance measure when comparing the prediction models. Accuracy is an appropriate evaluation metric in our setting, as the dataset contains an equal number of samples for both autistic and TD subjects. It is defined as

$$A = \frac{T_p}{T_s} \times 100,$$

where T_p represents the total number of correct predictions and T_s represents the total number of subjects.

We measure efficiency as the training and prediction runtime, in seconds, for each of the different models. Prediction time is given priority over training time, as training can be done offline but prediction must be executed online, in real time, and thus needs to be fast.

For each prediction or time series analysis method we tested, we tuned available hyper-parameters to obtain the highest possible effectiveness using that method. Due to lack of space, the details of the hyper-parameter tuning can be found in [18].

VII. Results And Discussion

A. Effectiveness Results

1) Base Models: We created eight different models, as described in Section V-B, one for each of the eight stimuli. The first six models, namely, DT, KNN, SVM, NB, RF and XGB, were built using the features extracted based on the two approaches mentioned in the EWP method described in Section V-A1, which splits the time series into a specified number of sections. We created different dataset representations with number of splits, n, ranging from 2 to 13, inclusive. For each value of n, after further splitting the training set into training and validation subsets, we trained different instances of all the six models using different combinations of hyper-parameters. Then, we chose the best model instance based on its validation accuracy. Finally, we re-trained the best model for each algorithm using the chosen best hyper-parameters and the entire original training set.

The DTWNN model utilizes the features extracted using the DTW approach mentioned in Section V-A2, which computes the Euclidean DTW distance between different subjects. Higher distance values imply lower similarity, and *vice versa*. For creating the overlapping splits, we chose $r = 10\%$. The SAXNN model was then built using the SAX feature construction method described in Section V-A3.

Fig. 5 (a) shows the comparison of the best base performing model instances for different algorithms, created using different feature extraction methods and using baseline stage data. We observed that, in almost all cases, the models created using SI features perform better than those created using MSD features. Also, among the two standard time series approaches, the models created using SAX features perform much better as compared to those based on DTW distances.

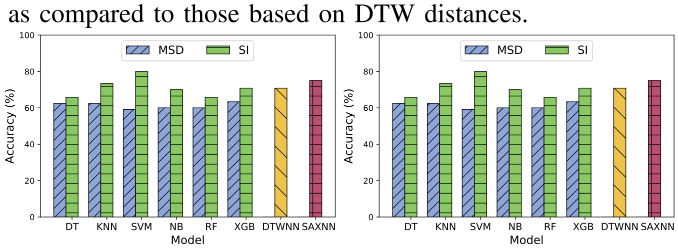

(a) (b)

Fig. 5. (a) Comparison of the best base models for the auditory (tones) stage. (b) Comparison of the best SGD ensemble models.

Table II shows the accuracy scores of the best models for each stimulus. Auditory (tones) and visual stimuli data result in the best performing models, with an accuracy of 80.00% (highlighted in bold). We also observed that two of the best performing models utilize both ECG and SC data for making predictions, showing that both types of sensor data are important in predicting autism.

TABLE II
Best Base Model Accuracy Values Using Each Stimulus

	Accuracy(%)	Model	Data Used
Baseline	75.83	SAXNN	SC
Auditory (Tones)	**80.00**	SVM	Both
Visual	**80.00**	XGB	SC
Auditory (Siren)	77.50	RF	ECG
Olfactory	77.50	SAXNN	SC
Tactile	74.17	SAXNN	SC
Vestibular	78.33	RF	Both
Recovery	73.33	SAXNN	Both

2) Ensemble Models: We combined the results from the models generated using different stimuli, presented in Section VII-A1, to create ensemble models. We compared the accuracy of the ensemble models with the best base models. Ensemble models were created using the three approaches described in Section V-B2.

Fig. 5 (b) shows the comparison of the best SGD ensemble models. We observed that models constructed from SI features outperformed those using MSD ones in almost all cases. The best performing model using SI features is an SGD ensemble XGB model that achieved an accuracy of 93.33%, which is 7.50% higher than the best performing model using MSD features.

As SI features performed better than the MSD ones, further comparisons with DTW and SAX-based approaches were done using only SI features. As mentioned in Sections V-A2 and V-A3, both DTW and SAX-based models are KNN models. Table III shows the best model accuracies for the different tested data processing and modeling methods. One can observe that all the models give the best accuracy using the SGD ensemble method. In this ensemble approach, as described in Section V-B2c, the SGD algorithm is applied on the training set to learn the weights of each stimulus towards making correct predictions.

TABLE III
Best Ensemble Model Accuracy Values

	Accuracy(%)	Ensemble Type	Data Used
DT	92.50	SGD	Both
KNN	81.67	SGD	SC
SVM	87.50	SGD	Both
NB	88.33	SGD	SC
RF	89.17	SGD	Both
XGB	**93.33**	**SGD**	**Both**
DTWNN	77.50	SGD	Both
SAXNN	92.50	SGD	ECG

The best overall performing model was the SGD ensemble XGB model, built using both ECG and SC data, which resulted in an accuracy of 93.33%. The value is approximately 4.16% greater than that achieved using either the majority vote or weighted prediction vote ensemble methods.

As the best accuracy is achieved using both ECG and SC data, we can infer that both types of sensors are important in accurately predicting autism. Additionally, we observed that the next best performing models were DT and SAXNN, which were built using either only ECG data or both ECG and SC data. This further highlights the importance of ECG data in predicting autism in children. In comparison to the best performing base model, the ensemble models performed much better in general. The best performing ensemble model (93.33%) had an accuracy that was 13.33% higher than the best performing base model (80.00%). Even ensemble models built using majority vote (89.17%) and weighted prediction (89.17%) decisions performed better than the base models.

Even though DTW is an important metric for comparing time series, we observed that classification models based on DTW failed to outperform other classification models in our problem. The best accuracy achieved by the DTWNN models was 77.50%, which is approximately 18% lower than that of the best performing model.

B. *Efficiency Results*

We measured the efficiency of the models based on the time taken to train and perform predictions. Fig. 6 shows the comparison of natural log transformed training and prediction times, in seconds. The log scaling in the figure is necessary due to the very wide range of values, which would otherwise hide most results in the graph.

The best performing model in terms of accuracy was the XGB model, which was the third slowest method, taking approximately 49,300 seconds to train and 1.23e$-$4 seconds to predict. On the other hand, the DTW-based model took approximately 4.40 times longer to train and 10^8 times longer to predict in comparison to the SAXNN model. The high execution time for training and prediction makes it difficult to utilize DTW-based models in real-world applications for our problem. On the other hand, the DT model achieved the second highest accuracy (92.50%) and predicts 7 times faster than the best performing XGB model.

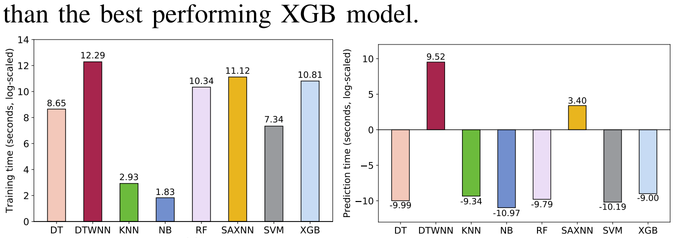

Fig. 6. Comparison of training time (left) and prediction time (right) for all methods.

VIII. Conclusions

In this paper, we described novel techniques we developed for analyzing very large time series of ECG and SC sensor data derived from a sensory trial administered to 50 autistic and TD children. Our analysis showed that autistic children are affected to a higher degree by some stimuli as compared to TD children and take longer to recover. Moreover, the feature extraction methods we developed were both effective and efficient in analyzing multivariate time series with over 2 million values. An XGB-based model trained on vectors constructed using a feature engineering method we developed (SI) achieved the best performance (93.33% accuracy) taking only a millisecond to predict samples.

References

[1] M. Norris and L. Lecavalier, “Screening accuracy of level 2 autism spectrum disorder rating scales: A review of selected instruments,” *Autism*, vol. 14, no. 4, pp. 263–284, 2010, doi:10.1177/1362361309348071.

[2] J. Constantino and C. Gruber, “Social responsive scale (srs) manual,” *Los Angeles, CA: Western Psychological Services*, 2005, doi:10.1177/1534508410380134.

[3] M. C. Chang, L. D. Parham, E. I. Blanche, A. Schell, C.-P. Chou, M. Dawson, and F. Clark, “Autonomic and behavioral responses of children with autism to auditory stimuli,” *American Journal of Occupational Therapy*, vol. 66, no. 5, pp. 567–576, 2012, doi:10.5014/ajot.2012.004242.

[4] T. Chaspari, M. Goodwin, O. Wilder-Smith, A. Gulsrud, C. A. Mucchetti, C. Kasari, and S. Narayanan, “A non-homogeneous poisson process model of skin conductance responses integrated with observed regulatory behaviors for autism intervention,” in *Acoustics, Speech and Signal Processing (ICASSP), 2014 IEEE International Conference on*. IEEE, 2014, pp. 1611–1615, doi:10.1109/ICASSP.2014.6853870.

[5] S. Chandler, T. Charman, G. Baird, E. Simonoff, T. Loucas, D. Meldrum, M. Scott, and A. Pickles, “Validation of the social communication questionnaire in a population cohort of children with autism spectrum disorders,” *Journal of the American Academy of Child & Adolescent Psychiatry*, vol. 46, no. 10, pp. 1324–1332, 2007, doi:10.1097/chi.0b013e31812f7d8d.

[6] L. Laufer and B. Németh, “Predicting user action from skin conductance,” in *Proceedings of the 13th international conference on Intelligent user interfaces*. ACM, 2008, pp. 357–360, doi:10.1145/1378773.1378829.

[7] C. Liu, K. Conn, N. Sarkar, and W. Stone, “Physiology-based affect recognition for computer-assisted intervention of children with autism spectrum disorder,” *International journal of human-computer studies*, vol. 66, no. 9, pp. 662–677, 2008, doi:10.1016/j.ijhsc.2008.04.003.

[8] M. Müller, “Dynamic time warping,” *Information retrieval for music and motion*, pp. 69–84, 2007, doi:10.1007/978-3-540-74048-3_4.

[9] L. Muda, M. Begam, and I. Elamvazuthi, “Voice recognition algorithms using mel frequency cepstral coefficient (mfcc) and dynamic time warping (dtw) techniques,” *arXiv preprint arXiv:1003.4083*, 2010. [Online]. Available: http://arxiv.org/abs/1003.4083

[10] B.-H. Juang, “On the hidden markov model and dynamic time warping for speech recognitiona unified view,” *Bell Labs Technical Journal*, vol. 63, no. 7, pp. 1213–1243, 1984, doi:10.1002/j.1538-7305.1984.tb00034.x.

[11] S. Salvador and P. Chan, “Toward accurate dynamic time warping in linear time and space,” *Intelligent Data Analysis*, vol. 11, no. 5, pp. 561–580, 2007. [Online]. Available: https://content.iospress.com/articles/intelligent-data-analysis/ida00303

[12] A. Mueen and E. Keogh, “Extracting optimal performance from dynamic time warping,” in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 2016, pp. 2129–2130.

[13] E. Keogh, S. Chu, D. Hart, and M. Pazzani, “Segmenting time series: A survey and novel approach,” in *Data mining in time series databases*. World Scientific, 2004, pp. 1–21.

[14] J. Lonardi and P. Patel, “Finding motifs in time series,” in *Proc. of the 2nd Workshop on Temporal Data Mining*, 2002, pp. 53–68.

[15] J. Lin, E. Keogh, L. Wei, and S. Lonardi, “Experiencing sax: a novel symbolic representation of time series,” *Data Mining and knowledge discovery*, vol. 15, no. 2, pp. 107–144, 2007.

[16] D. Anastasiu, A. Rashid, A. Tagarelli, and G. Karypis, “Understanding computer usage evolution,” in *2015 IEEE 31st International Conference on Data Engineering, ICDE 2015*, vol. 2015-May. IEEE Computer Society, 5 2015. doi: 10.1109/ICDE.2015.7113424 pp. 1549–1560.

[17] R. C. Schaaf, T. W. Benevides, E. Blanche, B. A. Brett-Green, J. Burke, E. Cohn, J. Koomar, S. J. Lane, L. J. Miller, T. A. May-Benson *et al.*, “Parasympathetic functions in children with sensory processing disorder,” *Frontiers in Integrative Neuroscience*, vol. 4, p. 4, 2010, doi:10.3389/fnint.2010.00004.

[18] M. Kapoor and D. C. Anastasiu, “A data-driven approach for detecting autism spectrum disorders,” San José State University, San José, CA, USA, Tech. Rep. 2019-1, 2019. [Online]. Available: pdf/papers/2019-Kapoor-asp-tr.pdf