PREDICITING SEX AND AGE USING SWIPE-GESTURE FORM A MOBILE DEVICE

by

Storm Davis

A Senior Honors Project Presented to the

Honors College

East Carolina University

In Partial Fulfillment of the

Requirements for

Graduation with Honors

by

Storm Davis

Greenville, NC

May, 2020

Approved by:

Dr. Nasseh Tabrizi

College of engineering and technology, department of computer science

Predicting Sex and Age Using Swipe-Gesture Data from a Mobile Device

Abstract

Swipe-gestures are by far the most common way to interact with mobile devices such as phones, tablets, and even some computers. As touch-screen technology has improved, the possibility of obtaining high-quality swipe-gesture data from touch-screen devices has become more and more prevalent. This has led to the exploration of its use in further improving authentication systems, and more recently, as the basis for soft biometrics prediction. This paper discusses the process of using swipe-gesture data for prediction of sex and age of individuals using mobile devices. The software used to obtain the data is presented, the features collected from the swipe data are detailed, and the machine learning classifiers are displayed in a way that the experiment can be replicated. During this experiment, a total of ten well-known classifiers have been used. The results of this analysis have further confirmed the possibility of predicting sex, obtaining an accuracy rate of 79% for a single classifier as well as a group average of almost 70%. Moreover, in the prediction of age category, the results are even more encouraging, obtaining an accuracy rate of nearly 80% on average as well as several of the classifiers performing well above the average.

1. Introduction

Traditionally, biometrics has been defined by [1] as the science of automatically recognizing people based on physical or behavioural characteristics. More recently, however, is the emergence of soft biometrics that uses traits such as height, weight, sex, and hair color as descriptive characteristics, but they cannot be used exclusively to identify a unique individual [2]. The fusion of soft-biometrics and traditional biometrics has been used to enhance a system's reliability as well as the overall performance by integrating typical physiological biometrics (iris, face, fingerprint) along with behavioural or soft biometrics (iris color, sex, height, age) traits of the same identity [3]. In such a fusion, soft biometrics are used as to complement the hard biometrics performance. However, the use of soft biometrics is not exclusive to biometrics scenarios. In the past, standard input devices such as mouse and keyboard have been used as biometric resources to help classify the emotional state of individuals based on their interaction [4]. In recent years, similar methods have been implemented using mobile devices and swipe-gestures as the basis for interaction. For example, the use of swipe-gestures on a mobile device in a recent study [5] has confirmed the possibility of sex prediction with a high rate of accuracy (78%) using data from two different directions. Predictions based on soft biometrics may soon be able to provide dynamic and enhanced interfaces that adapt to different individuals or emotional states [5].

The first instance of touch-screen technology used for the purposes of studying human-computer interaction dates back over 50 years to a paper published by [5,6]. In his paper, authors of [6] used wires connected to a CRT (cathode-ray tube) device that was sensitive to touch, simulating a touch-screen. Today, touch-screen technologies are everywhere, such as mobile devices, tablets, and even some computers that support a touch interface. With worldwide adoption of the touch-screen, these technologies have been getting continuously better, which enables the possibility to use data from touch input to predict user information [5]. In the recent past, these capabilities have led to more effective and secure means of authentication [7-10]. However, few works exist that use soft biometrics as a basis for prediction based on swipe-gestures on a touch-screen.

This paper aims to assess the possibility of predicting both the sex and age of individuals using swipe-gesture data from a game played on a mobile device. More specifically, we want to evaluate the accuracy of the prediction, as well as determining which classification algorithm best suits swipe-gesture data for soft biometric prediction.

2. Related Works

The use of soft biometrics and related swipe-gesture data have been proposed in several studies as a way of further enhancing authentication methods on mobile, touch-screen devices. One of the earliest studies [3], dating back to 2004, used the fusion of soft and hard biometrics to improve identity verification. Some years later, another related study emerged with the proposition of replacing text-based passwords with graphical patterns [7]. During this time (2010), touch-screen technologies were beginning to take off thanks to increasingly capable devices from Apple, Samsung, and other manufacturers.

With more capable touch-screen technologies, better data can be obtained through the soft biometrics of swipe-gestures. The study [10] was conducted on the soft biometrics data collected from users utilizing the Android lock pattern. For those unfamiliar with Android, the

Android lock pattern consists of a 3 x 3 grid of circles that users can use to create unique designs for authentication on Android devices. The authors [10] proved that security could be enhanced using the soft biometrics data from the swipe-gestures during the process of drawing the unlock pattern. The use of graphical passwords inputted through multi-touch-screens on tablets was also analyzed [11] for user authentication, where the authors' work yielded a result in which the equal rate error was 10% using a single multi-touch gesture on the tablet screen. Other studies like [12, 13] have analyzed the strength of such pattern authentication methods to determine the factors and complexity that make a particular pattern more secure.

The study [14] combined touch biometrics along with various sensors in a mobile device to verify the ownership. The biometric touch gestures observed were Tap, Scroll, and Fling. The features captured during each touch-gesture were: coordinates on the device screen, touch duration, as well as touch pressure across the use of three different applications. One hundred users interacted with the applications (Message, Album, and Twitter) in a static test. With ten interactions, the model demonstrated 80% accuracy for identifying a non-owner, whereas with six interactions, the model was nearly 100% accurate for the owner of the smartphone.

Even now, the use of touch-screen soft biometrics is still being used to identify users on mobile devices more securely. In [15], the authors proposed a method for enhancing traditional authentication systems through the incorporation of soft biometric information as a second level of user authentication. In this study, users draw each digit of a password rather than typing it in as usual authentication systems often use. The authors then analyzed the handwritten password, determining the discriminative power of each handwritten digit, and how the length of the passcode affected its robustness. The results of the study found that the handwriting method produced equal error rates of approximately 4.0% when an attacker knows the passcode. When compared to traditional methods of password authentication, the attacker would have a nearly 100% success rate when knowing the password under the same scenario (password obtained by looking over the shoulder of the user).

Up to this point, soft biometrics has nearly only been used to complement mobile authentication for more secure interaction with mobile devices. However, in [5] and [16-19] the authors used soft biometrics, not as a means for enhancing user authentication, but rather, as a way to predict specific characteristics of the user. In [16], the authors investigate sex prediction from iris images using a combination of geometric and texture features. The study yielded accuracies of up to 90% using the BioSecure Database for sex prediction. [17] used mouse biometrics for the classification of sex, incorporating kinematic and spatial analyses of 256 mouse movements performed by each user. The study yielded encouraging results that were further validated through binary logistic regressions.

[5] is more aligned to the purpose of this paper than previous studies on security and user authentication. In [5], swipe data were collected from 116 participants as they swiped through a set of jokes on a mobile device. The swipe data contained features such as total length, total time, width, height, area, as well as average thickness and pressure for each swipe. They used averages from this data as input for the selected machine learning classifiers (Decision Tree, Naïve Bayes, Support Vector-Machines, and Logistic Regression) to predict the sex of participants. According to [5] the results of the study yielded an encouraging 78% accuracy rate for the prediction of a user's sex. Also, in [18], the authors were able to obtain encouraging results for the prediction of age, sex, and operating hand through the use of keystroke extraction on a mobile device. Although this study does incorporate age in their prediction scheme, timing-based keystroke features were used rather than swipe-gestures. Moreover, in a very recent study, even greater performance is achieved for sex prediction on a mobile device using swipe-gestures [19].

As far as we are aware, [5] and [19] are the only previous works that have used swipe-gesture data as a basis for the prediction of sex, and there exist no other works that use swipe data (collected from a mobile game) to predict both the sex and age group of a user.

3. Methodology

This section describes our experimental methodology as well as the steps taken from the development of the mobile game to the swipe-gesture data acquisition and eventual sex and categorical age prediction. It will elaborate on how the data were collected, which features from the data were utilized, as well as how these features were selected and used with machine learning algorithms to obtain the prediction for a user's sex and categorical age.

3.1. Data Collection

The swipe-gesture data used in this experiment were collected manually. Rather than using an available dataset, the data were collected by asking approximately 48 participants to play the mobile game. Each participant was required to give their first name, age category (above or below 40), and their sex (male or female). In total, there were 27 males and 21 females with 18 participants being over the age of 40, and the rest below. While the participants played the game on a touch screen device, their swipe-gestures were being captured. The instrument used was an iPad 2 IOS tablet with a 9.70-inch touch display and a resolution of 1024 x 768 pixels (132 pixels per inch). No screen protectors or external casing was on the device during the participants' interaction during the study. Participants were instructed to hold the device in landscape orientation with their non-dominant hand while using the index finger on their dominant hand to perform the swipe- gestures necessary to play the game (see Figure 1).

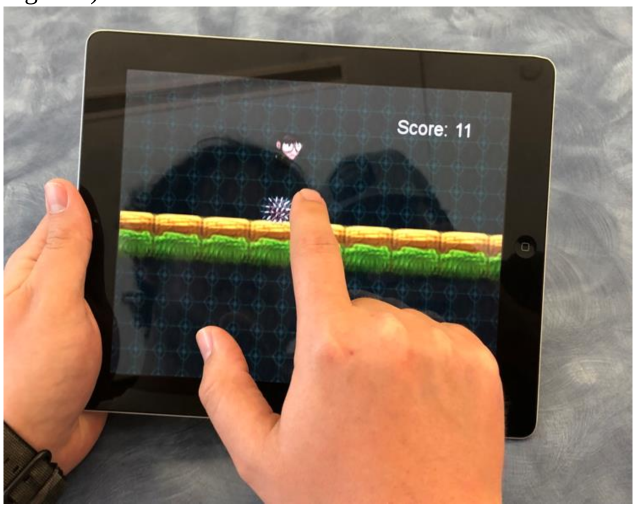

Figure 1: Mobile game used for swipe-gesture acquisition

Each participant played the game for approximately five minutes, continually restarting the game session until the time limit expired. The time limit of five minutes was determined to be best suitable for the manual collection process. Five minutes allows sufficient swipe-gesture data to be collected while also not exhausting a user's interest. The swipe-gestures of each participant were captured using a mobile game developed specifically for this study's purpose. The IOS game automatically captured the swipes of each participant as well as storing their personal information (age group, sex, name) with their swipe data locally on the iPad until it was uploaded to the server. The game itself consisted of a series of obstacles where the user would have to swipe up or down to avoid a collision. If the user's character collided with a barrier, the game would restart. Each participant averaged nearly 134 swipe-gestures during the span of five minutes with the game.

The swipe-gesture data collected consisted of angle from start to end, average acceleration, ending acceleration (vx, vy, and vz), staring acceleration (x, y, and vz, z), total length, total time, as well as width and maximum speed for each swipe.

3.2. Machine Learning Classifiers

The raw swipe data collected were used as inputs for a set of machine learning classifiers [20] in order to predict sex as well as age group. The data were initially tested using a set of 10 different classifiers in order to test the possibility of prediction. This set of classifiers consisted of the following: Nearest Neighbors [21], Linear and Radial Support-Vector Machines [22], Gaussian Process [23], Decision Tree and its ensemble Random Forest [24,25], Neural Network [26], Adaptive Boosting [27], as well as Naïve Bayes [28] and Quadratic Discriminant Analysis [29]. The collection of ten classifiers, as well as the parameters used, can be seen in Table 1. From the initial set of ten, the classifiers were narrowed down to three that best predicted a user's sex and age group based on the data collected, as well as those that produced poor predictions.

Table 1. Machine learning classifiers and parameters used.

Nearest Neighborhood	**n_neighbors=5, weights='uniform', algorithm='auto', leaf_size=30, p=2, metric='minkowski', metric_params=None, n_jobs=None, **kwargs**
Linear SVM	**C=0.025, kernel='linear', degree=3, gamma='auto_deprecated', oef0=0.0, shrinking=True, probability=False, tol=0.001, cache_size=200, class_weight=None, verbose=False, max_iter=-1, decision_function_shape='ovr', random_state=None**
RBF SVM	**C=1, kernel='rbf', degree=3, gamma=2, oef0=0.0, shrinking=True, probability=False, tol=0.001, cache_size=200, class_weight=None, verbose=False, max_iter=-1, decision_function_shape='ovr', random_state=None**
Gaussian Process	**kernel=1.0*RBF(1.0), optimizer='fmin_l_bfgs_b, n_restarts_optimizer=0, max_iter_predict=100, warm_start=False, copy_X_train=True, random_state=None, multiclass='one_vs_rest', n_jobs=None**
Decision Tree	**criterion='gini', splitter='best', max_depth=5, min_samples_split=2, min_samples_leaf=1, min_weight_fraction_leaf=0.0, max_features=None, random_state=None, max_leaf_nodes=None, min_impurity_decrease=0.0, min_impurtity_split=None, class_weight=None, presort=False**
Random Forest	**n_estimators=10, criterion='gini', max_depth=5, min_samples_split=2, min_samples_leaf=1, min_weight_fraction_leaf=0.0, max_features=1, random_state=None, max_leaf_nodes=None, min_impurity_decrease=0.0 min_impurity_split=None, class_weight=None, presort=False, bootstrap=True, oob_score=False, n_jobs=None, verbose=0, warm_start=False**
Neural Network	**hidden_layer_sizes=(100,), activation='relu',solver='adam', alpha=1, batch_size='auto', learning_rate='constant', learning_rate_init=0.001, power_t=0.5, max_iter=200, shuffle=True, random_state=None, tol=0.0001, verbose=False, momentum=0.9, nesterovs_momentum=True, early_stopping=False, validation_fraction=0.1, warm_start=False,**

	beta_1=0.9, beta_2=0.999, epsilon=1e-08, n_iter_no_change=10
Adaptive Boosting	**base_estimator=None, n_estimators=50, learning_rate=1.0, algoritm='SAMME.R', random_state=None**
Naïve Bayes	**priors=None, var_smoothing=1e-09**
Quadratic Discriminant Analysis	**priors=None, reg_param=0.0, store_covariance=False, tol=0.0001**

The classifiers used in this study are well-known machine learning algorithms that have been used successfully in various other machine learning studies [30-33]. All ten of the used classifiers were from the Scikit-learn repository [34] for python. Three different machine learning classifiers were selected for being the best candidates for predicting the sex and age group of a user. These classifiers range from Decision Trees (Random Forest) and ensemble boosting classifiers (AdaBoost) to probabilistic models (Gaussian Process). The next 2 sections (3.3 and 3.4) describe the three best classifiers, along with two classifiers that did not perform well when compared to other studies using the same classifiers.

3.3. Best Performing Classifiers

Decision Tree (Random Forest): Decision tree learning is undoubtedly one of the most popular algorithms for automatic machine learning. A Random Forest is not a Decision Tree, but rather, a collection of Decision Trees. A Random Forest is a meta estimator that fits several Decision Tree classifiers on various samples of the used dataset. Random Forests also use averaging to improve accuracy for a prediction as well as limit over-fitting. In short, it uses multiple Decision Trees to generate results that are typically more accurate than a simple Decision Tree algorithm is on its own [24].

Adaptive Boost (AdaBoost): AdaBoost is an ensemble classifier that combines several smaller classification algorithms to create a classifier that performs well and produces accurate predictions. The basic idea behind AdaBoost was presented by Freud and Schapire in 1997 [27]. Without delving too deep into how this classifier works, it combines several weak classifiers and uses the assignment of weights to produce a result that is an average of the weighted classifiers. This boosting method for classification has successfully been used for the prediction of subcellular localization in a previous study.

Gaussian Process: The Gaussian Process classifier is an algorithm for probabilistic classification in which predictions take the form of class probabilities. The specific classifier used is from Scikit-learn, but it is based on algorithms from [23] described by Rasmussen and Williams. This classifier is said to use "lazy learning" (A generalization of training data is delayed until a query is made.) and generates a measure of similarity between points to predict a specified value.

3.4. Poor Performing Classifiers

Although the above classifiers performed quite well, two classifiers did not produce expected results. Of the ten total classifiers used, three performed consistently well across both predictions for sex and age group (discussed in 3.2.1), two performed mediocrely, and the rest fell somewhere in between. The Naïve Bayes classifier, as well as the Quadratic Discriminant analysis (QDA), produced results that were below the average when compared to the eight other classifiers in both sex and age category predictions.

Naïve Bayes: The Naïve Bayes classifier [35,28] is based on a probabilistic Bayes' rule and is best suited for high dimensional inputs. The classifier assumes that the effect of a particular feature is independent of the effects of other predictors. This assumption is seemingly unrealistic, but it has performed quite well for additional studies in a range of uses.

Some possible reasons the classifier did not perform as well as expected could include limited dimensionality in the input data or unrealistic and simplified assumptions made by the algorithm.

Quadratic Discriminant Analysis: the Quadratic Discriminant Analysis (QDA) classifier uses a Quadratic Decision Boundary (QDB) that is generated by fitting class conditional densities to data and utilizing Bayes' rule. This classifier also has performed well in other predictive studies, but its results from this study were mediocre at best.
It is possible that this classifier performed similarly poor because of its reliance on Bayes' rule, a key component in the Naïve Bayes classifier that also produced below-average results.

3.5. Evaluation of Predictive Models

The machine learning models were created for all ten of the classifiers and the two categories of age group and sex. 75% of the acquired data were used as a training set, and the remaining 25% was used to test the predictive capability of the model. The evaluation of each model was carried out through a score value that each model produced. The score associated with each model was a percentage that quantified how accurate each model was in its prediction. Along with the score, the precision and recall were computed for each model to better analyze the results.

4. Results

In this section, the results obtained from the machine learning models are analyzed and discussed in detail. In *A*, the score of each classifier is presented for the prediction of sex. In the following section, the score of each model is given for the prediction of age category.

4.1. Sex Prediction

As can be seen in Figure 2, some machine learning models performed quite differently than others in terms of accuracy. The average rate of accuracy for the prediction of sex across the ten classifiers was approximately 70%. Some classifiers, like Decision Tree and Gaussian Process, performed well above the average in this case. Decision Tree showed an accuracy of 80% whereas Gaussian Process and AdaBoost were not far behind at 76% and 77% accuracy, respectively. However, other classifiers like Naïve Bayes and QDA failed to meet the average. Their scores were well below the average at 45% and 62%, respectively. The remaining six classifiers produced rates of accuracy that fell between the average and the higher accuracy rates of the Decision Tree and Gaussian Process. Without incorporating the poor results from Naïve Bayes and Quadratic Discriminant Analysis, the average for the remaining eight classifiers was an encouraging 74% accuracy rate.

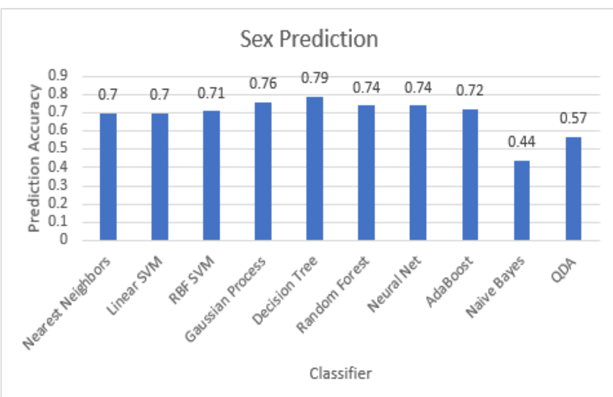

Figure 2: Classifier accuracy for sex prediction

However, accuracy alone is not enough to justify the results of the classifiers. Table 2 shows the precision and recall values computed from a confusion matrix generated for each of the classifiers when attempting to predict a user's sex. Although the accuracy of the Support-Vector Machine (SVM) classifier was around the group average of 70%, its confusion matrix illustrated that this model was not a good fit for the data. Linear SVM's matrix only showed values for the false positives and true positives sections in the matrix, meaning that the classifier solely predicted that every participant was a male. Similar results were produced by the Radial Basis Function (RBF) SVM classifier [36]. The confusion matrix for each classifier can be seen in Figure 3. This being said, the precision and high recall values for the SVM classifiers were not good representations of how they actually performed.

Table 2: Precision and recall values produced by each classifier for sex prediction

Classifier	Precision	Recall
Nearest Neighbors	**0.75**	**0.86**
Linear SVM	**0.7**	**1**
RBF SVM	**0.71**	**0.99**
Gaussian Process	**0.78**	**0.92**
Decision Tree	**0.8**	**0.92**
Random Forest	**0.76**	**0.99**
Neural Net	**0.75**	**0.97**
AdaBoost	**0.79**	**0.97**
Naïve Bayes	**0.79**	**0.28**
QDA	**0.71**	**0.65**

Most other classifiers produced reasonable confusion matrices, and the most accurate classifiers (Decision Tree, Gaussian Process, AdaBoost) demonstrated relatively high recall and precision values to match. Decision Tree and its ensemble Random Forest posted precision values of 0.80 and 0.76, respectively while their recall values were closer to 1. Similarly, the precision and recall values of the Gaussian Process and AdaBoost were also encouraging.

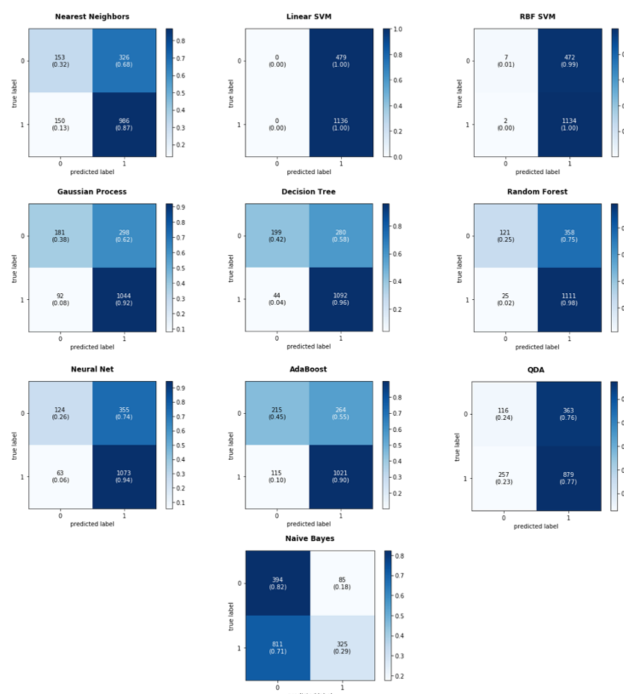

Figure 3: Confusion matrices produced by each classifier for sex prediction

The confusion matrices produced by Naïve Bayes and QDA seemed reasonable, and they each produced precision scores that were comparable to the other classifiers. However, the recall score from Naïve Bayes was even worse than its rate of accuracy. The recall of QDA, although quite better than that of Naïve Bayes, was still low when compared to the other classifiers.

4.2. Age Category Prediction

Figure 4 illustrates the performance, in terms of accuracy, for the prediction of age category for each of the ten classifiers. Keep in mind, age category is a general prediction for a user's age. In this study, there was an age category for those users below 40 years old as well as one for those above.

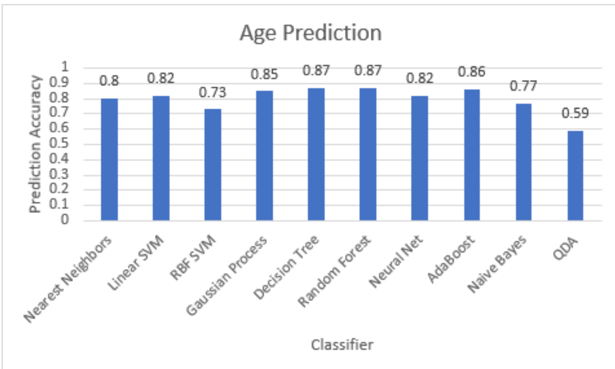

Figure 4: Classifier accuracy for age prediction

Similar to the results of the sex prediction, there are apparent performance differences between the classifiers used. However, the average rate of accuracy across the ten classifiers for age category prediction (80%) was significantly better than that of sex prediction (69%). Moreover, the results of each classifier, independently, were also mostly better than their predictions on the sex of a user. Naïve Bayes improved significantly, just barely performing below the average with an accuracy rate of 77%. Much like the results from sex prediction, Decision Tree, and Gaussian Process performed quite well. Decision Tree -as well as Random Forest- produced rates of accuracy near 87% while the Gaussian Process was not far behind at 85%. Another ensemble classifier, AdaBoost, boasted an 86% rate of accuracy, falling in between the performance of Decision Tree and Gaussian Process in terms of accuracy.

With the prediction of age group, the values of precision and recall for each classifier varied. Precision values for some classifiers were closer to 0.70 while others managed to achieve above 0.80. Table 3 shows the results of the precision and recall calculations performed from the confusion matrix of each classifier.

Table 3: Precision and recall values produced by each classifier for age prediction

Classifier	Precision	Recall
Nearest Neighbors	0.69	0.57
Linear SVM	0.81	0.49
RBF SVM	0.73	0.01
Gaussian Process	0.78	0.64
Decision Tree	0.82	0.69
Random Forest	0.83	0.62
Neural Net	0.81	0.56
AdaBoost	0.81	0.67
Naïve Bayes	0.72	0.3
QDA	0.38	0.73

More interestingly, the recall values of each classifier in the prediction of age category were quite low when compared to the values of sex prediction. Naïve Bayes obtained a similarly low recall value when compared to its results on sex prediction, but most others demonstrated significantly different values. While the recall values of the Gaussian Process, Decision Tree, and AdaBoost were between 0.90 and 1 in the prediction of sex, their recall values in the prediction of age group were between 0.60 and 0.70, a significant drop in recall. However, their precision values were quite comparable to the values produced for sex prediction.

The confusion matrix for each classifier can be seen in Figure 5. Similar to confusion matrix issues for sex prediction, the RBF SVM classifier's confusion matrix was not at all reasonable. As a result, it produced an inferior recall value of 0.01. However, all other classifiers produced rational confusion matrices with fairly balanced recall and precision scores to reflect.

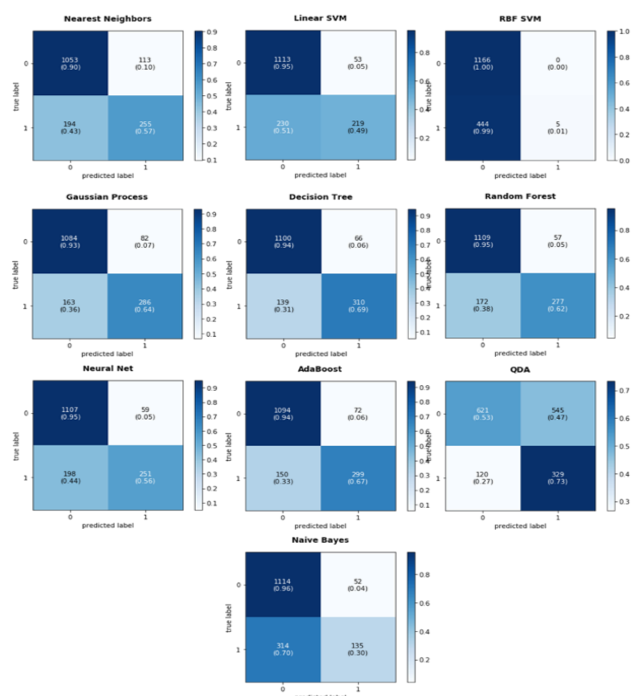

Figure 5: Confusion matrices produced by each classifier for age prediction

5. Conclusion

As of now, the primary form of interacting with a mobile device is through the use of a touch-screen. From this study, we have seen that swipe-gesture data are useful for the prediction of soft biometrics such as sex and age category. However, this opens the possibility for other soft biometrics predictions such as single- or two-handed usage, handedness, or even emotion [5]. Moreover, as touch-screen devices improve, it will be possible to obtain higher quality swipe-gesture data that will enhance the prediction of such soft biometrics.

The soft biometric traits discussed above can be utilized to further improve authentication systems on mobile devices, enhance the experience between mobile device and user, or even modify mobile advertisements or product suggestions to better suit the current user. This opens up the possibility of enhancements in the Human-Computer Interaction (HCI) field. With the ability to predict certain characteristics of a user (age, sex, etc.), interactions between systems and users can be more adaptive and tailored to the current user.

In short, this paper has explored and analyzed the possibility of sex and age category prediction using swipe-gesture data collected from a touch-screen mobile device. In this study, an IOS tablet (iPad 2) was used, but very similar swipe-gestures are used across all touch-screen devices, independent of their operating system or manufacturer.

The results of this exploration and analysis have further confirmed the possibility of sex prediction as first explored in [5] as well as demonstrating the potential of age category prediction. This study took a very different approach than [5], from the development of a mobile game and manual collection of data to the classifiers and data analysis techniques used. Moreover, the data used in similar, recent, predictive studies like [5,18,19] used heavy pre-processing of data whereas this study did not. For sex prediction, an average accuracy rate of 70% was demonstrated across ten classifiers, whereas the Decision tree alone produced an accuracy rate of 80% using the swipe-gesture data collected. Age category prediction produced even more encouraging results with an 80% group average accuracy rate across ten classifiers as well as Decision tree's high 87% rate of accuracy alone.

Regarding the most suitable machine learning classifier for the prediction of both sex and age category, the Decision Tree and its ensemble proved to be the best performers in both cases. Decision tree had the highest rate of accuracy in both predictive studies as well as demonstrating strong precision and recall values. Moreover, its ensemble, Random Forest, produced similarly high rates of accuracy as well as encouraging recall and precision values for both sex and age category prediction. Although each classifier utilized all the features from the swipe-gesture data, Decision Tree and its ensemble seemed to be the most capable of predicting sex and age category. Keep in mind that the swipe-gesture data used as inputs for the classifiers was raw, unprocessed data. More detailed analysis is needed to determine what features in the swipe-gesture data prompted enhanced prediction in the Decision Tree based models when compared to other classifiers.

Due to the small sample population used for this study, it is essential to acknowledge that general conclusions cannot be drawn due to this limitation. However, this study serves as valuable knowledge for the possibility of using swipe-gesture data for prediction. Future research on this topic of soft biometrics prediction could include exploring other traits for prediction like single or double handed use, handedness, or emotion. But also, future studies could explore the effect of large datasets on previous studies for prediction. More interestingly, however, are the possibilities of using similar predictive models to enhance Human-Computer Interaction (HCI).

6. Acknowledgements

This research is supported in part by grant #1560037 from the National Science Foundation.

References

[1] D. A. Reid, S. Samangooei, C. Chen, M. S. Nixon, and A. Ross, “Soft Biometrics for Surveillance: An Overview,” in Handbook of Statistics, vol. 31, Elsevier, 2013, pp. 327–352.

[2] S. Denman, C. Fookes, A. Bialkowski and S. Sridharan, "Soft-Biometrics: Unconstrained Authentication in a Surveillance Environment," 2009 Digital Image Computing: Techniques and Applications, Melbourne, VIC, 2009, pp. 196-203.

[3] R. Zewail, A. Elsafi, M. Saeb and N. Hamdy, "Soft and hard biometrics fusion for improved identity verification," The 2004 47th Midwest Symposium on Circuits and Systems, 2004. MWSCAS '04., Hiroshima, Japan, 2004, pp. I-225.

[4] A. Kołakowska, "A review of emotion recognition methods based on keystroke dynamics and mouse movements," 2013 6th International Conference on Human System Interactions (HSI), Sopot, 2013, pp. 548-555.

[5] O. Miguel-Hurtado, S. V. Stevenage, C. Bevan, and R. Guest, “Predicting sex as a soft-biometrics from device interaction swipe gestures,” Pattern Recognition Letters, vol. 79, pp. 44–51, Aug. 2016.

[6] E. A. Johnson, "Touch display--a novel input/output device for computers," in Electronics Letters, vol. 1, no. 8, pp. 219-220, October 1965.

[7] W. Hu, X. Wu and G. Wei, "The Security Analysis of Graphical Passwords," 2010 International Conference on Communications and Intelligence Information Security, Nanning, 2010, pp. 200-203.

[8] D. Crouse, H. Han, D. Chandra, B. Barbello, and A. K. Jain, “Continuous authentication of mobile user: Fusion of face image and inertial Measurement Unit data,” in 2015 International Conference on Biometrics (ICB), Phuket, Thailand, 2015, pp. 135–142.

[9] N. Sae-Bae, N. Memon and K. Isbister, "Investigating multi-touch gestures as a novel biometric modality," 2012 IEEE Fifth International Conference on Biometrics: Theory, Applications and Systems (BTAS), Arlington, VA, 2012, pp. 156-161.

[10] J. Angulo and E. Wästlund, “Exploring Touch-Screen Biometrics for User Identification on Smart Phones,” in Privacy and Identity Management for Life, vol. 375, J. Camenisch, B. Crispo, S. Fischer-Hübner, R. Leenes, and G. Russello, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 130–143.

[11] N. Sae-Bae, N. Memon and K. Isbister, "Investigating multi-touch gestures as a novel biometric modality," 2012 IEEE Fifth International Conference on Biometrics: Theory, Applications and Systems (BTAS), Arlington, VA, 2012, pp. 156-161.

[12] P. Andriotis, T. Tryfonas, and G. Oikonomou, “Complexity Metrics and User Strength Perceptions of the Pattern-Lock Graphical Authentication Method,” in Human Aspects of Information Security, Privacy, and Trust, vol. 8533, T. Tryfonas and I. Askoxylakis, Eds. Cham: Springer International Publishing, 2014, pp. 115–126.

[13] Andriotis, P., Tryfonas, T., Oikonomou, G., Yildiz, C.: A pilot study on the security of pattern screen-lock methods and soft side channel attacks. In: 6th ACM Conference on Security and Privacy in Wireless and Mobile Networks, WiSec 2013, pp. 1–6. ACM (2013)

[14] C. Bo, L. Zhang, X.-Y. Li, Q. Huang, and Y. Wang, “SilentSense: silent user identification via touch and movement behavioral biometrics,” in Proceedings of the 19th annual international conference on Mobile computing & networking - MobiCom ’13, Miami, Florida, USA, 2013, p. 187.

[15] R. Tolosana, R. Vera-Rodriguez and J. Fierrez, "BioTouchPass: Handwritten Passwords for Touchscreen Biometrics," in IEEE Transactions on Mobile Computing.

[16] M. Da Costa-Abreu, M. Fairhurst and M. Erbilek, "Exploring Sex Prediction from Iris Biometrics," 2015 International Conference of the Biometrics Special Interest Group (BIOSIG), Darmstadt, 2015, pp. 1-11.

[17] Van Balen, Nicolas & T. Ball, Christopher & Wang, Haining. (2017). A Behavioral Biometrics Based Approach to Online Sex Classification. 475-495. 10.1007/978-3-319-59608-2_27.

[18] A. Buriro, Z. Akhtar, B. Crispo and F. Del Frari, "Age, Sex and Operating-Hand Estimation on Smart Mobile Devices," 2016 International Conference of the Biometrics Special Interest Group (BIOSIG), Darmstadt, 2016, pp. 1-5.

[19] Ankita Jain, Vivek Kanhangad, Sex recognition in smartphones using touchscreen gestures, Pattern Recognition Letters, Volume 125, 2019, Pages 604-611, ISSN 0167-8655.

[20] S. B. Kotsiantis. 2007. Supervised Machine Learning: A Review of Classification Techniques. In Proceedings of the 2007 conference on Emerging Artificial Intelligence Applications in Computer Engineering: Real Word AI Systems with Applications in eHealth, HCI, Information Retrieval and Pervasive Technologies, Ilias Maglogiannis, Kostas Karpouzis, Manolis Wallace, and John Soldatos (Eds.). IOS Press, Amsterdam, The Netherlands, The Netherlands, 3-24.

[21] T. Cover and P. Hart, "Nearest neighbor pattern classification," in IEEE Transactions on Information Theory, vol. 13, no. 1, pp. 21-27, January 1967.

[22] C. Cortes and V. Vapnik, “Support-vector networks,” Mach Learn, vol. 20, no. 3, pp. 273–297, Sep. 1995.

[23] Williams, C., & Rasmussen, C. (1996). Gaussian Processes for Regression. In D. Touretzky, M. Mozer, & M. hasselmo (Eds.), Advances in Neural Processing Systems 8 (pp. 514-520). Cambridge, MA, USA: MIT Press.

[24] S. R. Safavian and D. Landgrebe, "A survey of decision tree classifier methodology," in IEEE Transactions on Systems, Man, and Cybernetics, vol. 21, no. 3, pp. 660-674, May-June 1991.

[25] L. Breiman, “Random Forests,” Machine Learning, vol. 45, no. 1, pp. 5–32, Oct. 2001.

[26] D. F. Specht, "A general regression neural network," in IEEE Transactions on Neural Networks, vol. 2, no. 6, pp. 568-576, Nov. 1991.

[27] Yoav Freund and Robert E Schapire. 1997. A Decision-Theoretic Generalization of On-Line Learning and an Application to Boosting. J. Comput. Syst. Sci. 55, 1 (August 1997), 119-139.

[28] Hand, D. J. and Yu, K. (2001), Idiot's Bayes—Not So Stupid After All?. International Statistical Review, 69: 385-398.

[29] P. A. Lachenbruch and M. Goldstein, “Discriminant Analysis,” Biometrics, vol. 35, no. 1, p. 69, Mar. 1979.

[30] V. P. B. felean, "Analysis and Predictions on Students' Behavior Using Decision Trees in Weka Environment," 2007

29th International Conference on Information Technology Interfaces, Cavtat, 2007, pp. 51-56.

[31] Hengl T, Heuvelink GBM, Kempen B, Leenaars JGB, Walsh MG, et al. (2015) Mapping Soil Properties of Africa at 250 m Resolution: Random Forests Significantly Improve Current Predictions. PLOS ONE 10(6): e0125814.

[32] Rogers, Alex, Maleki, Sasan, Ghosh, Siddhartha and Jennings, Nicholas R,(2011) Adaptive Home Heating Control Through Gaussian Process Prediction and Mathematical Programming. Second International Workshop on Agent Technology for Energy Systems (ATES 2011), Taiwan, Province of China. pp. 71-78 .

[33] Wen-Cong Lu, Yu-huan Jin, Bing Niu, Kai-Yan Feng, Yu-Dong Cai and Guo-Zheng Li, " Predicting Subcellular Localization with AdaBoost Learner", Protein & Peptide Letters (2008) 15: 286.

[34] Pedregosa, Fabian & Varoquaux, Gael & Gramfort, Alexandre & Michel, Vincent & Thirion, Bertrand & Grisel, Olivier & Blondel, Mathieu & Prettenhofer, Peter & Weiss, Ron & Dubourg, Vincent & Vanderplas, Jake & Passos, Alexandre & Cournapeau, David & Brucher, Matthieu & Perrot, Matthieu & Duchesnay, Edouard & Louppe, Gilles. (2012). Scikit-learn: Machine Learning in Python. Journal of Machine Learning Research. 12.

[35] D.D. Lewis, Naïve (Bayes) at forty: the independence assumption in information retrieval, Machine Learning: ECML-98, Springer, Berlin Heidelberg, 1998, pp. 4–15.

[36] B. Scholkopf et al., "Comparing support vector machines with Gaussian kernels to radial basis function classifiers," in IEEE Transactions on Signal Processing, vol. 45, no. 11, pp. 2758-2765, Nov. 1997.