

# Smartwatch-Based Biometric Gait Recognition

Andrew H. Johnston and Gary M. Weiss

Department of Computer and Information Science, WISDM Laboratory  
Fordham University

441 East Fordham Road, Bronx NY 10458

{ajohnston9, gaweiss}@fordham.edu

## Abstract

*The advent of commercial smartwatches provides an intriguing new platform for mobile biometrics. Like their smartphone counterparts, these mobile devices can perform gait-based biometric identification because they too contain an accelerometer and a gyroscope. However, smartwatches have several advantages over smartphones for biometric identification because users almost always wear their watch in the same location and orientation. This location (i.e. the wrist) tends to provide more information about a user's movements than the most common location for smartphones (pockets or handbags). In this paper we show the feasibility of using smartwatches for gait-based biometrics by demonstrating the high levels of accuracy that can result from smartwatch-based identification and authentication models. Applications of smartwatch-based biometrics range from a new authentication challenge for use in a multi-factor authentication system to automatic personalization by identifying the user of a shared device.*

## 1. Introduction

The modern smartwatch contains multiple sensors, substantial computing power, and the ability to communicate with smartphones and similar devices via Bluetooth. It is a relatively recent development; perhaps the first truly modern smartwatch, the Pebble, became available in 2013. Many additional smartwatches were released in 2014, most of which operate with Android phones and run the Android Wear subsystem. These watches include the LG G Watch, Moto 360, Samsung Gear Live, and Sony Smartwatch 3. While sales of these devices have been modest, the introduction of the Apple Watch in April 2015 has greatly increased the interest in such devices. While it is unclear whether smartwatches will become as ubiquitous as smartphones, market projections indicate that in the year 2020 nearly 400 million smartwatches will ship and thus represent a 25-fold increase over 2014 sales [1].

Like their smartphone counterparts, most modern smartwatches include both an accelerometer and gyroscope sensor, so that they may support similar capabilities and applications, such as health applications that require sensor-based activity recognition. Accelerometer and gyroscope sensors are ideal for gait-based biometrics; our prior work has shown that these sensors can enable smartphones to implement gait-based biometrics [2]. In this paper we show that smartwatches are similarly capable of performing gait-based identification and authentication. The identification task uses a single predictive model to identify a user within a group of users. The authentication task uses a per-user predictive model to determine if an unknown user is a “match” or is an imposter.

The study described in this paper utilizes smartwatches to collect and store sensor data which is ultimately sent to a server for processing. Android-based smartwatches and smartphones are used because they are currently the most readily-available smartwatches. More specifically, the experiments in this study utilize Samsung Galaxy S3 phones and LG G watches largely due to their low cost. Although our collection is limited to just one model of a watch, the data collected should be comparable to what would have been collected from any other smartwatch running the Android Wear subsystem.

Smartwatch-based biometrics is of great interest because of the relatively low cost of smartwatches, the fact that the device can be worn and carried with you (i.e. is mobile), and because it can transmit data and results to other devices via Bluetooth or, with the assistance of a paired smartphone, via the Internet. The smartwatch also has several important advantages over a smartphone and other mobile devices for gait-based biometrics. Primarily, smartwatches are almost always worn in the same location and in the same orientation (unlike smartphones, smartwatches running the Android Wear subsystem do not reorient the screen if the user holds it upside down, guaranteeing the user wears it in a specific orientation). This is a huge advantage over smartphones, where both the location and orientation of the de-

vice may vary depending on the user, what the user is wearing, and the activity that the user is performing. Changing the location of the device on the body will reduce the effectiveness of the biometric models and some locations simply may not generate a suitable biometric signature. The location issue for smartphones is especially prevalent with women, because unlike men they will often carry the phone “off-body”, such as in a pocket book. Further, in the case of the smartwatch, there is an added advantage resulting from the most common place on the body that the device will be carried: the wrist. If a user is walking, much more movement occurs at the wrist than at the most common smartphone location (upper thigh in a pant pocket) and hence it should prove superior for biometric identification.

Smartwatch-based gait biometrics can support several important applications. A smartwatch-based biometric system can serve as the foundation for a delegated authentication system. For example, as a user approaches his smart-home, his smartwatch could transmit its accelerometer and gyroscope sensor readings to the house, which would then compare the readings to past readings and open the door if they match. Gait-based biometrics can also be employed in a two-factor authentication scheme, thereby serving as a supplemental biometric mechanism to augment or replace traditional modalities of fingerprint and facial recognition.

This paper is organized as follows. Related work is introduced and discussed in Section 2. Section 3 describes the procedure for collecting the training data from the users, and how this data is transformed into a format suitable for conventional classification algorithms. The methodology for generating the identification and authentication models is then described in Section 4 and the results of these experiments are presented and analyzed in Section 5. Finally, in Section 6 we describe our conclusions and planned extensions to the research.

## 2. Related Work

Although this paper describes the first use of commercial smartwatches for gait-based biometrics, there is a substantial amount of prior work on gait-based biometrics in general. Most gait-based biometric work can be categorized as machine vision-based or wearable sensor-based, though there is some work on biometrics using floor-based sensors [3]. The work described in this paper falls under the wearable sensor-based category and hence we focus on work in that area. It should be pointed out that the machine vision-based approach [4, 5] has the advantage that it does not inherently impose requirements on the subject and in fact can be employed without the subject’s consent or even knowledge. Such systems are of particular interest in the field of surveillance, where video-based monitoring systems (e.g. CCTV) are already in place. Thus, for example, vision-based gait recognition can be used to supplement

face recognition biometrics in airports or other venues with special security requirements.

There are numerous wearable sensor-based systems for gait recognition, although most of these systems do not use commercially available devices. One system, which utilized 36 test subjects, produced good results using an accelerometer placed on the belt, at the subject’s back [6]. Another consisted of 21 test subjects and used a tri-axial accelerometer-based device attached to the user’s right lower leg [7]. Yet another system used two wireless sensors to collect tri-axial accelerometer and a bi-axial gyroscope readings from the ankles of four users [8].

More recent systems have utilized smartphones for gait recognition. An early effort placed a smartphone on the right hip of 6 test subjects and was able to achieve a 93.1% recognition rate [9]. However, this recognition rate was only achievable by taking specific steps to calibrate the orientation of the device with the user’s posture. Another similar study, which used the Google G1 phone, also placed the smartphone at the hip of 51 test subjects and was able to produce an Equal Error Rate of 20% [10]. Equal Error Rate (EER), a common metric for evaluating biometric performance, is the rate at which the false acceptance rate (i.e., the rate at which an imposter is incorrectly identified as the authorized user) equals the false rejection rate (i.e. the rate at which the authorized user is incorrectly identified as an imposter). A slightly more recent effort used Android smartphones with 36 test subjects, and employed continuous wavelet transforms to achieve excellent results with an EER of 1% [11].

A prior study from our lab, which employed Android phones, achieved strong results using only common descriptive statistical features and a 10-second sliding window [2]. This system, which employed 36 test subjects, achieved an identification rate of about 90% using a single 10 second sample of accelerometer data, but was able to perfectly identify all 36 users when using several 10 second samples. With respect to authentication, the system was able to achieve an average positive authentication rate of 85.9% and a negative authentication rate of 95.0% using a single 10 second sample; perfect authentication performance was achieved when using multiple 10 second samples.

The work in this study extends prior research by continuing the movement toward commercial devices for gait-based biometrics. This study employs more subjects than most other research studies and our results are competitive with, if not better than, most of the smartphone-based systems. Considering the advantages that we have identified of smartwatches over smartphones, certain measures taken in prior work (e.g. to calibrate the orientation of the phone with the user’s posture or clip it to the hip in a fixed position), which would not be realistic in everyday use, become unnecessary.

### 3. Data Collection and Feature Extraction

This section describes the process for collecting training data from the study participants, as well as the process for extracting useful features and transforming the time-series sensor data into examples that can be handled by conventional classifier induction programs (e.g. decision trees).

#### 3.1. Data Collection

The training data necessary for building the biometric identification and authentication models is a sample of each participant’s gait, as measured by the accelerometer and gyroscope on the smartwatch. The data collection process begins with participants enrolling in our study, which is approved by Fordham University’s Institutional Review Board, and granting written informed consent. This is necessary because we are technically “experimenting” on human subjects and there is a very small risk of injury (i.e. a participant could trip while walking). The participant then answers a few survey questions (e.g. age, gender, height, etc.), which are used to characterize our study population and can be used for more in-depth analyses in the future. The participant then fastens a smartwatch on the wrist of their non-dominant hand and places a Bluetooth-paired smartphone in their pocket. Both devices run a simple custom-designed application that controls the data collection process.

The application instructs the participant to input their name on the phone, turn off the phone’s screen, and then place the phone in their pocket. The participant is then instructed to walk for several minutes, using their normal gait, on a flat surface with relatively few turns. The smartphone instructs the smartwatch running our paired data collection app to collect the accelerometer and gyroscope data at 20Hz. Each sensor generates values for the  $x$ ,  $y$ , and  $z$  axes and appends a timestamp to the values. After 5 minutes the watch sends the data to the phone for transmission to our research server (a local copy is retained on the phone to preserve the data should transmission fail). After transmission, both the smartwatch and the smartphone vibrate to notify the user that the data collection process is complete and they can stop walking.

#### 3.2. Feature Extraction and Data Transformation

There are several ways to prepare the raw sensor data before using it for biometrics. Some gait-based biometric work utilizes the data within the time domain [12, 13, 14], but other successful systems map the time-series sensor data into examples using a sliding window approach, which permits the use of conventional classifier induction systems that cannot handle time-series data. This study utilizes the same sliding window approach employed in our prior smartphone-based study [2]. The transformation process

partitions the time-series sensor data into 10 second non-overlapping windows and then generates relatively simple features based on the data in the window. The accelerometer and gyroscope features are generated independently, but using the same feature encoding schemes. All features except one are based on the sensor values for a single axis, but 3 versions of each feature are generated corresponding to the 3 axes associated with the sensor data. Because the data is sampled at 20Hz and the window size is 10 seconds, there are 200 time-series values per axis per window, and 600 sensor values per window. This holds for both the accelerometer and the gyroscope sensor. Each of these 600 time-series values is transformed into 43 features using the feature encodings described below; they are also used in our prior smartphone-based biometric study [2]. The value in subscripts specifies how many features of the given type are generated.

- Average[3]: Average sensor value (each axis)
- Standard Deviation[3]: Standard deviation (each axis)
- Average Absolute Difference[3]: Average absolute difference between the 200 values and the mean of these values (each axis)
- Time Between Peaks[3]: Time between peaks in the sinusoidal waves associated formed by the data as determined by a simple algorithm (each axis)
- Binned Distribution[30]: The range of values is determined (maximum – minimum), 10 equal-sized bins are formed, and the fraction of the 200 values within each bin is recorded. (each axis)
- Average Resultant Acceleration[1]: For each of the 200 sensor samples in the window, take the square root of the sum of the squares of the  $x$ ,  $y$ , and  $z$  axis values, and then average them.

Each example, which represents 10 seconds of data, is appended with an numerical ID value that uniquely identifies each participant. This ID field serves as the class value for the identification task while it is mapped into a binary valued class variable for authentication tasks.

### 4. Experiment Methodology

This section describes the methodology used for running the biometric identification and authentication experiments.

#### 4.1. Dataset

The experiments utilize sensor data collected from 59 study participants, of which 57% are male and 43% are female. The participants range in age from 18 to 39, with a majority being college-aged (i.e. 18-23). As described in Section 3, the smartwatch accelerometer and gyroscope

sensor data were collected at a rate of 20Hz for both sensors. A single raw sensor record includes the data for the 3 axes for one sensor. Our raw data set contains 650,458 of these records (half for each sensor). Given that there are 200 of these records per 10 second window, this corresponds to 4.5 hours of data per sensor. This equates to 4.6 minutes of data per user (the average is less than 5 minutes because we only collected 2 minutes of data from the first few users before raising the limit). Note that the classifier induction algorithms do not operate on the raw data, but the transformed data.

## 4.2. Classifier Induction Algorithms

The WEKA data mining suite [15] is freely available and implements a large number of classifier induction algorithms. This study utilizes the following four WEKA algorithms: Multilayer Perceptron (MLP), Random Forest, Rotation Forest, and Naive Bayes. The Multilayer Perceptron algorithm is a neural network algorithm, Random Forest and Rotation Forest are ensembles of decision trees, and Naive Bayes is a probabilistic classifier based on Bayes' Theorem. These models are all suitable for real-time biometric identification because they can be generated and evaluated rapidly.

## 4.3. Identification Experiments

The biometric identification task is to identify a user from a pool of users based on a sample of their sensor data. This requires a sample of data from all users in the population that can be used in the training phase. The experiments for this task are quite simple. The transformed datasets associated with the accelerometer and gyroscope data are each used to train and evaluate biometric identification models, using 10-fold cross validation. For this task the class variable is the User ID and there are 59 class values, one per participant. A set of experiments is conducted for each of the 4 algorithms and 2 sensors, such that there are 8 (i.e.  $(4 \times 2)$ ) sets of experiments performed.

## 4.4. Authentication Experiments

In authentication, each user has their own authentication model (i.e. classifier), and when a sample of sensor data is provided, the task is determine if the sample belongs to the user or to an imposter. Authentication is a specialized case of identification, but rather than identifying the imposter, authentication seeks to distinguish the imposter from the legitimate user.

The authentication experiments create and evaluate a model for each of the 59 participants in the study. In each case, the first half of the user's data is used in the training set, and the second half is used in the testing set. Then eight random users are chosen and one minute of data is chosen from a random position for each of these users. Four of the

random users are placed into the training set and the other four are placed into the test set (as was done in a prior research study [11]). The authentication model is then built using each of the four algorithms mentioned; the process is then repeated for the other sensor.

## 5. Results

This section presents and analyzes the results for the experiments described in the previous section. The identification results are presented first, followed by the authentication results.

### 5.1. Identification Results

The identification experiments, as described in Section 4.3, involve building a single predictive model to identify a specific user from a set of users. At the lowest level, our results are based on identifying each user based on a single 10 second sample (referred to as an instance) of walking data. However, we can improve our results by using more than a single instance and then employing a majority voting scheme to identify the user (we call this the "Most Predicted User" scheme). In order to demonstrate how this scheme works, and to provide greater insight into the results, a confusion matrix generated by WEKA for this identification task is presented in Table 1. The actual confusion matrix is a  $59 \times 59$  matrix, but due to space limitations Table 1 shows the results only for the first ten users (i.e. the upper left quadrant of the matrix). The results in this table are based on an identification model generated from accelerometer data and using the Random Forest algorithm.

User	1	2	3	4	5	6	7	8	9	10
1	<b>7</b>	0	0	0	0	0	0	0	0	0
2	0	<b>17</b>	0	0	5	0	0	0	0	0
3	0	0	<b>27</b>	0	0	0	0	0	0	0
4	0	0	0	<b>9</b>	0	1	0	0	0	0
5	0	4	0	0	<b>20</b>	0	0	0	0	0
6	0	0	1	1	0	<b>26</b>	1	0	0	0
7	0	0	0	0	0	2	<b>23</b>	0	0	0
8	0	0	1	0	1	0	1	<b>20</b>	0	0
9	0	0	0	0	0	0	0	0	<b>11</b>	0
10	0	0	0	0	0	0	0	0	0	<b>28</b>

Table 1. Partial confusion matrix

The rows in Table 1 correspond to the actual users and the columns to the predicted users, so that the values in the diagonal (noted in boldface) correspond to correct identifications and all other values correspond to errors. The results clearly indicate that the model usually identifies the user (the number of total predictions varies because different amounts of labeled data were collected from different users). Based only on these partial results we could compute the accuracy for identifying a single user or the accuracy aggregated over all 10 users. For example, the accuracy

for identifying User 1 is 100% (7/7) while the accuracy for identifying User 2 is 77% (17/22). The overall accuracy would be simply the total number of correct predictions divided by the total number of predictions. These “raw accuracies” are so named as they are based off of a single 10 second instance. The accuracies computed solely from Table 1 are clearly optimistic estimates since the matrix is not complete and hence there will be errors that are not accounted for. Nonetheless, even from this partial matrix it is clear that the diagonal values tend to be the largest values.

In our identification scenario we assume that all of the sensor data from a device comes from the same user, so it is feasible to use multiple 10 second instances to make a prediction. Thus, we can use a simple strategy to improve the identification performance: set the identity of the user to the most frequently predicted user[2]. Based on our partial data in Table 1, for User 2 we make one prediction based on the 22 instances of data and identify the user as User 2. In this case the accuracy improves from 77% to 100% as the 5 errors no longer impact the final identification. This strategy requires a larger sample of data but yields dramatically improved performance.

We now turn to the full set of experimental results. Table 2 shows the raw accuracy results (i.e. using a single instance) for the two different sensors using the four classification algorithms. These results show that even a 10 second instance is sufficient to identify a user most of the time, especially if one uses the accelerometer data and a method other than Naive Bayes (note the accelerometer sensor data is clearly more informative than the gyroscope data). These results far outperform the “strawman” approach of predicting the most common class (i.e. the user with the greatest number of instances), which would yield an accuracy of only 1.96%.

Sensor	Naive Bayes	Random Forest	Rotation Forest	MLP	Avg.
Accel	66.8%	82.9%	84.0%	83.1%	79.2%
Gyro	52.4%	59.0%	66.4%	70.5%	62.1%

Table 2. Identification accuracy using a 10 second instance

The corresponding results interpreted with the Most Predicted User strategy are shown in Table 3. This strategy always leads to perfect results. Based on our visual inspection of the full confusion matrices, and based on the fact that there usually is not a second user who gets nearly as many “votes” as the actual user, we believe that for our population of 59 users, we could get perfect identification accuracy using fairly small samples of data.

## 5.2. Authentication Results

The results for the authentication experiments described in Section 4.4 are presented in this section. Table 4 provides the accuracy results for our authentication models. Recall

Sensor	Naive Bayes	Random Forest	Rotation Forest	MLP	Avg.
Accel	100%	100%	100%	100%	100%
Gyro	100%	100%	100%	100%	100%

Table 3. Identification accuracy using Most Predicted User

that these results are aggregated over all 59 of our authentication models (i.e. one per subject) and all authentication decisions presented here are based on a single 10 second instance of walking data.

Sensor	Naive Bayes	Random Forest	Rotation Forest	MLP	Avg.
Accel	94.9%	98.3%	98.3%	98.0%	97.2%
Gyro	92.9%	93.8%	94.8%	94.6%	93.8%

Table 4. Authentication accuracy using a 10 second instance

The results in Table 4 indicate that smartwatch-based authentication can be relatively accurate when using only a single 10 second instance of data. We see that, consistent with our identification results, the accelerometer sensor data is more useful than the gyroscope data and the Naive Bayes algorithm significantly underperforms the other algorithms.

Table 5 presents the EER results for authentication. The EER rates are quite good and perform competitively with the results from other biometrics studies, including several studies that used smartphones (see Section 2).

Sensor	Naive Bayes	Random Forest	Rotation Forest	MLP	Avg.
Accel	4.5%	1.4%	2.5%	2.0%	2.6%
Gyro	9.6%	9.6%	7.0%	6.3%	8.1%

Table 5. Authentication EER results

## 6. Conclusions and Future Work

This paper described an effective system for performing smartwatch-based biometric identification and authentication. It demonstrates that gait, as measured by commercial-grade smartwatch sensors, is sufficient to identify an individual with modest accuracy. Furthermore, a simple fixed-width sliding window approach is shown to be sufficient for representing the time-series sensor data. While the results are not necessarily sufficient to enable a smartwatch to serve as a single means of identification or authentication, the results are already strong enough for this technology to participate in a biometrics system that utilizes multiple identification or authentication mechanisms.

There are many ways to extend this work. This study demonstrated that the accelerometer data is more useful than the gyroscope data, but it is very possible that a fusion of the data from these two sensors will yield improved results. Similarly, the smartwatch sensor data can be fused

with the smartphone sensor data to see if having multiple sources of data improves biometric performance. We have also been experimenting with more sophisticated features, which capture specific elements of a user's gait, and plan to investigate if these features can yield additional improvements.

Another one of our goals for future work is to expand the evaluation of this technology, so that it is applied to more realistic situations. Thus, we plan to expand our user base significantly, increase the diversity of the users (especially with respect to age), and evaluate how the system operates when the training and test samples are collected over longer time frames. A key limitation of our current work is that the data for each user is collected on a single day; preliminary experiments indicate that our results degrade significantly when the training and test data are collected from different days. We are not quite sure about the specific reason for this, but our future work will focus on identifying and addressing this issue, since to be useful a biometric system must be able to function over reasonably long time intervals. One final goal of ours is to incorporate this biometrics technology into a real-time system.

## 7. Acknowledgements

The authors would like to thank Anthony Canicatti, Alexander Despotakis, Susanne George, Florian Shabanaj, and Cameron Walker for their assistance with this work, as well as the many test subjects who participated in this study.

## References

- [1] J. Siegal. (2013, sept) Smartwatch sales set to explode, expected to top 100m within four years. [Online]. Available: <http://bgr.com/2013/09/27/smartwatch-sales-forecast-2020/>
- [2] J. R. Kwapisz, G. M. Weiss, and S. A. Moore, "Cell phone-based biometric identification," in *Biometrics: Theory Applications and Systems (BTAS), 2010 Fourth IEEE International Conference on*. IEEE, 2010, pp. 1–7.
- [3] D. Gafurov, "A survey of biometric gait recognition: Approaches, security and challenges," in *Annual Norwegian Computer Science Conference*, 2007, pp. 19–21.
- [4] J. Wang, M. She, S. Nahavandi, and A. Kouzani, "A review of vision-based gait recognition methods for human identification," in *Digital Image Computing: Techniques and Applications (DICTA), 2010 International Conference on*. IEEE, 2010, pp. 320–327.
- [5] L. Wang, T. Tan, H. Ning, and W. Hu, "Silhouette analysis-based gait recognition for human identification," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 25, no. 12, pp. 1505–1518, 2003.
- [6] J. Mantjarvi, M. Lindholm, E. Vildjiounaite, S.-M. Makela, and H. Ailisto, "Identifying users of portable devices from gait pattern with accelerometers," in *Acoustics, Speech, and Signal Processing, 2005. Proceedings.(ICASSP'05). IEEE International Conference on*, vol. 2. IEEE, 2005, pp. 973–976.
- [7] D. Gafurov, K. Helkala, and T. Söndrol, "Biometric gait authentication using accelerometer sensor," *Journal of computers*, vol. 1, no. 7, pp. 51–59, 2006.
- [8] A. Annadhorai, E. Guenterberg, J. Barnes, K. Haraga, and R. Jafari, "Human identification by gait analysis," in *Proceedings of the 2Nd International Workshop on Systems and Networking Support for Health Care and Assisted Living Environments*, ser. HealthNet '08. New York, NY, USA: ACM, 2008, pp. 11:1–11:3. [Online]. Available: <http://doi.acm.org/10.1145/1515747.1515762>
- [9] S. Sprager and D. Zazula, "Gait identification using cumulants of accelerometer data," in *Proceedings of the 2nd WSEAS International Conference on Sensors, and Signals and Visualization, Imaging and Simulation and Materials Science*. World Scientific and Engineering Academy and Society (WSEAS), 2009, pp. 94–99.
- [10] M. O. Derawi, C. Nickel, P. Bours, and C. Busch, "Unobtrusive user-authentication on mobile phones using biometric gait recognition," in *Intelligent Information Hiding and Multimedia Signal Processing (IIH-MSP), 2010 Sixth International Conference on*. IEEE, 2010, pp. 306–311.
- [11] F. Juefei-Xu, C. Bhagavatula, A. Jaech, U. Prasad, and M. Savvides, "Gait-id on the move: Pace independent human identification using cell phone accelerometer dynamics," in *Biometrics: Theory, Applications and Systems (BTAS), 2012 IEEE Fifth International Conference on*. IEEE, 2012, pp. 8–15.
- [12] C. Nickel, H. Brandt, and C. Busch, "Classification of acceleration data for biometric gait recognition on mobile devices," *BIOSIG*, vol. 11, pp. 57–66, 2011.
- [13] M. O. Derawi, C. Nickel, P. Bours, and C. Busch, "Unobtrusive user-authentication on mobile phones using biometric gait recognition," in *Intelligent Information Hiding and Multimedia Signal Processing (IIH-MSP), 2010 Sixth International Conference on*. IEEE, 2010, pp. 306–311.
- [14] H. M. Thang, V. Q. Viet, N. D. Thuc, and D. Choi, "Gait identification using accelerometer on mobile phone," in *Control, Automation and Information Sciences (ICCAIS), 2012 International Conference on*. IEEE, 2012, pp. 344–348.
- [15] M. Hall, E. Frank, G. Holmes, B. Pfahringer, P. Reutemann, and I. H. Witten, "The weka data mining software: an update," *ACM SIGKDD explorations newsletter*, vol. 11, no. 1, pp. 10–18, 2009.