Acoustic Gait Analysis Using Support Vector Machines

Jasper Huang[1], Fabio Di Troia[1], and Mark Stamp[1]
[1]*Department of Computer Science, San Jose State University, San Jose, California*
jhuang821@student.fuhsd.org, fabioditroia@msn.com, mark.stamp@sjsu.edu

Keywords: Gait recognition; support vector machine; acoustic analysis; biometric.

Abstract: Gait analysis, defined as the study of human locomotion, can provide valuable information for low-cost analytic and classification applications in security, medical diagnostics, and biomechanics. In comparison to visual-based gait analysis, audio-based gait analysis offers robustness to clothing variations, visibility issues, and angle complications. Current acoustic techniques rely on frequency-based features that are sensitive to changes in footwear and floor surfaces. In this research, we consider an approach to surface-independent acoustic gait analysis based on time differences between consecutive steps. We employ support vector machines (SVMs) for classification. Our approach achieves good classification rates with high discriminative one-vs-all capabilities and we believe that our technique provides a promising avenue for future development.

1 Introduction

Biometrics—the analysis of human physiological and behavioral characteristics—has become a leading method of person identification, verification, and classification. Image-based biometric techniques include iris, face, and fingerprint analysis. These typically necessitate subject cooperation and a highly controlled environment. For example, current fingerprinting and hand geometry systems often require a specific placement or orientation of the finger or hand, while facial recognition algorithms can be complicated by non-frontal facial images and uncooperative subjects. One biometric that addresses many of these limitations is human gait, a person's manner of walking. Gait analysis is the systematic study of human walking, which can be defined as the method of body locomotion for support and propulsion (Whittle, 2014). In comparison to most other biometrics, gait analysis is less intrusive, and allows for passive monitoring.

Previous work has demonstrated that people can recognize subjects by observing illuminated moving joints (Johansson, 1973), and research has also shown that people are able to reliably identify co-workers by listening to their footstep sounds (Makela et al., 2003). Human gait analysis has also been applied to problems in medical diagnosis (Murray, 1967).

The majority of current gait analysis techniques are applied in the visual domain—the field of visual gait recognition has been an active area of research for at least the past 15 years (Lee, 2002). Although there has been significantly less focus on acoustic-based gait recognition, information from gait audio has been shown to be useful as well. Audio-based research has yielded promising results using frequency analysis of both footstep sounds and those produced by clothing movement and contact (Geiger et al., 2014). However, one issue that arises with frequency-related data is sensitivity to contact surfaces, as variations in footwear, clothing, floor surfaces, etc., produce different sounds that may adversely affect classification rates. Time-domain analysis would appears to be a promising avenue of research that could mitigate some of these issues. In this research, we consider temporal acoustic analysis for gait recognition.

Although gait information alone might not be sufficient to distinguish each individual in a large database, gait analysis could still be advantageous for low-cost pre-screening. Such technologies could be used, for example, as part of an anomaly detection system in a smart home, for access control in high-security buildings, and for surveillance in indoor environments. Our audio-based system can easily be combined with existing techniques for a multimodal gait recognition system.

Compared to visual systems, acoustic techniques are not sensitive to changes in illumination, visibility, and walking angle. The time-based features we analyze are also likely to be more robust with respect to changes in clothing, footwear, and floor surface. In addition, acoustic analysis requires only inexpensive sensors with minimal sensor density.

Our gait analysis method is also relevant in clinical and biomechanical applications. The fields of gait and human movement sciences have proven useful in the analysis, diagnosis, and treatment of various afflictions (Lai et al., 2009). For example, gait analysis has been applied to the diagnosis of cerebral palsy (Kamruzzaman and Begg, 2006) and age-related issues (Begg and Kamruzzaman, 2005; Begg et al., 2005). With minor modification, our technique has the potential for application in such domains.

The organization of the remainder of this paper is as follows. In Section 2, we provide background information, including a brief survey of relevant literature in gait and audio analysis. In Section 3, we outline our methodology and in Section 4 we present experimental results. We conclude the paper in Section 5 and provide suggestions for future work.

2 Related Work

Gait has generally been interpreted as a visual phenomenon, necessitating that the person be seen to be recognized or classified. Not surprisingly, the bulk of research in gait analysis has been visual-based, relying on video or image data. In authentication applications, gait analysis typically involves side-view silhouette features from the spatiotemporal domain (Wang et al., 2003). A variety of data analysis techniques have been employed, such as hidden Markov models on sequences of feature vectors corresponding to different postures (Kale et al., 2002; Sundaresan et al., 2003). In another comprehensive study (Man and Bhanu, 2006), a spatiotemporal "gait energy image" forms the basis for gait analysis. A review of human motion analysis in the field of computer vision appears in (Aggarwal and Cai, 1997).

Gait recognition methods that do not involve video or audio information include Doppler techniques (Kalgaonkar and Raj, 2007; Otero, 2005; Tahmoush and Silvious, 2009), floor pressure sensors (Qian et al., 2008; Middleton et al., 2005), and smartphone accelerometers (Sprager and Zazula, 2009; Nishiguchi et al., 2012; Mantyjarvi et al., 2005). In medical applications, many gait analysis techniques are also inherently image-based, relying on features derived from kinetics and kinematics (Aggarwal and Cai, 1997), such as hip and pelvis acceleration (Aminian et al., 2002). A survey of computational intelligence in gait research within the medical field can be found in (Lai et al., 2009).

An alternative approach to gait analysis is to capture the sounds of walking, that is, acoustic gait analysis. Current approaches typically treat the task as an instance of the more general problem of sound recognition, which is related to automatic speech recognition and speaker identification. Current general purpose speech recognition systems often apply hidden Markov models (HMMs) and other statistical techniques to n-dimensional real-valued vectors consisting of coefficients extracted from words or phonemes. Gaussian mixture models are commonly used for text-independent verification (Bimbot et al., 2004; Reynolds, 1995). Such models combine computational feasibility with statistical robustness.

Similar methods have been extended to acoustic gait recognition. Footstep detection and identification techniques are presented in (She, 2004) and (Shoji et al., 2004), respectively. In (Itai and Yasukawa, 2008), dynamic time warping and cepstral information extracted from acoustic gait data is used for classification, while audio data recorded on a staircase—for the purpose of recognizing home inhabitants—is considered in (Alpert and Allen, 2010). A recent study (Geiger et al., 2014) extracts mel-frequency cepstral coefficients (MFCCs) as audio features and uses HMMs with a cyclic topology for dynamic classification.

3 Methodology

The objective of this research is to consider a simple method for surface-independent audio based gait analysis. To address surface-independence, we use time domain analysis, as opposed to frequency analysis, which is likely to be much more vulnerable to surface variations.

During a gait cycle, there are multiple distinct visual stances that a subject transitions through. In a similar vein, it is likely that a subject's gait patterns varies in a way that affects the timing between consecutive footsteps. Therefore, we consider a sequence of time intervals between successive steps, which represents a subject's characteristic gait cadence. We extract statistical data from this set to obtain feature vectors based on footstep sound recordings. Given numerous labeled feature vectors, we employ a supervised learning technique for classification. Here, we provide preliminary results for the effectiveness of this approach, based on a small dataset.

Next, we describe the procedure by which we determine our input features. Then we briefly discuss support vector machines (SVM), which are used for classification in this research.

3.1 Input Features

To measure audio gait characteristics, a microphone is assumed to be in a static position on the floor, and only one person's footsteps are monitored at a time.[1] From the sound wave data, we detect footstep peaks as shown in Figure 1 and compile a sequence of times, denoted $t_1, t_2, \ldots, t_{n+1}$, at which a particular individual's footsteps are heard. Next, time intervals between successive steps are calculated and compiled into a second sequence $x_i = t_{i+1} - t_i$, yielding $x_1, x_2, \ldots, x_n$. We will refer to the x_i as the consecutive time interval (CTI) sequence, which represents time-based cadence and serves as the base for our analysis. Not surprisingly, the CTI sequence of each individual follows a (roughly) normal distribution, as can be seen in Figure 2.

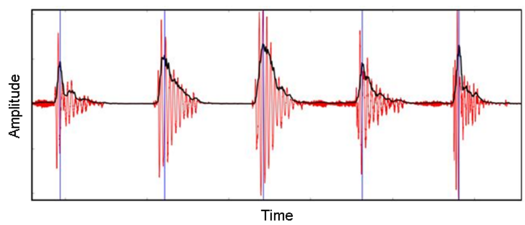

Figure 1: Five peaks corresponding to footstep impact sounds

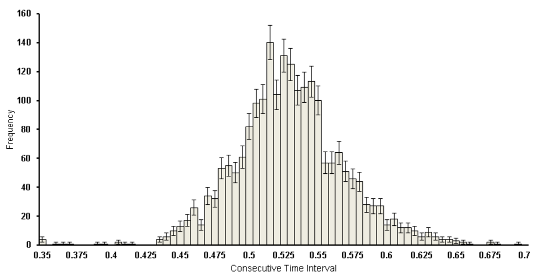

Figure 2: Sample CTI distribution with error bars

From the CTI sequence, the four central moments are calculated. That is, we compute each of the following.

- Mean — The center of the distribution will distinguish subjects that have shorter or longer strides on average. The mean is computed as

$$\mu = \frac{1}{n}\sum_{i=1}^{n} x_i = \frac{1}{n}\left(x_1 + x_2 + \cdots + x_n\right)$$

where the x_i are the CTI sequence and n is the length of this sequence.

- Standard deviation — The variance measures the spread of the data about the mean and will reveal the variability in stride times. The standard deviation, which is the square root of the variance, is computed as

$$\sigma = \sqrt{\frac{\sum_{i=1}^{n}(x_i - \mu)^2}{n-1}}$$

where μ is the mean of the CTI.

- Skewness — The symmetry of a distribution can help to distinguish subjects who tend to slow down or speed up. We compute the skew as

$$s = \frac{\sum_{i=1}^{n}(x_i - \mu)^3}{n\sigma^3} \tag{1}$$

where μ iand σ are the mean and standard deviation of the CTI, respectively.

- Kurtosis — The flatness or peakedness provides additional information about the tails of the CTI distribution. We compute the kurtosis as

$$k = \frac{\sum_{i=1}^{n}(x_i - \mu)^4}{n\sigma^4}. \tag{2}$$

The intuition here is that each of these characteristics of the distribution may reflect a particular aspect of a subject's walking patterns as observed over a period of time.

Each CTI in our dataset consists of 100 to 130 time differences for a specific individual. For each CTI, we compute a feature vector consisting of the first four central moments, (μ, σ, s, k), as discussed above.

3.2 Support Vector Classification

Originating from statistical learning theory (Vapnik and Vapnik, 1998), and first implemented in (Cortes and Vapnik, 1995), support vector machines (SVMs) are recognized as among the most efficient and powerful supervised machine learning algorithms (Byun and Lee, 2002). An SVM attempts to determine an optimal separating hyperplane between two sets of labeled training data.

Here, we provide a very brief summary of the SVM technique. For additional information on SVMs, see, for example, (Cortes and Vapnik, 1995) or (Stamp, 2017), or consult (Berwick, 2003) for a particularly engaging introduction to the subject.

An SVM is trained based on a set consisting of n samples, of the form $(\vec{x}_1, y_1), \ldots, (\vec{x}_n, y_n)$ where each

[1] In a noisy environment, a higher density of microphones would be needed, and additional pre-procession (e.g., noise suppression) would also be required.

vector $\vec{x}_i$ is m-dimensional, containing input features, and each y_i is labeled either $+1$ or -1 according to the class to which the vector $\vec{x}_i$ belongs. If possible, the SVM will find a hyperplane that divides the set of $\vec{x}_i$ labeled $y_i = +1$ from those labeled $y_i = -1$. The SVM-generated hyperplane will be optimal, in the sense that it will the maximize the margin, where "margin" is defined as the minimum distance between the hyperplane and any sample $\vec{x}_i$ in the training set.

In generally, the training data will not be linearly separable, that is, no separating hyperplane will exist in the input space. In such cases, we can use a nonlinear kernel function to transform the data to a higher dimensional feature space, where it is more likely to be linearly separable (or at least reduce the number of classification errors). The so-called kernel trick enables us to embed such a transformation within an SVM, without paying a significant penalty in terms of computational efficiency.

In our case, each input sample consists of a 4-dimensional feature vector and its corresponding label, while the desired output value is a label of 0, 1, or 2, which specifies the subject. After feature regularization, an optimal hyperplane boundary is determined by training an SVM classifier. We use a one-vs-one scheme for this multi-class classification problem, in which each class is individually compared to every other class, as opposed to a one-vs-all scheme.[2] Additional testing examples are classified to determine the accuracy of the resulting classification scheme.

In this research, the following three types of SVM kernels were tested and evaluated.

- Linear
$$k(\vec{x}_i, \vec{x}_j) = \vec{x}_i \bullet \vec{x}_j$$
- Quadratic
$$k(\vec{x}_i, \vec{x}_j) = (\vec{x}_i \bullet \vec{x}_j)^2$$
- Radial basis function (RBF)
$$k(\vec{x}_i, \vec{x}_j) = \exp\left(-\frac{\|\vec{x}_i - \vec{x}_j)\|^2}{2\sigma^2}\right)$$

4 Experimental Results

A three-person database was constructed, in which subjects were asked to walk normally at comfortable walking speed for 60 to 70 seconds around a circle 7 feet in diameter. This yielded 100 to 130 continuous strides per walking trial. The dataset contains 60 trials per subject collected over several days. The data was collected by an inexpensive microphone attached to a PC laptop, with the microphone placed on a hardwood floor at the center of the walking circle. Samples were recorded using Windows Sound Recorder and exported in single channel WAV file format with a bit rate of 1411 kbps and sampling rate of 44.1 kHz.

4.1 Data Analysis

The distributions for each of the four temporal features are shown in Figure 3. We observe that each feature contributes some discriminating information that should prove useful to the SVM. In Figure 3 (a), we notice that the individual denoted by blue may be easily separated. To distinguish the individuals denoted in red and green, we see that the standard deviation in Figure 3 (b) provides useful information. Figures 3 (c) and (d) show that skewness and kurtosis are also potentially discriminating features.

4.2 Training

The training set is comprised of 150 sound files, 50 for each of 3 subjects. The remaining 10 files per subject are reserved for testing. As discussed above, features vectors have been extracted from each file and input to an SVM for training and classification.

To visualize the general boundaries of each kernel, we first conduct kernel comparisons in lower-dimensional feature spaces. In Figure 4, we see that a combination of any two features at a time is inadequate to distinguish the subjects. It is interesting that the subject represented by dark blue has significantly lower means, corresponding to a shorter CTI that allows for easy discrimination, relative to the classes marked in light blue and red. There is considerable overlap in the other three features for all subjects.

4.3 Classification Results

We experimented with the number of samples used for training to obtain the results shown in Figure 5. Across each kernel, the learning curves show that increasing training set size generally decreases the training score while increasing the cross-validation score. Based on these results, we find that the linear kernel is superior for several reasons. First, we observe that the linear kernel yielded the greatest mean score of 78% with a 95% confidence interval. In addition, the training score curve for the linear kernel

[2] A one-vs-one scheme is less sensitive to imbalanced datasets, but it is computationally more expensive than one-vs-all approach. However, the computational cost is not a significant issue here, due to the small number of dimensions in our experiments.

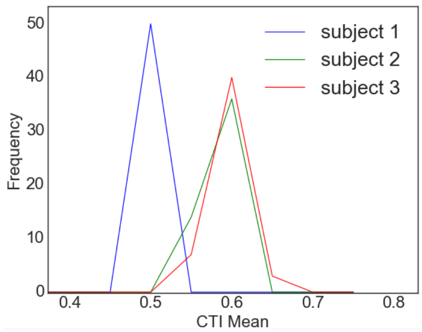

(a) Distributions of CTI mean

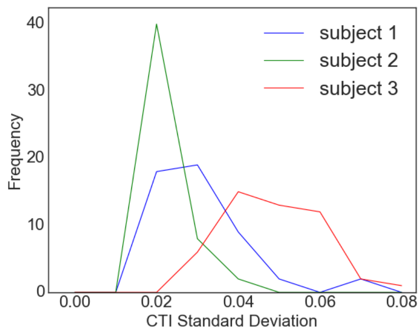

(b) Distributions of CTI standard deviation

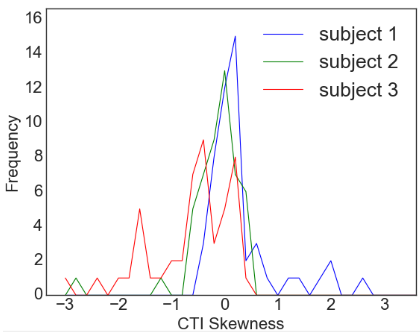

(c) Distributions of CTI skewness

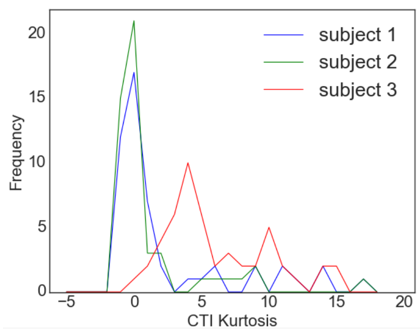

(d) Distributions of CTI kurtosis

Figure 3: Distributions for temporal features (different colors denote different individuals).

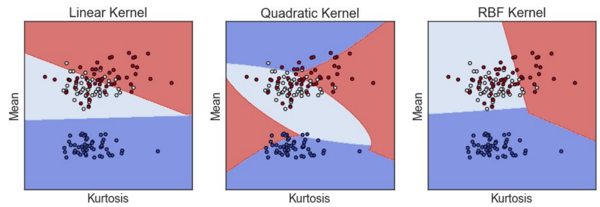

(a) Mean versus kurtosis

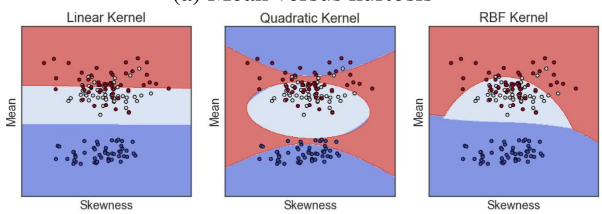

(b) Mean versus skewness

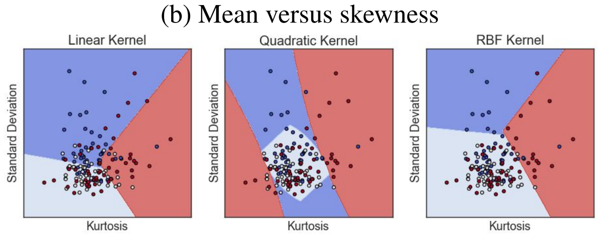

(c) Standard deviation versus kurtosis

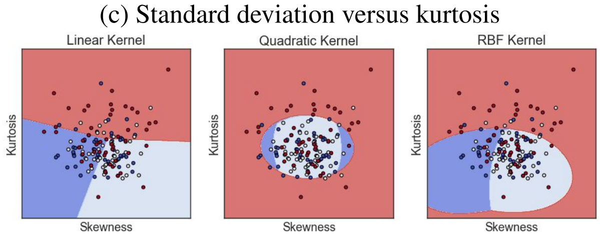

(d) Kurtosis versus skewness

Figure 4: SVM hyperplane boundaries in two-dimensional feature spaces

plateaus after 40 training examples, which implies that as few as 40 training examples is sufficient to achieve (essentially) optimal results. This offers a significant reduction in the training data requirement, as compared to the quadratic and RBF kernels. Finally, the training and cross-validation curves for the linear kernel approach similar scores in fewer training examples (again, as compared to the quadratic and RBF kernels), which indicates that the linear kernel is least likely to overfit the data.

In Figure 6 (a), we have given the ROC curve corresponding to a one-vs-all analysis, which represents the ability to discriminate one class from the other classes. Figure 6 (b) includes ROC curves for both the micro-averaged and macro-averaged cases. Micro-averaging calculates metrics globally, that is, we count all true positives, false negatives, and false positives, Note that in the micro-averaged case, larger classes have greater influence than smaller classes. In contrast, macro-averaging calculates metrics for each

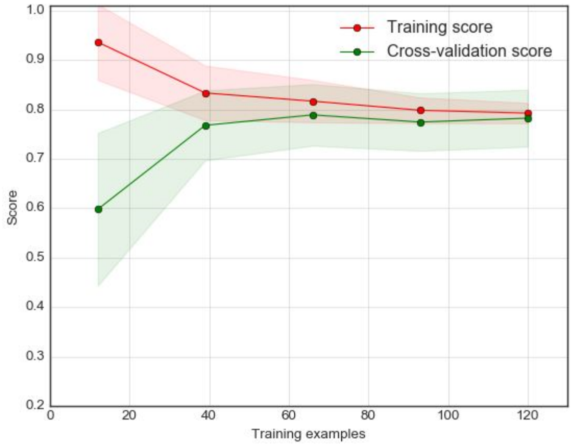

(a) Linear kernel

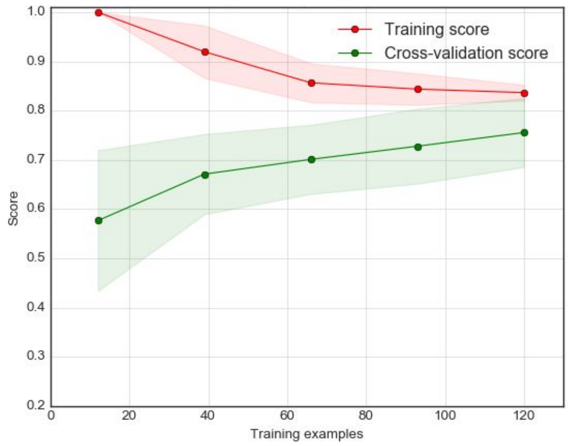

(b) Quadratic kernel

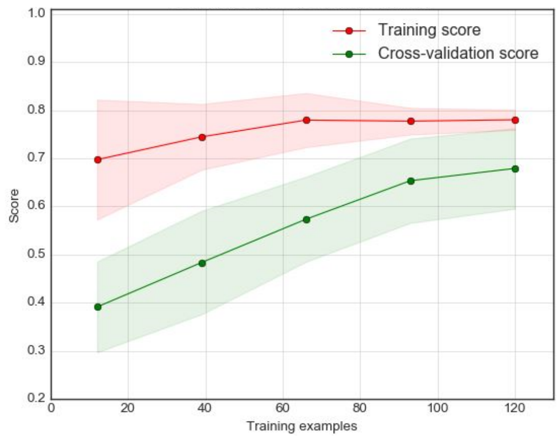

(c) RBF kernel

Figure 5: Learning curves for increasing number of training samples (95% confidence intervals)

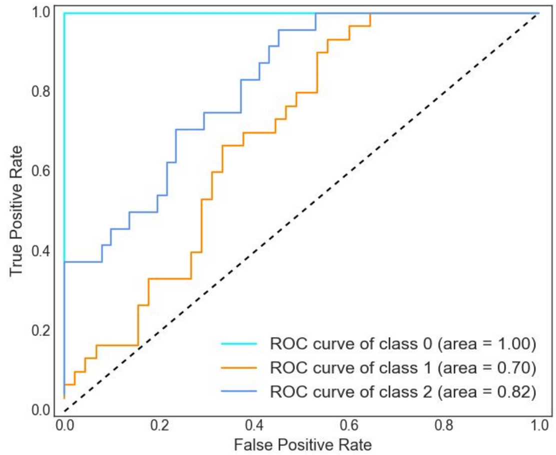

(a) One-vs-all ROC

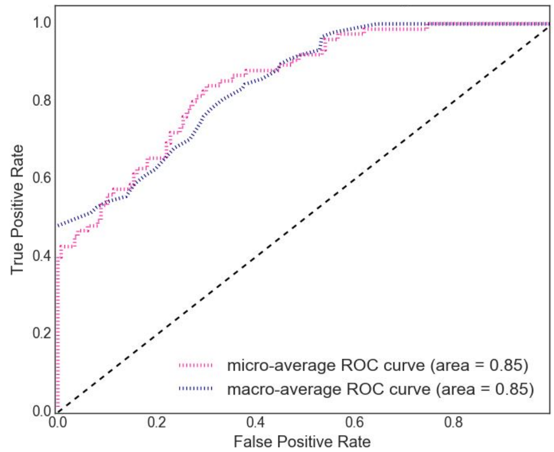

(b) Micro-averaged and macro-averaged ROC curves

Figure 6: ROC curves for linear kernel

individual category and computes their unweighted mean. Macro-averaging has the effect of weighting each class the same, regardless of any imbalances. Our data is not imbalanced, but we provide the macro-averaged results to facilitate comparison to imbalanced datasets.

In Figure 7 we give the results of these experiments in the form of a normalized confusion matrix. Finally, in Table 1 we give the accuracy results for this same set of experiments.

Table 1: Kernel accuracy comparison

	Kernel Function		
	Linear	Quadratic	RBF
Training set	0.787	0.827	0.773
Testing set	0.935	0.935	0.839

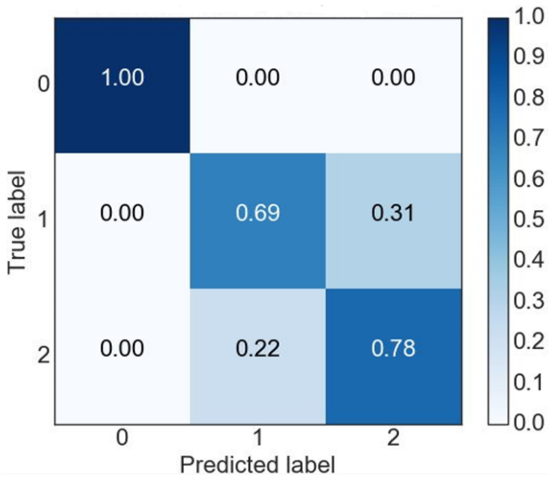

Figure 7: Normalized confusion matrix for linear SVM

5 Conclusions

In this paper, we have presented a computationally inexpensive method for gait analysis and classification based on audio information in the time-domain. Using an SVM, we were able classify samples from a small dataset with good accuracy.

The method considered here has various potential advantages and disadvantages, as compared to previous gait analysis techniques. One advantage is that our approach is computationally very inexpensive and requires minimal sensor density and sensitivity. Another advantage is that audio samples are inherently less sensitive to differences in clothing and illumination, and minor changes in footwear or walking angle. However, since our observation sequence is based on time, changes that affect a person's stride are likely to negatively affect performance. In addition, our analysis presumes a relatively quiet environment, in which only one person's footsteps are within sound range. However, in many cases, it would likely be possible to pre-process the data to extract footstep sounds from background noise.

For future work, we plan to conduct much larger scale experiments, with far more subjects and more test data for each individual. We also plan to test the robustness of our approach, with respect to changes in floor surface, differing footwear, background noise, and other practical environmental issues. In particular, the effects of microphone density and noise suppression will be carefully analyzed.

To deal with speed variations (e.g., walking hurriedly), we plan to study velocity-related features, such as the Doppler sensors discussed in (Kalgaonkar and Raj, 2007). In addition, we believe it will likely be profitable to test more advanced HMM-based techniques to investigate various timing transitions within gait-based audio datasets. For example, it has been demonstrated that spatiotemporal gait cycles can be successfully modeled as doubly stochastic processes with HMMs (Kale et al., 2002); we plan to apply this technique to the audio-based temporal gait cycles considered in this paper.

REFERENCES

Aggarwal, J. K. and Cai, Q. (1997). Human motion analysis: A review. In *Proceedings of IEEE Nonrigid and Articulated Motion Workshop, 1997*, pages 90–102. IEEE.

Alpert, D. T. and Allen, M. (2010). Acoustic gait recognition on a staircase. In *World Automation Congress, 2010*, WAC 2010, pages 1–6. IEEE.

Aminian, K., Najafi, B., Büla, C., Leyvraz, P.-F., and Robert, P. (2002). Spatio-temporal parameters of gait measured by an ambulatory system using miniature gyroscopes. *Journal of Biomechanics*, 35(5):689–699.

Begg, R. and Kamruzzaman, J. (2005). A machine learning approach for automated recognition of movement patterns using basic, kinetic and kinematic gait data. *Journal of Biomechanics*, 38(3):401–408.

Begg, R. K., Palaniswami, M., and Owen, B. (2005). Support vector machines for automated gait classification. *IEEE Transactions on Biomedical Engineering*, 52(5):828–838.

Berwick, R. (2003). An idiots guide to support vector machines (SVMs). http://www.svms.org/tutorials/Berwick2003.pdf.

Bimbot, F., Bonastre, J.-F., Fredouille, C., Gravier, G., Magrin-Chagnolleau, I., Meignier, S., Merlin, T., Ortega-García, J., Petrovska-Delacrétaz, D., and Reynolds, D. A. (2004). A tutorial on text-independent speaker verification. *EURASIP Journal on Applied Signal Processing*, 2004:430–451.

Byun, H. and Lee, S.-W. (2002). Applications of support vector machines for pattern recognition: A survey. In *Proceedings of the First International Workshop on Pattern Recognition with Support Vector Machines*, SVM '02, pages 213–236, London, UK, UK. Springer-Verlag.

Cortes, C. and Vapnik, V. (1995). Support-vector networks. *Machine Learning*, 20(3):273–297.

Geiger, J. T., Kneißl, M., Schuller, B. W., and Rigoll, G. (2014). Acoustic gait-based person identification using hidden markov models. In *Proceedings of the 2014 Workshop on Mapping Personality Traits Challenge and Workshop*, pages 25–30. ACM.

Itai, A. and Yasukawa, H. (2008). Footstep classification using simple speech recognition technique. In *IEEE International Symposium on Circuits and Systems, 2008*, ISCAS 2008, pages 3234–3237. IEEE.

Johansson, G. (1973). Visual perception of biological motion and a model for its analysis. *Perception & Psychophysics*, 14(2):201–211.

Kale, A., Rajagopalan, A., Cuntoor, N., and Kruger, V. (2002). Gait-based recognition of humans using continuous HMMs. In *Proceedings of Fifth IEEE International Conference on Automatic Face and Gesture Recognition, 2002*, pages 336–341. IEEE.

Kalgaonkar, K. and Raj, B. (2007). Acoustic doppler sonar for gait recogination. In *IEEE Conference on Advanced Video and Signal Based Surveillance, 2007*, AVSS 2007, pages 27–32. IEEE.

Kamruzzaman, J. and Begg, R. K. (2006). Support vector machines and other pattern recognition approaches to the diagnosis of cerebral palsy gait. *IEEE Transactions on Biomedical Engineering*, 53(12):2479–2490.

Lai, D. T., Begg, R. K., and Palaniswami, M. (2009). Computational intelligence in gait research: a perspective on current applications and future challenges. *IEEE Transactions on Information Technology in Biomedicine*, 13(5):687–702.

Lee, L. (2002). Gait analysis for recognition and classification. In *Proceedings of the IEEE Conference on Face and Gesture Recognition*, pages 155–161.

Makela, K., Hakulinen, J., and Turunen, M. (2003). The use of walking sounds in supporting awareness. In *Proceedings of the 2003 International Conference on Auditory Display*, ICAD'03, pages 1–4.

Man, J. and Bhanu, B. (2006). Individual recognition using gait energy image. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(2):316–322.

Mantyjarvi, J., Lindholm, M., Vildjiounaite, E., Makela, S.-M., and Ailisto, H. (2005). Identifying users of portable devices from gait pattern with accelerometers. In *Proceedings of IEEE International Conference on Acoustics, Speech, and Signal Processing, 2005*, ICASSP'05, pages 973–976. IEEE.

Middleton, L., Buss, A. A., Bazin, A., and Nixon, M. S. (2005). A floor sensor system for gait recognition. In *Fourth IEEE Workshop on Automatic Identification Advanced Technologies, 2005*, pages 171–176. IEEE.

Murray, M. P. (1967). Gait as a total pattern of movement: Including a bibliography on gait. *American Journal of Physical Medicine & Rehabilitation*, 46(1):290–333.

Nishiguchi, S., Yamada, M., Nagai, K., Mori, S., Kajiwara, Y., Sonoda, T., Yoshimura, K., Yoshitomi, H., Ito, H., Okamoto, K., et al. (2012). Reliability and validity of gait analysis by Android-based smartphone. *Telemedicine and e-Health*, 18(4):292–296.

Otero, M. (2005). Application of a continuous wave radar for human gait recognition. In Kadar, I., editor, *Signal Processing, Sensor Fusion, and Target Recognition XIV*, volume 5809, pages 538–548.

Qian, G., Zhang, J., and Kidané, A. (2008). People identification using gait via floor pressure sensing and analysis. In *Proceedings of the 3rd European Conference on Smart Sensing and Context*, EuroSSC '08, pages 83–98, Berlin, Heidelberg. Springer-Verlag.

Reynolds, D. A. (1995). Speaker identification and verification using gaussian mixture speaker models. *Speech Communication*, 17(1-2):91–108.

She, B. (2004). Framework of footstep detection in indoor environment. In *Proceedings of International Congress on Acoustics, Kyoto, Japan*, pages 715–718.

Shoji, Y., Takasuka, T., and Yasukawa, H. (2004). Personal identification using footstep detection. In *Proceedings of 2004 International Symposium on Intelligent Signal Processing and Communication Systems, 2004*, ISPACS 2004, pages 43–47. IEEE.

Sprager, S. and Zazula, D. (2009). A cumulant-based method for gait identification using accelerometer data with principal component analysis and support vector machine. *WSEAS Transactions on Signal Processing*, 5(11):369–378.

Stamp, M. (2017). *Introduction to Machine Learning with Applications in Information Security*. Chapman and Hall/CRC, Boca Raton.

Sundaresan, A., RoyChowdhury, A., and Chellappa, R. (2003). A hidden Markov model based framework for recognition of humans from gait sequences. In *Proceedings of 2003 International Conference on Image Processing*, volume 2 of *ICIP 2003*, pages 93–96. IEEE.

Tahmoush, D. and Silvious, J. (2009). Radar micro-doppler for long range front-view gait recognition. In *IEEE 3rd International Conference on Biometrics: Theory, Applications, and Systems, 2009*, BTAS'09, pages 1–6. IEEE.

Vapnik, V. N. and Vapnik, V. (1998). *Statistical Learning Theory*, volume 1. Wiley New York.

Wang, L., Tan, T., Ning, H., and Hu, W. (2003). Silhouette analysis-based gait recognition for human identification. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(12):1505–1518.

Whittle, M. W. (2014). *Gait Analysis: An Introduction*. Butterworth-Heinemann.