

A Gait Triaging Toolkit for Overlapping Acoustic Events in Indoor Home Environments

Kelvin Summoogum
Doctoral Researcher
University of Greenwich
Chatham, UK

kelvin.summoogum@greenwich.ac.uk

Debayan Das
Data Scientist
MiiCare Ltd
London, UK

debayan.das@miicare.co.uk

Parvati Jayakumar
Clinical Data Analyst
MiiCare Ltd
London, UK

parvati.jayakumar@miicare.co.uk

Ismael Essop
Associate Professor
University of Greenwich
Chatham, UK

I.A.Essop@greenwich.ac.uk

Abstract — Gait has been used in clinical and healthcare applications to assess the physical and cognitive health of older adults. Acoustic based gait detection is a promising approach to collect gait data of older adults passively and non-intrusively. However, there has been limited work in developing acoustic based gait detectors that can operate in noisy polyphonic acoustic scenes of homes and carehomes. We attribute this to the lack of good quality gait datasets from the real-world to train a gait detector on. In this paper, we put forward a novel machine-learning based filter which can triage gait audio samples suitable for training machine learning models for gait detection. The filter achieves this by eliminating noisy gait samples at an $f(1)$ score of 0.85 and prioritising gait samples with distinct spectral features and minimal noise. To demonstrate the effectiveness of the filter, we train and evaluate a deep learning model on gait datasets collected from older adults with and without applying the filter. The gait detector registers an increase of 25 points in its $f(1)$ score on unseen real-world gait data when trained with the filtered gait samples. The proposed filter will help automate the task of manual annotation of gait samples for training acoustic-based gait detection models for older adults in indoor environments.

Keywords — Acoustic Gait Analysis, Acoustic Gait Detection, Gait Sensing, Acoustic Event Detection, Geriatrics

I. INTRODUCTION

Gait has been extensively studied as a clinical biomarker for assessing deterioration of physical mobility [1] and cognitive health in older adults [2]. Detecting gait patterns of older adults has been explored using accelerometric sensors [3], pressure-sensitive floors [4], video feed [5], even audio recordings [6]. Acoustic based gait detection systems have the added benefit of being non-intrusive, portable, and economical when compared to those proposed in contemporary literature for gait analysis. However, there has been little to no progress in this field in the past decade. The advent of large-scale deep learning-based techniques have led to the development of state-of-art gait analysis using all the above data collection modalities except for that using audio.

Acoustic Gait Analysis (AGA) is an example of an Audio Event Detection (AED) task. Our analysis of the literature available for Audio Event Detection (AED) reveal the following findings:

1. Deep learning excels in achieving state of art [7] in AED only when the audio dataset used is clean and contains little to no background noise.
2. Majority of the available works in AED, including those for AGA use data captured from normal healthy individuals in controlled laboratories with clean acoustic scenes [8]. This is not representative of the

real-world acoustic scene in an indoor setting with an older adult living with one or more comorbidities. Real world gait recordings are noisy and have multiple overlapping acoustic events (i.e polyphony).

3. There is a very limited number of publicly available annotated acoustic datasets of gait data large enough to train deep learning models for gait detection in the real world.

Manual annotation of real-world gait data helps improve deep learning and machine learning based gait detection models in this regard. However, using manual labour to annotate gait data is time-consuming, not to mention the auditory expertise required to differentiate distinct clean gait samples from noisy gait samples captured in acoustic scenes with overlapping acoustic events. Figure 1 below shows the melspectrograms of two audio samples which a human annotator confidently labelled as “gait” samples despite the dominating presence of background noise and other acoustic events in one of them. Using distinct and relatively clean gait samples during training models will help them learn features better which will automatically translate to more robust detection of gait from noisy polyphonic recordings. For labelling purposes, we will call them “good” and “bad” gait samples respectively.




Figure 1 The two versions of gait as captured in real life.

In this paper, we explore the feasibility of using an automated “filter” which eliminates the “bad” gait samples from a population of human-annotated gait samples, leaving us with samples which are most likely to be “good”, i.e distinct and has minimal overlapping noise. We propose this specific approach because it is much easier to find noisy samples than clean samples from audio recordings captured in real world settings of the indoor home environment. Our objective is to develop a “pre-screening” technique that can be deployed in acoustic gait recording devices within the indoor home environment of older adults that will eliminate collection of “bad” noisy gait samples which is unsuitable for training gait analysis models, as much as possible directly at source.

This paper summarises the approach we took to develop this gait filter from a pre-annotated gait dataset of footsteps from older adults living with dementia (PLWD) with medium to high fall risks across carehomes in the UK. We present the

use of a Machine Learning (ML) model as a “gait sample filter” wherein we train the ML model on a set of gait samples which we consider as good gait samples, use it to classify a test set of “unfiltered” gait samples, eliminate bad gait samples, and then evaluate its “filtering” performance against human experts in gait analysis.

In Section II, we discuss the techniques presented in relevant acoustic based gait analysis literature for cleaning and feature extraction from gait samples. In Section III, we explain our process of extracting features to distinguish good gait samples from the bad noisy ones and passing these features into an ensemble model developing using AutoML [9] which maximises detection rate for bad gait samples. In Section IV, we reproduce the work done in [6] using a recurrent deep learning model for acoustic based gait detection in older adults and demonstrate how applying our ensemble model as an audio filter improved the detection performance of the deep learning model on real world gait data.

II. BACKGROUND

A. Feature Extraction for Gait

It is important to note the distinction between gait and footsteps. Gait is defined as the characteristics in the walking pattern of an individual. Gait “datasets” record sequence of multiple individual footstep events one after the other within a short time window. Most existing literature in AGA focuses on examination of the individual footstep audio event first and then applying it on all such events in a gait sample. However, there is no available literature (yet) to our knowledge that attempts differentiating clean from noisy gait sounds. The literature instead focuses on characteristics of the audio (“features”) which maximise the difference between gait and other non-gait sound like speech. We group these into two camps broadly speaking.

The first camp focuses on using statistical analysis of footsteps in the temporal and frequency domains. This includes extraction of features [10] like power, amplitude, ZCR, peaks, etc. from the waveform and FFT representations. Before the emergence of neural network-based feature extraction, such techniques enjoyed popularity across several works in machine learning and signal processing literature for gait detection and analysis. The common shortcoming among all these works using these types of features is that they use gait datasets collected in controlled acoustic scenes [8, 10] with little to minimal background noise. These features are highly susceptible to presence of noise and hence cannot be relied upon in a discussion on gait detection in the real world with noisy polyphonic acoustic scenes.

The second camp relies on spectrogram-based features to distinguish between footstep and other acoustic events. MFCC is the most popular spectrogram-based feature for deep learning models to learn from. Deep learning models exclusively use neural networks paired with pre-processing layers like a convolution layer (Convolutional Neural Networks) or a recurrent gate operation (Long Short-Term Memory and Gated Recurrent Units for instance) to automatically learn the best features to learn from the spectrogram representations. Contemporary AED have

achieved state of art [7] by applying deep learning models on spectrograms. This is the most promising among all approaches to build new feature extraction techniques for gait detection, as of the time of writing of this paper.

B. State of the Art in Gait Detection

Before the advent of deep neural networks, AGA relied on classifiers like Support Vector Machines (SVM) [10], Gaussian Mixture Models (GMM) [11], Hidden Markov Models (HMM) [12], Non-Negative Matrix Factorizations (NMF) [13] and so on. Annual competitions like DCASE [14] helped popularise the use of use of large deep learning models to achieve state of art in AED tasks like Acoustic Scene Classification, General AED in indoor environments, Audio Tagging, etc. The authors of [6] presented the use of synthetic soundscapes to combat the problem of unavailability of gait datasets and trained an Attention-based [15] Bidirectional LSTM model to perform gait detection in noisy overlapping audio streams. This approach was able to reproduce state of art achieved in DCASE 2021 with publicly available gait datasets like TUM GAID [16] and ESC50 [17] in the specific task of gait detection for older adults in indoor environments. However, we found that it failed to detect footsteps when we reproduced the work of [6] and applied it on real life gait recordings. In the upcoming section, we discuss what we how we developed our gait audio filter to make the model from [6] work in the noisy polyphonic acoustic scene of indoor homes of older adults.

III. METHODOLOGY

In this section, we present the problem we found when we attempted applying the gait detector reproduced from the work of [6] on a real-life gait dataset collected from older adults living in home and carehome settings and how we made it work by using a machine learning based gait sample filter to remove gait samples unfit for training the gait detector.

A. Dataset Collection

Our primary audio recording device for data collection from the homes of older adults was the MiiCube manufactured by MiiCare [18]. It is a ARM4 processor coupled with a 2-microphone setup from Seed Studio’s ReSpeaker Pi Hat [19] enclosed in a 3D-printed cubical box enclosure as shown in Figure 2 below. The microphones are located immediately beneath the top lid to ensure it is capturing the acoustic scene and all acoustic events without exposing the computing unit.



Figure 2 The MiiCube recording device from MiiCare.

Each MiiCube samples environmental sound at 16 KHz. We chose 16 KHz as the sampling rate because we have observed that gait do not have significant spectral presence beyond 8 KHz (Nyquist frequency). It records audio in batches

of 30s at any point of time and uploads it to a secure cloud storage. We deployed a total of 10 MiiCubes across different carehomes and home sites in the UK. 8 of them were set up in the rooms of older adults in carehomes living with advanced dementia and high fall risk. The remaining 2 were setup in the homes of two different older adults aged 80+ with mild cognitive impairment (MCI). We introduced an additional step of acoustic thresholding based on RMS of the incoming audio streams to determine when to start recording footsteps. This thresholding was done empirically by manual experimentation, so each MiiCube has its RMS threshold specific to its acoustic scene at the carehome or home where it is deployed. Details of this thresholding method are outside the scope of the contribution of this paper and will be discussed in future works.

B. Human Annotation

We employed external annotation services to annotate the collected recordings in cloud storage. For this work, we were able to get 6852 audio recordings annotated for 2 acoustic classes: “gait” and “non-gait”. For each file, the annotators produced sets of time segments where footsteps are heard through auditory inspection. The remaining segments of the file were marked as “non-gait”. These segments consisted of background noise from conversations, electrical appliances, door shuts, etc. Both the gait and non-gait segments had overlapping presence of multiple acoustic events. Majority of the gait segments had overlapping background noise and speech. Such quality of sound is to be expected for a real-life audio dataset collected from the indoors of homes where individuals lived unlike those in literature [8, 10] collected in controlled laboratory environments.

C. Motivation behind the filter

We reproduced the work of [6] on this newly annotated gait dataset. The pre-processing algorithm generated 3509 melspectrograms in total, with 1399 “gait” and 2110 “non-gait” melspectrograms. We applied an 80-20 split to create a training and test set which had 1119 and 280 gait melspectrograms respectively. The LSTM model achieved a $f(1)$ score of 0.58 on the 280 gait melspectrograms from the test set. We went through the test set ourselves to see how accurately we as humans could classify gait from non-gait. We observed that we missed more than half of the gait windows from the test set. We attribute this to the dominating presence of other acoustic events and background noise overlapping with footstep events in the melspectrograms which misled our judgments. Figure 1 shows two gait melspectrograms as examples in how noise and polyphony made it difficult for us to distinguish between the two classes correctly. This led us to realise that we need a technique to remove the noisy melspectrograms from the training set for the model.

D. Parameters for the filter

To study the patterns between suitable and unsuitable gait melspectrograms, we manually went through the 1399 gait melspectrograms and annotated them between what we considered as clean and unclean/noisy melspectrograms as human experts. We labelled a total of 422 of them as good while the remaining 977 as bad gait melspectrograms. Next,

we devised the following set of 3 parameters guided by our observations during this manual labelling process. Figure 3 below visualises the parameters from two melspectrograms.

1. **Average Peak Prominence:** Gait melspectrograms suitable for training purposes have distinct vertical peaks with little spectral noise around them. We capture this property using peak prominence in the energy signal representation of the melspectrogram. We define the energy signal from a melspectrogram as $E(i) = \sum_i F(i)$ for all frames in a melspectrogram where $F(i)$ is the i -th frame in the melspectrogram and $\sum_i F(i)$ is the sum of the amplitude magnitude values in each frame in decibels. Peak prominence measures how much a peak stands out in a signal due to its intrinsic height and its location relative to other peaks.
2. **Residuals to the RMS:** High amount of noise in a gait melspectrogram results in an energy signal E with less prominent peaks for the individual footstep events. To capture this property, we computed the residuals of $E(i)$ for all frames i around the RMS of $E(i)$. We define this operation as $r = \sum_i (E(i) - RMS(E))^2$ where r is our second parameter.
3. **Average Distance between Peaks:** Footstep events appear rhythmically in each gait melspectrogram. Non-gait melspectrograms which look like gait melspectrograms will not share this characteristic. We defined this parameter based on the average distance between all the peaks detected in the energy signal representation of the melspectrogram.

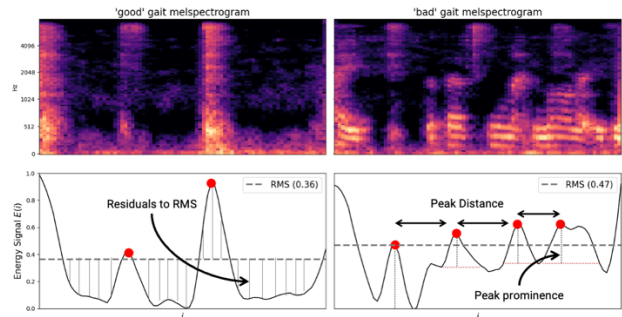


Figure 3 Visual representation of the three parameters as extracted from the energy signal representation $E(i)$ of the melspectrogram.




Figure 4 The parameters combined in 3D space achieve separability for ‘good’ and ‘bad’ gait melspectrograms.

We extracted these 3 parameters from all the 1399 spectrograms and visualised them in a third dimensional space using the TSNE algorithm to demonstrate that how it can

successfully separate the suitable from the unsuitable gait melspectrograms we found manually. Figure 4 above visualises the “feature separability” which we sought to conceptualise and exploit to automate our manual effort.

E. Training a machine learning algorithm as a filter

We used Microsoft FLAML’s AutoML [20] framework to automate finding the optimal ensemble of machine learning algorithms capable to identify bad gait melspectrograms from the good ones. Our focus was to maximise the predictive power of this model to detect the unsuitable bad melspectrograms as we have more of them in our training set. We took the original set of 1399 gait melspectrograms and our manual annotations and performed an 80-20 split again. The training set had 1119 melspectrograms with 338 labelled as “good” while the testing set had 280 melspectrograms with 84 labelled so. The AutoML model was trained on the 1119 melspectrograms with 10-fold cross-validation with the following parameters: ensemble=True, metric=“macro_f1”, eval_method=“cv”, n_splits=10, time_budget=300, early_stop=True, estimator_list=[“lgbm”, “xgboost”, “rf”, “extra_tree”, “kneighbour”], task = “classification”.

Table 1 below summarises the classification performance of the trained AutoML model as a filter on the test set of 280 melspectrograms against our manual annotations. Figure 5 shows the confusion matrix for the same.

TABLE I. Evaluating the AutoML filter against human expertise

Gait melspec label	Precision	Recall	$f(1)$	samples
“bad” gait melspecs	0.79	0.91	0.85	196
“good” gait melspecs	0.69	0.44	0.54	84
Average	0.76	0.77	0.75	

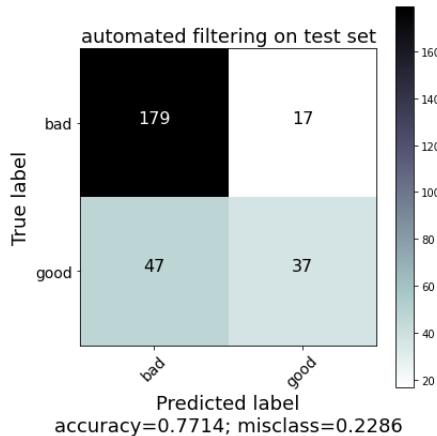


Figure 5 Confusion Matrix for Table 1 above.

IV. FILTER EFFECTIVENESS

We applied the newly trained AutoML model as a filter on the original “unfiltered” dataset of 3509 melspectrograms to get 2118 non-gait and 275 gait melspectrograms. To demonstrate the effectiveness of the filter on improving performance for gait detection models, we trained the LSTM model from [6] on the filtered melspectrogram dataset and compared the performance on the test sets from the model trained in Section III(D) earlier. Figure 6 below visualises the

improvements made by the new version of the model trained on the filtered gait melspectrograms across precision, recall and overall $f(1)$ scores.




Figure 6 Using the AutoML model as a filter on the training set improves performance metrics for gait and non-gait classification.

V. CONCLUSION

This paper has presented an AutoML based classifier as an audio based “triage” for gait data. We demonstrated how using the three parameters of average peak prominence, residuals around the RMS and the average peak distance together can help eliminate noisy gait recordings which are not suitable for training and developing robust acoustic based gait detection systems for older adults living in home and carehome settings. The proposed filter when applied to the training set of a gait detector immediately improved its overall performance in terms of precision, recall and $f(1)$ score in gait detection. As part of future works, we will investigate techniques to define good and bad gait samples using well-established features like signal-to-noise ratio (SNR) and explore techniques other than classification algorithms.

We believe this filter will help in two areas of developing acoustic based gait detection systems for AGA in geriatrics: One, it will speed up checking human annotations for large-scale gait datasets and help researchers focus on good quality gait recordings; Second, it will help eliminate bad quality gait recordings directly when applied at the source of gait data collection.

REFERENCES

- [1] Camicioli, R., Wang, Y., Powell, C., Mitnitski, A., & Rockwood, K. (2007). Gait and posture impairment, parkinsonism and cognitive decline in older people. *Journal of neural transmission*, 114(10), 1355-1361.
- [2] Kikkert, L. H., Vuillerme, N., van Campen, J. P., Hortobágyi, T., & Lamoth, C. J. (2016). Walking ability to predict future cognitive decline in old adults: A scoping review. *Ageing research reviews*, 27, 1-14.
- [3] Lakany, H. (2008). Extracting a diagnostic gait signature. *Pattern recognition*, 41(5), 1627-1637.
- [4] Qian, G., Zhang, J., & Kidané, A. (2008, October). People identification using gait via floor pressure sensing and analysis. In *European Conference on Smart Sensing and Context* (pp. 83-98). Springer, Berlin, Heidelberg.
- [5] Wang, F., Stone, E., Skubic, M., Keller, J. M., Abbott, C., & Rantz, M. (2013). Toward a passive low-cost in-home gait assessment system for older adults. *IEEE journal of biomedical and health informatics*, 17(2), 346-355.
- [6] Summoogum, K., & Palaniappan, R. (2021, November). Acoustic Based Footstep Detection in Pervasive Healthcare. In *2021 43rd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)* (pp. 4974-4977). IEEE.
- [7] Ebbers, Janek and Haeb-Umbach, Reinhold. Pre-Training And Self-Training For Sound Event Detection In Domestic Environments. Accessed at

- https://dcase.community/documents/challenge2022/technical_reports/DCA SE2022_Ebbers_125_t4.pdf
- [8] Altaf, M. U. B., Butko, T., & Juang, B. H. (2015). Acoustic gaits: Gait analysis with footstep sounds. *IEEE Transactions on Biomedical Engineering*, 62(8), 2001-2011.
- [9] He, X., Zhao, K., & Chu, X. (2021). AutoML: A survey of the state-of-the-art. *Knowledge-Based Systems*, 212, 106622.
- [10] Nakadai, K., Fujii, Y., & Sugano, S. (2013, July). Footstep detection and classification using distributed microphones. In *2013 14th International Workshop on Image Analysis for Multimedia Interactive Services (WIAMIS)* (pp. 1-4). IEEE.
- [11] X. Zhuang, X. Zhou, M. A. Hasegawa-Johnson, and T. S. Huang, "Real-world acoustic event detection," *Pattern Recognition Letters*, vol. 31, no. 12, pp. 1543–1551, 2010.
- [12] A. Mesaros, T. Heittola, A. Eronen, and T. Virtanen, "Acoustic event detection in real life recordings," in *Signal Processing Conference, 2010 18th European*. IEEE, 2010, pp. 1267–1271
- [13] Gemmeke, Jort F., et al. "An exemplar-based NMF approach to audio event detection." 2013 IEEE workshop on applications of signal processing to audio and acoustics. IEEE, 2013.
- [14] DCASE (Detection and Classification of Acoustic Scenes and Events). Accessed at <https://dcase.community>
- [15] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. (2017). Attention is all you need. *Advances in neural information processing systems*, 30.
- [16] Hofmann, M., Geiger, J., Bachmann, S., Schuller, B., & Rigoll, G. (2014). The TUM Gait from Audio, Image and Depth (GAID) database: Multimodal recognition of subjects and traits. *Journal of Visual Communication and Image Representation*, 25(1), 195-206.
- [17] Piczak, K. J. (2015, October). ESC: Dataset for environmental sound classification. In *Proceedings of the 23rd ACM international conference on Multimedia* (pp. 1015-1018).
- [18] MiiCare Ltd, UK. Company Number 11250727. Accessed at <https://miicare.co.uk>.
- [19] ReSpeaker 2-Mics Pi HAT. Accessed at <https://www.seeedstudio.com/ReSpeaker-2-Mics-Pi-HAT.html>
- [20] Wang, C., Wu, Q., Weimer, M., & Zhu, E. (2021). FLAML: A fast and lightweight auttml library. *Proceedings of Machine Learning and Systems*, 3, 434-447.