



HAL
open science

Scattering Features for Multimodal Gait Recognition

Srđan Kitić, Gilles Puy, Patrick Pérez, Philippe Gilberton

► **To cite this version:**

Srđan Kitić, Gilles Puy, Patrick Pérez, Philippe Gilberton. Scattering Features for Multimodal Gait Recognition. GlobalSIP 2017 - 5th IEEE Global Conference on Signal and Information Processing, Nov 2017, Montreal, Canada. hal-01523115

HAL Id: hal-01523115

<https://hal.science/hal-01523115>

Submitted on 16 May 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

SCATTERING FEATURES FOR MULTIMODAL GAIT RECOGNITION

Srđan Kitić, Gilles Puy, Patrick Pérez, Philippe Gilberton

Technicolor, Cesson-Sévigné, France

ABSTRACT

We consider the problem of identifying people on the basis of their walk (gait) pattern. Classical approaches to tackle this problem are based on, *e.g.*, video recordings or piezoelectric sensors embedded in the floor. In this work, we rely on acoustic and vibration measurements, obtained from a microphone and a geophone sensor, respectively. The contribution of this work is twofold. First, we propose a feature extraction method based on an (untrained) shallow scattering network, specially tailored for the gait signals. Second, we demonstrate that fusing the two modalities improves identification in the practically relevant open set scenario.

Index Terms— identification, walk, acoustic, vibration, scattering transform

1. INTRODUCTION

Identification lies at the heart of many user-defined services, ranging from movie recommendations to online banking. Due to its practical relevance, the problem of identifying people using various biometrics has triggered a significant amount of research in the signal processing and machine learning communities. Traditional means of identification, such as face [1] or speaker [2, 3] recognition, often require active participation in the recognition process, which may be intrusive in many applications. Therefore, a method that can reliably and *passively* identify people is advantageous in such a context.

In this work, we consider human gait as biometrics for identifying people present in a room. A number of approaches to gait-based identification have been proposed in the past, exploiting different signal modalities influenced by walk pattern, *e.g.*, based on video [4, 5], depth [5] or underfloor accelerometer measurements [6]. An appealing modality is structural (*e.g.*, floor) vibration induced by walking, and acquired through *geophones* [7], since it offers several practical advantages over other commonly used types of signals. One of them is increased security – it stems from the fact that there is no simple method (to the authors’ knowledge) that can accurately reproduce one’s gait in terms of the vibration signal. Another is preservation of privacy as vibration data is usually not considered a confidential, or sensitive information. Finally, the proposed setup is simple and cheap – typically one geophone is sufficient to monitor a medium-sized room.

Unfortunately, geophone measurements are not very rich in content, due to their very limited bandwidth. Currently, geophones are reliably measuring ground vibrations only in the very low frequency range [8], while the human footstep energy spans up to ultrasonic frequencies [9]. Hence, the loss of information is substantial.

In addition to vibrations (wave propagation in solids), a walking human also produces audible signals, which can be registered by standard microphones and used for identification [10, 5]. These have a much wider bandwidth, and, in addition to footsteps, they are also generated due to friction of the upper body (*i.e.*, due to leg and arm movements). However, modestly-priced microphones suffer from poor frequency response at very low frequencies, and the measured signals are susceptible to environmental noise, such as speech or music. Therefore, it seems that the vibration and acoustic modality complement each other: while the former is secure, robust and “senses” the low-frequency range, the latter carries more information, particularly at high frequencies. The goal of this work is to demonstrate that gait-based recognition using each of the modalities is a viable means of human identification, and that the two can be successfully coupled together in order to boost identification performance.

In the following section, we discuss the physical origin of acoustic and vibration gait measurements. Then, we introduce a feature extraction technique based on the *scattering transform* [11] and the specificities of the gait signal. In addition, we propose a simple feature fusion technique to enhance performance when bimodal measurements are available. Finally, we provide open set identification results, obtained from exhaustive experiments on a home-brewed dataset.

2. GAIT SIGNALS

A microphone and a geophone, placed (fixed) at the same location in a room, simultaneously acquire signals of a walking person. Their example outputs are shown in Fig. 1: while the two time series are markedly different, the envelope peaks (corresponding to footfalls) are obviously correlated. In fact, the two modalities are linked through latent physical quantity – (vertical) vibration particle velocity at the impact point – as described in the remainder of the section. Hereafter, \vec{r} denotes the coordinates of the impact (footfall) point relative to the position of the sensors, t denotes time and ω denotes the




Fig. 1. Vibration (top) and acoustic (bottom) recording of a person walking in silence.

angular frequency. The hat notation $\hat{\cdot}$ is used to denote the Fourier representation $\mathcal{F}(\cdot)$ of a signal.

Acoustic pressure signal $\hat{x}_a(\omega, \vec{r}) = \mathcal{F}(x_a(t, \vec{r}))$ can be related to the particle velocity $\hat{v}(\omega)$, as follows [12]:

$$\hat{x}_a(\omega, \vec{r}) = \hat{h}_a(\omega, \vec{r})\hat{v}(\omega) + \hat{e}_a(\omega) = \hat{g}_a(\omega, \vec{r})\frac{\hat{v}(\omega)}{\hat{z}(\omega)} + \hat{e}_a(\omega), \quad (1)$$

where $\hat{e}_a(\omega)$ is the additive noise of the microphone, and $\hat{h}_a(\omega, \vec{r})$ denotes the microphone transfer function. The latter comprises *specific acoustic impedance* $\hat{z}(\omega)$ (which is a material-related quantity of a medium [13]) at the impact point, and the (air) impulse response $\hat{g}_a(\omega, \vec{r})$, relating the impact point and the microphone location. While we may assume that the floor surface is homogeneous – thus $\hat{z}(\omega)$ does not change significantly with regards to \vec{r} – the impulse response $\hat{g}_a(\omega, \vec{r})$ is influenced to a larger extent by the change in position (this has been empirically verified in [14]).

Geophone measures the voltage corresponding to the velocity of the proof mass relative to the device case. In the prescribed frequency range, the velocity of the proof mass can be related to the impact point velocity $\hat{v}(\omega)$ [12] as

$$\hat{x}_g(\omega, \vec{r}) = \hat{h}_g(\omega, \vec{r})\hat{v}(\omega) + \hat{e}_g(\omega) = S_g\hat{g}_g(\omega, \vec{r})\hat{v}(\omega) + \hat{e}_g(\omega), \quad (2)$$

where $\hat{e}_g(\omega)$ is the additive noise of the geophone, and $\hat{h}_g(\omega, \vec{r})$ is the geophone transfer function. Furthermore, S_g denotes the sensitivity constant, while $\hat{g}_g(\omega, \vec{r})$ is the impulse response *within the floor* (hence different from $\hat{g}_a(\omega, \vec{r})$).

Transfer functions $\hat{h}_a(\omega, \vec{r})$ and $\hat{h}_g(\omega, \vec{r})$ (analogously, signals $\hat{x}_a(\omega, \vec{r})$ and $\hat{x}_g(\omega, \vec{r})$) are therefore, dependent on \vec{r} – the parameter we cannot control. This is the relative position between the walking person and the immobile sensors, which thus depends on time t , *i.e.*, $\vec{r} := \vec{r}(t)$. We assume that $\vec{r}(t)$ is a slowly varying function, *i.e.*, the impulse responses are (locally) stationary with respect to \vec{r} within a relatively short temporal window, and one can write

$$x_a(t, \vec{r}) \approx x_a(t) = h_a(t) * v(t) + e_a(t) \quad \text{and} \quad (3)$$

$$x_g(t, \vec{r}) \approx x_g(t) = h_g(t) * v(t) + e_g(t), \quad (4)$$

where $h_a(t)$ and $h_g(t)$ are time-domain representations of $\hat{h}_a(\omega) \approx \hat{h}_a(\omega, \vec{r})$ and $\hat{h}_g(\omega) \approx \hat{h}_g(\omega, \vec{r})$, respectively.

3. FEATURE EXTRACTION

The hypothesis is that the impact velocity $v(t)$ is sufficiently informative to discriminate people. The sensors, however, measure only the bandlimited convolution of $v(t)$ with the corresponding transfer functions. Fortunately, the local stationarity assumption enables us to exploit cancellation property of the so-called *normalized scattering* representation.

3.1. Scattering transform

Scattering transform is a novel feature extraction method, based on a cascade of wavelet transforms and modulus nonlinearities, bearing some resemblance to convolutional neural networks [15, 16]. An appealing property of scattering networks is that their filters are pre-defined, hence they require no training. Yet, classifiers using scattering features exhibit almost state-of-the-art performance on several problems, *e.g.*, [11, 15]. In the following, we briefly describe how the scattering transform is computed on the audio signal $x_a(t)$. The features from $x_g(t)$ are extracted in the same manner.

For a scattering of *order* p , the features are computed as $S_{\lambda_1 \dots \lambda_p}(x_a(t), t) = \phi_T(t) * U_{\lambda_1 \dots \lambda_p}(x_a(t), t)$. Here, ϕ_T denotes the real-valued lowpass filter of bandwidth $2\pi/T$ (where T is the targeted extent of time-invariance), and $U_{\lambda_1 \dots \lambda_p}(\cdot)$ is the so-called *wavelet propagator*¹:

$$U_{\lambda_1 \dots \lambda_p}(x_a) = |\psi_{\lambda_p} * |\psi_{\lambda_{p-1}} * \dots * |\psi_{\lambda_1} * x_a| \dots|, \quad (5)$$

where $\psi_{\lambda_i} := \psi_{\lambda_i}(t)$ is a complex analytic wavelet filterbank at $0 < i \leq p$. The set of scales $\lambda_i \in \Lambda_i$ is chosen such that the filterbank covers the frequency range $[\pi/T, \omega_a/2]$ (ω_a is the sampling frequency), possibly with certain redundancy. The expression above defines the recursion $U_{\lambda_1 \dots \lambda_p}(x_a) = |\psi_{\lambda_p} * U_{\lambda_1 \dots \lambda_{p-1}}(x_a)|$, with $U_{\emptyset}(x_a) = x_a$ at $i = 0$.

In [11], the authors further refine scattering features by making them nearly invariant to convolution by a filter h , when h is almost constant on the support of ψ_{λ_i} , which we assume to hold in our application. These normalized scattering coefficients are computed as component-wise division

$$\tilde{S}_{\lambda_1}(x_a) = \frac{S_{\lambda_1}(x_a)}{\phi_T * |x_a| + \varepsilon}, \quad \varepsilon > 0, \quad \text{for } i = 1, \quad (6)$$

$$\text{and } \tilde{S}_{\lambda_1 \dots \lambda_i}(x_a) = \frac{S_{\lambda_1 \dots \lambda_i}(x_a)}{S_{\lambda_1 \dots \lambda_{i-1}}(x_a)}, \quad \text{for } i > 1. \quad (7)$$

The zero-order coefficients $\tilde{S}_{\emptyset}(x_a) := S_{\emptyset}(x_a) = \phi_T * x_a$ remain unchanged.

Hence, if we independently consider signal segments of duration τ for which our local stationarity assumption holds, the normalized scattering features should be nearly invariant to filtering by h_a (accordingly, filtering by h_g for the geophone signal), and would mostly reflect the behavior of the fingerprint function v in a given bandwidth.

¹For notational convenience, the variables t and ω are dropped when the dependence is obvious.




Fig. 2. Normalized ($p = 1$) scattering features for the vibration (top) and audio (bottom) modality, representing the same person, at different time instances (left/right).

3.2. Feature fusion

When bimodal (microphone and geophone) measurements are available, one can exploit the fact that their effective bandwidths – ranges for which the magnitudes of frequency responses are large – are somewhat complementary. Excluding $\tilde{S}_\emptyset(\cdot)$, their respective normalized scattering representations should be complementary as well: the most informative coefficients of each modality appear at scales that do not overlap with one another, except perhaps within a narrow band. Indeed, while the vibration signal x_g has a very low and narrow frequency range, the audio x_a is a wideband signal.

This intuition can be verified in Fig. 2, where dark color indicates low magnitude coefficients, and vice-versa. For simplicity, the geophone signal x_g is upsampled to match the length of the audio signal x_a , thus the feature matrices have the same size. This suggests a simple fusion technique: since the coefficients are nonnegative, one can simply compute a weighted average of the two modalities to obtain a more informative representation, whose (implicit) effective bandwidth is extended. We remark that this is not a pure heuristics, as normalized scattering approach described before places the two representations in the same “impact velocity feature space”.

The fused scattering $S^{(f)}$ at orders $i \geq 1$ is given as

$$S_{\lambda_1, \dots, \lambda_i}^{(f)} = \alpha(x_a) \tilde{S}_{\lambda_1, \dots, \lambda_i}(x_a) + \alpha(x_g) \tilde{S}_{\lambda_1, \dots, \lambda_i}(x_g), \quad (8)$$

with a weight $\alpha(\cdot)$ defined as

$$\alpha(\cdot) = \left(\max_{\{\lambda_i \in \Lambda_i\}_{1 \leq i \leq p}} \tilde{S}_{\lambda_1, \dots, \lambda_i}(\cdot) \right)^{-1}, \quad (9)$$

to account for the magnitude disparity among the modalities.

The rows corresponding to zero-order coefficients $\tilde{S}_\emptyset(x_a)$ and $\tilde{S}_\emptyset(x_g)$ are simply concatenated with the fused ones.

3.3. Feature postprocessing

The lowpass filtering by ϕ_T makes the output invariant to translations smaller than T . It was shown [16] that the information loss induced by lowpass filtering is compensated by the higher-order scattering coefficients, with the scattering order p predominantly driven by the signal content [15, 11]. The rule of thumb is that the larger T is, the higher order the scattering transform should be. Unfortunately, this significantly increases computational complexity: scattering transform yields a tree-like representation (*cf.* Fig. 2 in [15]), where each “path” $\{\lambda_1, \lambda_2, \dots, \lambda_p\}_{\lambda_1 \in \Lambda_1, \lambda_2 \in \Lambda_2, \dots, \lambda_p \in \Lambda_p}$ needs to be traversed (*i.e.*, a full sequence of convolutions needs to be performed) to reach a leaf node.

As applications enabled by person identification often require real-time processing, our aim is to reduce the computational burden and compute normalized scattering features only up to $p = 1$ order (“shallow” scattering network), which implies that T cannot be large. However, features computed from very short signal segments cannot capture temporal dynamics of the gait, which we deem useful for identification. Indeed, the average period of normal walk is about 1.22s (two footfalls with the same leg) [17], and computing sufficiently informative scattering features with T that large is computationally prohibitive. While sophisticated classifiers, such as those based on Hidden Markov Models [10], could be used to model the temporal dynamics between successive feature vectors, we opted for a simpler alternative. By inspecting two scattering feature matrices, with the same label but computed at different time instances (Fig. 2 left and right), one may notice that the main source of variability is due to global temporal offset. This can be easily suppressed by computing the Fourier transform of the scattering matrix across temporal direction, and applying the modulus operator, *i.e.*, by discarding the phase. Thus, we extract a segment of duration $\tau > 1.22s \gg T$, and then postprocess the obtained feature matrix by applying the Fourier modulus row-wise. Since very long segments violate the local stationarity assumption, we set $\tau \approx 1.5s$. Hence, the postprocessing phase introduces an additional layer to the first order scattering network.

As suggested in [11], to separate multiplicative signal components and reduce dimensionality, we apply logarithm and *PCA* (*Principal Component Analysis*) – or its approximation through *DCT* (*Discrete Cosine Transform*) – to postprocessed the feature matrix. The features are standardized (centered and scaled to unit variance) before *PCA* (*DCT*).

4. RESULTS

While gait recognition attracted considerable amount of research, vibration- and audio-based bimodal identification has not been investigated so far, to the best knowledge of the authors. This led us to build our own dataset, by simultaneously recording signals using one IONTM SM-24 geophone (sam-

| $T \backslash N$ | 30 | 50 | 100 | 150 |
|------------------|---------------|--------|--------|---------------|
| 0.046s | 23.18% | 20.00% | 16.13% | 15.08% |
| 0.093s | 16.56% | 15.48% | 15.96% | 13.35% |
| 0.186s | 16.67% | 15.96% | 15.08% | 16.13% |
| 0.371s | 20.00% | 19.34% | 19.35% | 21.64% |
| 0.046s | 26.67% | 25.81% | 27.92% | 26.80% |
| 0.093s | 25.49% | 23.33% | 23.47% | 31.94% |
| 0.186s | 20.78% | 24.44% | 29.51% | 30.03% |
| 0.371s | 23.87% | 29.22% | 28.89% | 29.03% |
| 0.046s | 19.05% | 16.67% | 12.86% | 12.09% |
| 0.093s | 16.13% | 12.79% | 12.38% | 10.00% |
| 0.186s | 15.96% | 14.74% | 13.33% | 13.33% |
| 0.371s | 19.15% | 16.67% | 18.21% | 19.68% |

Table 1. EER performance of the audio (top), vibration (middle) and fused (bottom) features (lower is better).

pling rate \approx 1kHz), and one Samson Meteor[®] microphone (44.1kHz). The recordings involved 8 male and 4 female participants, each recorded during three days, and asked to wear the same type of shoes on (at least) two different days. All recordings were taken in the same room with carpet floor covering. The participants walked the same route 10 times per day: starting \sim 6m away, they approached the sensors, and returned to the initial point.

Open set identification refers to the case when classes not seen during training may appear in the test phase, and the recognition system needs to label them as “unknowns”. This type of problem is common in speaker recognition, which shares many traits with gait-based identification (interestingly, in the referenced literature, we found no connection between the two). The gist of current state-of-technology in speaker recognition are variants of *GMM-UBM* (*Gaussian Mixture Model - Universal Background Model*) framework – an interested reader may consult, *e.g.*, [2, 3] – which we here apply to gait identification. The gait dataset is divided into the “training” and “test” sets, such that the “training” set contains recordings taken on those two days when the participants wore different type of shoes. In this way, we ensure that the training data is sufficiently diverse. The data recorded on the third day constitutes the test set.

We split the training dataset such that the recordings of 6 randomly chosen individuals are used for training the UBM, and the training data of 3 among the remaining 6 (also randomly chosen), is used for the enrolment, *cf.* [3]. The test data of these 6 participants is used in the evaluation phase (thus, there are 3 unknown persons). This random partitioning is repeated 100 times, to verify that the results are consistent.

Normalized scattering, with a redundant Morlet wavelet filterbank, is computed on overlapping signal segments of duration τ (stepsize = 0.25s), using the Scatnet toolbox [18]. The GMM-UBM system [19], with 64 Gaussian components, is then fed with the postprocessed scattering features.




Fig. 3. Best performance for each feature type.

Series of experiments is performed by varying the hyperparameters T and N (the number of retained DCT coefficients), for each random partition. The median results, in terms of *EER* (*Equal Error Rate*) [2], are presented in Table 1. Overall, as expected, with the geophone-only features the recognition is somewhat poor. The audio modality performs better, while the fused features perform best, regardless of parameterization. Boxplots for the best-performing parameterizations (boldface values in the table), in Fig. 3, show that the EERs of the fused representation have the smallest variance. Concerning the choice of time-invariance parameter T , the optimal value is between 0.093s and 0.186s, which is consistent with average duration of the footfall impact event [17]. The preferred number of features seems to be modality-dependent (*e.g.*, richer representations favor larger N), and may be related to the preset number of GMM components.

5. CONCLUSION

We have presented a novel feature extraction approach for person identification based on audio and vibration gait measurements. In a low ambient noise environment, using the audio modality increases recognition accuracy, as demonstrated by the exhaustive experimentation on our bimodal signal dataset. Additionally, we have shown that the two modalities can be fused together to further improve recognition performance. Future work will focus on recognition in adverse conditions, *e.g.*, in the presence of auditory noise, and/or several people walking. For the latter, we feel that a body of work on speaker diarization [20] could be exploited to target such problems. Moreover, bimodal data may offer distinct advantages, both in terms of “walker diarization”, but also in terms of robustness to ambient noise, since the two modalities are usually not simultaneously affected by the same noise source. Finally, in this work we opted for (deterministic) scattering feature extraction, due to the size of our training dataset. If this is not a limiting factor, recent trends in machine learning suggest that a deep neural network may achieve superior performance.

6. REFERENCES

- [1] A. Jain and S. Li, *Handbook of face recognition*, Springer, 2011.
- [2] J. Hansen and T. Hasan, “Speaker recognition by machines and humans: A tutorial review,” *IEEE Signal processing magazine*, vol. 32, no. 6, pp. 74–99, 2015.
- [3] D. Reynolds and W. Campbell, “Text-independent speaker recognition,” in *Springer Handbook of Speech Processing*, pp. 763–782. Springer, 2008.
- [4] J. Phillips, S. Sarkar, I. Robledo, P. Grother, and K. Bowyer, “The gait identification challenge problem: Data sets and baseline algorithm,” in *16th International Conference on Pattern Recognition, 2002*. IEEE, 2002, vol. 1, pp. 385–388.
- [5] M. Hofmann, J. Geiger, S. Bachmann, B. Schuller, and G. Rigoll, “The tum gait from audio, image and depth (gaid) database: Multimodal recognition of subjects and traits,” *Journal of Visual Communication and Image Representation*, vol. 25, no. 1, pp. 195–206, 2014.
- [6] D. Bales, P. Tarazaga, M. Kasarda, D. Batra, A. Woolard, J. Poston, and V. Malladi, “Gender classification of walkers via underfloor accelerometer measurements,” *IEEE Internet of Things Journal*, 2016.
- [7] S. Pan, N. Wang, Y. Qian, I. Velibeyoglu, H. Y. Noh, and P. Zhang, “Indoor person identification through footstep induced structural vibration,” in *Proceedings of the 16th International Workshop on Mobile Computing Systems and Applications*. ACM, 2015, pp. 81–86.
- [8] M. Hons and R. Stewart, “Transfer functions of geophones and accelerometers and their effects on frequency content and wavelets,” *CREWES Res. Rep.*, vol. 18, pp. 1–18, 2006.
- [9] A. Ekimov and J. Sabatier, “Ultrasonic wave generation due to human footsteps on the ground,” *The Journal of the Acoustical Society of America*, vol. 121, no. 3, pp. EL114–EL119, 2007.
- [10] J. Geiger, M. Kneißl, B. Schuller, and G. Rigoll, “Acoustic gait-based person identification using hidden Markov models,” in *Proceedings of the 2014 Workshop on Mapping Personality Traits Challenge and Workshop*. ACM, 2014, pp. 25–30.
- [11] J. Andén and S. Mallat, “Deep scattering spectrum,” *IEEE Transactions on Signal Processing*, vol. 62, no. 16, pp. 4114–4128, 2014.
- [12] A. Ekimov and J. Sabatier, “Vibration and sound signatures of human footsteps in buildings,” *The Journal of the Acoustical Society of America*, vol. 118, no. 3, pp. 762–768, 2006.
- [13] F. Fahy, *Foundations of engineering acoustics*, Academic press, 2000.
- [14] D. Bard, J. Sonnerup, and G. Sandberg, “Human footsteps induced floor vibration,” *Journal of the Acoustical Society of America*, vol. 123, no. 5, pp. 3356, 2008.
- [15] J. Bruna and S. Mallat, “Invariant scattering convolution networks,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 35, no. 8, pp. 1872–1886, 2013.
- [16] S. Mallat, “Understanding deep convolutional networks,” *Phil. Trans. R. Soc. A*, vol. 374, no. 2065, pp. 20150203, 2016.
- [17] A. Ekimov and J. Sabatier, “Rhythm analysis of orthogonal signals from human walking,” *The Journal of the Acoustical Society of America*, vol. 129, no. 3, pp. 1306–1314, 2011.
- [18] L. Sifre, M. Kapoko, E. Oyallon, and V. Lostanlen, “Scatnet: a MATLAB toolbox for scattering networks,” 2013.
- [19] S. Sadjadi, M. Slaney, and L. Heck, “MSR identity toolbox v1.0: A matlab toolbox for speaker-recognition research,” *Speech and Language Processing Technical Committee Newsletter*, vol. 1, no. 4, 2013.
- [20] X. Anguera, S. Bozonnet, N. Evans, C. Fredouille, G. Friedland, and O. Vinyals, “Speaker diarization: A review of recent research,” *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 20, no. 2, pp. 356–370, 2012.