



The bag-of-frames approach: a not so sufficient model for urban soundscapes

Mathieu Lagrange, Grégoire Lafay, Boris Defreville, Jean-Julien Aucouturier

► **To cite this version:**

Mathieu Lagrange, Grégoire Lafay, Boris Defreville, Jean-Julien Aucouturier. The bag-of-frames approach: a not so sufficient model for urban soundscapes. 2014. <hal-01082501>

HAL Id: hal-01082501

<https://hal.archives-ouvertes.fr/hal-01082501>

Submitted on 13 Nov 2014

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.

The bag-of-frames approach: a not so sufficient model for urban soundscapes

Mathieu Lagrange, Grégoire Lafay, Boris Defreville, Jean-Julien
Aucouturier
mathieu.lagrange@cnrs.fr

Abstract:

The “bag-of-frames” approach (BOF), which encodes audio signals as the long-term statistical distribution of short-term spectral features, is commonly regarded as an effective and sufficient way to represent environmental sound recordings (soundscapes) since its introduction in an influential 2007 article. The present paper describes a conceptual replication of this seminal article using several new soundscape datasets, with results strongly questioning the adequacy of the BOF approach for the task. We show that the good accuracy originally reported with BOF likely result from a particularly thankful dataset with low within-class variability, and that for more realistic datasets, BOF in fact does not perform significantly better than a mere one-point average of the signal’s features. Soundscape modeling, therefore, may not be the closed case it was once thought to be. Progress, we argue, could lie in reconsidering the problem of considering individual acoustical events within each soundscape.

PACS numbers: 43.60.Cg, 43.50.Rq, 43.60.Lq, 43.66.Jh, 43.66.Ba

1. Introduction

In this paper, we study a particularly common pattern-recognition paradigm, the “bag-of-frames” approach (BOF), which represents audio signals as the long-term statistical distribution of their short-term spectral features. Typically implemented as a Gaussian mixture model (GMM) of Mel-Frequency Cepstrum Coefficients (MFCCs), the BOF approach is commonly applied to model music [18] and soundscapes [4] and is a standard component in several audio pattern recognition software packages [13]. Further, recent psychoacoustical evidence suggest the approach bears some resemblance with human auditory processing for sound textures [11, 12].

In an influential 2007 article, Aucouturier, Defreville & Pachet [2] applied a BOF model to categorize both polyphonic music and soundscapes. Their results showed that, while BOF was a meriting model for their polyphonic music dataset, it was spectacularly effective for soundscapes, reaching accuracies of 96%. The contrast, they interpreted, lied in differences in the temporal structure of both types of stimuli, with music being more strongly organized and soundscapes more easily summarized by statistics. In a later companion study [1], they showed that soundscapes could be time-shuffled without changing much of their perceptive qualities, a process which however impaired most of music’s perception. While more work was needed for music, the authors therefore concluded that BOF was a sufficient model to approximate human perception for soundscapes, practically ruling out the need to recognize the local acoustic events in a texture in order to identify it.

That article, which is commonly regarded as the birth act of BOF modeling for audio, was influential in the soundscape community (118 citations. Source: Google Scholar, Oct. 2014). BOF found applications in audio sound scene recognition [4], audio event identification [7] and soundscape quality assessment [14]. Because of its original comparison with more-demanding music stimuli, soundscape BOF modeling was also taken as a benchmark against which new music models were developed, and standardized soundscape datasets were assembled to conduct algorithm evaluations [7]. Though, to the best of our knowledge, the adequacy of the BOF approach for soundscapes, although mostly derived from a single seminal study, was never challenged. Not only did the original BOF article introduced a new task (soundscape modeling) to the world of pattern recognition, it also appeared to close the case in the same effort.

In this paper, we conduct a conceptual replication of the 2007 BOF article, using not only the original dataset, but three others, more recent soundscape datasets with less within-class variability. On these four datasets, we compare the original BOF model proposed by Aucouturier, Defreville and Pachet (2007) to the simple one-point averaging of sound features, on a standard task of categorization. We show that, while the excellent soundscape accuracy of BOF is indeed replicated on the original dataset, it does not hold, by far, for more varied and realistic data. More surprisingly perhaps, for all datasets, BOF did not perform significantly better than the much simpler one-point average approach. In other words, soundscape modeling may not be not the closed case it was once thought to be.

2. Experimental Protocol

2.1. Algorithms

We implemented the same BOF algorithm that was originally proposed in [2], namely Gaussian Mixture Models (GMMs) of Mel-Frequency Cepstrum Coefficients (MFCCs). MFCCs are computed as coefficients from the Fourier transform along frequency of the signal's log spectrogram. By extracting MFCCs on every successive short-time window of a signal, it is transformed into a time-series of feature points in an euclidean space. Here, we extracted $d=20$ MFCC coefficients from 2048-point time windows using the AUDITORY toolbox implementation [15].

GMMs are statistical models used to estimate a probability distribution $\mathcal{P}(x)$ as the weighted sum of M gaussian distributions $\mathcal{N}_i, \forall i < M$, each parameterized by a mean μ_i and covariance matrix Σ_i ,

$$\mathcal{P}(x) = \sum_i^M \pi_i \mathcal{N}_i(x, \mu_i, \Sigma_i) \quad (1)$$

where π_i is the weight of gaussian distribution \mathcal{N}_i . Given a collection of MFCC feature vectors, viewed as samples from a random variable, the parameters $\pi_i, \mu_i, \Sigma_i, \forall i < M$ of a GMM that maximizes the likelihood of the data can be estimated by the Expectation Maximization (EM) algorithm. Here, we took $M=50$.

In order to compare two audio signals p and q , we estimate the parameters of a GMM for each signal's collection of MFCC vectors $p[n]$ and $q[m]$, and then compare the two GMMs \mathcal{P}_p and \mathcal{P}_q using their Kullback Leibler (KL) divergence:

$$d_{KL}(p, q) = \int \mathcal{P}_p(x) \log \frac{\mathcal{P}_q(x)}{\mathcal{P}_p(x)} \quad (2)$$

As a comparison for BOF, we also evaluated a simpler matching algorithm in which a signal's MFCC series $p[n]$ is averaged along time to a single mean MFCC vector

\tilde{p} of size $d=20$:

$$\tilde{p} = \frac{1}{N_T} \sum_{t=1}^{t=N_T} p[t], \quad (3)$$

where N_T is the number of measured time points in the original representation. Two average MFCC vectors can then be compared using the simple euclidean distance, defined as

$$d_e(\tilde{p}, \tilde{q}) = \sqrt{\sum_i (\tilde{p}_i - \tilde{q}_i)^2} \quad (4)$$

where \tilde{p}_i and \tilde{q}_i are the i^{th} coordinate of average vectors \tilde{p} and \tilde{q} .

2.2. Datasets

We collected four datasets from the previous work of different laboratories (Table 1):

1. **AucoDefr07**: Used in the original BOF article, this dataset was recorded by Boris Defreville in Paris (France) and consists of 16 recordings distributed in 4 classes: “avenue”, “street”, “market” and “park”. The first class consists of 5 recordings of 3 avenues, the second of 3 recordings of 2 streets, the third of 5 recordings of 2 markets and the fourth of 3 recordings of 1 park. For the work of [2], these recordings were segmented into 3-minute units, giving 78 units with about 20 units per class. We are publishing this previously unreleased dataset as part of this paper: https://archive.org/details/defreville-Aucouturier_urbanDb.
2. **Guastavino**: provided by Catherine Guastavino and recorded as material for psychoacoustics experiments [9]. The dataset consists of 16 recordings of less than 3 minutes, distributed in 5 classes: “vehicles”, “vehicles and pedestrians”, “cafes”, “markets” and “parks”. Each class consist of 3 recordings of 2 places, except for the fourth which consists of 3 recordings of 3 places.
3. **Tardieu**: recorded by Julien Tardieu in Paris (France) for the purpose of psychoacoustics experiments [19], and consists of 66 recordings in 6 different French train stations distributed in 6 classes: “Platforms”, “Halls”, “Corridors”, “Waiting room”, “Ticket booths” and “Shops”. Classes are well balanced in terms of number of recordings and number of different locations.
4. **QMUL**: recorded in a wide variety of Greater London (UK) locations over Summer and Autumn 2012 by 3 members of Queen Mary University of London [6]. The dataset consists of 100 30-second recordings equally distributed among 10 classes: “bus”, “busy street”, “office”, “open air market”, “park”, “quiet street”, “restaurant”, “supermarket”, “tube”, and “tube station”. To ensure that no systematic variation in the recordings covaried with scene type, all recordings were made in moderate weather conditions, at varying times of day and week, and each operator recorded occurrences of every scene type. The public part of the dataset considered in this study is available on the C4DM Research Data Repository: <http://c4dm.eecs.qmul.ac.uk/rdr/handle/123456789/29>.

These four datasets were not created equal in their potential to validate computational approaches. The former two appear limited in terms of variety of recording locations (4 and 5, resp.) and number of recordings per class (4 and 3, resp.). Results obtained on these datasets shall therefore be considered with care. The latter two, while still rather small for a classification task, boast a more realistic variety of recording locations (6 and 10, resp.) and number of recording per class (11 and 10, resp.).

dataBase	number	recordings		p@5	human	
		number per class	duration		map	accuracy
AucoDefr07	16	4±1	947±597	-	-	-
Guastavino	16	3±0	15±1	96	99	-
Tardieu	66	11±4	17±2	72	63	-
QMUL	100	10±0	30±0	-	-	75

Table 1. Morphology of the soundscape datasets used in this study. Where available, we also report categorization performance of humans (%).

2.3. Metrics

We followed the same evaluation paradigm as [2]. Algorithms were evaluated on their ability to retrieve, for a given seed item, nearest neighbors that belong to the same category in a leave-one-out fashion. This ability was measured with two standard metrics: 1) the precision at rank 5 (p@5), *i.e.* the average number of items of the same class among the 5 closest items and 2) the Mean Average Precision (MAP), *i.e.* the average precision across all possible ranks. For both metrics, the target rank is bounded by the cardinality of the relevant class. Statistical significance was tested using paired-sample t-test at the 5 % significance level.

3. Results

We tested for the impact of a few algorithmic design variants within the general guidelines provided in [2]: whether audio recordings are amplitude-normalized prior to encoding, whether the first MFCC coefficient (related to the energy at the frame level) is discarded before modeling, and whether models are compared using the Monte Carlo sampling (MCMC) method or by the direct marginalization on the features. We found that none of the factors significantly affected either metrics. In the following, we used the configuration that gave the best result on the QMUL dataset, that is: audio normalization, including first MFCC coefficient, and MCMC sampling.

dataBase	chance	Average	BOF
AucoDefr07	44 / 37	96 / 82	97* / 83*
Guastavino	28 / 58	39 / 67	40* / 69*
Tardieu	33 / 33	49* / 43*	48 / 42
QMUL	28 / 25	47 / 42*	48* / 38

Table 2. Categorization performance in terms of precision-at-rank-5 and mean-average-precision (P@5 / MAP) in percents. Performance scores in bold are tested as non significantly worse than the best performance per dataset (depicted with *).

Table 2 show categorization precision scores broken by dataset. Scores appearing in bold are non-significantly worse than the best score per dataset. The evaluated computational approaches perform better than random, showing that some information is indeed retained by the models. Further, MAP scores do not strongly differ from p@5, suggesting coherent performance across different ranks (the large MAP values achieved with the Guastavino dataset are most probably due to its small size).

While performance scores in our tests are consistent across the Guastavino, Tardieu and QMUL datasets, performance for the original AucoDefr07 dataset clearly stands out. We believe that this result can be explained by two factors. First, in order to replicate results discussed in [2], the recordings were divided into 3-minute units and these units were next categorized. This means that units that are close to one

another will most probably come from the same recording. If this processing step is removed and each recording is categorized as a non-separable whole, we found precision dropped to 71 %. Second, the average number of recording locations per class is only 2, which may have the same effect of biasing the results to favor recordings from identical locations.

The average precision over the two most realistic datasets in this study, Tardieu and QMUL, was only 48 %. We believe that this figure is a more appropriate benchmark of the adequacy of BOF for soundscape categorization than the 96% obtained on the overly thankful AucoDefr07 dataset. This figure can be compared to that obtained by BOF for musical genre classification (77%) in [2].

Finally, for all four datasets, no statistical difference was found between the BOF approach and a simple one-point feature average. The computational complexity involved in computing and comparing GMM models therefore seems an unnecessary burden: not only is there no added advantage in considering several, rather than one, gaussian components per item, but there also does not seem to be any discriminative information encoded in the standard deviation, rather than in the sole mean value of the feature distribution. Furthermore, with simple averaging, recordings are mapped directly in a 20-dimensional Euclidean space, which allows faster matching and indexing methods than the GMM-based metric space.

4. Discussion

These new results challenge the commonly-held view that the BOF approach is a sufficient model to categorize soundscapes. While we replicated the excellent performance of BOF described in [2], it appears that this figure likely results from an overly thankful dataset with abnormally low within-class variability: a majority of exemplars in each class were different extracts of the same recordings, with plausibly strongly similar acoustic characteristics that would not generalize across different recordings. For the more realistic (yet not overly complex) tasks associated with the three alternative datasets tested here, BOF performance was severely degraded (48%). In fact, it did not score significantly better than a mere one-point average of the signal's features.

These results suggest that, while meaningful information for soundscape categorization is preserved in their summary statistics, this information by no means allows a sufficient level of generalization to handle categories with strong intrinsic variability. While recent results showed some improvements using more advanced statistical techniques than the basic BOF implementation tests here [4], it remains doubtful whether purely textural models can solve the problem.

Far from an anti-BOF epiphany, we believe these results should be taken to indicate that soundscape categorization is, and remains, a difficult computational problem when considered at realistic levels of within- and between-class variability. A more promising approach would likely combine textural modeling with prior recognition of individual sources. To this aim, efforts could be made at developing ecologically- and cognitively-grounded context models, linking events to textures (e.g. car horns, traffic sounds, pedestrian chatter to a “busy street” texture). Such efforts will need to confront such hard questions as what prior information can be reliably assumed about an event's distribution across soundscapes and about an event's acoustic signature, and what distinguishes an event from its background model. These are questions routinely and efficiently solved by the human auditory system [3, 8, 10, 12]; the present results show that our best computational models are not there yet.

From a methodological point of view, the problem found here with the AucoDefr07 dataset is reminiscent of a similar fallacy identified in the music information retrieval community for the task of musical genre classification. In early datasets contemporaneous with AucoDefr07, exemplars from a given genre category often consisted, by ease of construction, in several songs by the same artist or the same album

- which artificially boosted the classification accuracy [16]. Guidelines to avoid such biases were only recently introduced [5], and debate is still ongoing as to what constitutes a dataset apt to evaluate the ability to recognize the concept of genre, rather than an incidental correlate thereof [17].

More generally, this paper's critical re-examination of previously published results is one more argument in favor of direct and conceptual replications in science in general, and computational science in particular. This research was only made possible by our peers' willingness to share their research datasets or to publish them as part of evaluation campaigns[6]. These practices are making science a better place, and we were grateful to benefit from them here.

Acknowledgments

Research project partly funded by ANR-11-JS03-005-01(HOULE) and ERC Grant 335536 (CREAM). The authors would like to thank Catherine Guastavino, Damien Tardieu for sharing their datasets.

References and links

- [1] Jean-Julien Aucouturier and Boris Defreville. Judging the similarity of soundscapes does not require categorization: Evidence from spliced stimuli. *The Journal of the Acoustical Society of America*, 125(4):2155–2161, 2009.
- [2] Jean-Julien Aucouturier, Boris Defreville, and François Pachet. The bag-of-frames approach to audio pattern recognition: a sufficient model for urban soundscapes but not for polyphonic music. *The Journal of the Acoustical Society of America*, 122(2):881–891, 2007.
- [3] James A Ballas and Timothy Mullins. Effects of context on the identification of everyday sounds. *Human performance*, 4(3):199–219, 1991.
- [4] D. Barchiesi, D. Giannoulis, D. Stowell, and M. D. Plumbley. Acoustic scene classification. *IEEE Signal Processing Magazine*, in press.
- [5] Arthur Flexer. A closer look on artist filters for musical genre classification. In *Proc. of the 2007 International Conference on Music Information Retrieval (ISMIR)*, 2007.
- [6] Dimitrios Giannoulis, Dan Stowell, Emmanouil Benetos, Mathias Rossignol, Mathieu Lagrange, and Mark D Plumbley. A database and challenge for acoustic scene classification and event detection. In *Proceedings EUSIPCO*, pages 1–5, 2013.
- [7] Dimitrios Giannoulis, Dan Stowell, Emmanouil Benetos, Mathias Rossignol, Mathieu Lagrange, and Mark D Plumbley. Detection and classification of acoustic scenes and events: an IEEE AASP challenge. In *WASPAA*, 2013.
- [8] C. Guastavino, B. Katz, J.D. Polack, D.J. Levitin, and D. Dubois. Ecological validity of soundscape reproduction. *Acta Acustica united with Acustica*, 91:333–341, 2005.
- [9] Catherine Guastavino. Categorization of environmental sounds. *Canadian journal of experimental psychology = Revue canadienne de psychologie expérimentale*, 61(1):54–63, March 2007.
- [10] Robert Leech, Brian Gygi, Jennifer Aydelott, and Frederic Dick. Informational factors in identifying environmental sounds in natural auditory scenes. *The Journal of the Acoustical Society of America*, 126(6):3147–3155, 2009.
- [11] Josh H McDermott, Michael Schemitsch, and Eero P Simoncelli. Summary statistics in auditory perception. *Nature neuroscience*, 16(4):493–8, April 2013.
- [12] Israel Nelken and Alain de Cheveigné. An ear for statistics. *Nature neuroscience*, 16(4):381–2, April 2013.
- [13] Kevin R Page, Ben Fields, David De Roure, Tim Crawford, and J Stephen Downie. Capturing the workflows of music information retrieval for repeatability and reuse. *Journal of Intelligent Information Systems*, 41(3):435–459, 2013.
- [14] Tae Hong Park, Johnathan Turner, Michael Musick, Jun Hee Lee, Christopher Jacoby, Charlie Mydlarz, and Justin Salamon. Sensing urban soundscapes. In *EDBT/ICDT Workshops*, pages 375–382, 2014.
- [15] Malcolm Slaney. "auditory toolbox,". *Interval Research Technical Report*, 1998.
- [16] Bob L Sturm. An analysis of the gtzan music genre dataset. In *Proceedings of the second international ACM workshop on Music information retrieval with user-centered and multimodal strategies*, pages 7–12, 2012.
- [17] Bob L Sturm. Classification accuracy is not enough. *Journal of Intelligent Information Systems*, 41(3):371–406, 2013.
- [18] L. Su, C. Yeh, J. Liu, J. Wang, and Y. Yang. A systematic evaluation of the bag-of-frames representation for music information retrieval. *IEEE Transactions on Multimedia*, 16(5):1188–1200, 2014.
- [19] Julien Tardieu, Patrick Susini, Franck Poisson, Pauline Lazareff, and Stephen McAdams. Perceptual study of soundscapes in train stations. *Applied Acoustics*, 69(12):1224–1239, December 2008.