



This is a repository copy of *Background-tracking acoustic features for genre identification of broadcast shows*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/92449/>

Version: Accepted Version

---

**Proceedings Paper:**

Saz, O., Doulaty, M. and Hain, T. (2014) Background-tracking acoustic features for genre identification of broadcast shows. In: Spoken Language Technology Workshop (SLT), 2014 IEEE. Spoken Language Technology Workshop (SLT), 07-10 Dec 2014, South Lake Tahoe, NV. IEEE , 118 - 123. ISBN 9781479971299

<https://doi.org/10.1109/SLT.2014.7078560>

---

**Reuse**

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

**Takedown**

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing [eprints@whiterose.ac.uk](mailto:eprints@whiterose.ac.uk) including the URL of the record and the reason for the withdrawal request.



[eprints@whiterose.ac.uk](mailto:eprints@whiterose.ac.uk)  
<https://eprints.whiterose.ac.uk/>

# BACKGROUND-TRACKING ACOUSTIC FEATURES FOR GENRE IDENTIFICATION OF BROADCAST SHOWS

*Oscar Saz, Mortaza Doulaty, Thomas Hain*

Speech and Hearing Group, Department of Computer Science, University of Sheffield, UK

## ABSTRACT

This paper presents a novel method for extracting acoustic features that characterise the background environment in audio recordings. These features are based on the output of an alignment that fits multiple parallel background-based Constrained Maximum Likelihood Linear Regression transformations asynchronously to the input audio signal. With this setup, the resulting features can track changes in the audio background like appearance and disappearance of music, applause or laughter, independently of the speakers in the foreground of the audio. The ability to provide this type of acoustic description in audiovisual data has many potential applications, including automatic classification of broadcast archives or improving automatic transcription and subtitling. In this paper, the performance of these features in a genre identification task in a set of 332 BBC shows is explored. The proposed background-tracking features outperform short-term Perceptual Linear Prediction features in this task using Gaussian Mixture Model classifiers (62% vs 72% accuracy). The use of more complex classifiers, Hidden Markov Models and Support Vector Machines, increases the performance of the system with the novel background-tracking features to 79% and 81% in accuracy respectively.

**Index Terms**— Acoustic background, genre identification, broadcast data.

## 1. INTRODUCTION

The media domain presents many opportunities for the application of speech technologies. With audiovisual data growing larger and larger every day due to digital television, social media and on-line streaming there is a great need for performing automatic processing of this type of data. Possible applications include automatic transcription and subtitling, classification of audiovisual archives and acoustic information retrieval. Further research in this area is being also pushed by initiatives like the MediaEval Benchmarking for Multimedia Evaluation [1], which covers several of these tasks in the multimedia domain. The technologies required cover the whole

range of speech technologies: Automatic Speech Recognition (ASR); speaker identification; diarisation; identification of acoustic events; etc.

The ability of automatically detecting the genre of a broadcast show falls within the set of potential applications of speech technologies that could become useful within the media domain. While genres are subjective divisions usually defined depending on the content of the show, shows belonging to the same genre will share similar acoustic conditions that can be detected using automatic speech processing. In this context, multimodal approaches, merging features from audio and video processing, have been very commonly used [2, 3, 4, 5] and have consistently provided results above 90-95% accuracy. Regarding the type of acoustic features used, from the early works the focus of research has been on the use of short-term features [6], including Mel-Frequency Cepstral Coefficient (MFCC) features [7]. A full evaluation of the use of MFCCs and Gaussian Mixture Model (GMM) classifiers across 3 different test sets achieved 86% in a RAI dataset, 78% in a Quaero dataset and 58% in a YouTube dataset [4]. Using also MFCCs and GMMs, other authors achieved 94% accuracy in the RAI dataset when processing whole shows and 82% on segments as short as 6 seconds [8].

The different performances across sets indicate that short-term spectral features present solid classification abilities, but are not robust in heterogeneous and complex datasets. MFCCs, as well as Perceptual Linear Prediction (PLP) features [9], represent the short-term characteristics of speech, like the spectral properties of phonemes and speakers, but are not designed to characterise long-term properties of audio. This could explain why, in homogeneous datasets, where shows and speakers might often recur, like episodes from the same TV series or broadcast news programmes, MFCCs performed outstandingly. A solution to this was proposed using Factor Analysis (FA) to extract factors related to the genre, achieving 50% improvement over the use of MFCC features on Internet videos [10].

Other approaches to this task [11, 12] aim to identify specific audiovisual events that can be used as semantic blocks to understand the narrative of the overall show or video. However this is a more complex task, due to the need to identify very subtle events, and its performance still does not match the works previously mentioned in genre identification.

---

This work was supported by the EPSRC Programme Grant EP/I031022/1 (Natural Speech Technology).

The work in this paper aims to provide a novel set of long-term background-tracking features that can perform a more natural description of the type of acoustic background present, also tracking its temporal variations. In order to have robust genre classification abilities, these features should be able to represent different background conditions that can characterise shows, like studio recordings, outdoor noises, applause, laughter, different types of music, etc. On the other hand, to ensure generalisation in the genre classification task, the features should factor out the influence of the speaker and the foreground. The proposal explored in this work arises from the output of an asynchronous factorisation of background and speaker with feature transformations, previously used in an ASR task [13].

This paper is organised as follows: Section 2 will present the audio processing system used to extract the background-tracking features from audio files. Section 3 will describe the experimental setup designed to perform genre identification in a set of broadcast shows from the BBC. Finally, Sections 4 and 5 will present the results and conclusions of this work.

## 2. BACKGROUND-TRACKING FEATURES

In [13], a novel method was presented to perform asynchronous factorisation of background and speaker in ASR tasks. This method relied in using a set of Constrained Maximum Likelihood Linear Regression (CMLLR) transformations [14] characterising different possible background conditions that were switched asynchronously in the training and decoding process. As a byproduct, applying this set of background transformations asynchronously on a given audio segment will yield a sequence of states that will indicate which CMLLR transform was applied in each frame and, hence, which corresponding background was considered to be more likely.

The first step in order to extract the proposed background-tracking features is to use a previously trained Hidden Markov Model (HMM) to align the input audio data to its transcription, or to the output of a previous decoding if the transcription is not available, using a set of asynchronous CMLLR transformations trained to represent different background conditions. The sequence of transformations applied in the best path from the alignment can be written into a vector  $x = \{x(0), x(1), \dots, x(n), \dots, x(N - 1)\}$ , with  $N$  being the length of the input audio signal in frames and each value  $x(n)$  given by the index assigned to each background CMLLR transformation from a fixed set of values  $\{0, 1, \dots, t, \dots, T - 1\}$ , where  $T$  is the total number of background CMLLR transforms. Indicator functions  $c_t(n)$ , as defined in Equation 1, can be used to identify whether the value of  $x(n)$  is  $t$  or not.

$$c_t(n) = \begin{cases} 1 & \text{if } x(n) = t \\ 0 & \text{otherwise} \end{cases} \quad (1)$$




Fig. 1. Background-tracking feature extraction.

The new feature vector proposed in this work is denoted as  $v(m)$  and can be calculated as the moving average of the indicator functions  $c_t(n)$  over a span of  $P$  of the original frames. This new vector has a length of  $M = N/P$  and a dimension of  $T$ , being formed by all the values  $v_t(m)$ , computed as in Equation 2, generating  $v(m) = \{v_0(m), v_1(m), \dots, v_t(m), \dots, v_{T-1}(m)\}$ .

$$v_t(m) = \frac{1}{P} \sum_{p=0}^{P-1} c_t(m * P + p) \quad (2)$$

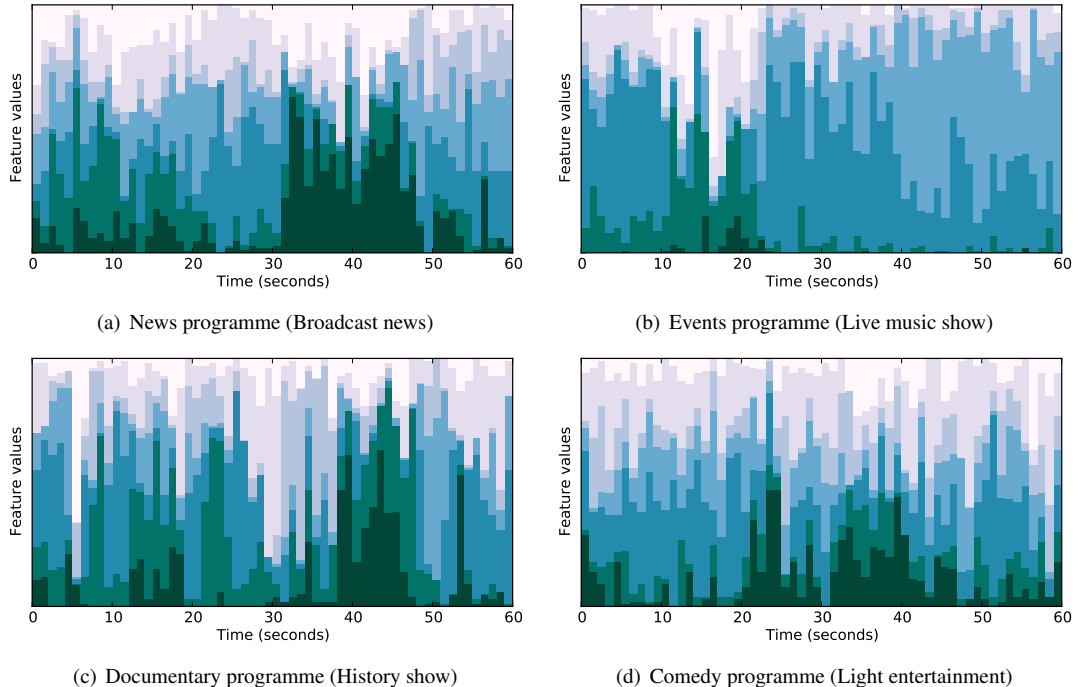
A graphical description of how this process is done can be seen in Figure 1. In this example, there are  $T = 4$  possible background transformations, and values are aggregated every  $P = 12$  frames of the original input vector  $x(n)$  generated as output of the asynchronous alignment.

## 3. EXPERIMENTAL SETUP

The experiments for the evaluation of the proposed background-tracking features were done in a set of 332 shows, totalling 231 hours, broadcast by the BBC during the first week of May in 2008. These programmes were divided into the following 8 genres according to an internal BBC classification:

- Advice: Consumer, DIY and property shows.
- Children’s: Including cartoons and educational shows.
- Comedy: Sit-coms and light entertainment shows.
- Competition: Quiz shows and other contest shows.
- Documentary: Including fly-on-the-wall shows.
- Drama: Soap operas and other serialised dramas.
- Events: Live events, sports and concerts.
- News: Broadcast news and current affair shows.

These genres are very heterogeneous, as the BBC classifies a large number of subgenres. For instance, the “Events”



**Fig. 2.** 1-minute samples of background-tracking features for different shows

genre covers music shows as well as live sports; or the “Documentary” genre covers nature documentaries as well as fly-on-the-wall shows. Since the dataset contains all the BBC broadcasts from a single week, covering all genres, it is a very complete scenario for the evaluation of background characterisation and genre identification techniques.

For the experiments, 285 shows were used for training and 47 shows were used for testing. The number of shows and amount of time covered by each genre is presented in Table 1. The selection of the test set was done with the idea of providing equal coverage of genres and subgenres, with each genre represented by around 3 hours of broadcast time, except for documentaries that have a larger representation due to the multiple subgenres existing. Shows in the test set were also classified depending whether a previous instalment of the same show appeared in the training set, as this indicated whether speakers and environments appearing in the test set also appeared in the training set. 28 of the 47 shows had previous instalments in the training set, with the remaining 19 shows being unique instances of a show in the whole set.

While full transcriptions were not available for these shows, the close captioning subtitles that were broadcast with the shows were available in order to train HMMs for ASR in an Hidden Markov Model Toolkit (HTK) setup [15] using a lightly supervised training process [16]. 7 CMLLR asynchronous transformations were originally trained on a modified version of the WSJCAM0 corpus [17], used in adaptation experiments [18], containing 7 types of acoustic backgrounds: clean speech; classical music; contemporary music;

**Table 1.** Distribution of shows by genre.

| Genre       | Train |         | Test  |        |
|-------------|-------|---------|-------|--------|
|             | Shows | Time    | Shows | Time   |
| Advice      | 34    | 24.5h.  | 4     | 3.0h.  |
| Children’s  | 45    | 18.5h.  | 8     | 3.0h.  |
| Comedy      | 20    | 9.7h.   | 6     | 3.2h.  |
| Competition | 37    | 25.9h.  | 6     | 3.3h.  |
| Documentary | 41    | 29.8h.  | 9     | 6.8h.  |
| Drama       | 19    | 14.4h.  | 4     | 2.7h.  |
| Events      | 23    | 29.8h.  | 5     | 4.3h.  |
| News        | 66    | 50.3h.  | 5     | 2.0h.  |
| Total       | 285   | 203.0h. | 47    | 28.3h. |

applause; cocktail party noise; traffic noise and wildlife noise. These transformations were asynchronously retrained in the BBC dataset to represent the different acoustic disturbances present in this data. The asynchronous alignment required to extract background-tracking features was also performed based on the existing subtitles with  $T = 7$  and  $P = 100$ , so a 7-dimensional feature vector was extracted from each second of the input audio.

An illustration of the output of the background-tracking feature extraction can be seen in the images in Figure 2. These images visualise the 7-dimensional feature vectors extracted, as explained earlier in the Section. These samples represent 4 periods of one minute (60 frames) from 4 different shows. The values of each of the 7 dimensions are represented by


the size of the 7 coloured bars in each frame. Figure 2(a) is one minute in a broadcast news programme, where the background changes from music to street noise to clean studio and ends with street noise. Figure 2(b) is one minute in a music event show, where the music changes from rock music to solo singing and then to instrumental rock music. Figure 2(c) is one minute in a historical documentary show, that starts with bell sounds, followed by a period of music, another period of clean speech and finishes with sounds of seaside and birds. Finally, Figure 2(d) is one minute in a light entertainment show that mixes speech with long bursts of laughter.

#### 4. RESULTS

The first set of experiments were designed to evaluate the performance of the proposed background-tracking features compared to short-term features in the genre identification task. Genre-based GMMs were trained with the feature vectors extracted from all the shows in the training set belonging to each genre. A set of GMMs was trained with 13-dimensional PLP features extracted every 10 ms. and another set with 7-dimensional background-tracking features extracted every second. First and second derivatives were also computed and added to the feature vectors, for a total of 39 dimensions in the PLP features and 21 dimensions in the background-tracking features. The background-tracking features were tested on two conditions, the first one assuming that the subtitles of the shows in the test set were available for the alignment, and the second one using the transcription provided by the ASR system to do the alignment. The classification of the genre for each show in the test set was done by selecting the GMM that maximised the overall likelihood of all the input frames in the test show. The results in terms of accuracy (number of correctly classified shows divided by the total number of test shows) for different number of Gaussians in the GMMs for both types of features are presented in Figure 3.

Background-tracking features outperformed PLPs in this task. While the proposed features achieved up to 72.4% accuracy, PLPs only reached 61.7% accuracy. In further analysis, PLPs required a higher number of Gaussians (up to 1,024 and 2,048) to achieve their best performance, while background-tracking features required less model complexity. This was due to the long-term nature of the background-tracking features, which were extracted every second, instead of every 10 milliseconds. While a total of 73,528,233 frames were available for training the PLP GMMs, only 730,621 were available with the background-tracking features. Figure 3 also shows that there was little difference between using the subtitles or the decoding transcripts to extract the background-tracking features in the test shows, indicating that the feature extraction process was robust to the use of noisy transcriptions in the asynchronous alignment. Following this, all further experiments were based on the alignment to the subtitles.

The final element for analysis is presented in thinner lines




**Fig. 3.** Accuracy in genre identification PLP and background-tracking features with GMM classifiers (Thicker lines represent global accuracy, thinner lines represent accuracy for repeated and non-repeated shows).

around the main lines in Figure 3. These lines mark the accuracies achieved in shows that have previous instalments in the training set and the accuracy achieved in the rest of the shows. PLP features presented a larger spread (represented by the shaded area in the Figure) between these two types of shows, 15% to 20% difference in absolute accuracy across most of the range of GMM sizes, while background-tracking exhibited lower difference, 5% to 10% maximum. For shows with previous episodes in the training set, PLP features achieved 67.8% accuracy, narrowing the gap to the 75.0% obtained with background-tracking features for the same shows. However, for the rest of the shows, PLPs only reached 52.6% accuracy, while background-tracking features reached a more robust 68.4%. This pointed out how short-term features were more sensitive to the presence of known speakers and environments in the training set.

Afterwards, more advanced classifiers were evaluated using background-tracking features. Two experiments were set to study two aspects of classification: Modelling of temporal changes and discriminative methods. The first classifier used were HMMs, which are generative classifiers like GMMs, but, unlike GMMs, they also model temporal transitions among hidden states existing in the input data. For these experiments, HMMs with 8 states were found to provide the best performance and were, subsequently, used. The Gaussian components in each state and the transition probabilities among states were learnt using a Maximum Likelihood (ML) [19] approach from all the input feature vectors from the shows in the training data. The selection of the genre for each test show was also done maximising the likelihood.

The second classifier used at this stage were Support Vector Machines (SVM) [20]. SVMs are widely used discriminative classifiers and had been previously used in the genre identification task [4]. In these experiments, the inputs to the SVM classifier were supervectors obtained by concatenating




**Fig. 4.** Accuracy in genre identification of GMM, HMM and SVM classifiers using background-tracking features.

the Gaussian means of show-based GMMs trained via Maximum A Posteriori (MAP) adaptation [21]. Gaussian-kernel SVMs were trained [22] for each genre to classify whether shows belonged or not to that genre. The final decision for each test show was made for the genre whose SVM gave the best score from all the genre-based SVMs.

The results of the GMM, HMM and SVM classifiers are shown in Figure 4 for different values of model complexity. They showed that both HMMs and SVMs outperformed GMMs. The best result for HMMs, 78.7% accuracy, was achieved with a total model complexity of 256 Gaussians (8 states with 32 Gaussians each); while the best result for SVMs, 80.9% accuracy, was achieved with a lower model complexity, only 16 Gaussians.

To evaluate the identification abilities of the proposed systems, the F-measure of the two best HMM and SVM systems are presented in Figure 5 for each genre. The F-measure, defined as the harmonic mean of precision and recall for each class, allows to evaluate the accuracy and specificity of a classifier. Figure 5 shows that SVMs performed better identifying the “Advice”, “Children’s”, “Events” and “News” genres, while HMMs outperformed SVMs in the “Comedy”, “Competition” and “Drama” genres.

Finally, system combination based on the confidence scores given by the best HMM and SVM systems was performed [23]. System combination has traditionally been proposed as a solid way of exploiting the outputs of different classifiers with different properties; in this task, the modelling of dynamics given by HMMs and the discriminative modelling provided by SVMs. The confidence of the HMM classifier was based on the likelihood score of the decided HMM; while the confidence score of the SVM classifier was based on the distance score provided by the decided SVM, both normalised to the range of  $[0, 1]$ . When both systems provided the same hypothesis, this was accepted straight-away; but when they disagreed, the output of the system with highest confidence was selected. The result of the combination of both systems in terms of global accuracy was 83.0%.



**Fig. 5.** F-measure for HMM and SVM classifiers.

## 5. CONCLUSIONS

The proposed background-tracking features have shown, through a range of different classifiers, that they can provide robust results in the task of genre identification of broadcast shows. While, in absolute terms, the use of acoustic and video features has been reported to provide better performance [2, 3, 4], the results are very promising when compared with previous results using only acoustic features. Furthermore, some types of broadcasts, such as radio or podcasts, do not have video and rely only on the acoustics for classification. Future work will have to see these novel acoustic features merged with state-of-the-art video features to compare with the best performing systems in this task.

The experiments have also shown that the use of long-term features outperforms usual short-term features in tasks that require an acoustic characterisation of the background. Features like PLPs or MFCCs have great classification capabilities in speech but fail to generalise well, as shown by [4] in their comparison of different datasets, because they mostly describe the phonemes or speakers in the audio. Long-term background-based features provide a more comprehensive description of the acoustic conditions of broadcasts, and are less sensitive to the recurring presence or not of the same speakers and environments.

There are many other tasks where the background-tracking features could be exploited. In the future, these features can be used to automatically split complete shows or videos into homogeneous segments with a similar acoustic background. These segments could be clustered by similarity and then used to let users browse and link segments with a similar acoustic background. From the point of view of speech technologies, it is needed to explore how these features can be used in ASR tasks in noisy conditions. Background-tracking features could be used to adapt or compensate to background noises and disturbances, even in the case when the background changes asynchronously, enhancing ASR performance.

## 6. REFERENCES

- [1] Larson, M., Anguera, X., Reuter, T., Jones, G.J.F., Ionescu, B., Schedl, M., Piatrik, T., Hauff, C. and Soleymani, M. (eds.), “Proceedings of the MediaEval 2013 Multimedia Benchmark Workshop”, Barcelona, Spain, 2013.
- [2] Montagnuolo, M., and Messina, A., “TV Genre Classification Using Multimodal Information and Multilayer Perceptron”, in *Proceedings of AI\*IA 2007: Artificial Intelligence and Human-Oriented Computing*, pp. 730–741, Rome, Italy, 2007.
- [3] Montagnuolo, M., and Messina, A., “Parallel Neural Networks for Multimodal Video Genre Classification”, *Multimedia Tools and Applications*, 41, pp. 125–129, 2009.
- [4] Ekenel, H.K., and Semela, T., “Multimodal Genre Classification of TV Programs and YouTube Videos”, *Multimedia Tools and Applications*, 63, pp. 547–567, 2013.
- [5] Mironica, I., Ionescu, B., Knees, P. and Lambert, P., “An In-depth Evaluation of Multimodal Video Genre Categorization”, in *Processings of the 11th International Workshop on Content-based Multimedia Indexing (CBMI)*, pp. 11–16, Veszprem, Hungary, 2013.
- [6] Liu, Z., Huang, J. and Wang, Y., “Classification of TV Programs Based on Audio Information Using Hidden Markov Models”, in *Proceedings of the IEEE Second Workshop on Multimedia Signal Processing*, pp. 27–32, Redondo Beach, CA, 1998.
- [7] Roach, M. and Mason, J., “Classification of Video Genre using Audio”, in *Proceedings of Interspeech 2001*, pp. 2693–2696, Aalborg, Denmark, 2001.
- [8] Kim, S., Georgiu, P., and Narayanan, S., “On-line Genre Classification of TV Programs Using Audio Content”, in *Proceedings of ICASSP 2013*, pp. 798–802, Vancouver, Canada, 2013.
- [9] Hermansky, H., “Perceptual Linear Predictive (PLP) Analysis of Speech,” *Journal of the Acoustic Society of America*, 87(4), pp. 1738–1752, 1990.
- [10] Rouvier, M., Matrouf, D. and Linares, G., “Factor Analysis for Audio-based Video Genre Classification”, in *Proceedings of Interspeech 2009*, pp. 1155–1158, Brighton, UK, 2009.
- [11] Lee, K., and Ellis D.P.W., “Audio-Based Semantic Concept Classification for Consumer Video”, *IEEE Transactions on Audio, Speech and Language Processing*, 18(6), pp. 1406–1416, 2010.
- [12] Castán, D., and Akbacak, M., “Indexing Multimedia Documents with Acoustic Concept Recognition Lattices”, in *Proceedings of Interspeech 2013*, pp. 2643–2647, Lyon France, 2013.
- [13] Saz, O., and Hain, T., “Asynchronous Factorisation of Speaker and Background with Feature Transforms in Speech Recognition”, in *Proceedings of Interspeech 2013*, pp. 1238–1242, Lyon, France, 2013.
- [14] Gales, M. J. F., and Woodland, P. C., “Mean and Variance Adaptation within the MLLR Framework”, *Computer, Speech and Language*, 10(4), pp. 249–264, 1996.
- [15] Young, S. J., Evermann, G., Gales, M. J. F., Hain, T., Kershaw, D., Moore, G., Odell, J. J., Ollason, D. G., Povey, D., Valtchev, V., and Woodland, P. C., “The HTK Book version 3.4”, Cambridge University Engineering Department, Cambridge, UK, 2006.
- [16] Lanchantin, P., Bell, P., Gales, M., Hain, T., Liu, X., Long, Y., Quinnell, J., Renals, S., Saz, O., and Seigel, M., “Automatic Transcription of Multi-genre Media Archives”, in *Proceedings of First Workshop on Speech, Language and Audio in Multimedia*, Marseille, France, 2013.
- [17] Robinson, T., Fransen, J., Pye, D., Foote, J. and Renals, S., “WSJCAM0: A British English Speech Corpus for Large Vocabulary Continuous Speech Recognition”, in *Proceedings of ICASSP 1995*, pp. 81–84, Detroit MI, USA.
- [18] Saz, O., and Hain, T., “Using Contextual Information in Joint Factor Eigenspace MLLR for Speech Recognition in Diverse Scenarios”, in *Proceedings of ICASSP 2014*, pp. 6314–6318, Florence, Italy, 2014.
- [19] Dempster, A.P., Laird N.M. and Rubin, D.B., “Maximum Likelihood from Incomplete Data via the EM Algorithm”, *Journal of the Royal Statistical Society. Series B (Methodological)*, 39(1), pp. 1–38, 1977.
- [20] Cortes, C., and Vapnik, V., “Support-Vector Networks”, *Machine Learning* 20(3), 273–297, 1995.
- [21] Gauvain, J.L., and Lee, C., “Maximum A Posteriori Estimation for Multivariate Gaussian Mixture Observations of Markov Chains”, *IEEE Transactions on Speech and Audio Processing*, 2, pp. 291–298, 1994.
- [22] Joachims T., “Making Large-Scale SVM Learning Practical”, in *Advances in Kernel Methods: Support Vector Learning*, Chapter 11, pp. 169–184, MIT Press: Cambridge, MA, 1999.
- [23] Silva, C., Ribeiro, B., “Inductive Inference for Large Scale Text Classification: Kernel Approaches and Techniques”, Springer: Berlin, 2010.