

INCREMENTAL TRANSFER LEARNING IN TWO-PASS INFORMATION BOTTLENECK BASED SPEAKER DIARIZATION SYSTEM FOR MEETINGS

Nauman Dawalatabad¹, Srikanth Madikeri², C Chandra Sekhar¹, Hema A Murthy¹

¹Indian Institute of Technology Madras, India

²Idiap Research Institute, Martigny, Switzerland

ABSTRACT

The two-pass information bottleneck (TPIB) based speaker diarization system operates independently on different conversational recordings. TPIB system does not consider previously learned speaker discriminative information while diarizing new conversations. Hence, the real time factor (RTF) of TPIB system is high owing to the training time required for the artificial neural network (ANN). This paper attempts to improve the RTF of the TPIB system using an incremental transfer learning approach where the parameters learned by the ANN from other conversations are updated using current conversation rather than learning parameters from scratch. This reduces the RTF significantly. The effectiveness of the proposed approach compared to the baseline IB and the TPIB systems is demonstrated on standard NIST and AMI conversational meeting datasets. With a minor degradation in performance, the proposed system shows a significant improvement of 33.07% and 24.45% in RTF with respect to TPIB system on the NIST RT-04Eval and AMI-1 datasets, respectively.

Index Terms— Speaker diarization, transfer learning, information bottleneck

1. INTRODUCTION

Speaker diarization systems determine “*who spoke when?*” in a conversational recording [1, 2]. Diarization has applications in telephone conversations, meetings, broadcast news, and TV shows. Diarizing conversational meetings is challenging owing to the dynamic nature of speaker switching. The diarization output is often used as a front-end in speech applications like automatic speech recognition [3] and speaker linking in a large corpus [4, 5]. Agglomerative hierarchical clustering is the most used approach for diarization, where short segments of speech are clustered in a bottom-up manner. Speaker diarization systems based on hidden Markov model/Gaussian mixture model (HMM/GMM) [6] and information bottleneck (IB) approach [7] are the most popular.

Real Time Factor (RTF) is defined as the ratio of the time taken by an algorithm to the duration of the input. IB based diarization systems are significantly fast and are suitable for

real time applications [7, 8]. To enable better speaker discrimination, there has been numerous efforts in finding the deep neural network (DNN) based speaker discriminative embeddings [9–15]. Different network architectures and/or loss functions are used to obtain speaker discriminative representations. However, in all cases, the DNN needs to be trained on huge amounts of labeled/unlabelled data. The two-pass IB based speaker diarization (TPIB) [16] system is a recently proposed unsupervised approach that does not use any separate training data to learn speaker discriminative characteristics. In this approach, the speaker discriminating features are automatically learned from the current conversational recording to be diarized. Although the error rates reported by TPIB system are better than the IB system, the RTF is high mainly due to the increased time for training the artificial neural network (ANN).

In this paper, we attempt to improve the RTF of the TPIB system. All the unsupervised methods (including the TPIB system) operate independently on different audio recordings under consideration. The information learned during diarizing one recording can be used to improve the diarization on another recording. The system should remember the previous speaker discriminating knowledge, learn new information and then transfer this information incrementally for diarizing other audio recordings. This “*Remember–Learn–Transfer*” approach is the natural way for humans. We propose to use transfer learning across recordings in an incremental fashion to learn speaker discriminating features continuously.

The paper is organized as follows. Section 2 briefly explains the IB and TPIB systems. In Section 3 we describe the proposed system. Section 4 presents the results of all systems on different datasets. Finally, Section 5 concludes the paper.

2. INFORMATION BOTTLENECK BASED SYSTEMS

This section briefly describes the IB and TPIB systems.

2.1. Baseline Information Bottleneck System

In agglomerative information bottleneck algorithm [7], short segments \mathbf{X} of speech are clustered in a bottom-up manner. A GMM is modeled using \mathbf{X} where the GMM components are denoted by \mathbf{Y} , and represent relevant information. The IB

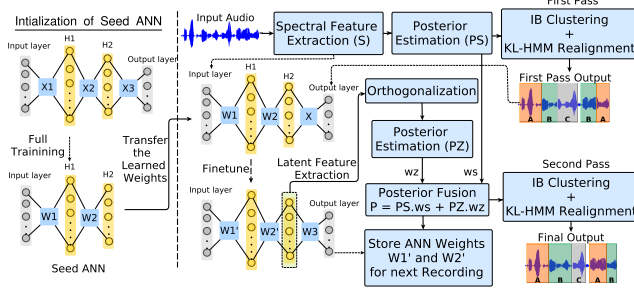


Fig. 1: Block diagram of the proposed system (TPIB-ITL).

based approach clusters the set of segments \mathbf{X} into a set of clusters \mathbf{C} such that most of the relevant information about \mathbf{Y} is captured. The objective function \mathcal{F} is given by

$$\mathcal{F} = I(\mathbf{Y}, \mathbf{C}) - \frac{1}{\beta} I(\mathbf{C}, \mathbf{X}) \quad (1)$$

where $I(\cdot)$ denotes the mutual information and β is a Lagrange multiplier. The normalized mutual information (NMI) [7] is used as a criterion for terminating the clustering process.

2.2. Two-pass IB based Diarization System

Two-pass IB (TPIB) based speaker diarization system [16] consists of 2 passes of the IB diarization as given below.

- **First pass:** In the first pass, standard IB based diarization is performed followed by Kullback-Leibler hidden Markov model (KL-HMM) based realignment [17].

- **ANN Training, Feature Extraction and Orthogonalization:** ANN initialized with random weights is trained on the output boundary labels and the spectral features obtained in the first pass. Latent features are extracted from the penultimate layer of the trained ANN. Principal component analysis (PCA) is then applied to whiten the features.

- **Second pass:** The latent features are used along with the spectral features in the second pass of IB clustering which is then followed by a KL-HMM based realignment.

3. TPIB WITH INCREMENTAL TRANSFER LEARNING

Transfer learning is the improvement of learning in a new task through the transfer of knowledge from a related task that has already been learned [18]. In TPIB system, the ANN learns the parameters from scratch for each conversation independent of the other. It is the primary bottleneck in terms of run time. We attempt to solve this problem by retaining the knowledge learned by the model from other audio recordings. New discriminative information from the current recording is used to learn more speaker discriminative features. As shown in Fig. 1, the following steps are involved in the proposed TPIB system with incremental transfer learning (TPIB-ITL).

- **Step 0 - Training seedANN:** A *seedANN* is trained from the first audio to be diarized by the system. We first perform IB based diarization followed by KL-HMM based realignment to get the relative speaker labels. We initialize all layers' weights ($X1$, $X2$, and $X3$) of *seedANN* using Xavier initialization technique as described by Glorot and Bengio in [19]. The obtained speaker labels and the spectral features are used to train the *seedANN*. This *seedANN* is used for the subsequent recordings to be diarized. It is important to note that no separate data is used for training the *seedANN*.

- **Step 1 - First pass:** In the first pass, we perform IB based diarization that is followed by KL-HMM based realignment.

- **Step 2 - Transfer of knowledge and fine-tuning:** Unlike random initialization of ANN weights in TPIB, we initialize the parameters of the current ANN with the parameters of the *seedANN*. However, the number of speaker labels in the first-pass output of the current recording can be different from the number of output neurons in the *seedANN*. To handle this mismatch, the parameters $W1$ and $W2$ (as shown in Fig. 1) are initialized from the *seedANN*. The final layer weights alone are initialized using Xavier initialization (X). After initialization, fine-tuning of the ANN is performed for a small number of epochs. Fine-tuning is performed using the first pass output labels and the spectral features of the current recording to be diarized. The fine-tuned ANN weights ($W1'$ and $W2'$) are stored as ANN^* .

- **Step 3 - Latent Feature Extraction and Orthogonalization:** Once the ANN is fine-tuned for the current recording, the output of the penultimate layer is used as the latent feature (LF) representation. The output features are subjected to PCA for orthogonalization.

- **Step 4 - Posterior Merging:** The latent space features (LSF) are merged with the spectral features [7] as,

$$P(\mathbf{y}|\mathbf{f}_t^s, \mathbf{f}_t^z) = P(\mathbf{y}|\mathbf{f}_t^s).w_s + P(\mathbf{y}|\mathbf{f}_t^z).w_z \quad (2)$$

where \mathbf{f}_t^s and \mathbf{f}_t^z are feature vectors at time t from spectral feature stream s and latent feature stream z , respectively. Here, w_s and w_z are the weights assigned to the feature streams s and z respectively, such that $w_s + w_z = 1$.

- **Step 5 - Second pass:** Finally, the second pass of IB clustering is performed on $P(\mathbf{y}|\mathbf{f}_t^s, \mathbf{f}_t^z)$ which is followed by KL-HMM based realignment. This is the final diarization output.

For every new recording, the parameters of the ANN is initialized with the ANN^* . Notice that ANN^* gets updated after every fine-tuning step (Step 2) for each recording.

3.1. Remember-Learn-Transfer

The neural network primarily learns to discriminate between speakers in a recording. The characteristics can therefore be similar across the recordings. The discrimination learned by the ANN on one conversational recording can be useful for training an ANN on another recording. This is missing in the TPIB framework where the already learned speaker discriminative information is never utilized for training any fu-

Table 1: AMI meeting datasets.

AMI-1	ES2008c, ES2013a, ES2013c, ES2014d, ES2015a, IS1001c, IS1007a, IS1008c, IS1008d, IS1009c
AMI-2	ES2010b, ES2013b, ES2014c, ES2015b, ES2015c, IS1004b, IS1006c, IS1007c, IS1008a, IS1009d

ture model. Transfer learning enables the new model to start learning from a better point than starting from random initialization, which ultimately leads to faster convergence.

The fine-tuning of ANN is essential for the current recording as we are interested in discriminating the speakers in that recording. Moreover, this also helps to get rid of problems of unseen speakers (or classes). The fine-tuning is performed on every new recording, which can lead to additional speaker discriminative information. Hence, the ANN model gets better at discrimination as it continuously learns from previously diarized conversations.

4. EXPERIMENTAL SETUP AND RESULTS

This section explains the datasets, experimental setup and results obtained.

4.1. Dataset, Features and Evaluation Measure

Experiments are performed on standard NIST-RT (RT-04Dev, RT-04Eval, RT-05Eval) [20] datasets and on the subsets of Augmented Multi-Party Interaction (AMI) corpus [21]. The list of meeting IDs from AMI datasets recorded at Idiap (IS) and Edinburgh (ES) is given in Table 1. These are the same AMI meetings selected randomly in [16] for evaluation. The number of speakers in the NIST dataset ranges from 3 to 8 whereas there are four speakers in each of the AMI recordings. Most recordings have a different set of speakers. RT-04Dev was used as the development dataset and remaining datasets were used for the testing purpose.

Mel frequency cepstral coefficients (MFCC) of 19 dimensions extracted with 10 ms shift from 26 filterbanks are used as input spectral features. Diarization error rate (DER) is the sum of missed speech (MS), false alarm (FA) and speaker error rate (SER). The MS and FA depend on errors in speech activity detection whereas the SER is due to the speaker mismatch. As the focus of the paper is primarily on clustering, speech/non-speech hypotheses were obtained from the ground truth and SER is used as the evaluation metric. As the proposed system does not use any separate training data (labelled/unlabelled) to obtain speaker discriminative features, we compare the performance of the proposed system to the baseline IB and the TPIB systems. We kept a forgiveness collar of 0.25 sec and included overlapped speech in our evaluations.

4.2. Models and Experimental Setup

A multi-layer feed-forward neural network (MLFFNN) with 2 hidden layers was used for both TPIB and TPIB-ITL sys-

Table 2: Speaker Error Rate (SER) on different systems are mentioned. The feature fusing weights are mentioned in parentheses. Avg. denotes the average SER over all fusing weights combination. Best SER on both systems for each dataset is indicated in bold font.

System	Feature(s)	Dev. Set	Test Sets			
		RT-04Dev	RT-04Eval	RT-05Eval	AMI-1	AMI-2
IB	MFCC	15.1	13.5	16.4	17.9	23.5
TPIB	LSF	15.1	11.6	14.2	17.5	21.3
	MFCC+LSF (0.8, 0.2)	13.1	12.5	16.6	16.4	22.7
	MFCC+LSF (Avg.)	14.9	12.6	15.3	17.8	22.4
Proposed System						
TPIB-ITL	LSF	15.5	12.5	15.1	17.5	22
	MFCC+LSF (0.1, 0.9)	15.2	12.2	15	18	21.2
	MFCC+LSF (Avg.)	15.8	12.5	15.4	17.8	21.9
TPIB-ITL (Dev.)	LSF	15.5	12.9	14.8	17.5	22.1
	MFCC+LSF (0.1, 0.9)	15.2	12.5	15	17.5	22
	MFCC+LSF (Avg.)	15.8	13.3	15.6	17.8	22.5

tems. The first layer has 30 neurons with tanh activation function and the second layer has 16 neurons with linear function. We use TensorFlow’s implementation of stochastic gradient descent with cross-entropy loss function to train/fine-tune the ANN models.

All the weights were initialized using Xavier initialization. The early stopping criterion based on cross-entropy error was set to achieve the best results on the development set in terms of both, SER and RTF. This implementation of TPIB system gave lower RTF for TPIB than that reported in [16].

The results reported in Table 2 for TPIB-ITL are when the system is executed in the standard chronological order of the meeting IDs within the dataset. The sequence of the datasets used is also the same as shown (left to right) in Table 2. This standard sequence is maintained just for reproducibility of the results. However, similar trends were observed when the TPIB-ITL system was run on 10 different random permutations of the meeting sequences with different recording for *seedANN*. As we will show in the next section, the order of meetings does not influence the overall performance of TPIB-ITL. As the system uses information from previous recordings, an alternative experiment was also conducted. Here, the incremental transfer learning is performed only on development data. The parameters of the ANN for each test conversation are then fine-tuned from the trained ANN independent of the other conversations. The trained ANN is used for all the test recordings independently. The SERs for this experiment is given in TPIB-ITL (Dev.) (i.e., last row of Table 2). The number of fine-tuning epochs was set to 50 for both TPIB-ITL and TPIB-ITL (Dev.).

We use open source IB toolkit [22] in all the experiments. The values of NMI and β were set to 0.4 and 10, respectively. All RTFs are calculated on 2.6 GHz CPU with 2 threads. The reported RTFs are calculated by averaging the RTFs across 10 independent runs. All the results reported in this paper are based on the best performing parameters tuned on the development dataset.

Table 3: The RTF on different systems for different datasets are mentioned. Impr. denotes the relative improvement in RTF with respect to TPIB system. The observed deviation in RTFs is between 0.002–0.004 on NIST datasets and 0.003–0.007 on AMI datasets.

Sys/Dataset	RT-04Dev	RT-04Eval	RT-05Eval	AMI-1	AMI-2
IB	0.070	0.081	0.086	0.241	0.304
TPIB	0.248	0.257	0.254	0.642	0.740
TPIB-ITL	0.175	0.172	0.180	0.485	0.605
Impr. (%)	29.44	33.07	29.13	24.45	18.24

4.3. Results and Discussion

The SERs for different approaches to diarization are given in Table 2. It can be seen that both the proposed system (TPIB-ITL) and the TPIB system outperform the baseline IB system in most cases. The SERs of TPIB and TPIB-ITL systems are comparable. The TPIB-ITL system showed an absolute improvement of 2.3% with respect to the baseline IB system on AMI-2 dataset. Though the best feature combination tuned on development set gives better SER than the IB system, we observed that it may not always show improvement over LSF. This happens when the extracted LSF itself is better at speaker discrimination than the combination. Hence to check the overall behaviour of the system under different feature fusing weights, we also provide the averaged SER across all the feature fusing weights for each dataset. It gives a measure of robustness of the system to the feature fusing weights. It can be seen from Table 2 that the average SER of the proposed system show absolute 1.6% improvement on AMI-2 dataset. This is significant, as it is calculated over the average of all possible feature fusing weight combinations. It confirms the robustness of the TPIB-ITL system to the feature fusing weight combinations.

The SER on development data for TPIB-ITL is slightly higher than other systems. This is expected as the ANN has learned speaker discrimination from only a few recordings. Its performance is better on the test data, as incremental transfer learning happens. The improvement in SER can be observed on all test datasets. Average SER (avg.) of 21.9% is observed for TPIB-ITL system on AMI-2 (last dataset in sequence) whereas it is 22.4% for TPIB system. This confirms that the ANN model in TPIB-ITL gets better over time.

To ensure that the results are not biased towards ordering of meetings, experiments were conducted for random permutations with random *seedANN*. The overall trend in the result was observed to be similar. For example, the average SER for TPIB-ITL over all the feature fusing weights on each dataset for one of the random permutations of the meeting sequence different from the sequence shown in Table 2 is RT-04Dev=15.5, RT-04Eval=12.6, RT-05Eval=14.9, AMI-1=17.7, AMI-2=22.2. This confirms that the sequence of recordings does not influence the overall performance of the proposed system.

The results where the training of ANN is used incremen-

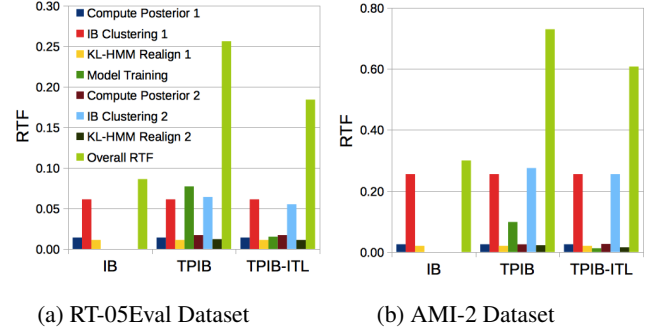


Fig. 2: RTFs of individual modules for all systems.

tally in TPIB-ITL using only development data is given in TPIB-ITL(Dev.) (last row of the Table 2). The system outperforms the baseline IB system and is comparable to TPIB system. Notice that the system does not use any labeled information from the development dataset (RT-04Dev). Moreover, it should also be noted that the RT-04Dev is a very small dataset comprising of only 8 meeting recordings.

A significant improvement in RTF is obtained over the TPIB system. It can be seen from Table 3 that the best case improvement of 33.07% in RTF is observed on RT-04Eval dataset for TPIB-ITL system. The RTFs for all NIST-RT datasets are in a similar range. As the meeting duration in AMI datasets is longer than that in NIST-RT datasets, the RTF of the IB system itself is more for AMI datasets. Nevertheless, we observe improvements in RTF of 24.45% and 18.24% for AMI-1 and AMI-2 datasets, respectively. Though TPIB-ITL system seems inherently sequential, it can be easily run in parallel on different batches of recordings. The IB system showed an average (over all datasets) error of 2.2 in estimating a correct number of speakers, whereas the TPIB and TPIB-ITL showed a similar average error of 1.15 speakers.

Fig. 2 shows module-wise RTF for all the systems on two datasets; NIST RT-05Eval and AMI-2. Other respective NIST and AMI datasets also show a similar trend. As the orthogonalization step takes negligible time, we do not report it. From Fig. 2, it can be seen that considerable time is consumed in the ANN model training stage which is reduced significantly by TPIB-ITL system on both NIST and AMI datasets.

5. CONCLUSION

The proposed TPIB-ITL system uses “Remember-Learn-Transfer” principle to incrementally learn speaker discriminative characteristics. The ANN model gets better at speaker discrimination over time even with a small number of fine-tuning epochs. With a minor degradation in performance, the proposed system show a significant improvement of 33.07% and 24.45% in RTF on RT-04Eval and AMI-1 datasets over TPIB system, respectively. Since the proposed system follows the TPIB framework, it also inherits the advantages of the TPIB framework, i.e., (i) No separate training data is needed, and (ii) It incorporates recording-specific speaker discriminative features during the diarization process.

6. REFERENCES

- [1] Xavier Anguera Miro, Simon Bozonnet, Nicholas Evans, Corrine Fredouille, Gerald Friedland, and Oriols Vinyals, "Speaker Diarization: A Review of Recent Research," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 20, no. 2, pp. 356–370, 2012.
- [2] Sue E. Tranter and Douglas Reynolds, "An Overview of Automatic Speaker Diarization Systems," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 14, no. 5, pp. 1557–1565, 2006.
- [3] Dimitrios Dimitriadis and Petr Fousek, "Developing On-Line Speaker Diarization System," in *Proceedings of INTERSPEECH*, 2017, pp. 2739–2743.
- [4] Douglas E Sturim and William M Campbell, "Speaker Linking and Applications Using Non-Parametric Hashing Methods.," in *Proceedings of INTERSPEECH*, 2016, pp. 2170–2174.
- [5] Marc Ferras, Srikanth Madikeri, and Hervé Bourlard, "Speaker Diarization and Linking of Meeting Data," *IEEE/ACM Transactions on Audio, Speech & Language Processing*, vol. 24, no. 11, pp. 1935–1945, 2016.
- [6] Jitendra Ajmera and Charles Wooters, "A Robust Speaker Clustering Algorithm," in *IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, 2003, pp. 411–416.
- [7] Deepu Vijayasenan, Fabio Valente, and Herve Bourlard, "An Information Theoretic Approach to Speaker Diarization of Meeting Data," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 17, no. 7, pp. 1382–1393, 2009.
- [8] Srikanth Madikeri, David Imseng, and Hervé Bourlard, "Improving Real Time Factor of Information Bottleneck-based Speaker Diarization System," , no. Idiap-RR-18-2015, 2015.
- [9] Sree H. Yella, Andreas Stolcke, and Malcolm Slaney, "Artificial Neural Network Features for Speaker Diarization," in *IEEE Spoken Language Technology Workshop (SLT)*, 2014, pp. 402–406.
- [10] Gaël Le Lan, Delphine Charlet, Anthony Larcher, and Sylvain Meignier, "A Triplet Ranking-based Neural Network for Speaker Diarization and Linking," in *Proceedings of INTERSPEECH*, 2017, pp. 3572–3576.
- [11] Hervé Bredin, "Tristounet: Triplet loss for speaker turn embedding," *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 5430–5434, 2017.
- [12] Daniel Garcia-Romero, David Snyder, Gregory Sell, Daniel Povey, and Alan McCree, "Speaker Diarization using Deep Neural Network Embeddings," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2017, pp. 4930–4934.
- [13] Quan Wang, Carlton Downey, Li Wan, Philip Andrew Mansfield, and Ignacio Lopez Moreno, "Speaker Diarization with LSTM," *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 5239–5243, 2018.
- [14] David Snyder, Daniel Garcia-Romero, Gregory Sell, Daniel Povey, and Sanjeev Khudanpur, "X-Vectors: Robust DNN Embeddings for Speaker Recognition," *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 5329–5333, 2018.
- [15] Arindam Jati and Panayiotis Georgiou, "Neural Predictive Coding using Convolutional Neural Networks towards Unsupervised Learning of Speaker Characteristics," *ArXiv e-prints*, Feb. 2018.
- [16] Nauman Dawalatabad, Srikanth Madikeri, C. Chandra Sekhar, and Hema A. Murthy, "Two-Pass IB Based Speaker Diarization System Using Meeting-Specific ANN Based Features," in *Proceedings of INTERSPEECH*, 2016, pp. 2199–2203.
- [17] Deepu Vijayasenan, Fabio Valente, and Hervé Bourlard, "KL Realignment for Speaker Diarization with Multiple Feature Streams," in *Proceedings of INTERSPEECH*, Sept 2009, pp. 1059–1062.
- [18] Emilio Soria Olivas et al., *Handbook Of Research On Machine Learning Applications and Trends: Algorithms, Methods and Techniques - 2 Volumes*, Information Science Reference - Imprint of: IGI Publishing, Hershey, PA, 2009.
- [19] Xavier Glorot and Yoshua Bengio, "Understanding the difficulty of training deep feedforward neural networks," in *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, 2010, pp. 249–256.
- [20] "The NIST Rich Transcription," <https://www.nist.gov/itl>.
- [21] Jean Carletta et al., "The AMI Meeting Corpus: A Pre-announcement," in *Proceedings of the Second International Conference on Machine Learning for Multimodal Interaction*, 2006, MLMI'05, pp. 28–39.
- [22] Deepu Vijayasenan and Fabio Valente, "DiarTk: An Open Source Toolkit for Research in Multistream Speaker Diarization and its Application to Meetings Recordings," in *Proceedings of INTERSPEECH*, 2012.