

Contaminated speech training methods for robust DNN-HMM distant speech recognition

Mirco Ravanelli, Maurizio Omologo

Fondazione Bruno Kessler, Trento, Italy

mravanelli@fbk.eu, omologo@fbk.eu

Abstract

Despite the significant progress made in the last years, state-of-the-art speech recognition technologies provide a satisfactory performance only in the close-talking condition. Robustness of distant speech recognition in adverse acoustic conditions, on the other hand, remains a crucial open issue for future applications of human-machine interaction. To this end, several advances in speech enhancement, acoustic scene analysis as well as acoustic modeling, have recently contributed to improve the state-of-the-art in the field. One of the most effective approaches to derive a robust acoustic modeling is based on using contaminated speech, which proved helpful in reducing the acoustic mismatch between training and testing conditions.

In this paper, we revise this classical approach in the context of modern DNN-HMM systems, and propose the adoption of three methods, namely, asymmetric context windowing, close-talk based supervision, and close-talk based pre-training. The experimental results, obtained using both real and simulated data, show a significant advantage in using these three methods, overall providing a 15% error rate reduction compared to the baseline systems. The same trend in performance is confirmed either using a high-quality training set of small size, and a large one.

Index Terms: distant speech recognition, multi-condition training, DNN.

1. Introduction

During the last decade, much research has been devoted to improving Automatic Speech Recognition (ASR) performance. Nevertheless, most state-of-the-art systems are still based on close-talking solutions, forcing the user to speak very close to a microphone-equipped device. There are, however, various real-life situations where Distant Speech Recognition (DSR) is more natural, convenient and attractive [1]. An emerging application is, for instance, speech-based domestic control, where users might prefer to freely interact with their home appliances without wearing or even handling any microphone-equipped device. This scenario was addressed under the EU DIRHA project¹, which had the ultimate goal of developing real-time systems for motor-impaired users. Unfortunately, current DSR technologies still exhibit a significant lack of robustness and flexibility, due to the adverse acoustic conditions originated by non-stationary noises and acoustic reverberation [2].

Several efforts have been devoted during the last years to improve DSR, as also witnessed by some recent international challenges such as REVERB [3], CHIME [4, 5] and ASPIRE.

¹This work was partially funded by EU's 7th Framework Programme under grant agreement n. 288121 DIRHA. More details can be found in <http://dirha.fbk.eu/>.

Considerable progresses were fostered by substantial advances in spatial filtering [6, 7], microphone selection [8], source separation [9], speech dereverberation [10] as well as speaker localization [11], acoustic event detection [12] and acoustic modeling. Concerning the latter field, the research community is currently experiencing a revolution due to the introduction of Deep Neural Networks (DNNs) [13], which have consistently outperformed previous Gaussian Mixture Models (GMMs) for both close [14] and distant-talking scenarios [15, 16, 17, 18, 19]. One of the most effective and straightforward approach so far applied to DSR for acoustic modeling is based on multi-condition training with contaminated speech corpora [20, 21, 22]. Contaminated speech is generated by convolving close-talking signals with an Impulse Response (IR) usually measured in the targeted environment. Background noise is then typically added to the convoluted signals, in order to make the resulting corpus more realistic and better matching real-world conditions. This approach has also been adopted recently in [3, 4, 23, 24].

This paper deals with the use of contaminated speech to train a DNN-HMM speech recognition system. Training such very complex systems is typically difficult and tricky. The basic strategy to run it, and the way one performs optimization steps, often play a crucial role in order to get best performance. Our work, which relies on the Kaldi toolkit [25], has concerned all those optimization efforts, e.g., including an optimal choice of number of hidden layers and of units. However, in this paper we focus on two specific aspects, about which we believe the experimental results demonstrate novel and interesting evidences.

First, we propose some methods able to exploit information from both distant and close-talking datasets to derive robust acoustic models. In particular, one proposed approach consists in inheriting the high-quality labels that might be generated from close-talk datasets to train distant-talking DNNs. Another novel method consists in replacing a standard Restricted Boltzmann Machines (RBMs) based pre-training [26] with a supervised close-talk pre-training, leading to a smarter initialization of the distant-talking DNN.

As a second contribution, we propose to replace the traditional symmetric context windows, which are typically adopted to gather several consecutive features in the DNN framework, with an Asymmetric Context Window (ACW). Interestingly, this approach has proved effective in counteracting the harmful effects of the acoustic reverberation, due to a better use of the contextual information.

The experimental validation is carried out in a living-room of a real apartment using both simulated and real test data. Similarly to [23], the evaluation is performed on a phone-loop task, in order to avoid any non-linear influence and artifacts possibly originated by a language model.

The rest of the paper is organized as follows. In Sec. 2 the

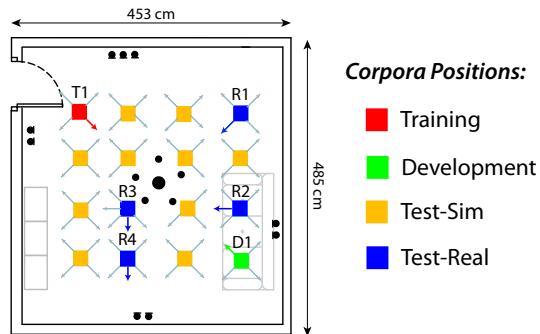


Figure 1: The reference DIRHA living-room environment. Squares and arrows represent the speaker positions and orientations adopted for the real and simulated corpora.

contamination approach is outlined, while Sec. 3 and 4 describe the adopted corpora and the ASR system, respectively. Finally, the proposed methods are experimentally validated in Sec. 5, while Sec. 6 draws the conclusions.

2. Speech contamination

The speech contamination process is represented by the following basic relationship:

$$y(t) = x(t) * h(t) + \alpha n(t) \quad (1)$$

where $x(t)$ and $y(t)$ are the close and distant-talking signals, respectively, while $n(t)$ consists in recorded environmental noise controlled by the gain α in order to obtain different SNRs. $h(t)$ is an IR that describes the acoustic reverberation effects of a generic source-microphone path. In particular, if one assumes to deal with a linear time-invariant acoustical transmission system, the IR provides a complete description of the changes a sound signal undergoes when it travels from a particular position in space to a given microphone [27]. The method adopted here for estimating IR consists in diffusing in the target environment a known Exponential Sine Sweep (ESS) signal, and recording it by a microphone properly placed in space [28]. This method of measuring the IR has a remarkable impact on the speech recognition performance, as discussed in [29].

3. Simulated and Real Corpora

The contamination process described in Sec. 2, can be exploited to generate simulated data starting from different close-talking datasets. In this work, simulated data are used for training and development purposes only, while test is performed also on real data. The sampling rate is 16 kHz for all the datasets. In the following sections, the reference environment, the close-talking corpora, and the simulated and real datasets are described.

3.1. Acoustic Environment and Microphone Set-up

The reference environment considered for this study is the living-room of a real apartment, which was available for experiments under the DIRHA project. The original microphone network includes several microphones installed on the walls and on the ceiling. However, due to the purpose of this work, in the following we will consider only one microphone, i.e., the central one of the ceiling array. As shown in Fig. 1, an IR measurement session exploring a large number of positions and

orientations of the sound source has been conducted. The mean reverberation time T_{60} measured in the room is 0.7 s.

3.2. Close-talking Corpora

In this work, two close-talk datasets are used for training, namely APASCI and Euronews. APASCI [30] consists of more than 5200 phonetically-rich sentences uttered by 163 speakers and recorded in a professional studio, overall corresponding to about 6 hours. For comparison purposes, the Italian part of the Euronews database [31], consisting of about 100 hours of TV news, is also considered for training. The Euronews signals are frequently corrupted by cross-talk, background music, environmental noise as well as spontaneous speech phenomena.

A small development set is also adopted to tune the free-parameters of the DSR systems, as discussed in Sec. 5.1. To this purpose, a portion of the DIRHA corpus [32], consisting of 220 high-quality phonetically-rich sentences uttered by 10 speakers is employed.

For testing purposes, a different portion of the DIRHA corpus, composed of 430 phonetically-rich sentences uttered by 20 speakers is adopted.

3.3. Simulated Corpora

For each of the close-talking corpora outlined in Sec. 3.2, two contaminated versions were generated. The first version is based on a convolution of the close-talking sentences with impulse responses measured in the reference living-room environment. Such simulated datasets are intended to study a scenario where only reverberation acts as a source of disturbance. Furthermore, a second version of the contaminated databases was directly derived from the reverberated datasets by adding noisy background sequences recorded in the target environment, ensuring a mean SNR of about 10 dB. This version is intended to study more challenging scenarios where both reverberation and noises act as a source of disturbance. For the training corpora a single impulse response (from T1 in Fig.1) was considered, while for development a different IR (D1) was adopted. As shown in [23], note that a single impulse response is sufficient to model room reverberation effects for DSR purposes.

3.4. DIRHA Real Corpus

Besides the simulated test-sets, a corpus of real phonetically-rich sentences was recorded in the reference living-room. These real recordings involved 18 native Italian speakers (9 males and 9 females), who read a list of 40 phonetically-rich sentences in four different positions (R1-4 in Fig.1). In total, 2880 sentences, corresponding to more than two hours of speech material were acquired. These recordings were conducted in quiet acoustic conditions, where reverberation was the main source of disturbance. The estimated SNR is about 22 dB.

4. DNN based ASR

In this work, we use a Context-Dependent DNN-HMM speech recognizer, where every unit is modeled by a three state left-to-right HMM, and the tied-state observation probabilities are estimated through a DNN.

Feature extraction is based on blocking the signal into frames of 25 ms with 10 ms overlapping. For each frame, 13 MFCCs plus pitch and Probability of Voicing (PoV) are extracted. The pitch and PoV are estimated through the normalized autocorrelation method discussed in [33]. The resulting




Figure 2: PER(%) obtained by DNN context windows which progressively integrate past or future frames.

features, together with their first and second order derivatives, are then arranged into a single observation vector of 45 components. Finally, a context window gathering several consecutive frames followed by a mean and variance normalization of the feature space are applied before feeding the DNN.

The DNN, trained with the Kaldi toolkit (Karel’s recipe) [25], is composed of sigmoid-based hidden neurons, while the output layer is based on softmax activation functions. The pre-training phase is carried out by stacking Restricted Boltzmann Machines (RBM) [26] to form a deep belief network, while the fine-tuning is performed by a stochastic gradient descent optimizing cross-entropy loss function. In the latter phase, the initial learning rate is kept fixed as long as the increment of the frame accuracy on the dev-set is higher than 0.5%. For the following epochs, the learning rate is halved until the increment of frame accuracy is less than the stopping threshold of 0.1%.

The decoding is performed by adopting a phone-loop based grammar. As previously outlined, even though the use of more complex grammars or language models is certainly helpful in increasing the recognition performance, the adoption of a simple phone-loop is due to the need of an experimental evidence not biased by a LM. In this study, as in [23], a set of 26 phone units of the Italian language was chosen for evaluation purposes.

5. Experiments

In the following section, the proposed techniques are better described and an experimental validation is provided². The experiments are conducted in three different acoustic scenarios of increasing complexity: close-talking, distant-talking with reverberation (Rev), and distant-talking with both noise and reverberation (Rev&Noise). The reported performance is obtained when a matching acoustic condition between test and training is met (e.g., for the reverberated speech test, the reverberated training is used). The evaluation is performed on both simulated and real data, using either APASCI or Euronews training-sets.

5.1. Baselines

This section reports the results obtained with a standard DNN-HMM system, after optimizing the main ASR free-parameters on the dev-set. The best performance is reached with a DNN composed of 6 hidden layers with 1500 neurons per layer, adopting a symmetric context window of 17 frames (8 before and 8 after the current frame) and using a learning rate of 0.008. Note that such free-parameters are used for both the training corpora and for all the addressed acoustic scenarios, due to a

²Part of the experiments are conducted with a Tesla K40 donated by the NVIDIA Corporation.

similar outcome of the optimization step. Only a difference concerning the optimal number of tied-states is observed, leading to 2.5k and 5k tied-states for APASCI and Euronews, respectively. Table 1 shows the performance achieved on the test-sets.

| <i>Test</i> \ <i>Train</i> | APASCI (6 h) | Euronews (100 h) |
|----------------------------|--------------|------------------|
| Close-talk | 14.0 | 14.5 |
| Sim-Rev | 37.6 | 37.0 |
| Real-Rev | 41.0 | 40.1 |
| Sim-Rev&Noise | 52.3 | 51.8 |

Table 1: Phone Error Rate (PER%) of the baseline system.

Table 1 clearly highlights that the performance in distant-talking scenarios is dramatically reduced, if compared to a close-talking case. This loss of performance further confirms how challenging DSR is under adverse acoustic conditions and without any help from grammars or LMs. It is also highly relevant that with only 6 hours of high-quality speech material a performance similar to that of a 100 hours corpus is obtained.

5.2. Asymmetric Context Window (ACW)

In distant-talking scenarios, the acoustic reverberation introduces a long-term redundancy in the signal. This disturbance can be modeled as a long causal FIR filter, whose taps introduce a forward memory in the speech sequence. As a consequence, since the length of the window adopted for features extraction (i.e., 25 ms) is much smaller than the reverberation time (i.e., 700 ms), a progressive concatenation of future frames tends to embed information correlated with the central frame contents. On the other hand, a concatenation of past frames seems to be more helpful since, on average, different and complementary information can potentially be analyzed to perform a frame-level prediction. Based on this concept, we explore the use of an Asymmetric Context Window (ACW), which analyzes more past frames than future frames.

Fig. 2 reports a first experiment involving context windows which progressively integrate only past (negative x-axis) or future frames (positive x-axis). In this section, only results obtained using APASCI as training-set are reported. However, a similar trend is observed with Euronews.

From Fig. 2, it is clear that in a close-talking scenario a symmetric behaviour is attained, meaning that past and future information provides a similar contribution in improving the system performance. Conversely, the role of the past information seems significantly more important in distant-talking scenarios, since a faster decrease of the PER(%) is progres-

sively obtained when past frames are concatenated. Along this line, another experiment, reported in Table 2, shows the results achieved when the symmetric context window of 17 frames introduced in Sec. 5.1 is replaced by different asymmetric context windows.

| <i>Test</i> | <i>Train</i> | Context Window | | | | |
|---------------|--------------|----------------|-------------|-------------|--------|--------|
| | | P16-F0 | P10-F6 | P8-F8 | P6-F10 | P0-F16 |
| Close-talk | | 15.5 | 14.2 | 14.0 | 14.4 | 16.3 |
| Sim-Rev | | 38.0 | 37.0 | 37.6 | 38.3 | 43.6 |
| Real-Rev | | 41.6 | 40.2 | 41.0 | 41.6 | 46.8 |
| Sim-Rev&Noise | | 53.8 | 51.8 | 52.3 | 53.0 | 57.0 |

Table 2: PER(%) obtained with Symmetric and Asymmetric context windows composed of 17 frames. ‘‘P’’ refers to the number of past frames, while ‘‘F’’ refers to future frames.

Consistently with the previous experiment, in a close-talking scenario the typical symmetric window performs better than any asymmetric window. Differently, an asymmetric context window which embeds more past information (10 frames) than future information (6 frames) performs slightly better in distant-talking conditions. Note that this is consistent for both real and simulated test-sets. Moreover, results obtained with different context window lengths (e.g., 9, 11, 13, 21, 27 frames) further confirms that ACW is a better choice in distant-talking scenarios, thanks to an improved robustness against reverberation. Note also that the use of an ACW does not imply any additional computational effort and is particularly suitable for on-line systems [34, 35], since the latency originated by waiting for the future frames is minimized.

5.3. Close-talking labels for Distant-talking DNN

In standard ASR, the labels for DNN training are derived by a forced-alignment of the training corpus over the tied-states. Although some GMM-free solutions have been proposed [36], this alignment is typically performed using a standard CD-GMM-HMM system. However, this phase can be very critical because a precise alignment could be difficult to reach, especially in challenging acoustic scenarios characterized by noise and reverberation. As a consequence, the DNN learning process might be more critical due to a poor supervision.

In the contaminated speech training framework, however, a more precise supervision can be obtained from the close-talking dataset. For this reason, we propose to inherit the high-quality labels that can be generated from the close-talking data to train the distant-talking DNNs. This requires to train, with the original clean datasets, a standard CD-GMM-HMM system, and exploit it to generate a precise tied-state forced alignment over the close-talking training corpus, later inherited as supervision for the distant-talking DNN. Table 3 reports the results obtained with the standard approach, which is based on labels derived from distant-talking signals, and the proposed solution, which directly inherits labels derived from close-talking signals. The ACW of 17 frames (P10-F6) discussed in Sec. 5.2 is still used for the following experiments.

Results show that the proposed approach provides a substantial improvement in the performance, over both real and simulated data. Specifically, a relative improvement of 10% and 12% is achieved for APASCI and Euronews, respectively. Besides this significant performance improvement, another interesting aspect is the faster convergence of the iterative learning procedure, due to a better supervision provided to the DNN. In

| <i>Test</i> | <i>Train</i> | APASCI (6 h) | | Euronews (100 h) | |
|---------------|--------------|--------------|--------|------------------|--------|
| | | Standard | CT-lab | Standard | CT-lab |
| Sim-Rev | | 37.0 | 33.0 | 36.1 | 32.1 |
| Real-Rev | | 40.2 | 35.9 | 39.3 | 34.3 |
| Sim-Rev&Noise | | 51.8 | 47.3 | 50.1 | 46.4 |

Table 3: PER(%) obtained in distant-talking scenarios with the standard and with the proposed technique based on close-talking labels (CT-lab).

these experiments, 15 epochs are needed to converge with the standard solution, while only 12 are sufficient with the proposed approach, reducing the training time of 20%.

5.4. Close-talking Pre-Training for Distant-talking DNN

In this section, we propose a further way of taking advantage of the rich information embedded in the close-talking dataset. In particular, instead of adopting a standard unsupervised RBM pre-training of the DNN, we propose to use a supervised pre-training method based on the close-talking data. More precisely, the idea is to train a close-talking DNN and inherit its parameters for initializing the distant-talking DNN. A subsequent fine tuning phase is then carried out on distant-talking data using a slightly reduced learning rate (LR=0.005 is used here). Table 4 reports the results obtained using such approach on real and simulated datasets. The ACW (P10-F6) discussed in Sec. 5.2 and the close-talking labels proposed in Sec. 5.3 are still used for this experiment.

| <i>Test</i> | <i>Train</i> | APASCI (6 h) | | Euronews (100 h) | |
|---------------|--------------|--------------|-------|------------------|-------|
| | | Stand-PT | CT-PT | Stand-PT | CT-PT |
| Sim-Rev | | 33.0 | 31.4 | 32.1 | 31.2 |
| Real-Rev | | 35.9 | 34.7 | 34.3 | 33.0 |
| Sim-Rev&Noise | | 47.3 | 46.2 | 46.4 | 45.1 |

Table 4: PER(%) obtained in distant-talking scenarios with a standard RBM pre-training (Stand-PT) and with the proposed supervised close-talking pre-training (CT-PT).

Results clearly show an improvement deriving from this approach, which is consistent over both simulated and real data. This suggests that a close-talk pre-training is a smart way of initializing the DNN, which somehow first learns the speech characteristics and only at a later stage learns how to counteract the adverse acoustic conditions.

6. Conclusions and Future Work

In this paper, some novel methods for taking advantage of contaminated speech training in modern DNN-HMM systems are proposed. In particular, we explored approaches based on an asymmetric context window, a close-talk supervision and a supervised close-talk pre-training. Experimental results are very promising with a relative performance improvement of more than 15% over the baseline system on both real and simulated data.

Along this line, our next activities will concern the exploration of different architectures, such as convolutional and recurrent NNs. We also plan to consider different acoustic environments, languages and tasks and, at a later stage, to address a multi-microphone case.

7. References

- [1] M. Wölfel and J. McDonough, *Distant Speech Recognition*. Wiley, 2009.
- [2] E. Hänsler and G. Schmidt, *Speech and Audio Processing in Adverse Environments*. Springer, 2008.
- [3] K. Kinoshita, M. Delcroix, T. Yoshioka, T. Nakatani, E. Habets, R. Haeb-Umbach, V. Leutnant, A. Sehr, W. Kellermann, R. Maas, S. Gannot, and B. Raj, “The reverb challenge: A common evaluation framework for dereverberation and recognition of reverberant speech,” in *WASPAA 2013*, 2013, pp. 1–4.
- [4] J. Barker, E. Vincent, N. Ma, H. Christensen, and P. Green, “The PASCAL CHiME speech separation and recognition challenge,” *Computer Speech and Language*, vol. 27, no. 3, pp. 621–633, 2013.
- [5] J. Barker, R. Marxer, E. Vincent, and S. Watanabe, “The third CHiME Speech Separation and Recognition Challenge: Dataset, task and baselines,” in *Submitted to ASRU 2015*, 2015.
- [6] M. Brandstein and D. Ward, *Microphone arrays*. Springer, Berlin, 2000.
- [7] W. Kellermann, *Beamforming for Speech and Audio Signals*. in *HandBook of Signal Processing in Acoustics*, Springer, 2008.
- [8] M. Wolf and C. Nadeu, “Channel selection measures for multi-microphone speech recognition,” *Speech Communication*, vol. 57, pp. 170–180, Feb. 2014.
- [9] S. Makino, T. Lee, and H. Sawada, *Blind Speech Separation*. Springer, 2010.
- [10] P. A. Naylor and N. D. Gaubitch, *Speech Dereverberation*. Springer, 2010.
- [11] A. Brutti, M. Ravanelli, P. Svaizer, and M. Omologo, “A speech event detection/localization task for multiroom environments,” in *Proc. of HSCMA*, 2014, pp. 157–161.
- [12] A. Temko, C. Nadeu, D. Macho, R. Malkin, C. Zieger, and M. Omologo, “Acoustic event detection and classification,” in *Computers in the Human Interaction Loop*. Springer London, 2009, pp. 61–73.
- [13] D. Yu and L. Deng, *Automatic Speech Recognition - A Deep Learning Approach*. Springer, 2015.
- [14] G. Dahl, D. Yu, L. Deng, and A. Acero, “Context-dependent pre-trained deep neural networks for large vocabulary speech recognition,” *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 20, no. 1, pp. 30–42, 2012.
- [15] P. Swietojanski, A. Ghoshal, and S. Renals, “Hybrid acoustic models for distant and multichannel large vocabulary speech recognition,” in *Proc. of ASRU 2013*, 2013, pp. 285–290.
- [16] Y. Liu, P. Zhang, and T. Hain, “Using neural network front-ends on far field multiple microphones based speech recognition,” in *Proc. of ICASSP 2014*, 2014, pp. 5542–5546.
- [17] F. Weninger, S. Watanabe, J. Le Roux, J. Hershey, Y. Tachioka, J. Geiger, B. Schuller, and G. Rigoll, “The MERL/MELCO/TUM System for the REVERB Challenge Using Deep Recurrent Neural Network Feature Enhancement,” in *IEEE REVERB Workshop*, 2014.
- [18] S. S. Masato Mimura and T. Kawahara, “Reverberant speech recognition combining deep neural networks and deep autoencoders,” in *IEEE REVERB Workshop*, 2014.
- [19] A. Schwarz, C. Huemmer, R. Maas, and W. Kellermann, “Spatial diffuseness features for dnn-based speech recognition in noisy and reverberant environments,” in *Accepted at ICASSP 2015*, 2015.
- [20] M. Matassoni, M. Omologo, D. Giuliani, and P. Svaizer, “Hidden Markov model training with contaminated speech material for distant-talking speech recognition,” *Computer Speech & Language*, vol. 16, no. 2, pp. 205–223, 2002.
- [21] L. Couvreur, C. Couvreur, and C. Ris, “A corpus-based approach for robust ASR in reverberant environments,” in *Proc. of INTERSPEECH*, 2000, pp. 397–400.
- [22] T. Haderlein, E. Nth, W. Herbordt, W. Kellermann, and H. Niemann, “Using Artificially Reverberated Training Data in Distant-Talking ASR,” ser. *Lecture Notes in Computer Science*, vol. 3658. Springer, 2005, pp. 226–233.
- [23] M. Ravanelli and M. Omologo, “On the selection of the impulse responses for distant-speech recognition based on contaminated speech training,” in *Proc. of INTERSPEECH 2014*, 2014, pp. 1028–1032.
- [24] A. Brutti and M. Matassoni, “On the use of Early-to-Late Reverberation Ratio for ASR in reverberant environments,” in *Proc. of ICASSP 2014*, 2014, pp. 4670–4675.
- [25] D. Povey, A. Ghoshal, G. Boulianne, L. Burget, O. Glembek, N. Goel, M. Hannemann, P. Motlicek, Y. Qian, P. Schwarz, J. Silovsky, G. Stemmer, and K. Vesely, “The Kaldi Speech Recognition Toolkit,” in *ASRU 2011*, 2011.
- [26] O. S. T. Y. Hinton, G.E., “A fast learning algorithm for deep belief nets,” *Neural Computation*, vol. 18, no. 1, pp. 1527–1554., 2006.
- [27] H. Kuttruff, “Room acoustic-fifth edition,” in *Spon Press*, 2009.
- [28] A. Farina, “Simultaneous measurement of impulse response and distortion with a swept-sine technique,” in *Proc. of the 108th AES Convention*, 2000, pp. 18–22.
- [29] M. Ravanelli, A. Sosi, P. Svaizer, and M. Omologo, “Impulse response estimation for robust speech recognition in a reverberant environment,” in *Proc. of EUSIPCO*, 2012, pp. 1668–1672.
- [30] B. Angelini, F. Brugnara, D. Falavigna, D. Giuliani, R. Gretter, and M. Omologo, “Speaker independent continuous speech recognition using an acoustic-phonetic Italian corpus,” in *Proc. of IC-SLP*, 1994, pp. 1391–1394.
- [31] R. Gretter, “Euronews: a multilingual speech corpus for ASR,” in *Proc. of LREC 2014*, 2014.
- [32] L. Cristoforetti, M. Ravanelli, M. Omologo, A. Sosi, A. Abad, M. Hagmueller, and P. Maragos, “The DIRHA simulated corpus,” in *Proc. of LREC 2014*, 2014, pp. 2629–2634.
- [33] P. Ghahremani, B. BabaAli, D. Povey, K. Riedhammer, J. Trmal, and S. Khudanpur, “A pitch extraction algorithm tuned for automatic speech recognition,” in *Proc. of ICASSP 2014*, 2014, pp. 2494–2498.
- [34] X. Lei, A. Senior, A. Gruenstein, and J. Sorensen, “Accurate and compact large vocabulary speech recognition on mobile devices,” in *Proc. of INTERSPEECH 2013*, 2013, pp. 662–665.
- [35] G. Chen, C. Parada, and G. Heigold, “Small-footprint keyword spotting using deep neural networks,” in *ICASSP 2014*, 2014, pp. 4087–4091.
- [36] A. W. Senior, G. Heigold, M. Bacchiani, and H. Liao, “GMM-free DNN acoustic model training,” in *ICASSP 2014*, 2014, pp. 5602–5606.