A NETWORK OF DEEP NEURAL NETWORKS FOR DISTANT SPEECH RECOGNITION

Mirco Ravanelli^{§,*}, Philemon Brakel[†], Maurizio Omologo[§], Yoshua Bengio[†]

[§]Fondazione Bruno Kessler, Trento, Italy [†]Université de Montréal, Montréal, Canada

ABSTRACT

Despite the remarkable progress recently made in distant speech recognition, state-of-the-art technology still suffers from a lack of robustness, especially when adverse acoustic conditions characterized by non-stationary noises and reverberation are met.

A prominent limitation of current systems lies in the lack of matching and communication between the various technologies involved in the distant speech recognition process. The speech enhancement and speech recognition modules are, for instance, often trained independently. Moreover, the speech enhancement normally helps the speech recognizer, but the output of the latter is not commonly used, in turn, to improve the speech enhancement.

To address both concerns, we propose a novel architecture based on a network of deep neural networks, where all the components are jointly trained and better cooperate with each other thanks to a full communication scheme between them. Experiments, conducted using different datasets, tasks and acoustic conditions, revealed that the proposed framework can overtake other competitive solutions, including recent joint training approaches.

Index Terms— speech recognition, speech enhancement, joint training, deep neural networks

1. INTRODUCTION

Distant Speech Recognition (DSR) is a technology of fundamental importance towards more flexible and effective human-machine interfaces. Such a technology might eventually allow users to access speech recognition services even when adverse acoustic conditions are met or when a distant-talking (far-field) interaction with a machine is required. A crucial role in improving current solutions is being played by deep learning [1], which has recently contributed to outperforming previous HMM-GMM speech recognizers [2]. The progress in the field was also fostered by the considerable success of some international challenges such as CHiME [3] and REVERB [4].

Despite the great efforts of the past years, state-of-the-art techniques still exhibit a significant lack of robustness to acoustic conditions characterized by non-stationary noises and acoustic reverberation [5]. To counteract such adversities, most DSR systems must rely on a combination of several interconnected technologies [6], including methods for speech enhancement [7], speech separation [8], acoustic event classification [9, 10], speaker identification [11], just to name a few.

A significant limitation of most current techniques lies in the lack of matching between the various modules being combined. For example, speech enhancement and speech recognition are often designed independently and, in several cases, the enhancement part is tuned according to metrics which are not directly correlated with the final speech recognition performance. Another potential limitation is the lack of communication between the various modules of the DSR system. Most systems are based on a unidirectional information flow across the DSR pipeline. Speech enhancement tries to help the speech recognizer, but the latter does not contribute, in turn, to improve the performance of the speech enhancement module. We argue that establishing this missing link can nevertheless be very useful, since a hint on the recognized phone sequence might help the speech enhancement in performing its task. A fruitful integration between the various systems, however, was very difficult for many years, mainly due to the different nature of the technologies involved at the various steps. Nevertheless, the recent success of deep learning has not only largely contributed to the substantial improvement of the speech recognition part of a DSR system [12-14], but has also enabled the development of competitive DNN-based speech enhancement solutions [15, 16], making an effective integration between such modules easier.

Within the DNN framework, the adoption of a joint training approach between speech enhancement and speech recognition DNNs has recently been proposed to mitigate the lack-of-matching issue [17–24]. The core idea is to pipeline such DNNs and to jointly update their parameters as if they were within a single bigger network. Despite the considerable effectiveness of joint training, most of these methods suffer from a lack of communication, since the standard DSR pipeline based on a unidirectional communication between the speech enhancement and speech recognition is still employed.

In this paper, we attempt to evolve standard joint training approaches, by proposing a novel DNN architecture, potentially able to address at the same time both the lack-of-matching and the lack-of-communication arising in current systems. Such an architecture replaces the standard DSR pipeline with a *network of deep neural networks* in which all the modules are jointly trained and better cooperate with each other thanks to a full communication between them. The architecture, after being properly unrolled, is trained with a variation of the standard back-propagation algorithm, which is based on a back-propagation of both speech recognition and speech enhancement gradients through the network of DNNs.

The preliminary results reported in this paper confirm the effectiveness of this approach, showing that the proposed paradigm is able to overtake even recently proposed joint training methods. The experimental validation has been carried out in a distant-talking scenario considering different training datasets, tasks and acoustic conditions.

2. A NETWORK OF DEEP NEURAL NETWORKS

The proposed network of deep neural networks is depicted in Fig. 1. The upper part of the figure highlights the full communication between the speech enhancement and speech recognition DNNs. How-

^{*}This work was done while the author was visiting the Montreal Institute for Learning Algorithms (MILA) and was supported by the FBK mobility programme.



Fig. 1. The proposed network of deep neural networks.

ever, such a communication modality inherently entails a chickenand-egg problem. This is caused by the fact that the speech recognizer is fed by the speech enhancement, which requires, in turn, also the speech recognition output itself to generate the enhanced speech. We can circumvent this issue by unfolding the proposed architecture to an arbitrary number of levels L, as shown in Fig. 1. The resulting computational graph is built by concatenating several speech enhancement and speech recognition DNNs, which form various interaction levels. At the first level ($\ell = 0$) the two networks are independent, but a full communication is established and continuously refined by progressively adding more levels to the architecture.

An effective communication, however, can be impaired by the high dimensionality of the speech recognizer output $\hat{y}_{SR_{\ell}}$. This dimensionality derives from the number of considered context-dependent states, which typically ranges from 1000 to 4000 (depending on the phonetic decision tree and on the dataset). To avoid feeding the speech enhancement with such a high dimensional input, we jointly estimate the monophone targets (which are only some dozens). Similarly to [25], this is realized by adding an additional softmax classifier on the top of the last hidden layer of each speech recognition DNN.

The proposed architecture is jointly trained with the backpropagation through network algorithm described in Alg. 1, where x are the input features, L the number of levels, N the number of samples in the minibatch, g the gradients and θ the DNN parameters. This algorithm is repeated for all the minibatches and iterated for several epoch until convergence. The basic idea is to perform a forward pass, compute the loss functions at the output of each DNN (mean-squared error MSE for speech enhancement and negative multinomial log-likelihood NLL for speech recognition), compute the corresponding gradients, and back-propagate them. In particular, following an approach inherited by recent joint training works, the gradient of each DNN is back-propagated through all the connected lower-level DNNs. Therefore, the speech enhancement parameter updates not only depend on the speech enhancement cost, but also on the higher-level speech recognition loss (line 11 of Alg. 1). In this way the enhancement process is in part guided by the speech recognition cost function and it would hopefully be able to provide enhanced speech which is more suitable for the subsequent speech recognition task. Similarly, the updates of the ℓ -level speech

Algorithm 1 back-propagation through network algorithm				
1: Forward Pass:				
2: Starting from input x do a forward pass through the DNNs.				
3: Compute Cost Functions ($\ell \in 0,, L$):				
4: $MSE_{\ell} = \frac{1}{N} \sum_{n=1}^{N} (\hat{x}_{SE_{\ell}} - x_{clean})^2$				
5: $NLL_{\ell} = -\frac{1}{N} \sum_{n=1}^{N} y_{SR} log(\hat{y}_{SR_{\ell}})$				
6: Gradient Computation ($\ell \in 0,, L$):				
7: $g_{SE_{\ell}} = \frac{\partial MSE_{\ell}}{\partial \theta_{SE_{\ell}}}, g_{SR_{\ell}} = \frac{\partial NLL_{\ell}}{\partial \theta_{SR_{\ell}}}$				
8: Back-propagate $g_{SE_{\ell}}, g_{SR_{\ell}}$ through $\ell - 1$ level:				
9: $g_{SE_{\ell} \to SR_{\ell-1}} = \frac{\partial MSE_{\ell}}{\partial \theta_{SR_{\ell-1}}}, g_{SR_{\ell} \to SE_{\ell-1}} = \frac{\partial NLL_{\ell}}{\partial \theta_{SE_{\ell-1}}}$				
10: Parameter Updates ($\ell \in 0,, L$):				
11: $\theta_{SE_{\ell}} \leftarrow \theta_{SE_{\ell}} - \eta[(1-\lambda)g_{SE_{\ell}} + \lambda g_{SR_{\ell+1} \to SE_{\ell}}]$				
12: $\theta_{SR_{\ell}} \leftarrow \theta_{SR_{\ell}} - \eta[(1-\lambda)g_{SR_{\ell}} + \lambda g_{SE_{\ell+1} \to SR_{\ell}}]$				

recognizer also depend on the $\ell + 1$ speech enhancement cost function. In general, as we discussed in [24], the integration of different gradients coming from the higher levels produces a regularization effect, which could significantly help the training of the system. According to this vision, the parameter λ , which weights the gradient coming from the higher levels, can be regarded as a regularization hyperparameter, whose optimal value can be determined on the development-set.

The joint training of the network of DNNs, however, can be complicated by the fact that the output distribution of the speech enhancement and speech recognition systems may change substantially during the optimization procedure. The speech recognition and speech enhancement modules would have to deal with an input distribution that is non-stationary and unnormalized. To mitigate this issue, we suggest to couple the proposed architecture with batch normalization. Batch normalization [26], which has been recently proposed in the machine learning community, addresses this concern (known as internal covariate shift) by normalizing the mean and the variance of each layer for each training mini-batch, and back-propagating through the normalization step. Similarly to what we observed in [24], batch normalization resulted crucial to significantly achieve better performance, to improve convergence of the proposed training algorithm, and to avoid any time-consuming pretraining steps.

3. RELATED WORK

The idea of using more than one neural network in the speech recognition process has long been explored. Examples of multi-DNN systems were, for instance, the so-called hierarchical bottleneck DNNs [27–29], which considered a cascade between a short-term and a long-term DNN to embedding longer-term information in the speech recognition process.

More recently, the joint training methods outlined in Sec.1 have gained considerable attention. This work can be considered as an evolution of such approaches, in which we employ a more advanced architecture based on a full communication between the DNNs. Similarly to this work, an iterative pipeline based on feeding the speech recognition output into a speech enhancement DNN has recently been proposed in [19, 30]. The main difference with our approach is that the later circumvent the chicken-and-egg problem by simply feeding the speech enhancement with the speech recognition alignments generated at the previous iteration, while our solution faces this issue by adopting the unrolling procedure over different interaction levels previously discussed. Our paradigm has also some similarities with traditional multitasking techniques [31]. The main difference is that the latter are based on sharing some hidden layers across the tasks, while our method relies on exchanging DNN outputs at various interaction levels. Finally, the proposed training algorithm has some aspects in common with the back-propagation through structure originally proposed for the parsing problem [32]. The main difference is that the latter back-propagates the gradient through a tree structure, while the proposed variation back-propagates it on a less constrained network of components. Another difference is that in the original algorithm the same neural network is used across all the levels of the tree, while in this work different types of DNNs (i.e., speech enhancement and speech recognition) are involved.

4. EXPERIMENTAL SETUP

4.1. Corpora and tasks

To provide an accurate evaluation of the proposed technique, the experimental validation has been conducted using different training datasets, different tasks and various environmental conditions. In particular, a set of experiments with TIMIT has been performed to test the proposed paradigm in low-resources conditions. To validate on a more realistic task, the proposed technique has also been evaluated on a WSJ task¹.

The experiments with TIMIT are based on a phoneme recognition task (aligned with the Kaldi s5 recipe). The original training dataset has been contaminated with a set of impulse responses measured in a real apartment. The reverberation time (T_{60}) of the considered room is about 0.7 seconds. Development and test data have been simulated with the same approach, but considering a different set of impulse responses.

The WSJ experiments are based on the popular wsj5k task (aligned with the CHiME 3 [3] task) and are conducted under two different acoustic conditions. For the *WSJ rev* case, the training set is contaminated with the same set of impulse responses adopted for TIMIT. For the *WSJ rev+noise* case, we also added non-stationary noises recorded in a domestic context (the average SNR is about 10 dB). The test phase is carried out with the DIRHA-English corpus (real-data part), consisting of 409 WSJ sentences uttered by six native American speakers in the above mentioned apartment. More details on this corpus and on the impulse responses adopted in this work can be found in [33, 34].

4.2. System details

The features considered in this work are standard 39 Mel-Cepstral Coefficients (MFCCs) computed every 10 ms with a frame length of 25 ms. The speech enhancement DNNs are fed with a context of 21 consecutive frames and predict (every 10 ms) 11 consecutive frames of enhanced MFCC features. The speech recognition DNNs are fed by such 11 speech enhanced frames and predict both context-dependent and monophone targets at their output. All the layers used Rectified Linear Units (ReLU), except for the output of the speech enhancement DNNs (linear) and the output of the speech recognition modules (softmax). Batch normalization [26] and dropout [35] are employed for all the hidden layers. The labels for the speech

enhancement DNN (denoted as x_{clean} in Fig. 1) are the MFCC features of the original clean datasets. The labels for the speech recognition DNN (denoted as y_{SR} in Fig. 1) are derived by performing a forced alignment procedure on the original training datasets. See the standard s5 recipe of Kaldi for more details [36].

The weights of the networks are initialized according to the *Glorot* initialization [37], while biases are initialized to zero. Training is based on a standard Stochastic Gradient Descend (SGD) optimization with mini-batches N of size 128. The performance on the development set is monitored after each epoch and the learning rate η is halved when the performance improvement is below a certain threshold. The training ends when no significant improvements have been observed for more than four consecutive epochs.

The main hyperparameters of the system (i.e., learning rate η , number of hidden layers, hidden neurons per layer, dropout factor, gradient weighting factor λ and number of unfolding levels L) have been optimized on the development set. As a result, speech enhancement and speech recognition DNNs with 4 hidden layers of 1024 neurons and DNNs with 6 hidden layers of 2048 neurons are employed for TIMIT and WSJ tasks, respectively. The initial learning rate is 0.08, the dropout factor is 0.2 and the considered number of levels is 3 (l = 0, ..., 2). Similarly to [24], λ is fixed to 0.1.

The proposed system, which has been implemented with Theano [38], has been coupled with the Kaldi toolkit [36] to form a context-dependent DNN-HMM speech recognizer.

5. RESULTS

5.1. Close-talking baselines

The Phoneme Error Rate (PER%) obtained by decoding the original test sentences of TIMIT is 19.5% (using DNN models trained with the original dataset). The Word Error Rate (WER%) obtained by decoding the close-talking DIRHA-English WSJ sentences is 3.3%. It is worth noting that, under such favorable acoustic conditions, the DNN model leads to a very accurate sentence transcription, especially when coupled with a language model.

5.2. Network of DNNs performance

The proposed network of DNNs approach is compared in Table 1 with other competitive systems. The first line reports the results obtained with a single neural network. In this case, only the speech recognition labels are used and the DNN is not forced to perform any speech enhancement task. The second line shows the performance obtained when the single DNN is coupled with a traditional multi-task learning, in which a speech enhancement and a speech recognition task are simultaneously considered. This multi-task architecture shares the first half of the hidden layers across the tasks, while the second half of the architecture is task-dependent. This approach aims to discover (within the shared layers) more general and robust features which can be exploited to better solve both correlated tasks. The third line reports the performance achieved with the joint training approach recently proposed in [24]. In this case a bigger DNN composed of a cascade of a speech enhancement and a speech recognition DNNs is jointly trained by back-propagating the speech recognition gradient also into the speech enhancement DNN. The last line finally shows the performance achieved with the proposed network of deep neural network approach. To allow a fair comparison, batch normalization is adopted for all the considered systems.

Table 1 highlights that the proposed approach significantly outperforms all the single DNN systems. For instance, a relative

¹The DIRHA-English dataset will be publicly distributed through the Linguistic Data Consortium (LDC). The scripts to generate the contaminated version of TIMIT will be available at https://github.com/mravanelli

Dataset	TIMIT	WSJ	WSJ
System	rev	rev	rev+noise
Single DNN	31.9	8.1	14.3
Single DNN +multitask	31.4	8.1	13.8
Joint SE-SR training	29.1	7.8	12.7
Network of DNNs	28.7	7.6	12.3

Table 1. Performance of the proposed network of DNN approach compared with other competitive DNN-based systems (PER% for TIMIT, WER% for WSJ).

Dataset	Level 0	Level 1	Level 2
TIMIT rev	31.4	29.1	28.7
WSJ rev	8.0	7.7	7.6
WSJ rev+noise	14.3	12.7	12.3

Table 2. Performance of the proposed network of DNN achieved at various levels of the architecture.

improvement of about 14% over the single DNN baseline is obtained for the *WSJ rev+noise* case. The network of deep neural networks also outperforms the considered joint training method. This result suggests that the improved cooperation between the networks achieved with our full communication scheme can overtake the standard DSR pipeline based on a partial and unidirectional information flow (which is still considered in the context of joint training approaches).

Table 2 shows the results obtained by decoding the speech recognition output at the various levels of the proposed architecture (denoted as \hat{y}_{SR_0} , \hat{y}_{SR_1} , \hat{y}_{SR_2} in Fig. 1). One can note that the performance become progressively better as the level of the network of DNNs increases. As expected, the first level speech recognizer (*SR DNN-0*) performs similarly to the single DNN baseline. The second level (*SR DNN-1*) is based on a simple cascade between a speech enhancement and a speech recognition DNNs, and thus provides results similar to that obtained with standard joint training. The third level (*SR DNN-2*) achieves the best performance, confirming that the progressive interaction of the DNNs involved in the DSR process helps in improving the system performance. No additional benefits have been observed for the considered tasks by adding more than 3 levels.

5.3. Alternative architectures

The architecture depicted in Fig. 1 represents only one of the possible ways to implement the proposed network of DNN paradigm. An alternative solution would be to share the parameters across the various speech enhancement and speech recognition DNNs. However, results (not reported here) show no benefits from this approach. Another possible modification is represented by pre-training each DNN before performing the back-propagation through network. As observed in [24], when batch normalization is adopted, no benefits from any pre-training methods have been observed. Although a more detailed exploration of alternative solutions is under study, in this section we report some preliminary results obtained by evolving the proposed method with an architecture partly-inspired by residual networks (ResNets) [39]. ResNets have recently achieved state-ofthe-art performance in computer vision and are based on the idea that, instead of learning a distribution directly, we can learn more easily the residual functions with reference to the input layer. In this work, as shown in Fig. 2 and Eq.1, we employ a ResNet-inspired



Fig. 2. The architectural variation inspired by residual networks.

Dataset	TIMIT	WSJ	WSJ
System	rev	rev	rev+noise
Original Arch.	28.7	7.6	12.3
ResNet-inspired	28.4	7.5	12.0

 Table 3. Performance with the ResNet-inspired architecture.

architecture for the speech enhancement DNNs:

$$R_{SE_l} = \hat{x}_{SE_{l-1}} - \hat{x}_{SE_l} \ . \tag{1}$$

This choice is motivated by the fact that we progressively expect less variation between the input and the output of the speech enhancement DNN as the levels increase. Directly learning the residual can thus make the training of the higher levels easier. The results achieved with the ResNet inspired architecture are reported in Tab. 3 and confirm the effectiveness of this approach, showing a coherent improvement over all the considered tasks.

6. CONCLUSION AND FUTURE WORKS

In this paper, we proposed a novel architecture for distant speech recognition based on a network of deep neural networks. The experiments, conducted considering different tasks, datasets and acoustic conditions, show that our method is effectively able to take advantage of a full communication between a speech enhancement and a speech recognizer, leading to a performance significantly better than that obtained with more standard DSR pipelines.

This work, however, represents only a first step towards more advanced architectures able to better cooperate and communicate each other to achieve a common goal. Our future research efforts will be thus focused on improving the current paradigm. In this paper, for instance, we only considered speech enhancement and speech recognition DNNs. Nevertheless, the proposed approach is a general framework that can be extended in a straightforward way by involving other modules, eventually including DNNs for acoustic scene classification, voice activity detection and speaker identification. Moreover, a more detailed exploration of alternative architectures and the natural extension of this paradigm to RNNs will also be considered as future research.

7. REFERENCES

- I. Goodfellow, Y. Bengio, and A. Courville, "Deep learning," Book in preparation for MIT Press, 2016.
- [2] D. Yu and L. Deng, Automatic Speech Recognition A Deep Learning Approach, Springer, 2015.

- [3] J. Barker, R. Marxer, E. Vincent, and S. Watanabe, "The third CHiME Speech Separation and Recognition Challenge: Dataset, task and baselines," in *Proc. of ASRU*, 2015.
- [4] K. Kinoshita et al., "The reverb challenge: A Common Evaluation Framework for Dereverberation and Recognition of Reverberant Speech," in *Proc. of WASPAA 2013*, pp. 1–4.
- [5] E. Hänsler and G. Schmidt, Speech and Audio Processing in Adverse Environments, Springer, 2008.
- [6] T. Yoshioka et al., "The NTT CHiME-3 system: Advances in speech enhancement and recognition for mobile multimicrophone devices," in *Proc. ASRU*, 2015, pp. 436–443.
- [7] M. Brandstein and D. Ward, *Microphone arrays*, Springer, Berlin, 2000.
- [8] S. Makino, T. Lee, and H. Sawada, *Blind Speech Separation*, Springer, 2010.
- [9] 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*, pp. 61–73. Springer London, 2009.
- [10] M. Ravanelli, B. Elizalde, K. Ni, and G. Friedland, "Audio concept classification with hierarchical deep neural networks," in *Proc. of EUSIPCO*, 2014.
- [11] H. Beigi, *Fundamentals of Speaker Recognition*, Springer, 2011.
- [12] P. Swietojanski, A. Ghoshal, and S. Renals, "Hybrid acoustic models for distant and multichannel large vocabulary speech recognition," in *Proc. of ASRU*, 2013, pp. 285–290.
- [13] Y. Liu, P. Zhang, and T. Hain, "Using neural network frontends on far field multiple microphones based speech recognition," in *Proc. of ICASSP*, 2014, pp. 5542–5546.
- [14] M. Ravanelli and M. Omologo, "Contaminated speech training methods for robust DNN-HMM distant speech recognition," in *Proc. of Interspeech*, 2015.
- [15] Y. Xu, J. Du, L. R. Dai, and C. H. Lee, "A regression approach to speech enhancement based on deep neural networks," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 23, no. 1, pp. 7–19, Jan 2015.
- [16] F. Weninger, H. Erdogan, S. Watanabe, E. Vincent, J. Le Roux, J. R. Hershey, and B. W. Schuller, "Speech enhancement with LSTM recurrent neural networks and its application to noiserobust ASR," in *Proc. of LVA/ICA*, 2015, pp. 91–99.
- [17] A. Narayanan and D. Wang, "Joint noise adaptive training for robust automatic speech recognition," in *Proc. of ICASSP*, 2014, pp. 4380–4384.
- [18] T. Gao, J. Du, L. R. Dai, and C. H. Lee, "Joint training of front-end and back-end deep neural networks for robust speech recognition," in *Proc. of ICASSP*, 2015, pp. 4375–4379.
- [19] Z. Chen, S. Watanabe, H. Erdogan, and J. Hershey, "Speech enhancement and recognition using multi-task learning of long short-term memory recurrent neural networks," in *Proc. of Interspeech*, 2015, pp. 3274–3278.
- [20] T. N. Sainath, R. J. Weiss, K. W. Wilson, A. Narayanan, M. Bacchiani, and A. Senior, "Speaker localization and microphone spacing invariant acoustic modeling from raw multichannel waveforms," in *Proc. of ASRU*, 2015.

- [21] M. Mimura, S. Sakai, and T. Kawahara, "Joint optimization of denoising autoencoder and dnn acoustic model based on multitarget learning for noisy speech recognition," in *Proc. of Interspeech*, 2016.
- [22] T. Higuchi, T. Yoshioka, and T. Nakatani, "Optimization of speech enhancement front-end with speech recognition-level criterion," in *Proc. of Interspeech*, 2016.
- [23] X. Xiao et al., "Deep beamforming networks for multi-channel speech recognition," in *Proc. of ICASSP*, 2016.
- [24] M. Ravanelli, P. Brakel, M. Omologo, and Y. Bengio, "Batchnormalized joint training for dnn-based distant speech recognition," in *Proc. of SLT*, 2016.
- [25] M. Seltzer and J. Droppo, "Multi-task learning in deep neural networks for improved phoneme recognition," in *Proc. of ICASSP*, 2014, pp. 6965–6969.
- [26] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in *Proc. of ICML*, 2015, pp. 448–456.
- [27] F. Grézl and M. Karafiát, "Hierarchical Neural Net Architectures for Feature Extraction in ASR," in *Proc. of Interspeech*, 2010, pp. 1201–1204.
- [28] C. Plahl, R. Schlüter, and H. Ney, "Hierarchical Bottle-Neck Features for LVCSR," in *Proc. of Interspeech*, 2010.
- [29] M. Ravanelli, V.H. Do, and A. Janin, "TANDEM-Bottleneck Feature Combination using Hierarchical Deep Neural Networks," in *Proc. of ISCSLP 2014*, 2014, pp. 113–117.
- [30] H. Erdogan, J. R. Hershey, S. Watanabe, and J. Le Roux, "Phase-sensitive and recognition-boosted speech separation using deep recurrent neural networks," in *Proc. of ICASSP*, 2015, pp. 708–712.
- [31] G. Pironkov, S. Dupont, and T. Dutoit, "Multi-Task Learning for Speech Recognition: An Overview," in *Proc. of ESANN*, 2016.
- [32] C. Goller and A. Kuchler, "Learning task-dependent distributed representations by backpropagation through structure," in *Proc. of ICNN*, 1996, pp. 347–352.
- [33] M. Ravanelli, L. Cristoforetti, R. Gretter, M. Pellin, A. Sosi, and M. Omologo, "The DIRHA-English corpus and related tasks for distant-speech recognition in domestic environments," in *Proc. of ASRU*, 2015, pp. 275–282.
- [34] M. Ravanelli, P. Svaizer, and M. Omologo, "Realistic multimicrophone data simulation for distant speech recognition," in *Proc. of Interspeech*, 2016.
- [35] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: A simple way to prevent neural networks from overfitting," *Journal of Machine Learning Research*, vol. 15, pp. 1929–1958, 2014.
- [36] D. Povey et al., "The Kaldi Speech Recognition Toolkit," in *Proc. of ASRU*, 2011.
- [37] X. Glorot and Y. Bengio, "Understanding the difficulty of training deep feedforward neural networks," in *Proc. of AISTATS*, 2010, pp. 249–256.
- [38] Theano Development Team, "Theano: A Python framework for fast computation of mathematical expressions," *arXiv eprints*, vol. abs/1605.02688, May 2016.
- [39] K. He, X. Zhang, S. Ren, and Jian Sun, "Deep Residual Learning for Image Recognition," in *Proc. of CVPR*, 2016.