

END-TO-END TRAINING OF TIME DOMAIN AUDIO SEPARATION AND RECOGNITION

Thilo von Neumann^{1,2} Keisuke Kinoshita¹ Lukas Drude² Christoph Boeddeker²
Marc Delcroix¹ Tomohiro Nakatani¹ Reinhold Haeb-Umbach²

¹NTT Communication Science Laboratories, NTT Corporation, Kyoto, Japan

²Paderborn University, Department of Communications Engineering, Paderborn, Germany

ABSTRACT

The rising interest in single-channel multi-speaker speech separation sparked development of End-to-End (E2E) approaches to multi-speaker speech recognition. However, up until now, state-of-the-art neural network-based time domain source separation has not yet been combined with E2E speech recognition. We here demonstrate how to combine a separation module based on a Convolutional Time domain Audio Separation Network (Conv-TasNet) with an E2E speech recognizer and how to train such a model jointly by distributing it over multiple GPUs or by approximating truncated back-propagation for the convolutional front-end. To put this work into perspective and illustrate the complexity of the design space, we provide a compact overview of single-channel multi-speaker recognition systems. Our experiments show a word error rate of 11.0% on WSJ0-2mix and indicate that our joint time domain model can yield substantial improvements over cascade DNN-HMM and monolithic E2E frequency domain systems proposed so far.

Index Terms— End-to-end speech recognition, speech separation, multi-speaker speech recognition, time domain, joint training

1. INTRODUCTION

Automatic Speech Recognition (ASR) is a key technology for the task of automatic analysis of any kind of spoken speech, e.g., phone calls or meetings. For scenarios of relatively clean speech, e.g., recordings of telephone speech or audio books, ASR technologies have improved drastically over the recent years [1]. More realistic scenarios like spontaneous speech or meetings with multiple participants in many cases require the ASR system to recognize the speech of multiple speakers simultaneously. In meeting scenarios for example, the overlap is in the range of 5% to 10%¹ and can easily exceed 20% in informal get-togethers². Thus, there has been a growing interest in source separation systems and multi-speaker ASR. A special focus lies on processing of single-channel recordings, as this is not only important in scenarios where only a single-channel is available (e.g., telephone conference recordings), but as well for multi-channel recordings where conventional multi-channel processing methods, e.g., beamforming, cannot separate the speakers well enough in case, e.g., they are spatially too close to each other.

The topic of single-channel source separation has been examined extensively over the last few years, trying to solve the cocktail party problem with techniques such as Deep Clustering (DPCL) [3], Permutation Invariant Training (PIT) [4] and TasNet [5, 6]. In DPCL, a neural network is trained to map each time-frequency bin to an embedding vector in a way that embedding vectors of the same speaker form a cluster in the embedding space. These clusters can be

found by a clustering algorithm and be used for constructing a mask for separation in frequency domain. Concurrently, PIT has been developed which trains a simple neural network with multiple outputs to estimate a mask for each speaker with a permutation invariant training criterion. The reconstruction loss is calculated for each possible assignment of training targets to estimations for a mixture, and the permutation that minimizes the loss is then used for training. Both DPCL and PIT show a good separation performance in time-frequency domain. The permutation-invariant training scheme was adopted to time domain source separation with a Time domain Audio Separation Network (TasNet) which replaces the commonly used Short-Time Fourier Transform (STFT) with a learnable transformation and directly works on the raw waveform. TasNet achieves a Signal-to-Distortion Ratio (SDR) gain of more than 15 dB on WSJ0-2mix, even outperforming oracle masking in frequency domain.

Based on these source separation techniques, multi-speaker ASR systems have been constructed. DPCL and PIT have been used as frequency domain source separation front-ends for a state-of-the-art single-speaker ASR system and extended to jointly trained E2E or hybrid systems [7, 8, 9, 10]. They showed that joint (re-)training can improve the performance of these models over a simple cascade system. The effectiveness of TasNet as a time domain front-end for ASR was investigated in [11], showing an improvement over frequency domain processing for both source separation and ASR results. However, TasNet was not yet optimized jointly with an ASR system, possibly due to the intricacies of dealing with the high memory consumption or the novelty of the TasNet method.

In this paper, we combine a state-of-the-art front-end, i.e., Conv-TasNet [4], with an E2E CTC/attention [12, 13, 14] ASR system to form an E2E multi-speaker ASR system that directly operates on raw waveform features. We try to answer the questions whether it is possible to jointly train a time domain source separation system like Conv-TasNet with an E2E ASR system and whether the performance can be improved by joint fine-tuning. Going further on the investigations from [11], we retrain pre-trained front- and back-end models jointly and show by evaluating on the WSJ0-2mix database that a simple combination of an independently trained Conv-TasNet and ASR system already provides competitive performance compared to other E2E approaches, while a joint fine-tuning of both modules in the style of an E2E system can further improve the performance by a large margin. We enable joint training by distributing the model over multiple GPUs and show that an approximation of truncated back-propagation [15] through time for convolutional networks enables joint training even on a single GPU by significantly reducing the memory usage while still providing a good performance.

We finally put this work into perspective by providing a compact overview of single-channel multi-speaker ASR system and illustrate the complexity of the design space.

¹Measured on the AMI meeting corpus [2].

²Measured on the CHiME-5 database.




Fig. 1. Architecture of the joint E2E ASR model. Sources are separated by a Conv-TasNet and separated audio streams are processed by a single-speaker ASR system. During training, the permutation problem is solved based on the signal level loss with π_{sig} .

2. RELATION TO PRIOR WORK

Other works already studied the effectiveness of frequency domain source separation techniques as a front-end for ASR. DPCL and PIT have been efficiently used for this purpose, and it was shown that joint retraining for fine-tuning can improve performance [7, 8, 10]. E2E systems for single-channel multi-speaker ASR have been proposed that no longer consist of individual parts dedicated for source separation and speech recognition, but combine these functionalities into one large monolithic neural network. They extend the encoder of a CTC/attention-based E2E ASR system to separate the encoded speech features and let one or multiple attention decoders generate an output sequence for each speaker [16, 17]. These models show promising performance, but they are not on par with hybrid cascade systems yet. Drawbacks of these monolithic E2E models compared to cascade systems include that they cannot make use of parallel and single-speaker data and that they do not allow pre-training of individual system parts. The impact of using raw waveform features directly for the task of multi-speaker ASR has only been investigated for a combination of TasNet and a single-speaker ASR system [11], but not yet jointly trained.

3. SOURCE SEPARATION AND SPEECH RECOGNITION

3.1. Time domain source separation with Conv-TasNet

Conv-TasNet [6] is a single-channel source separating front-end which can produce waveforms for a fixed number of speakers from a mixture waveform. It is a variant of [4], replacing the feature extraction by a learnable transformation and the separation network by a convolutional architecture. It outputs an estimated audio stream in time domain for each speaker present in the input signal \mathbf{x} :

$$[\mathbf{x}_1^{(\text{enh})}, \mathbf{x}_2^{(\text{enh})}] = \text{ConvTasNet}(\mathbf{x}). \quad (1)$$

The model directly works on the raw waveform instead of STFT frequency domain features which makes it possible both to easily model and reconstruct phase information and to propagate gradients through the feature extraction and signal reconstruction parts. While propagating from raw waveform at the output to raw waveform at the input, it is possible to directly optimize a loss on the time domain signals, such as the Scale-Invariant-Signal-to-Distortion Ratio (SI-SNR) loss, which we call the front-end loss $\mathcal{L}^{(\text{FE})}$ here. This loss is optimized in a permutation-invariant manner by picking the assignment π_{sig} of estimations to targets that minimizes the loss.

Since the Conv-TasNet is built upon a convolutional architecture, it can be heavily parallelized on GPUs as compared to RNN-based models, but has a limited receptive field. When optimized for source separation only, the limited length of the receptive field is actively exploited by performing a chunk-level training on chunks of 4 s length randomly cut from the training examples, both increasing the variability of data within one minibatch and simplifying implementation, since the length of all training examples is fixed.

3.2. End-to end CTC/attention speech recognition

As a speech recognizer we use a CTC/attention-based ASR system. We use an architecture similar to [13] with an implementation included in the ESPnet framework [18], but we replace the original filterbank and pitch feature extraction by log-mel features implemented in a way such that gradients can be propagated through. This way, gradients can flow from the ASR system to the front-end.

The multi-target loss for the ASR system $\mathcal{L}^{(\text{ASR})}$ is composed of a CTC and an attention loss, as in [13] Sections II B and C,

$$\mathcal{L}^{(\text{ASR})} = \lambda \mathcal{L}^{(\text{CTC})} + (1 - \lambda) \mathcal{L}^{(\text{att})} \quad (2)$$

with a weight λ that controls the interaction of both loss terms. During training, teacher forcing using the ground truth transcription labels is employed for the attention decoder.

4. JOINT END-TO-END MULTI-SPEAKER ASR

We propose to combine a Conv-TasNet as the source separation front-end with a CTC/attention speech recognizer as displayed in Fig. 1. The input mixture \mathbf{x} is separated by the front-end and the separated audio streams are processed by a single-speaker ASR back-end. Although multi-speaker source separation can already be performed by combining independently trained front- and back-end systems, the source separator produces artifacts unknown to the ASR system which disturb its performance. According to [19], and as also shown in [8, 10], such a mismatch can be mitigated by jointly fine-tuning the whole model at once.

We here compare three different variants of joint fine-tuning: (a) fine-tuning just the ASR system on the enhanced signals, (b) fine-tuning just the front-end by propagating gradients through the ASR system but only updating the front-end parameters and (c) jointly fine-tuning both systems. The losses for the front- and back-end are combined as

$$\mathcal{L} = \alpha \mathcal{L}^{(\text{FE})} + \beta \mathcal{L}^{(\text{ASR})}, \quad (3)$$

where α and β are manually chosen weights for the front-end and ASR losses. For (a) α is set to 0 and β to 1, for (b) α is set to 1 and β to 0, and for (c) β is set to 1 and α is set to 0.5 without carefully optimizing them. The system does not seem to be very sensitive to the choices of α and β .

In order to choose the transcription for teacher forcing and loss computation a permutation problem needs to be solved. Two possible options are to use the permutation π_{CTC} that minimizes the CTC loss as in [16], or the permutation π_{sig} that minimizes the signal level loss, as in [8]. While using π_{CTC} has the advantage of not requiring parallel data, permutation assignment based on π_{sig} works more reliably in our experiments and we use π_{sig} for all fine-tuning experiments even when the front-end is not optimized.

4.1. Approximated truncated back-propagation through time

One-dimensional Convolutional Neural Networks (1D-CNNs) over time, as they are used in the Conv-TasNet, can be seen as an alternative to Recurrent Neural Network (RNN) architectures. Similar to RNNs, this can lead to enormous memory consumption when a sufficiently long time series is used for back-propagation. For example, we here fine-tune the Conv-TasNet with the E2E ASR model jointly on single mixtures. Although we constrain ourselves to a batch size of one, this requires us to split the model onto four GPUs, three GPUs for the front-end and one GPU for the back-end, by placing individual layers on different devices.

This memory consumption can be addressed by generalizing truncated back-propagation through time (TBPTT) to 1D-CNN architectures. TBPTT for 1D-CNNs can in theory be realized by back-propagating the gradients for a part of the output only. While moving back towards the input, the gradients reaching over the borders of this part are ignored. In practice, however, this is difficult to implement and we here approximate TBPTT for the Conv-TasNet front-end by ignoring the left and right context of the block the gradients are computed for. We first compute the forward step on the whole mixture without building the backward graph to obtain an output estimation for the whole signal. Note that this only requires to store the output signal and no persistent data for the backward computation. We then compute the forward step again with enabled backward graph construction, but only for a chunk randomly cut from the input signal. The approximated output for the whole utterance is formed by overwriting the corresponding part of the full forward output with the approximated chunk output. This full output is passed to the back-end and gradients reaching the front-end are only back-propagated through the approximated chunk. This technique allows to run the joint training on a single GPU in our case, but even with larger GPU memory this permits increasing the batch size which in general speeds up training and produces a smoother gradient.

5. EXPERIMENTS

We carry out experiments on the WSJ database and the commonly used WSJ0-2mix dataset first proposed in [3]. The data in WSJ0-2mix is generated by linearly mixing clean utterances from WJS0 (si.tr.s for training and si.et.05 for testing) at ratios randomly chosen from 0 dB to 5 dB. It consists of two different datasets, namely the min and max datasets. The min dataset was designed for source separation and is formed by truncating the longer one of the two mixed recordings, so that it only contains fully overlapped speech. We use this dataset for pre-training the Conv-TasNet, but it is not suitable for joint training where the audio data needs to match the full transcription. For this need, we use the max dataset that does

not truncate any recordings. We use a sampling frequency of 8 kHz for both the front- and back-end to speed up the training process. We remove any labels marked as noisy, i.e., special tokens such as "lip smack" or "door slam", from the training transcriptions since these cannot be assigned to one speaker based on speech information by the front-end which makes their estimation ambiguous.

We evaluate our experiments in terms of Word Error Rate (WER), Character Error Rate (CER) and, where applicable, by SDR as supplied by the BSS-EVAL toolbox [20] and SI-SNR [21]. For the experiments on mixed speech, the metrics are computed for all possible combinations of predictions and ground truth transcriptions and the metrics for the permutation that minimizes WER is reported.

5.1. Conv-TasNet time domain source separation

We use the best performing architecture according to [6] and optimize it with the ADAM optimizer [22].³ In particular, following the hyper-parameter notations in the original paper, we set $N = 512$, $L = 16$, $B = 128$, $H = 512$, $P = 3$, $X = 8$ and $R = 3$ with global layer normalization. For distributing over multiple GPUs, we split between the three repeating convolutional blocks. Table 1 lists SDR and SI-SNR performance for our Conv-TasNet model, comparing the min and max subsets of WSJ0-2mix. It can be seen that our implementation of the Conv-TasNet reaches a comparable performance on the min dataset when compared to the original paper. There is a slight degradation in performance on the max dataset caused by the mismatch of training and test data because the model never saw long single-speaker regions during training and learned to always output a speech signal on both outputs, while these regions are present in the max dataset.

Table 1. SDR and SI-SNR in dB for the min and max test (tt) datasets of the WSJ0-2mix database.

Dataset	SDR	SI-SNR
WSJ0-2mix min [6]	15.6	15.3
WSJ0-2mix min (ours)	14.3	13.8
WSJ0-2mix max (ours)	13.8	13.4

5.2. CTC/attention ASR model

We use a configuration similar to [16] without the speaker dependent layers for the speech recognizer. This results in a model with two CNN layers followed by two BLSTMP layers with 1024 units each for the encoder, one LSTM layer with 300 units for the decoder and a feature dimension of 80. The multi-task learning weight was set to $\lambda = 0.2$. We use a location-aware attention mechanism and ADADELTA [23] as optimizer. All decoding is performed with an additional word-level RNN language model. Our ASR model achieves a WER of 6.4 % on the WSJ eval92 set.

5.3. Joint finetuning

The results of the different fine-tuning variants are listed for comparison in Table 2. It is notable that combining the independently trained models (Conv-TasNet + RNN) already gives a competitive performance compared other methods (see Section 5.4 and Table 3). Fine-tuning just the ASR system (+ fine-tune ASR) can further cut the WER almost in half from 22.9 % to 11.7 %.

Joint fine-tuning without a signal level loss (+ fine-tune joint), when the system is no longer constrained to transport meaningful

³Our implementation is based on <https://github.com/funcwj/conv-tasnet>.

Table 2. CER and WER on max test (tt) set of WSJ0-2mix for different variants of fine-tuning. All models are pre-trained.

Model	fine-tune		Joint training type	additional SI-SNR loss	CER	WER	SDR	SI-SNR
	front-end	back-end						
Conv-TasNet + RNN	—	—	—	—	13.9	22.9	13.8	13.4
+ fine-tune ASR	—	✓	single GPU	—	6.7	11.7	”	”
+ fine-tune TasNet	✓	—	multi GPU	—	7.7	14.2	10.5	9.5
+ SI-SNR loss	✓	—	multi GPU	✓	7.7	14.3	12.5	12.1
+ fine-tune joint	✓	✓	multi GPU	—	6.2	11.7	9.8	8.4
+ SI-SNR loss	✓	✓	multi GPU	✓	6.0	11.1	13.8	13.5
+ fine-tune joint TBPTT	✓	✓	single GPU (TBPTT)	—	6.1	11.0	11.7	11.5
+ SI-SNR loss	✓	✓	single GPU (TBPTT)	✓	18.0	23.9	12.4	12.1

Table 3. Comparison of single-channel multi-speaker ASR systems. They differ heavily in their used architecture, training data and technique. *: The models are evaluated on mixtures based on WSJ [16, 17] and the WERs are not comparable to the other models.

Model	structure	pre-training	joint training	signal reconstr.	no paral. data req.	data		CER	WER
						train	eval		
DPCL				✓	—	WSJ0-2mix			
+ DNN-HMM [7]	hybrid	✓	—	✓	—	WSJ0	WSJ0-2mix	—	16.5
+ CTC/attention [8]	E2E	✓	—	✓	—	WSJ0	WSJ0-2mix	23.1	—
+ joint fine-tuning [8]	E2E	✓	✓	✓	—	WSJ0-2mix	WSJ0-2mix	13.2	—
PIT-ASR (best) [10, 17]	hybrid	—	✓	—	✓	WSJ0-2mix	WSJ0-2mix	—	28.2
E2E ASR [17]	E2E	—	✓	—	✓	WSJ0-2mix	WSJ0-2mix	—	25.4
<i>E2E ASR* [16]</i>	<i>E2E</i>	(✓)	✓	—	✓	<i>WSJ-2mix</i>	<i>WSJ-2mix</i>	<i>13.2</i>	<i>28.2</i>
<i>E2E ASR* [17]</i>	<i>E2E</i>	—	✓	—	✓	<i>WSJ-2mix</i>	<i>WSJ-2mix</i>	<i>10.9</i>	<i>18.4</i>
joint TasNet (our best)	E2E	✓	✓	✓	—	WSJ & WSJ0-2mix	WSJ0-2mix	6.1	11.0

speech between front- and back-end, can not improve much over just fine-tuning the ASR system and significantly lowers the source separation performance. This indicates that there is enough information available for reliable speech recognition in the separated signals (i.e., retraining of the front-end is not required), but that not all information required to reconstruct speech is required for ASR.

Using a signal-level loss (+ fine-tune joint + SI-SNR loss) can further improve the WER to 11.1%. In this case, the source separation performance stays comparable to the Conv-TasNet model. A signal-level loss helps the model to better separate the speech.

The performance for just fine-tuning the front-end (+ fine-tune TasNet) cannot reach the performance of fine-tuning the back-end. This means that it is easier to mitigate the mismatch for the ASR back-end (i.e., learn to ignore the artifacts produced by the front-end) than it is for the front-end (i.e., learn to suppress the artifacts).

Comparing the results of the chunk-based fine-tuning (+ fine-tune joint TBPTT) as an approximation of TBPTT with the full joint fine-tuning (+ fine-tune joint), it can be seen that even though the TBPTT-based approach is just an approximation, its performance is comparable to the full joint model if no signal-level loss is used. It even performs slightly better, possibly because TBPTT allowed to use a larger batch size. The slightly odd degradation in performance for the case with a signal level loss (+ fine-tune joint TBPTT + SI-SNR loss) might be caused by the signal-level loss penalizing the approximation heavily, while the gradient propagated through the ASR system is less harmful. This case was not evaluated further.

5.4. Comparison with related work

This section compares the performance of the different related works presented in Section 2. Their major differences and performance in terms of CER and WER are listed in Table 3. While these compar-

isons are not fair because the presented works differ greatly in their overall model structure, training methods and data, the numbers are meant to give a rough indication of how these methods compare and how complex the design space is.

Among the related systems, the hybrid DNN-HMM system still outperforms all monolithic approaches although this system is not fine-tuned jointly. On the other hand, the E2E ASR model can outperform the jointly optimized hybrid PIT-ASR on the same dataset. Their WERs, however, are far from the joint DPCL model. Although not directly comparable, the best results in this table were produced by cascade models that allow reconstruction of the enhanced separated signals (DPCL + ASR, joint TasNet), which suggests that having dedicated parts for source separation and speech recognition is helpful, while joint fine-tuning improves the performance. Our time domain approach gives the best result in this comparison.

6. CONCLUSIONS

We propose to use a time domain source separation system like Conv-TasNet as a front-end for a single-speaker E2E ASR system to form a multi-speaker E2E speech recognizer. We show that independently training the front- and back-end already gives a competitive performance and that joint fine-tuning can drastically improve the performance. Fine-tuning can be performed jointly with the whole model distributed over multiple GPUs, but can as well be sped up roughly by a factor of 2 on a single GPU by approximating TPBTT for convolutional neural networks, while keeping the performance comparable. The results suggest that retraining the ASR part can much better compensate the mismatch between front-end and back-end than a fine-tuned front-end could.

7. REFERENCES

- [1] W. Xiong, L. Wu, F. Allewa, J. Droppo, X. Huang, and A. Stolcke, "The microsoft 2017 conversational speech recognition system," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, April 2018, pp. 5934–5938.
- [2] J. Carletta, S. Ashby, S. Bourban, M. Flynn, M. Guillemot, T. Hain, J. Kadlec, V. Karaiskos, W. Kraaij, M. Kronenthal, G. Lathoud, M. Lincoln, A. Lisowska, I. McCowan, W. Post, D. Reidsma, and P. Wellner, "The ami meeting corpus: A pre-announcement," in *Machine Learning for Multimodal Interaction*, S. Renals and S. Bengio, Eds., Berlin, Heidelberg, 2006, pp. 28–39, Springer Berlin Heidelberg.
- [3] Y. Isik, J. L. Roux, Z. Chen, S. Watanabe, and J. R. Hershey, "Single-channel multi-speaker separation using deep clustering," *arXiv preprint arXiv:1607.02173*, 2016.
- [4] D. Yu, M. Kolbæk, Z.-H. Tan, and J. Jensen, "Permutation invariant training of deep models for speaker-independent multi-talker speech separation," in *The 42nd IEEE International Conference on Acoustics, Speech and Signal Processing/IEEE International Conference on Acoustics, Speech and Signal Processing*. IEEE, 2017, pp. 241–245.
- [5] Y. Luo and N. Mesgarani, "Tasnet: time-domain audio separation network for real-time, single-channel speech separation," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 696–700.
- [6] Y. Luo and N. Mesgarani, "Conv-tasnet: Surpassing ideal time–frequency magnitude masking for speech separation," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 27, no. 8, pp. 1256–1266, 2019.
- [7] T. Menne, I. Sklyar, R. Schlüter, and H. Ney, "Analysis of deep clustering as preprocessing for automatic speech recognition of sparsely overlapping speech," *arXiv preprint arXiv:1905.03500*, 2019.
- [8] S. Settle, J. Le Roux, T. Hori, S. Watanabe, and J. R. Hershey, "End-to-end multi-speaker speech recognition," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 4819–4823.
- [9] D. Yu, X. Chang, and Y. Qian, "Recognizing multi-talker speech with permutation invariant training," *arXiv preprint arXiv:1704.01985*, 2017.
- [10] Y. Qian, X. Chang, and D. Yu, "Single-channel multi-talker speech recognition with permutation invariant training," *Speech Communication*, vol. 104, pp. 1–11, 2018.
- [11] F. Bahmaninezhad, J. Wu, R. Gu, S.-X. Zhang, Y. Xu, M. Yu, and D. Yu, "A comprehensive study of speech separation: spectrogram vs waveform separation," *arXiv preprint arXiv:1905.07497*, 2019.
- [12] S. Kim, T. Hori, and S. Watanabe, "Joint CTC-attention based end-to-end speech recognition using multi-task learning," in *2017 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2017, pp. 4835–4839.
- [13] S. Watanabe, T. Hori, S. Kim, J. R. Hershey, and T. Hayashi, "Hybrid CTC/attention architecture for end-to-end speech recognition," *IEEE Journal of Selected Topics in Signal Processing*, vol. 11, no. 8, pp. 1240–1253, 2017.
- [14] W. Chan, N. Jaitly, Q. Le, and O. Vinyals, "Listen, attend and spell: A neural network for large vocabulary conversational speech recognition," in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2016, pp. 4960–4964.
- [15] P. J. Werbos, "Backpropagation through time: what it does and how to do it," *Proceedings of the IEEE*, vol. 78, no. 10, pp. 1550–1560, 1990.
- [16] H. Seki, T. Hori, S. Watanabe, J. L. Roux, and J. R. Hershey, "A purely end-to-end system for multi-speaker speech recognition," *arXiv preprint arXiv:1805.05826*, 2018.
- [17] X. Chang, Y. Qian, K. Yu, and S. Watanabe, "End-to-end monaural multi-speaker asr system without pretraining," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 6256–6260.
- [18] S. Watanabe, T. Hori, S. Karita, T. Hayashi, J. Nishitoba, Y. Unno, N. E. Y. Soplin, J. Heymann, M. Wiesner, N. Chen, et al., "Espnet: End-to-end speech processing toolkit," *arXiv preprint arXiv:1804.00015*, 2018.
- [19] J. Heymann, L. Drude, C. Boeddeker, P. Hanebrink, and R. Haeb-Umbach, "Beamnet: End-to-end training of a beamformer-supported multi-channel asr system," in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017, pp. 5325–5329.
- [20] E. Vincent, "BSSEval. a toolbox for performance measurement in (blind) source separation," *lúnea*. Disponible: http://bassdb.gforge.inria.fr/bss_eval/. [Último acceso: 29 Marzo 2017], 2005.
- [21] J. Le Roux, S. Wisdom, H. Erdogan, and J. R. Hershey, "SDR—half-baked or well done?," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 626–630.
- [22] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *CoRR*, vol. abs/1412.6980, 2014.
- [23] M. D. Zeiler, "Adadelata: an adaptive learning rate method," *arXiv preprint arXiv:1212.5701*, 2012.