



Revisiting joint decoding based multi-talker speech recognition with DNN acoustic model

Martin Kocour* Kateřina Žmolíková* Lucas Ondel*¹ *Ján Švec*
Marc Delcroix[†] Tsubasa Ochiai[†] Lukáš Burget* Jan “Honza” Černocký*

* Brno University of Technology, Faculty of Information Technology, Speech@FIT, Czechia

[†] NTT Corporation, Japan

ikocour@fit.vut.cz

Abstract

In typical multi-talker speech recognition systems, a neural network-based acoustic model predicts senone state posteriors for each speaker. These are later used by a single-talker decoder which is applied on each speaker-specific output stream separately. In this work, we argue that such a scheme is sub-optimal and propose a principled solution that decodes all speakers jointly. We modify the acoustic model to predict joint state posteriors for all speakers, enabling the network to express uncertainty about the attribution of parts of the speech signal to the speakers. We employ a joint decoder that can make use of this uncertainty together with higher-level language information. For this, we revisit decoding algorithms used in factorial generative models in early multi-talker speech recognition systems. In contrast with these early works, we replace the GMM acoustic model with DNN, which provides greater modeling power and simplifies part of the inference. We demonstrate the advantage of joint decoding in proof of concept experiments on a mixed-TIDIGITS dataset.

Index Terms: Multi-talker speech recognition, Permutation invariant training, Factorial Hidden Markov models

1. Introduction

Although automatic speech recognition (ASR) systems today achieve remarkable performance, they degrade in presence of interference. Notably, recognizing speech in multi-talker environments is still a challenge, as demonstrated e.g. in recent CHiME evaluations [1, 2]. The problem is even more pronounced in the single-channel scenario, where the recording of only one microphone is available. Significant amount of research addressing single-channel multi-talker speech recognition has been conducted in the past [3, 4, 5], including early approaches using factorial generative models [6, 7] and later neural network-based acoustic models [4, 5].

For multi-talker speech recognition, neural network-based systems are dominant nowadays [4, 5]. One of the most commonly used approaches is a multi-talker acoustic model trained with permutation invariant training (PIT) [8, 4]. In this case, the neural network constituting the acoustic model has separate

outputs predicting the senone state posteriors for each speaker. The ambiguity in the order of speakers at the outputs and the labels is addressed by considering all possible permutations of the outputs during the training. In the inference time, the ASR decoder is applied separately on each sequence of posteriors of each speaker. PIT-ASR systems have shown to be promising for multi-talker ASR even for complex large vocabulary tasks [4, 9].

However, in the PIT-ASR approach, we expect the acoustic model to fully solve the “separation” of the speakers, i.e. fully attribute parts of the mixed speech signal to the different outputs of the network. This separation is based only on the acoustic signal and without higher level language information. This may be challenging especially when the voices of the speakers are very similar. While the decoder could help to solve the “separation” with the use of the dictionary and grammar, in the current scheme, there is no interaction between the decoders of the individual speakers. This can lead to duplicity where the same phoneme or word is attributed to multiple speakers [10]. In this paper we investigate a principled way to overcome this issue by *decoding all speakers jointly*.

In pre-deep learning works, joint decoding of multiple speakers was proposed in the framework of factorial GMM-HMM (f-GMM-HMM) [6]. This model showed success at recognizing speech in the monaural speech separation and recognition challenge [3], where it even outperformed human listeners [11]. In the model, features of each speaker are modeled by GMM-HMMs with states corresponding to parts of speech (such as phonemes). Combining the per-speaker models using a speaker-interaction model gives rise to the f-GMM-HMM which can explain the mixed speech of multiple speakers. By doing inference of the most probable sequence of states, we can transcribe the speech of multiple speakers at once. This joint decoding of all speakers can take full advantage of the available dictionary and grammar, which can significantly improve the resulting hypothesis. The f-GMM-HMM has been however applied only to very constrained tasks with known speakers and is difficult to scale up to open speaker case, in part because of the iterative acoustic inference caused by the interaction model [6]. Besides, the GMM-HMM ASR systems have been long outperformed by neural network-based systems [12, 4].

In this work, we propose to combine the advantage of neural network-based approaches with the joint decoding used in factorial models. To do this, we modify the neural network in the acoustic model to predict joint state posteriors of multiple speakers given the mixed speech. By doing this, we enable the neural network to provide uncertainty about the attribution of the parts of speech to the speakers, which can be later resolved in the decoding stage using higher level language information.

Lucas Ondel is now affiliated with LISN, CNRS, Université Paris-Saclay, France.

The work was partly supported by European Union’s Horizon 2020 projects No. 864702 - ATCO2 and No. 884287 HAAWAIL, and by Czech Ministry of Education, Youth and Sports from project no. LTAIN19087 “Multi-linguality in speech technologies”. Part of high-performance computation run on IT4I supercomputer and was supported by the Ministry of Education, Youth and Sports of the Czech Republic from the Large Infrastructures for Research, Experimental Development and Innovations project e-Infrastructure CZ – LM2018140.

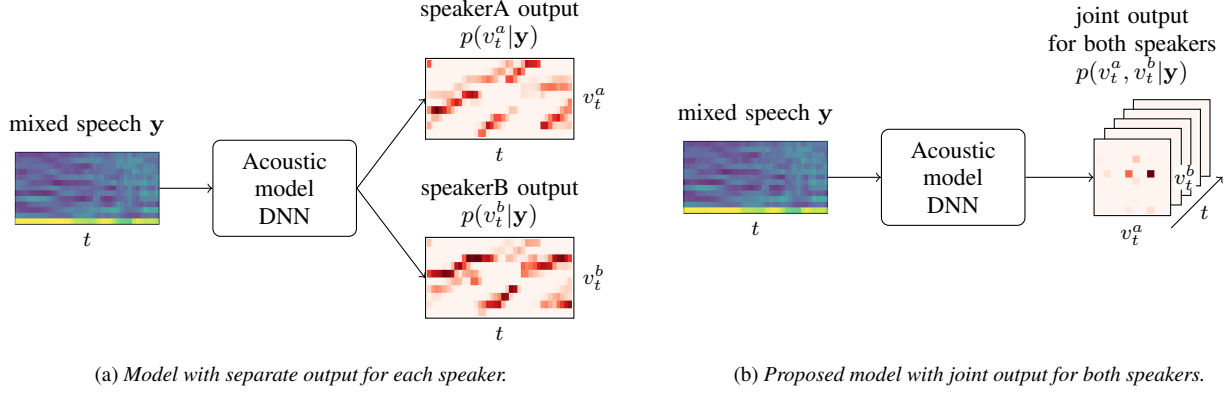


Figure 1: Comparison of neural network architectures for single channel multi-talker speech recognition.

The joint state posteriors can be leveraged by the joint decoder based on factorial HMM. To make the inference feasible, we employ the loopy belief propagation algorithm as proposed in f-GMM-HMM works [6, 13]. Contrary to f-GMM-HMM, our system can take the advantage of the strong modeling power of neural networks. The neural network also simplifies the acoustic inference. In a proof of concept experiment, we demonstrate that 1) decoding each speaker separately in the PIT-ASR approach is not optimal, and 2) using the factorial HMM framework in combination with the neural network acoustic model, we can improve the performance by joint decoding of multiple speakers.

2. Multi-talker speech recognition

In this work, we assume that the observed single-channel signal y is a combination of speech x^a and x^b from two speakers¹ as

$$y_t = x_t^a + x_t^b. \quad (1)$$

The goal is to recognize the speech of both speakers in the mixture. This can be expressed as

$$\hat{\mathbf{v}}^a, \hat{\mathbf{v}}^b = \arg \max_{\mathbf{v}^a, \mathbf{v}^b} p(\mathbf{y} | \mathbf{v}^a, \mathbf{v}^b) P(\mathbf{v}^a) P(\mathbf{v}^b), \quad (2)$$

where $\mathbf{v}^a = [v_1^a, \dots, v_T^a]$ is a hidden state sequence (which is mapped to a word sequence) for speaker a , $p(\mathbf{y} | \mathbf{v}^a, \mathbf{v}^b)$ is a likelihood function of observed mixed speech and $P(\mathbf{v}^a)$ and $P(\mathbf{v}^b)$ are prior probabilities estimated by language model.

Conventional multi-talker approaches such as PIT-ASR decode each speaker k separately as $\hat{\mathbf{v}}^k = \arg \max_{\mathbf{v}^k} p(\mathbf{y} | \mathbf{v}^k) p(\mathbf{v}^k)$ [8, 4]. This rule can be obtained by assuming conditional independence of the state sequences of the speakers given the observation, i.e. $p(\mathbf{v}^a, \mathbf{v}^b | \mathbf{y}) = p(\mathbf{v}^a | \mathbf{y}) p(\mathbf{v}^b | \mathbf{y})$. However, this is a strong assumption that may not be valid in practice. We propose instead to consider the dependencies that should eventually lead to better modeling.

In this paper, we focus on DNN-HMM hybrid systems, where a DNN predicts senone states posteriors and HMM is used to model state transitions. In the following, we briefly review the conventional scheme for PIT-ASR with separate decoding and then introduce our proposed joint decoding with factorial DNN-HMM.

¹To simplify the derivation, we restrict ourselves to the two speaker case and discuss extension to more speakers in Section 5.

2.1. Conventional PIT-ASR with separate decoding

In PIT-ASR, the acoustic model predicts the posterior probabilities for each speaker separately

$$[p(v_t^a | y_t), p(v_t^b | y_t)] = f_{\text{NN}}(y_t), \quad (3)$$

where $p(v_t^a | y_t)$ and $p(v_t^b | y_t)$ are posterior probabilities of hidden states given observed frame y_t , $f_{\text{NN}}(y_t)$ is the DNN forward function. Such model is depicted in Figure 1a. Then, the ASR decoding is performed for each speaker independently using e.g. Viterbi algorithm. Below, we briefly explain the decoding considering the case of speaker a .

The joint probability $p(\mathbf{x}^a, \mathbf{v}^a)$ of source \mathbf{x}^a and sequence of hidden states \mathbf{v}^a is modeled by HMM

$$p(\mathbf{x}^a, \mathbf{v}^a) = \prod_t p(x_t^a | v_t^a) p(v_t^a | v_{t-1}^a) \quad (4)$$

$$p(x_t^a | v_t^a) = \frac{p(v_t^a | x_t^a) p(x_t^a)}{p(v_t^a)} \quad (5)$$

$$p(v_t^a | x_t^a) \approx p(v_t^a | y_t), \quad (6)$$

where $p(v_t^a | y_t)$ is a posterior modeled by neural network (3). Note that the evidence $p(x_t^a)$ does not affect the decoding results, since it is constant. That is why we use pseudo-likelihoods $\bar{p}(x_t^a | v_t^a)$

$$\bar{p}(x_t^a | v_t^a) = \frac{p(v_t^a | y_t)}{p(v_t^a)} \quad (7)$$

However, in this work, we found out that normalizing by prior $p(v_t^a)$ does not really help either, so we kept $\bar{p}(x_t^a | v_t^a) \approx p(v_t^a | y_t)$.

Recognizing speech using an HMM involves finding the most likely sequence $\hat{\mathbf{v}}^a$ given observed data \mathbf{x}^a (i.e. the maximum a posteriori (MAP) state sequence). The MAP state sequence is obtained by Viterbi algorithm, where messages $m(v_t)$ are defined as

$$m(v_{t+1}^a) = \max_{v_t^a} p(v_{t+1}^a | v_t^a) m(v_t^a) \bar{p}(x_t^a | v_t^a) \quad (8)$$

$$\tilde{v}_t(v_{t+1}^a) = \arg \max_{v_t^a} p(v_{t+1}^a | v_t^a) m(v_t^a) \bar{p}(x_t^a | v_t^a). \quad (9)$$

The MAP state sequence $\hat{\mathbf{v}} = [\hat{v}_1, \dots, \hat{v}_T]$ is recovered by backtracking $\hat{v}_t = \tilde{v}_t(\hat{v}_{t+1})$, where $\hat{v}_T = \arg \max_{v_T} m(v_T)$ initiates the recursion.




Figure 2: Factorial HMM for speaker a and b with observed sequence \mathbf{y} and hidden state sequences \mathbf{v}^a and \mathbf{v}^b . Green dashed lines show message passing in loopy belief propagation.

2.2. Proposed PIT-ASR with joint decoding

To take account of the dependencies, we model the joint probability of mixed speech \mathbf{y} and hidden state sequences $\mathbf{v}^a, \mathbf{v}^b$ in Factorial HMM framework [14]. Figure 2 shows dynamic Bayesian network representing this model, where the joint probability is

$$p(\mathbf{y}, \mathbf{v}^a, \mathbf{v}^b) = \prod_t p(y_t | v_t^a, v_t^b) p(v_t^a | v_{t-1}^a) p(v_t^b | v_{t-1}^b), \quad (10)$$

where $p(y_t | v_t^a, v_t^b)$ is derived from a neural network that predicts the posterior probabilities of a tuple of states (v_t^a, v_t^b) as

$$p(v_t^a, v_t^b | y_t) = g_{\text{NN}}(y_t), \quad (11)$$

where $g_{\text{NN}}(y_t)$ is a forward function of proposed neural network shown in Figure 1b (also referred as joint network in follow-up text). Similar to the separate decoding, we use pseudo-likelihood $\tilde{p}(y_t | v_t^a, v_t^b) \approx p(v_t^a, v_t^b | y_t)$.

Note that our model is similar to Factorial HMM proposed in [6], where the authors model sources \mathbf{x}^a and \mathbf{x}^b explicitly. They introduce interaction function $p(y_t | \mathbf{x}^a, \mathbf{x}^b)$, which models the coupling. However, we instead use a more direct and simpler approach where a DNN is trained to directly predict hidden state tuples given mixed speech \mathbf{y} .

Standard Viterbi algorithm can be used to exactly infer the MAP hidden state sequences in FHMM. However, its time complexity is $O(TKV^{K+1})$ for K speakers and HMM with V states [6] (i.e. it scales exponentially w.r.t. number of speakers). Fortunately, there are algorithms, that can approximate the inference in a more efficient way (e.g. variational methods [14] or loopy belief propagation [6]). In this work we use the loopy belief propagation (LBP).

In LBP framework, messages are passed between variables, which share common factors [15] according to a predefined schedule. Our message passing schedule was inspired by [6] and has the following form for one speaker:

$$\tilde{p}(y_t | v_t^a) = \max_{v_t^b} \tilde{p}(y_t | v_t^a, v_t^b) \tilde{p}_{fw}(v_t^b) \tilde{p}_{bw}(v_t^b) \quad (m1)$$

$$\tilde{p}_{fw}(v_t^a) = \max_{v_{t-1}^a} p(v_t^a | v_{t-1}^a) \tilde{p}_{fw}(v_{t-1}^a) \tilde{p}(y_{t-1} | v_{t-1}^a) \quad (m2)$$

$$\tilde{p}_{bw}(v_t^a) = \max_{v_{t+1}^a} p(v_{t+1}^a | v_t^a) \tilde{p}_{bw}(v_{t+1}^a) \tilde{p}(y_{t+1} | v_{t+1}^a) \quad (m3)$$

The message passing is also depicted in Figure 2. Green dashed lines depict the message passing for speaker a , while messages of speaker b are fixed.

When all messages for the first speaker are computed, the process is repeated for the next speaker, while messages of other speakers are fixed. Similar to Viterbi, we store also maximizing arguments \tilde{v}_t for all messages. These are used to recover the MAP state sequences in the manner analogous to the Viterbi algorithm. The whole process is repeated until it converges or some number of iterations is reached. Although there is no guarantee that the LBP finds the global MAP configuration, it is known to work well in practice [16]. The time complexity of proposed LBP inference is $O(TKV^2)$, and thus it scales linearly w.r.t. number of speakers.

3. Experiments

3.1. TIDIGITS 2 mixture data

To evaluate the performance of proposed architectures, we conducted experiments on artificially mixed TIDIGITS dataset [17]. Each mixture consists of exactly two speakers randomly selected from original TIDIGITS. Each utterance is a sequence of pronounced digits, the length may vary from 1 to 7 digits. The training data contains 6 328 different pairs of speakers, 52.5 hours of speech. The evaluation data comprises of 4 821 different speaker pairs, in total 5.3 hours of audio. More details about the dataset are available on GitHub². We use this simple dataset to allow for easy analysis of the method and fast experimentation. For discussion on scaling up to larger data, see Section 5.

3.2. Acoustic model

Both the PIT-ASR with separate decoding and the proposed ASR with joint decoding employ a DNN-based acoustic model. Both models are trained in PyTorch [18] using permutation invariant training (PIT)[8, 4] with cross entropy (CE) loss function. We can use PIT-CE objective function for both approaches, since, for every frame, both neural networks predict probabilities of being in HMM state v_t^a and v_t^b although, the state probabilities are modeled in different manner.

Both networks shares the same architecture: TDNN [19] layers with 384 units followed by batch normalization and ReLU activation function. We compared neural networks with 5 or 10 TDNN layers. The neural network for separate decoding contains 2 output layers with size of 62 each. On the other hand, the neural network used with joint decoding consist of single output layer with the size of 62×62 , i.e. 3844. Note that the size of joint network output layer grows exponentially w.r.t. number of speakers. We discuss how this issue can be addressed in Section 5. The input consists of 40-dimensional MFCCs.

The frame-level labels (i.e. state index for each time frame) for each source l_t^a and l_t^b are obtained by applying the force-alignment on original sources \mathbf{x}^a and \mathbf{x}^b . We generate the labels by mono-phone GMM-HMM models, taken from Kaldi [20] recipe for TIDIGITS dataset. For the conventional separate PIT model, the CE loss target probabilities are $[\hat{p}(v_t^a | y_t), \hat{p}(v_t^b | y_t)]$ where $\hat{p}(v_t^a = i | y_t) = \delta_{l_t^a, i}$, where $\delta_{l_t^a, i} = 1$ if training label $l_t^a = i$ and is zero otherwise. For the joint decoding model, we derive the target probabilities as $\hat{p}(v_t^a = i, v_t^b = j | y_t) = \delta_{l_t^a, i} \delta_{l_t^b, j}$.

²https://github.com/MartinKocour/TIDIGITS_mix

3.3. Decoding network

The recognition network is similar to WFST used in Kaldi TIDIGITS setup³ for monophone GMM-HMM ASR system. The network is based on unigram LM, where each word, i.e. digit, is equally likely. Pronunciations are taken from CMU Dictionary⁴. We model silence with 5 HMM states and the other 19 phones with 3 HMM states. The HMM states do not share emission probabilities, i.e. we use 62 distinct PDFs in total. The network is implemented in Julia [21] using the MarkovModels toolkit⁵.

4. Results

We evaluate the speech recognition systems in all experiments using Word error rate (WER). To pair the multi-talker hypotheses with the references, we use the oracle permutation with minimal WER.

4.1. Comparison of joint and separate decoding

We evaluate the performance of separate and joint decoding on 5.3 hours of 2-talker mixed speech. We use two different sizes of TDNN with 5 layers and 10 layers. In Table 1, we compare the results obtained with both models using *separate* and *joint* decoding. For a more fair comparison, we include third *separate-marginal* method marked with Σ -symbol, which combines the AM predicting joint posteriors (Sec. 2.2) with separate decoding (Sec. 2.1). For this, we marginalize the joint posteriors for each speaker as $p(v_i^a | y_i) = \sum_{v_i^b} p(v_i^a, v_i^b | y_i)$. This allows us to separately evaluate the benefit of the joint posterior output (which also induces increased number of parameters) and the benefit of the joint decoding itself.

By comparing the results of *separate* decoding in rows 1 and 2 for the smaller model, and 3 and 4 for the larger model, we can see the advantage of the joint posteriors on the output of the acoustic model. More importantly, comparing the *separate-marginal* decoding with the *joint* decoding, we can see improvements for both smaller and larger models, demonstrating the benefit of decoding both speakers jointly in the factorial framework.

We also compare the performance of Kaldi decoder with our implementation of separate decoding. The results slightly differ when marginalized joint posteriors are used, which we believe is caused by pruning. In our approach, we do not use any kind of pruning.

4.2. Analysis of 2-talker speech recognition results

We further analyzed the WER differences between separate and joint decoder on the speech mixtures of same or opposite genders. The results are shown in Table 2. We observe that the proposed joint decoding greatly improves performance on same gender mixtures, where it may be particularly difficult to distinguish speakers. We hypothesize that thanks to employing dictionary and grammar (in recognition network) to “separate” speech, the proposed method can better discriminate the speakers than the separate decoding scheme performing “separation” only based on the acoustic information.

³The recipe is available on GitHub <https://github.com/kaldi-asr/kaldi/tree/master/egs/tidigits>

⁴The dictionary is available on <https://github.com/cmuspinx/cmudict>

⁵The toolkit is available on GitHub <https://github.com/lucasondel/MarkovModels.jl>

Table 1: Comparison of separate and joint decoding in terms of WER. The results marked with Σ -symbol are computed from marginalized joint posteriors.

Arch	Output [dim]	#params	Separate	Joint	Kaldi
5L-TDNN	[62, 62]	1.9 M	26.09	-	25.99
5L-TDNN	[3844]	3.3 M	17.55 Σ	15.79	17.83 Σ
10L-TDNN	[62, 62]	4.1 M	18.68.	-	18.66
10L-TDNN	[3844]	5.5 M	16.36 Σ	14.70	16.97 Σ

Table 2: Comparison of joint and separate decoder on the mixed speech of two female speakers ($F + F$), two male speakers ($M + M$), same ($F + F \vee M + M$) and opposite gender speakers ($F + M$).

Genders	Separate [%WER]	Joint [%WER]
F + F	30.54	21.45
M + M	32.61	27.12
same	31.56	24.26
opposite	6.17	5.42

5. Discussion and Future work

In this study we analyzed the performance on rather simple task as a proof-of-concept of the joint decoding scheme. The results suggest that considering joint decoding could greatly improve recognition performance of multi-speaker ASR. These results come at the expense of increased computational complexity. For example, in our experiments we observed that the real time factor of the decoding on a CPU increased from approximately 0.3 with separate decoding to 0.5 with joint decoding. However, the decoding time would further increase when dealing with larger vocabulary tasks where the number of senones and the computational graph is larger. In future work, we plan to reduce the computational complexity by investigating factorized model for the joint posteriors and implementing fast decoder using pruning, sparse matrices and GPU-based computations[22], which would allow tackling larger tasks.

6. Conclusion

This work investigates the potential of joint-decoding for multi-talker speech recognition, in contrast with separate decoding in conventional approaches. The proposed method is based on early multi-talker speech recognition works, which use generative factorial approaches. Our work extends this ideas by replacing the acoustic model with DNN. The proof-of-concept experiments showed that the approach has potential to improve performance in challenging conditions where it may be difficult to achieve high separation by simply relying on the acoustic information (e.g. potentially when dealing with noisy and reverberant mixtures). We hope that the findings of this study will promote the importance of joint decoding for future research in multi-talker ASR.

7. References

- [1] J. Barker, S. Watanabe, E. Vincent, and J. Trmal, "The Fifth CHiME Speech Separation and Recognition Challenge: Dataset, Task and Baselines," in *Proc. Interspeech 2018*, 2018, pp. 1561–1565. [Online]. Available: <http://dx.doi.org/10.21437/Interspeech.2018-1768>
- [2] S. Watanabe, M. Mandel, J. Barker, E. Vincent, A. Arora, X. Chang, S. Khudanpur, V. Manohar, D. Povey, D. Raj, D. Snyder, A. S. Subramanian, J. Trmal, B. B. Yair, C. Boeddeker, Z. Ni, Y. Fujita, S. Horiguchi, N. Kanda, T. Yoshioka, and N. Ryant, "CHiME-6 Challenge: Tackling Multispeaker Speech Recognition for Unsegmented Recordings," in *Proc. The 6th International Workshop on Speech Processing in Everyday Environments (CHiME 2020)*, 2020, pp. 1–7. [Online]. Available: <http://dx.doi.org/10.21437/CHiME.2020-1>
- [3] M. Cooke, J. R. Hershey, and S. J. Rennie, "Monaural speech separation and recognition challenge," *Computer Speech & Language*, vol. 24, no. 1, pp. 1–15, 2010.
- [4] 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. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0167639317304193>
- [5] C. Weng, D. Yu, M. L. Seltzer, and J. Droppo, "Deep neural networks for single-channel multi-talker speech recognition," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 23, no. 10, pp. 1670–1679, 2015.
- [6] S. J. Rennie, J. R. Hershey, and P. A. Olsen, "Single-channel multi-talker speech recognition," *IEEE Signal Processing Magazine*, vol. 27, no. 6, pp. 66–80, 2010.
- [7] J. Hershey, T. Kristjansson, S. Rennie, and P. A. Olsen, "Single channel speech separation using factorial dynamics," in *Advances in Neural Information Processing Systems*, B. Schölkopf, J. Platt, and T. Hoffman, Eds., vol. 19. MIT Press, 2007. [Online]. Available: <https://proceedings.neurips.cc/paper/2006/file/6788076842014c83cedadbe6b0ba0314-Paper.pdf>
- [8] M. Kolbæk, D. Yu, Z.-H. Tan, and J. Jensen, "Multitalker speech separation with utterance-level permutation invariant training of deep recurrent neural networks," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 25, no. 10, pp. 1901–1913, 2017.
- [9] X. Chang, Y. Qian, and D. Yu, "Monaural multi-talker speech recognition with attention mechanism and gated convolutional networks," in *INTERSPEECH*, 2018, pp. 1586–1590.
- [10] N. Kanda, R. Ikeshita, S. Horiguchi, Y. Fujita, K. Nagamatsu, X. Wang, V. Manohar, N. E. Y. Soplín, M. Maciejewski, S.-J. Chen *et al.*, "The Hitachi/JHU CHiME-5 system: Advances in speech recognition for everyday home environments using multiple microphone arrays," *Proc. CHiME-5*, pp. 6–10, 2018.
- [11] J. R. Hershey, S. J. Rennie, P. A. Olsen, and T. T. Kristjansson, "Super-human multi-talker speech recognition: A graphical modeling approach," *Computer Speech & Language*, vol. 24, no. 1, pp. 45–66, 2010.
- [12] G. Hinton, L. Deng, D. Yu, G. E. Dahl, A.-r. Mohamed, N. Jaitly, A. Senior, V. Vanhoucke, P. Nguyen, T. N. Sainath *et al.*, "Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups," *IEEE Signal processing magazine*, vol. 29, no. 6, pp. 82–97, 2012.
- [13] S. J. Rennie, J. R. Hershey, and P. A. Olsen, "Hierarchical variational loopy belief propagation for multi-talker speech recognition," in *2009 IEEE Workshop on Automatic Speech Recognition & Understanding*. IEEE, 2009, pp. 176–181.
- [14] Z. Ghahramani and M. I. Jordan, "Factorial Hidden Markov Models," *Machine Learning*, vol. 29, no. 2, pp. 245–273, 1997. [Online]. Available: <https://doi.org/10.1023/A:1007425814087>
- [15] C. M. Bishop, *Pattern Recognition and Machine Learning*. Springer, 2006.
- [16] Y. Weiss and W. Freeman, "On the optimality of solutions of the max-product belief-propagation algorithm in arbitrary graphs," *IEEE Transactions on Information Theory*, vol. 47, no. 2, pp. 736–744, 2001.
- [17] R. Leonard, "A database for speaker-independent digit recognition," in *ICASSP '84. IEEE International Conference on Acoustics, Speech, and Signal Processing*, vol. 9, 1984, pp. 328–331.
- [18] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, "Pytorch: An imperative style, high-performance deep learning library," in *Advances in Neural Information Processing Systems 32*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, Eds. Curran Associates, Inc., 2019, pp. 8024–8035. [Online]. Available: <http://papers.nips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>
- [19] V. Peddinti, D. Povey, and S. Khudanpur, "A time delay neural network architecture for efficient modeling of long temporal contexts," in *Proc. Interspeech 2015*, 2015, pp. 3214–3218.
- [20] 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 *IEEE 2011 Workshop on Automatic Speech Recognition and Understanding*. IEEE Signal Processing Society, Dec. 2011, iEEE Catalog No.: CFP11SRW-USB.
- [21] J. Bezanson, A. Edelman, S. Karpinski, and V. B. Shah, "Julia: A fresh approach to numerical computing," *SIAM review*, vol. 59, no. 1, pp. 65–98, 2017. [Online]. Available: <https://doi.org/10.1137/141000671>
- [22] L. Ondel, L.-M. Lam-Yee-Mui, M. Kocour, C. F. Corro, and L. Burget, "GPU-accelerated forward-backward algorithm with application to lattice-free MMI," in *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022, pp. 8417–8421.