



Speaker-Targeted Audio-Visual Models for Speech Recognition in Cocktail-Party Environments

Guan-Lin Chao¹, William Chan¹, Ian Lane^{1,2}

Carnegie Mellon University

¹Electrical and Computer Engineering, ²Language Technologies Institute

{guanlinchao, williamchan, lane}@cmu.edu

Abstract

Speech recognition in cocktail-party environments remains a significant challenge for state-of-the-art speech recognition systems, as it is extremely difficult to extract an acoustic signal of an individual speaker from a background of overlapping speech with similar frequency and temporal characteristics. We propose the use of speaker-targeted acoustic and audio-visual models for this task. We complement the acoustic features in a hybrid DNN-HMM model with information of the target speaker's identity as well as visual features from the mouth region of the target speaker. Experimentation was performed using simulated cocktail-party data generated from the GRID audio-visual corpus by overlapping two speakers's speech on a single acoustic channel. Our audio-only baseline achieved a WER of 26.3%. The audio-visual model improved the WER to 4.4%. Introducing speaker identity information had an even more pronounced effect, improving the WER to 3.6%. Combining both approaches, however, did not significantly improve performance further. Our work demonstrates that speaker-targeted models can significantly improve the speech recognition in cocktail-party environments.

Index Terms: Automatic Speech Recognition, Speaker-Specific Modelling, Audio-Visual Processing, Cocktail Party Effect

1. Introduction

Automatic Speech Recognition (ASR) in cocktail-party environments aims to recognize the speech of an individual speaker from a background containing many concurrent voices, and has attracted researchers for decades [1, 2]. Current ASR systems can decode clear speech well in relatively noiseless environments. However, in a cocktail-party environment, their performance is severely degraded in the presence of loud noise or interfering speech signals, especially when the acoustic signal of the speaker of interest and the background share similar frequency and temporal characteristics [3]. Some previous approaches to this problem can be: *multimodal robust features* and *blind signal separation*, or a hybrid of both.

In ASR systems, it is common to adapt a well-trained, general acoustic model to new users or environmental conditions. [4] proposed to supply speaker identity vectors, *i-vectors* as input features to a deep neural network (DNN) along with acoustic features. [5] extended [4] by factorizing *i-vectors* to represent speaker as well as acoustic environment. [6] trained speaker-specific parameters jointly with acoustic features in an adaptive DNN-hidden Markov model (DNN-HMM) for word recognition. [7, 8, 9] proposed training speaker-specific discriminant features (referred to as *speaker codes* and *bottle-*

neck features) for fast DNN-HMM speaker adaptation in speech recognition. [10] extended the speaker codes approach to convolutional neural network-HMM (CNN-HMM) systems. [11] investigated different NN architectures of learning *i-vectors* for input feature mapping.

Inspired by humans' ability to use other sensory information like visual cues and knowledge about the environment to recognize speech, research in audio-visual ASR has also demonstrated the advantage of using audio-visual features over audio-only features in robust speech recognition. The McGurk effect was introduced in [12], which illustrates that visual information can affect human's interpretation of audio signals. In [1], low dimensional lip movement vectors, *eigenlips*, were used to complement acoustic features for ASR. In [13], generalized versions of HMMs, factorial HMM and the coupled HMM, were used to fuse auditory and visual information, in which the HMM parameters were able to be trained with dynamic Bayesian networks. In [14], the authors proposed a DNN-based approach to learning multimodal features and a shared representation between modalities. In [15], the authors presented a deep neural network that used a bilinear softmax layer to account for class specific correlations between modalities. In [16], a deep learning architecture with multi-stream HMM model was proposed. Using noise-robust acoustic features extracted by autoencoders and mouth region of interest (ROI) image features extracted by CNNs, this approach achieved higher word recognition rate than the use of non-denoised features or normal HMMs. [17] proposed an active appearance model-based approach to extracting visual features of jaw and lip ROI on four image streams which were then combined with acoustic features for in-car audio-visual ASR.

Traditional cocktail-party ASR methods suggest performing blind signal separation prior to auditory speech recognition of individual signals. Blind signal separation aims at estimating multiple unknown sources from the sensor signals. When there is only a single-channel signal available, source separation on the cocktail-party problem becomes even more difficult [18]. A main assumption in the signal separation is that speech signals from different sources are statistically independent [3]. Another common assumption of signal separation is that all the sources have zero-mean and unit variance for the convenience of performing Independent Component Analysis [19, 20]. However, these two assumptions are not always correct in practice. Therefore, we try to lift these assumptions by directly recognizing single-channel signals of overlapping speech in this work.

In this paper, we propose a speaker-targeted audio-visual ASR model of multi-speaker acoustic input signals in cocktail-party environments without the use of blind signal separation. With the term *speaker-targeted* model, we refer to a speaker-

independent model with speaker identity information input. We complement the acoustic features with information of the target speaker's identity in embeddings similar to i-vectors in [4], along with raw pixels of the target speaker's mouth ROI images, to supply multimodal input features to a hybrid DNN-HMM for speech recognition in cocktail-party environments.

2. Model

In this work, we focus on the cocktail-party problem with overlapping speech from two speakers. We approach this problem using DNN acoustic models with different combinations of additional modalities: visual features and speaker identity information. The acoustic features are filterbank features extracted from the audio signals where two speakers' speech is mixed on a single acoustic channel. The visual features are raw pixel values of the mouth ROI images of the target speaker whose speech the system is expected to recognize. The speaker identity information is represented by the target speaker's ID-embedding. Details about feature extraction are described in the Experiments section.

DNN acoustic models have been widely and successfully used in ASR [21]. Let \mathbf{x} be a window of acoustic frames (i.e., context of filterbanks), the standard DNN acoustic models model the posterior probability:

$$p(y|\mathbf{x}) = \text{DNN}(\mathbf{x}) \quad (1)$$

where y is a phoneme label or alignment (i.e., from GMM-HMM) and DNN is a deep neural network with softmax outputs. The DNN is typically trained to maximize the log probability of the phoneme alignment or minimize the cross-entropy error. However, this optimization problem is difficult when $\mathbf{x} = \mathbf{x}_1 + \mathbf{x}_2$ is a superposition of two signals \mathbf{x}_1 and \mathbf{x}_2 (i.e., cocktail party).

In this work, we extend the previous traditional DNN acoustic models to leverage additional information in order to model our phonemes. By leveraging combinations of the visual features and speaker identity information, the standard DNN acoustic model is extended to have multimodal inputs.

We train the DNN acoustic models with four possible combinations of input features, audio-only, audio-visual, speaker-targeted audio-only, and speaker-targeted audio-visual, in two steps: *speaker-independent models training* followed by *speaker-targeted models training*. The details of the two steps are described in the following sub-sections.

2.1. Two-Speaker Speaker-Independent Models

First, we leverage the visual information in conjunction with the acoustic features. The standard DNN acoustic model:

$$p(y|\mathbf{x}) = \text{DNN}_A(\mathbf{x}) \quad (2)$$

with additional input of visual features becomes:

$$p(y|\mathbf{x}, \mathbf{w}) = \text{DNN}_{AV}(\mathbf{x}, \mathbf{w}) \quad (3)$$

where \mathbf{w} are the visual features. In this step, speaker-independent audio-only model DNN_A and audio-visual model DNN_{AV} are trained for the two-speaker cocktail-party problems. The acoustic and visual features are concatenated directly as DNN inputs for the audio-visual model. The speaker-independent models are illustrated in Figure 1, where the figure without the dashed arrow represents the audio-only model, and the figure with the dashed arrow represents the audio-visual model.




Figure 1: *Speaker-Independent Models.* This figure illustrates a DNN architecture, where the phoneme labels are modeled in the output layer. The arrow connecting the visual features (mouth ROI pixels) and the input layer is a dashed arrow. The speaker-independent audio-only model is illustrated without the dashed arrow. The speaker-independent audio-visual model is illustrated with the dashed arrow, where the acoustic features (filterbank features) and video features are concatenated as DNN inputs.

2.2. Two-Speaker Speaker-Targeted Models

Secondly, we try to leverage the speaker identity information to extend the previous models, DNN_A and DNN_{AV} . DNN_A is extended to:

$$p(y|\mathbf{x}, \mathbf{z}) = \text{DNN}_{AI}(\mathbf{x}, \mathbf{z}) \quad (4)$$

and DNN_{AV} is extended to:

$$p(y|\mathbf{x}, \mathbf{w}, \mathbf{z}) = \text{DNN}_{AVI}(\mathbf{x}, \mathbf{w}, \mathbf{z}) \quad (5)$$

where \mathbf{z} are the speaker identity information. In this step, we adapt the audio-only and audio-visual speaker-independent models to speaker-targeted models respectively (i.e., from DNN_A to DNN_{AI} and from DNN_{AV} to DNN_{AVI}) by hinting the network which target speaker to attend to by supplying speaker identity information as input. The speaker identity information is represented by an embedding that corresponds to the target speaker's ID. We investigate three ways to fuse the audio-visual features with the speaker identity information:

- (A) Concatenating the speaker identity directly with audio-only and audio-visual features.
- (B) Mapping speaker identity into a compact but presumably more discriminative embedding and then concatenating the compact embedding with audio-only and audio-visual features.
- (C) Connecting the speaker identity to a later layer than audio-only and audio-visual features.

The three fusion techniques introduce the three variants (A), (B) and (C) of both the speaker-targeted models DNN_{AI} and DNN_{AVI} . The speaker-targeted models of the three invariants are shown in Figure 2, where the figures without the dashed arrow represent the audio-only models, and the figures with the dashed arrow represent the audio-visual models.

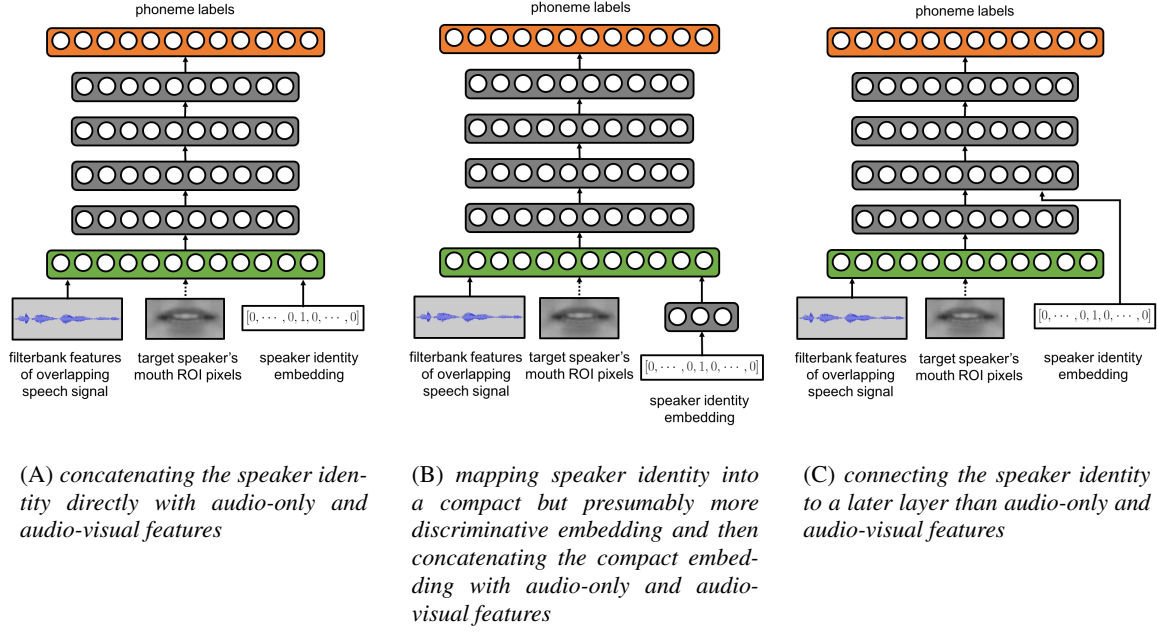


Figure 2: *Three Variants of Speaker-Targeted Models.* These figures illustrate three fusion techniques of audio-visual features with speaker identity information in a DNN architecture, where the phoneme labels are modeled in the output layer. The arrows connecting the visual features and the input layers are dashed arrows. The speaker-targeted audio-only models are illustrated without the dashed arrows. The speaker-targeted audio-visual models are illustrated with the dashed arrows, where the acoustic and video features are concatenated as DNN inputs.

Moreover, we train single-speaker speaker-independent models in comparison with the two-speaker speaker-independent models. We also train 6 randomly-selected speaker's speaker-dependent models (adapted from speaker-independent models as well) to compare with the speaker-targeted models.

3. Experiments

3.1. Dataset

The GRID corpus [22] is a multi-speaker audio-visual corpus. This corpus consists of high-quality audio and video recordings of 34 speakers in quiet and low-noise conditions. Each of the speakers read 1000 sentences which are simple six-word commands obeying the following syntax:

\$command \$color \$preposition \$letter \$digit \$adverb

We use the utterances of 31 speakers (16 males and 15 females) from the GRID corpus, excluding speaker 2, 21 and 28 and part of the utterances of the remaining 31 speakers due to the availability of mouth ROI image data. In the one-speaker datasets, there are 15395 utterances in the training set, 548 in the validation set, and 540 in the testing set, following the convention of CHiME Challenge [23]. The GRID corpus utterances that don't belong to the one-speaker datasets are termed background utterance set. To simulate the overlapping speech audios for the two-speaker datasets, we mix the target speaker and a background speaker's utterances with equal weights on a single acoustic channel using SoX software [24]. The background speaker's utterances are randomly selected from the background utterance set excluding the utterances of the target speaker. The resulting mixed audio's length is as long as the length of the target speaker's utterance.

3.2. Feature Extraction

3.2.1. Audio Features

Log-mel filterbank features with 40 bins are extracted, and a context of ± 5 frames was used for audio input features (i.e., $440 = 40 * (5 + 1 + 5)$ dimensions per acoustic feature \mathbf{x}).

3.2.2. Visual Features

We use target speaker's mouth ROI images' pixel values as visual features. The facial landmarks are first extracted by IntraFace software [25], and each video frame is cropped into a 60 pixel * 30 pixel mouth ROI image [26] according to the mouth region landmarks (i.e., $1800 = 60 * 30$ dimensions per visual feature \mathbf{w}). The gray-scale pixel values are then concatenated with audio features to form audio-visual features.

3.2.3. Speaker Identity Information

Speaker identity information is represented by the target speaker's ID-embedding, which is simply a one-hot vector of thirty-three 0s and a single 1, $[0, \dots, 0, 1, 0, \dots, 0]$, in which the entry of 1 corresponds to the target speaker's ID (i.e., 34 dimensions per speaker identity embedding \mathbf{z}).

3.3. Acoustic Model

Here we describe the architecture of our DNNs. The number of hidden layers for audio-only models and speaker-independent audio-visual models is 4, while it is 5 for speaker-targeted and speaker-dependent audio-visual models. Each hidden layer contains 2048 nodes. Rectified linear function (ReLU) is used for activation in each hidden layer. The output layer has a softmax of 2371 phoneme labels. We use stochastic gradient descent




Figure 3: *WER Comparisons of Two-Speaker models for Individual Speakers.* WER of two-speaker models for individual speakers are illustrated. The dashed line is plotted on the right vertical axis which represents the speaker-independent audio-only model. The solid lines and markers are plotted on the left vertical axis. Speaker-dependent models for speaker 1, 17, 22, 24, 25 and 30 are plotted in markers. The chart demonstrates a similar trend between different models’ performance on individual speakers.

with a batch size of 128 frames and a learning rate of 0.01.

3.4. Results

Table 1: *WER Comparisons of Single-Speaker Models*

	audio-only	audio-visual
speaker-independent	0.3%	0.4%

Table 2: *WER Comparisons of Two-Speaker Models*

	audio-only	audio-visual
speaker-independent	26.3%	4.4%
speaker-targeted A	4.0%	3.6%
speaker-targeted B	3.6%	3.9%
speaker-targeted C	4.4%	4.4%
speaker-dependent	3.9%	3.4%

The aforementioned single-speaker models are used to decode the single-speaker testing dataset, while the two-speaker models are used to decode the two-speaker testing dataset. Table 1 shows the WER of single-speaker models. Table 2 shows the WER of two-speaker models. The audio-only baseline for two-speaker cocktail-party problem is 26.3%. The results of speaker-independent models for single-speaker and two-speaker suggest that automatic speech recognizers’ performance degrades severely in cocktail-party environments compared to low-noise conditions. It is also demonstrated that the introduction of visual information to acoustic features can reduce WER significantly in cocktail-party environments, improving the WER to 4.4%, although it may not help when the environmental noise is low. WER comparisons between two-speaker’s audio-only speaker-independent and speaker-targeted models suggest that using speaker identity information in conjunction with acoustic features achieves a better improvement on WER, reducing WER up to 3.6%.

The results of two-speaker’s speaker-targeted models A, B, and C suggest a weak tendency that providing speaker information in earlier layers of the network seems to have advantage. WER comparisons between two-speaker speaker-dependent and speaker-targeted models suggest an intuitive result that a speaker-dependent ASR system which is optimized for one specific speaker performs better than a speaker-targeted ASR system which is optimized for multiple speakers simulta-

neously. We also find the introduction of visual information improves the WER of speaker-dependent acoustic models while it doesn’t improve the speaker-targeted acoustic models. We subscribe this finding to the limitation of the capacity of the neural network architecture that we use for both models, that it is able to optimize for one specific speaker’s visual information in a speaker-dependent model, but not powerful enough to learn a unified optimization for all 31 speakers’ visual information in a single speaker-targeted model. Figure 3 illustrates the WER of the individual speakers. A similar trend between different models’ performance on individual speakers is demonstrated.

4. Conclusions

A speaker-targeted audio-visual DNN-HMM model for speech recognition in cocktail-party environments is proposed in this work. Different combinations of acoustic and visual features and speaker identity information as DNN inputs are presented. Experimental results suggest that the audio-visual model achieves significant improvement over the audio-only model. Introducing speaker identity information introduces an even more pronounced improvement. Combining both approaches, however, does not significantly improve performance further.

Future work will aim to investigate better representations in multimodal data space to incorporate audio, visual and speaker identity information with the objective to improve the speech recognition performance in cocktail-party environments. More complex architectures can be explored such as CNNs for modeling image structures and recurrent neural networks (RNNs) and Long-Short Term Memory (LSTM) models for modeling variable time-sequence inputs in order to achieve a better unified optimization for the speaker-targeted audio-visual models.

5. Acknowledgements

We would like to thank Niv Zehngut for preparing the image dataset, Srikanth Kallakuri for technical support, Benjamin Elizalde and Akshay Chandrashekar for proofreading and suggestions.

6. References

- [1] C. Bregler and Y. Konig, "eigenlips for robust speech recognition," in *Acoustics, Speech, and Signal Processing, 1994. ICASSP-94., 1994 IEEE International Conference on*, vol. 2. IEEE, 1994, pp. II-669.
- [2] J. H. McDermott, "The cocktail party problem," *Current Biology*, vol. 19, no. 22, pp. R1024–R1027, 2009.
- [3] S. Choi, H. Hong, H. Glotin, and F. Berthommier, "Multichannel signal separation for cocktail party speech recognition: A dynamic recurrent network," *Neurocomputing*, vol. 49, no. 1, pp. 299–314, 2002.
- [4] G. Saon, H. Soltan, D. Nahamoo, and M. Picheny, "Speaker adaptation of neural network acoustic models using i-vectors," in *ASRU*, 2013, pp. 55–59.
- [5] P. Karanasou, Y. Wang, M. J. Gales, and P. C. Woodland, "Adaptation of deep neural network acoustic models using factorised i-vectors," in *INTERSPEECH*, 2014, pp. 2180–2184.
- [6] J. S. Bridle and S. J. Cox, "Recnorm: Simultaneous normalisation and classification applied to speech recognition," in *NIPS*, 1990, pp. 234–240.
- [7] O. Abdel-Hamid and H. Jiang, "Fast speaker adaptation of hybrid nn/hmm model for speech recognition based on discriminative learning of speaker code," in *Acoustics, Speech and Signal Processing (ICASSP), 2013 IEEE International Conference on*. IEEE, 2013, pp. 7942–7946.
- [8] S. Xue, O. Abdel-Hamid, H. Jiang, L. Dai, and Q. Liu, "Fast adaptation of deep neural network based on discriminant codes for speech recognition," *Audio, Speech, and Language Processing, IEEE/ACM Transactions on*, vol. 22, no. 12, pp. 1713–1725, 2014.
- [9] R. Doddipatla, M. Hasan, and T. Hain, "Speaker dependent bottleneck layer training for speaker adaptation in automatic speech recognition," in *INTERSPEECH*, 2014, pp. 2199–2203.
- [10] O. Abdel-Hamid and H. Jiang, "Rapid and effective speaker adaptation of convolutional neural network based models for speech recognition," in *INTERSPEECH*, 2013, pp. 1248–1252.
- [11] Y. Miao, H. Zhang, and F. Metze, "Speaker adaptive training of deep neural network acoustic models using i-vectors," *Audio, Speech, and Language Processing, IEEE/ACM Transactions on*, vol. 23, no. 11, pp. 1938–1949, 2015.
- [12] H. McGurk and J. MacDonald, "Hearing lips and seeing voices," *Nature*, vol. 264, pp. 746–748, 1976.
- [13] A. V. Nefian, L. Liang, X. Pi, X. Liu, and K. Murphy, "Dynamic bayesian networks for audio-visual speech recognition," *EURASIP Journal on Advances in Signal Processing*, vol. 2002, no. 11, pp. 1–15, 2002.
- [14] J. Ngiam, A. Khosla, M. Kim, J. Nam, H. Lee, and A. Y. Ng, "Multimodal deep learning," in *Proceedings of the 28th international conference on machine learning (ICML-11)*, 2011, pp. 689–696.
- [15] Y. Mroueh, E. Marcheret, and V. Goel, "Deep multimodal learning for audio-visual speech recognition," in *Acoustics, Speech and Signal Processing (ICASSP), 2015 IEEE International Conference on*. IEEE, 2015, pp. 2130–2134.
- [16] K. Noda, Y. Yamaguchi, K. Nakadai, H. G. Okuno, and T. Ogata, "Audio-visual speech recognition using deep learning," *Applied Intelligence*, vol. 42, no. 4, pp. 722–737, 2015.
- [17] A. Biswas, P. Sahu, and M. Chandra, "Multiple cameras audio visual speech recognition using active appearance model visual features in car environment," *International Journal of Speech Technology*, vol. 19, no. 1, pp. 159–171, 2016.
- [18] M. N. Schmidt and R. K. Olsson, "Single-channel speech separation using sparse non-negative matrix factorization," in *Spoken Language Processing, ISCA International Conference on (INTERSPEECH)*, 2006.
- [19] P. Comon, "Independent component analysis, a new concept?" *Signal processing*, vol. 36, no. 3, pp. 287–314, 1994.
- [20] G.-J. Jang and T.-W. Lee, "A probabilistic approach to single channel blind signal separation," in *Advances in neural information processing systems*, 2002, pp. 1173–1180.
- [21] G. Hinton, L. Deng, D. Yu, G. Dahl, A. rahman Mohamed, N. Jaitly, A. Senior, V. Vanhoucke, P. Nguyen, T. Sainath, and B. Kingsbury, "Deep Neural Networks for Acoustic Modeling in Speech Recognition," November 2012.
- [22] M. Cooke, J. Barker, S. Cunningham, and X. Shao, "An audio-visual corpus for speech perception and automatic speech recognition," *The Journal of the Acoustical Society of America*, vol. 120, no. 5, pp. 2421–2424, 2006.
- [23] J. Barker, E. Vincent, N. Ma, H. Christensen, and P. Green, "The pascal chime speech separation and recognition challenge," *Computer Speech & Language*, vol. 27, no. 3, pp. 621–633, 2013.
- [24] "Sox - sound exchange," online; accessed 30-Mar-2016. [Online]. Available: <http://sox.sourceforge.net>
- [25] F. de la Torre, W.-S. Chu, X. Xiong, F. Vicente, X. Ding, and J. Cohn, "Intraface," in *Automatic Face and Gesture Recognition (FG), 2015 11th IEEE International Conference and Workshops on*, vol. 1. IEEE, 2015, pp. 1–8.
- [26] N. Zehngut, "Audio visual speech recognition using facial landmark based video frontalization," 2015, unpublished technical report.