

# Speaker disentanglement in video-to-speech conversion

Dan Oneață

University POLITEHNICA of Bucharest  
Bucharest, Romania  
dan.oneata@speed.pub.ro

Adriana Stan

Technical University of Cluj-Napoca  
Cluj-Napoca, Romania  
adriana.stan@com.utcluj.ro

Horia Cucu

University POLITEHNICA of Bucharest  
Bucharest, Romania  
horia.cucu@upb.ro

**Abstract**—The task of video-to-speech aims to translate silent video of lip movement to its corresponding audio signal. Previous approaches to this task are generally limited to the case of a single speaker, but a method that accounts for multiple speakers is desirable as it allows to (i) leverage datasets with multiple speakers or few samples per speaker; and (ii) control speaker identity at inference time. In this paper, we introduce a new video-to-speech architecture and explore ways of extending it to the multi-speaker scenario: we augment the network with an additional speaker-related input, through which we feed either a discrete identity or a speaker embedding. Interestingly, we observe that the visual encoder of the network is capable of learning the speaker identity from the lip region of the face alone. To better disentangle the two inputs—linguistic content and speaker identity—we add adversarial losses that dispel the identity from the video embeddings. To the best of our knowledge, the proposed method is the first to provide important functionalities such as (i) control of the target voice and (ii) speech synthesis for unseen identities over the state-of-the-art, while still maintaining the intelligibility of the spoken output.

**Index Terms**—video-to-speech conversion, speech synthesis, speaker disentanglement, multispeaker

## I. INTRODUCTION

In this paper we are concerned with learning a mapping from silent lip video of a person talking to its corresponding audio speech signal, while also controlling the speaker identity of the output speech. A compelling application of this task is to enable people who lost their ability to speak to interact with a speech synthesizer in a more personalised, fast and natural way. A fair number of video-to-speech systems have been already proposed [1]–[5], but none of them explicitly models the speaker identity and they can synthesise speech only in a predefined speaker’s voice. When trained on multiple speakers at once, the results usually suffer, especially when attempting to synthesise the voices of unseen speakers [5]. To counter these limitations, we propose ways of explicitly incorporating speaker information. The benefits of this approach are (i) the ability to use larger sets of training data by allowing multiple speakers present in the dataset and (ii) the possibility to control the speaker identity at inference time.

In text-to-speech synthesis (TTS) systems controlling the speaker identity has been extensively explored [6], [7]. The

most common way of conditioning the output on a speaker involves a simple augmentation of the internal representations with a speaker embedding, be it learnt [6] or transferred from an external embedding network [7]. However, whereas the text input to the TTS system is devoid of any speaker information, the same cannot be said about video. Our initial experiments revealed that a speaker-independent network is able to maintain the correct voice for each video independent of any speaker identity conditioning. As a result, the difficulty of our video-to-speech task lies in the fact that the two sources of information—linguistic content and speaker identity—are entangled. Ideally, we want to be able to separately specify the content (*what is being said*) from the video input and the speaker information (*who is speaking*). One simple workaround would be to apply a lip-tracking algorithm, which discards any other visual input. However, by limiting the input to a small set of data points, other essential information is lost as well (e.g. tongue or teeth positions and movement), and it is also prone to additional modelling errors.

Therefore, we attempt to strip the entire speaker information from the video stream neural encoding, and control the speaker through a separate input to the network. Our approach takes inspiration from work on supervised learning of disentangled representations [8], [9]. Our main contributions are: (i) a new video-to-speech architecture that leverages state-of-the-art components; (ii) a network adaptation to incorporate the speaker identity; (iii) a means to improve the network’s disentanglement capabilities for the content and speaker information. Code and samples are available at <https://speed.pub.ro/xts/>.

## II. RELATED WORK

The task of video-to-speech conversion has initially been approached by using rich sensory information (ultra-sound of articulatory movements) [10] and hand-crafted features [11]. Recent literature shifted towards deep end-to-end networks that operate directly at pixel level [1], [3]–[5]. A common aspect of all these works is the heavy use of convolutional layers, and the notable differences rely in the output representation—the main choices to encode the audio are LPC coefficients [1], spectrogram [4], or raw audio [5]. Recently, both Vougioukas *et al.* [5] and Salik *et al.* [12] evaluated the video-to-speech task in a speaker-independent scenario, but without adapting their architectures to explicitly model the speaker identity. The

This work was supported in part by a grant of the Romanian Ministry of Education and Research, CNCS - UEFISCDI, project number PN-III-P1-1.1-PD-2019-0918, within PNCDI III.

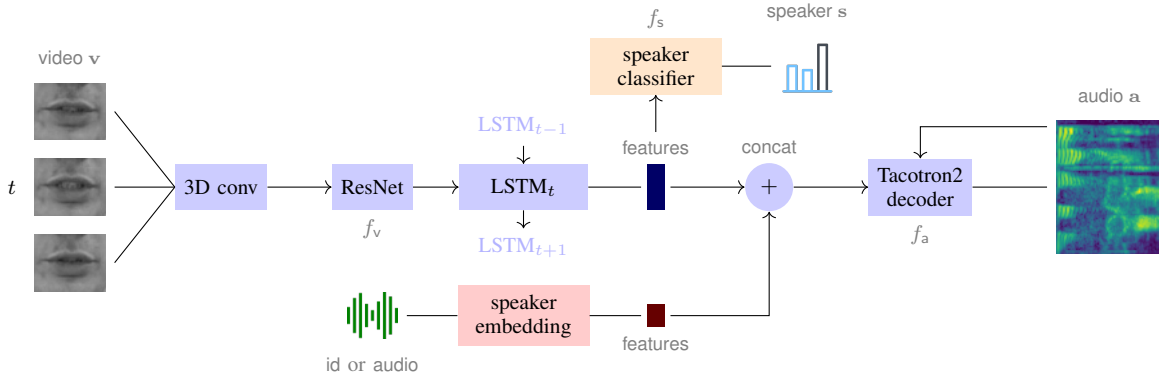


Fig. 1. The proposed video-to-speech architecture. The baseline model (blue) consists of a visual processing front-end (3D conv, ResNet, LSTM) and an audio decoder (Tacotron2). This backbone is augmented with a speaker embedding component (red), which injects speaker information into the decoder (either an identity or an audio sample), and a speaker classifier (orange), which removes speaker information from the visual features and disentangles content and identity.

concurrent work of [13] proposes a multi-speaker architecture that incorporates the speaker embedding, but differently from us they do not employ disentanglement techniques and consider a scenario in which the speaker embeddings are extracted from test utterances to generalize to unseen speakers.

Two other related and well-studied tasks are lip reading and text-to-speech synthesis. In our work we leverage state-of-the-art systems from both domains in order to build the backbone of our architecture. More exactly, for the visual encoder we draw inspiration from the deep lip reading front-end [14] and for the speech decoder we borrow ideas from the Tacotron2 architecture [15]. Both of these systems use words as either output or input, while our work avoids the explicit textual representation, resulting in a simpler, attention-free architecture (since video and audio are time-aligned).

### III. METHOD DESCRIPTION

This section describes our methodology for the video-to-speech task. We first describe a baseline model which we then augment with an auxiliary branch to incorporate speaker information, and an extra loss function to encourage the disentanglement of the speaker identity from the visual features.

**Baseline model.** The base model (B) combines the ResNet architecture [16] (for visual processing) with the Tacotron2 decoder [15] (for speech synthesis). As the visual and audio signals are aligned, we do not use attention modules, but simply upscale temporarily the visual features to match the number of audio frames. To make the model more suitable for video processing, we insert a layer of 3D convolutions and add an LSTM layer over the ResNet features. The architecture is illustrated in Figure 1. The network generates reduced Mel spectrograms which are inverted to audio using the Spectrogram Super-resolution Network (SSRN) [17] combined with a Griffin-Lim [18] vocoder.

**Modeling speaker identity.** In order to control the voice of the synthesized audio, we concatenate the speaker identity to the visual features (obtained after the LSTM) and the augmented vector is then passed to the audio decoder. The speaker information is either: (i) a learned speaker-specific vector based on the discrete identity of each speaker (SI) or

(ii) an utterance-level embedding obtained from a pre-trained speaker recognition network (SE). For the SE method, we extract embeddings using the state-of-the-art network of Xie *et al.* [19]. We freeze the speaker recognition network and project the speaker embeddings from 512 dimensions to 32 using a learned linear transformation. The dimensionality reduction step ensures that the SE and SI architectures have a comparable number of parameters. We found that the scale of the speaker embeddings affects the behaviour of the network. In particular, for SE, the original scale of the embeddings is small and the model tends to ignore the auxiliary speaker input, relying only on the information available in the visual input. Hence, we experiment with a variant (SE-norm), which standardizes the speaker embedding (normalize to zero mean and unit variance).

**Disentangling identity from content.** We hypothesize that we can control the speaker identity better if the intermediate embedding (fed to the audio decoder) contains the content and speaker information disentangled: ideally, the visual features should reflect only the content of what is being said, while the auxiliary speaker embedding should contain only speaker information. However, as mentioned in the introduction, we observed that the visual front-end consistently extracts speaker information. We change the architecture to *dispel* the speaker information from the visual features. The idea is to train a speaker classifier on top of the features extracted from the visual processing network, and use the classifier’s predictions to update the extraction process. We consider two approaches. The first approach is based on adversarial learning: a discriminator learns to classify speakers based on visual features, while the generator changes the visual features to fool the discriminator and still be able to reconstruct the original audio. Concretely, we optimize the following two losses:

$$L_d(f_s) = H(s, (f_s \circ f_v)(\mathbf{v})) \quad (1)$$

$$L_g(f_a, f_v) = \|\mathbf{a} - (f_a \circ f_v)(\mathbf{v})\|_2^2 - \lambda H((f_s \circ f_v)(\mathbf{v})) \quad (2)$$

where  $\mathbf{v}$  denotes the input video,  $\mathbf{a}$  the target audio, and  $s$  the speaker identity;  $f$  are the neural network functions, with the subscripts denoting the different components:  $f_v$  the video processing net,  $f_a$  the audio decoder net, and  $f_s$  the speaker

classifier.  $H$  denotes the cross-entropy (Eq. 1) or entropy (Eq. 2) and  $\lambda$  represents the weight of the speaker classifier loss. Similar losses have been proposed for face generation [8], but our generator maximizes the entropy of the predictions instead of minimizing the cross-entropy with respect to the true speaker identity.

The second approach to disentanglement uses the idea of gradient reversal [20]. We optimize for an objective function involving audio reconstruction and speaker classification:

$$L(f_a, f_s, f_v) = \|a - (f_a \circ f_v)(v)\|_2^2 + \lambda H(s, (f_s \circ f_v)(v)) \quad (3)$$

where the notations are the same as in Eqs. 1 and 2. The gradient reversal step updates the visual features such that they become indistinguishable in terms of speaker identity.

Both approaches rely on a speaker classifier  $f_s$  to predict the speaker identity for the entire input video sequence. We used two speaker classification variants, which differ in terms of complexity. The simpler variant (denoted as *linear*) performs average pooling across the temporal sequence of visual features and then applies a linear classifier. The more complex variant (denoted as *MLP*) uses an additional two-layer network before average pooling.

#### IV. EXPERIMENTAL RESULTS

**Dataset.** We carried our experiments on the GRID corpus [21], the test bed for the video-to-speech task [1], [3], [5]. The dataset consists of 34,000 video-audio samples coming from 34 different speakers. The vocabulary is constrained to 52 words, but no two samples contain the exact same sequence of words.

**Evaluation metrics.** We use multiple metrics to assess the desired properties of video-to-speech systems: *quality*, *intelligibility*, and *speaker identity*. To measure the *quality* of the generated signal, we employ Mel-cepstral distortion (MCD) and perceptual evaluation of speech quality (PESQ). For *intelligibility*, we report the short-term objective intelligibility (STOI) and word error rate (WER) between the reference transcription and the output of an automatic speech recognition (ASR) system on the synthesised audio. The ASR is implemented as a time-delay neural network [22] in Kaldi and it was trained on the TED-LIUM 2 dataset [23]; the language model uses a finite state grammar derived from the word sequences present in the GRID dataset. The performance of the ASR system is around 2.8% WER on the natural speech samples. Finally, to measure the *identity* of the speaker we compute the equal error rate (EER) on pairs of audio—natural and synthesized.

**Implementation details.** To improve the conditioning on the input and make the model less dependent on the auto-regressive signal, we experimented with a dropout-like mechanism inspired by [24]. The idea is to randomly replace a fraction of the audio frames on which we condition (20% in our case) with a fixed frame—the mean of all the frames over the training set. The weight  $\lambda$  for the disentanglement losses was set to  $10^{-4}$ . Other hyper-parameters (learning rate, parameters for the

TABLE I  
EVALUATION ON SEEN SPEAKERS FROM THE GRID CORPUS FOR THREE OF OUR METHODS: BASELINE TRAINED ON ALL SPEAKERS (B), BASELINE TRAINED PER SPEAKER (B-SPK), SPEAKER INDEPENDENT MODEL (SI). THE RESULTS MARKED WITH † DENOTE THAT THEY WERE RECOMPUTED. ARROWS INDICATE THE DIRECTION OF BETTER PERFORMANCE FOR EACH METRIC.

	STOI ↑	PESQ ↑	MCD ↓	WER ↓
Lip2AudSpec [3]	0.446	1.82	38.14	32.5
V2S GAN [5]	0.518	1.71	22.29	26.6
V2S GAN [5]†	<b>0.525</b>	1.72	<b>22.02</b>	27.1
B	0.470	<b>1.88</b>	32.28	21.8
B-spK	0.452	1.82	32.42	<b>17.8</b>
SI	0.468	1.85	32.08	19.9

3D convolution and LSTM layers) were chosen by a random search procedure on the validation split. The SSRN vocoder was trained on 17 hours of speech data collected from multiple speakers, but not including the GRID speakers.

##### A. Evaluation of video-to-speech synthesis

We evaluate our proposed methods for the task of video-to-speech and compare to previously proposed methods [3], [5]. We consider the speaker dependent setup of [5], which consists of four speakers, each with 900 samples for training, 50 for validation and 50 for testing. We evaluate three variants of our methods: a speaker-independent baseline trained on all four speakers at once (B), a speaker-dependent baseline trained for each speaker separately (B-spK), and a model trained on all four speakers at once, but which explicitly incorporates the speaker identity (SI). The quantitative results are presented in Table I, while qualitative samples can be found online.<sup>1</sup>

Compared to previous work, our methods obtain competitive results with respect to both speech quality and content intelligibility metrics, yielding best results in terms of PESQ and WER. The performance of our proposed methods is similar, but they come with a different set of trade-offs. The B-spK and SI variants model the speaker explicitly, allowing for the additional functionality of controlling the speaker identity of the generated voice. The SI method has the additional benefit of being more efficient than B-spK, which scales linearly in the number of speakers as it trains a separate model for each speaker in the dataset.

Surprisingly, even if the B model is completely oblivious to the explicit speaker identity, it can still produce results on par with the speaker-dependent variants (B-spK and SI). Indeed, the left image of Figure 2 shows that network is able to implicitly model the speakers, due to the visual cues.

##### B. Controlling the speaker identity

In this subsection we present experiments on the control of the speaker identity for the generated audio. The input consists of a video and a target identity, and the desired output is an utterance consisting of the words pronounced in the video, but

<sup>1</sup><https://speed.pub.ro/xts/>

in the voice of the selected identity. As a qualitative example, we show in Figure 2 (right) the speaker embeddings extracted from the audio generated by running the SI method with a fixed embedding (the average speaker embedding across all speakers). As expected, even if the videos belong to different persons, the synthesised audio roughly pertains to the same speaker, which means that the average embedding can indeed control the voice identity.

For the quantitative evaluation, we consider a setup where we pair videos of *unseen* speakers with target voice identities of *seen* speakers. This scenario measures the controllability of the output speech in terms of speaker identity, as well as the network’s capability to encode the linguistic content from the video. Because of the random pairing of video and speaker identities, MCD, STOI or PESQ metrics cannot be computed as all three rely on a reference audio signal. Instead, we focus on measuring the content intelligibility and speaker identity. The intelligibility is computed as the WER of the previously described ASR system over the output utterance. The speaker identity of the synthesized speech is evaluated in terms of equal error rate (EER). The EER operating point relies on the false acceptance and rejection rates given by varying a threshold on the cosine distance between the embeddings of the synthesised and natural samples; the embeddings are obtained from the same speaker embedding network used in the training process. This evaluation protocol measures that the embedding of the synthesized speech for a given speaker is (i) close to all embeddings of natural samples for that speaker, and (ii) far from embeddings of any other speaker. We randomly chose 20k pairs of embeddings out of which 1460 are positive (synthesized and natural embeddings pertain to the same speaker), and the rest are negative (the embeddings are from different speakers). A naïve method, which assigns a random distance between 0 and 1 to any pair, would obtain an EER of around 50%.

The results of this evaluation scenario are presented in Table II. Previous work (row 1) and our baseline (B; row 2) generate significantly less intelligible speech for unseen speakers (WER 41.9%) than for seen speakers (Table I: WER 22–27%). Moreover, these two methods are not able to control the speaker identity, while the rest of the architectures presented in Table II can (rows 3–20). In terms of speaker control, SI yields the best results on average, but this architecture is only able to produce speech in the voices seen during training. The other architectures (SE and SE-norm) are capable of producing speech based on a speaker *embedding*. Among these two, the variant which uses normalized speaker embeddings (SE-norm) offers better speaker control, because it ensures that the speaker embeddings have a comparable scale to the visual features.

Analyzing the variants of each of the three speaker-dependent architectures, we observe that dropping frames consistently improves intelligibility, but this improvement is generally obtained at the expense of speaker control (rows 4, 10, 16 in Table II). If we factor in the disentanglement losses (dispel and reverse gradient), we improve the speaker control, as expected, but lose in terms of WER. Among the speaker classifiers




Fig. 2. Speaker embeddings for synthesised audio generated from videos of 14 different speakers (denoted by a different color). On the left is a speaker-independent model (B) trained on 14 speakers, which maintains the voice of the person in the video without using the explicit speaker identity. On the right is a speaker-dependent model (SI), which controls the identity through an input speaker embedding—in the figure all samples are generated in the same voice using a fixed speaker embedding (the mean of all training embeddings).

TABLE II  
WORD ERROR RATE (WER) AND EQUAL ERROR RATE (EER) RESULTS FOR THE SPEAKER IDENTITY CONTROL ON UNSEEN SPEAKERS. ARROWS INDICATE THE DIRECTION OF BETTER PERFORMANCE FOR EACH METRIC.

	Architecture	Drop	Disentanglement	WER ↓	EER ↓
1	V2S GAN [5]	—	—	41.9	N/A
2	B	no	—	41.9	N/A
3	SI	no	—	43.7	6.9
4		yes	—	43.8	7.1
5		yes	dispel	50.2	7.5
6		yes	dispel	43.7	<b>6.8</b>
7		yes	rev. grad.	45.2	6.9
8		yes	rev. grad.	<b>42.7</b>	7.3
9	SE	no	—	36.5	18.0
10		yes	—	<b>31.2</b>	48.6
11		yes	dispel	41.9	<b>7.1</b>
12		yes	dispel	35.5	12.7
13		yes	rev. grad.	37.7	8.9
14		yes	rev. grad.	36.1	13.6
15	SE-norm	no	—	40.6	11.7
16		yes	—	<b>38.7</b>	12.5
17		yes	dispel	49.6	7.8
18		yes	dispel	40.1	10.6
19		yes	rev. grad.	41.5	<b>7.6</b>
20		yes	rev. grad.	38.9	11.9

used for disentanglement, the MLP generally shows better speaker control results, while the linear classifier yields better intelligibility results. Overall, we noticed an on-going trade-off between the two goals (intelligibility versus speaker control), but compared to the baseline model these results hold promise: we are able to maintain the content intelligibility (WER 42.7%), while providing very good speaker control (EER 7.3%), see row 8 in Table II, or we can significantly improve the results in terms of content intelligibility (WER 35.5%), while still obtaining decent speaker control (EER 12.7%), see row 12 in Table II.

**Listening test.** To obtain subjective measures of speaker similarity (corresponding to the objective EER) and intelligibility (corresponding to the objective WER) we conducted a listening test; the test was carried out on eighteen volunteers and using samples from six methods from Table II (which are




Fig. 3. Listening test results for six of our methods (identified by the row number in Table II) and reference samples. We report speaker similarity (SS; evaluated on a MuSHRA scale, 0–100) and intelligibility (WA; evaluated in terms of word accuracy). For both metrics higher values are better.

identified by their row number). The results are presented in Figure 3 and show that the relative ordering of the methods is similar for both the subjective and objective measures. In particular, we observe that the best performing system with respect to each of the two objective measures is best evaluated in the listening test, as well.

## V. CONCLUSIONS

This paper addressed the task of multi-speaker video-to-speech conversion. We highlighted the challenge of entangled content and speaker identity in the video features and proposed ways to mitigate this issue and, consequently, allow for better speaker control. Our methodology extended a baseline video-to-speech architecture with speaker information by inputting speaker identities or speaker embeddings. All these variants maintain or improve the content intelligibility, while allowing the control of the speaker identity. From a practical point of view, this is a highly-desirable functionality, as an impaired user can choose a voice—either an existing voice from the database or even discover new voices by interpolating in the embedding space. We have conducted a realistic and difficult evaluation in which we assumed we encounter new (unseen) speakers at test time. Our ablation study showed that dropping frames consistently improved intelligibility at the expense of speaker control and that the disentanglement losses had a reverse effect. While we observed that intelligibility and speaker control are often at odds, there are still good compromises to be achieved (for example, the variants based on speaker embeddings with disentanglement) and the proposed methods fare well when compared to the speaker-independent model.

## REFERENCES

- [1] Ariel Ephrat and Shmuel Peleg, “Vid2speech: Speech reconstruction from silent video,” in *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2017, pp. 5095–5099.
- [2] Yaman Kumar, Mayank Aggarwal, Pratham Nawal, Shin’ichi Satoh, Rajiv Ratn Shah, and Roger Zimmermann, “Harnessing AI for speech reconstruction using multi-view silent video feed,” in *ACM International Conference on Multimedia*, 2018, pp. 1976–1983.
- [3] Hassan Akbari, Himani Arora, Liangliang Cao, and Nima Mesgarani, “Lip2AudSpec: Speech reconstruction from silent lip movements video,” in *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2018, pp. 2516–2520.
- [4] Yuki Takashima, Tetsuya Takiguchi, and Yasuo Ariki, “Exemplar-based lip-to-speech synthesis using convolutional neural networks,” in *International Workshop on Frontiers of Computer Vision*, 2019.

- [5] Konstantinos Vougioukas, Pingchuan Ma, Stavros Petridis, and Maja Pantic, “Video-driven speech reconstruction using generative adversarial networks,” in *Interspeech*, 2019.
- [6] Andrew Gibiansky, Sercan Arik, Gregory Diamos, John Miller, Kainan Peng, Wei Ping, Jonathan Raiman, and Yanqi Zhou, “Deep voice 2: Multi-speaker neural text-to-speech,” in *Advances in Neural Information Processing Systems*, 2017, pp. 2962–2970.
- [7] Ye Jia, Yu Zhang, Ron Weiss, Quan Wang, Jonathan Shen, Fei Ren, Patrick Nguyen, Ruoming Pang, Ignacio Lopez Moreno, Yonghui Wu, et al., “Transfer learning from speaker verification to multispeaker text-to-speech synthesis,” in *Advances in Neural Information Processing Systems*, 2018, pp. 4480–4490.
- [8] Yu Liu, Fangyin Wei, Jing Shao, Lu Sheng, Junjie Yan, and Xiaogang Wang, “Exploring disentangled feature representation beyond face identification,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 2080–2089.
- [9] Wei-Ning Hsu, Yu Zhang, Ron J Weiss, Yu-An Chung, Yuxuan Wang, Yonghui Wu, and James Glass, “Disentangling correlated speaker and noise for speech synthesis via data augmentation and adversarial factorization,” in *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2019, pp. 5901–5905.
- [10] Thomas Hueber and Gérard Bailly, “Statistical conversion of silent articulation into audible speech using full-covariance HMM,” *Computer Speech & Language*, vol. 36, pp. 274–293, 2016.
- [11] Thomas Le Cornu and Ben Milner, “Reconstructing intelligible audio speech from visual speech features,” in *Interspeech*, 2015.
- [12] Khwaja Mohd Salik, Swati Aggarwal, Yaman Kumar, Rajiv Ratn Shah, Rohit Jain, and Roger Zimmermann, “LipNet: Speaker independent speech synthesis using multi-view lipreading,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, 2019, vol. 33, pp. 10023–10024.
- [13] KR Prajwal, Rudrabha Mukhopadhyay, Vinay P Namboodiri, and CV Jawahar, “Learning individual speaking styles for accurate lip to speech synthesis,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2020, pp. 13796–13805.
- [14] Triantafyllos Afouras, Joon Son Chung, and Andrew Zisserman, “Deep lip reading: A comparison of models and an online application,” in *Interspeech*, 2018, pp. 3514–3518.
- [15] Jonathan Shen, Ruoming Pang, Ron J Weiss, Mike Schuster, Navdeep Jaitly, Zongheng Yang, Zhifeng Chen, Yu Zhang, Yuxuan Wang, Rj Skerrv-Ryan, et al., “Natural TTS synthesis by conditioning WaveNet on Mel spectrogram predictions,” in *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2018, pp. 4779–4783.
- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, “Deep residual learning for image recognition,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 770–778.
- [17] Hideyuki Tachibana, Katsuya Uenoyama, and Shunsuke Aihara, “Efficiently trainable text-to-speech system based on deep convolutional networks with guided attention,” *CoRR*, vol. abs/1710.08969, 2017.
- [18] D. Griffin and Jae Lim, “Signal estimation from modified short-time Fourier transform,” *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 32, no. 2, pp. 236–243, 1984.
- [19] Weidi Xie, Arsha Nagrani, Joon Son Chung, and Andrew Zisserman, “Utterance-level aggregation for speaker recognition in the wild,” in *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2019, pp. 5791–5795.
- [20] Yaroslav Ganin and Victor Lempitsky, “Unsupervised domain adaptation by backpropagation,” in *International Conference on Machine Learning*, 2015, pp. 1180–1189.
- [21] Martin Cooke, Jon Barker, Stuart Cunningham, and Xu 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.
- [22] Vijayaditya Poddinti, Daniel Povey, and Sanjeev Khudanpur, “A time delay neural network architecture for efficient modeling of long temporal contexts,” in *Interspeech*, 2015, pp. 3214–3218.
- [23] Anthony Rousseau, Paul Deléglise, and Yannick Estève, “Enhancing the TED-LIUM corpus with selected data for language modeling and more TED talks,” in *LREC*, 2014, pp. 3935–3939.
- [24] Peng Liu, Xixin Wu, Shiyin Kang, Guangzhi Li, Dan Su, and Dong Yu, “Maximizing mutual information for Tacotron,” *arXiv preprint arXiv:1909.01145*, 2019.