



# DubWise: Video-Guided Speech Duration Control in Multimodal LLM-based Text-to-Speech for Dubbing

Neha Sahipjohn<sup>1</sup>, Ashishkumar Gudmalwar<sup>1</sup>, Nirmesh Shah<sup>1</sup>, Pankaj Wasnik<sup>1</sup>, Rajiv Ratn Shah<sup>2</sup>

<sup>1</sup>Sony Research India Pvt. Ltd., Bangalore, India

<sup>2</sup>Indraprastha Institute of Information Technology (IIIT), Delhi, India

{ashish.gudmalwar1,nirmesh.shah,pankaj.wasnik}@sony.com, rajivratan@iiitd.ac.in

## Abstract

Audio-visual alignment after dubbing is a challenging research problem. To this end, we propose a novel method, DubWise: Multi-modal Large Language Model (LLM)-based Text-to-Speech (TTS), which can control the speech duration of synthesized speech in such a way that it aligns well with the speaker's lip movements given in the reference video even when the spoken text is different or in a different language. To accomplish this, we propose to utilize cross-modal attention techniques in a pre-trained GPT-based TTS. We combine linguistic tokens from text, speaker identity tokens via a voice cloning network, and video tokens via a proposed duration controller network. We demonstrate the effectiveness of our system on the Lip2Wav-Chemistry and LRS2 datasets. Also, the proposed method achieves improved lip sync and naturalness compared to the SOTAs for the same language but different text (i.e., non-parallel) and the different language, different text (i.e., cross-lingual) scenarios.

**Index Terms:** Multimodal TTS, Duration Controllable, LLM-based TTS

## 1. Introduction

Recently, AI-based dubbing technologies have become popular with the emergence of LLM-based models like GPT-2 [1–4] and XTTS [5, 6], resulting in high-quality speech synthesis and voice cloning capabilities [7]. The AI-based dubbing technologies first generate the subtitles in the source language with the help of the Automatic Speech Recognition (ASR) service. Then, the Neural Machine Translation (NMT) service converts these subtitles from a source language to the target language. Finally, Text-to-Speech (TTS) technology generates a speech signal in the target language [8]. The critical issue with these approaches is that speech generated via TTS in the target language have a different length than the corresponding segment in the source language [9, 10]. Hence, dubbed audio does not correctly align with the source video content. Hence, the dubbing output look unnatural from an audio-visual synchronization perspective.

To tackle these alignment issues, one can efficiently utilize time scale modification-based signal processing algorithms, such as WSOLA and others [11–13], to alter the length of synthesized speech. Another possible workaround could be to alter the length of the source video either by downsampling or up-sampling video frames via interpolation-based methods [14]. Some approaches have also proposed utilizing phoneme duration predictor networks in TTS to regulate the synthesized speech's overall duration via a multiplication factor [15, 16]. While these methods effectively achieve global duration control, they fail to align audio with reference video. In addition,

these methods can alter or control the length of an audio to a certain extent; beyond that, they will significantly deteriorate the quality and intelligibility of speech and, hence, the final dubbing output.

Traditionally, the goal of multimodal TTS systems is to control emotion or speaker-specific information from video modalities [17–21]. On the contrary, this paper proposes DubWise, a video-guided speech duration-controllable multimodal TTS for dubbing applications. Our method utilizes visual cues extracted from the video to achieve duration controllability in GPT-based TTS while maintaining intelligibility and speech quality. Videos provide more reliable guidance than audio for alignment in noisy settings. In particular, the proposed architecture utilizes a cross-modal attention technique to combine the video tokens learned via the proposed duration controller network, the audio tokens learned via the voice cloning network, and the linguistic tokens from the input text. The effectiveness of the proposed method has been shown to be superior to the SOTA methods on the Lip2wav Chemistry and LRS2 dataset.

In summary, our work makes the following contributions:

- To the best of the authors' knowledge, this is the first attempt of its kind that utilizes video-based modality for achieving duration controllability in autoregressive (AR) Large Language Model (LLM)-based multimodal TTS.
- Propose a novel duration loss within the auto-regressive model to enable further control over the speech duration.
- Propose a cross-modal attention technique to integrate video modality against traditional concatenation-based strategies.
- Propose a lightweight, faster training strategy by only training the randomly initialized cross-attention layers and transpose convolutions layers of the models.
- Propose to utilize VideoCLIP feature integration to understand global scene context to achieve improved zero-shot voice cloning capability.

## 2. Related Work

**Automatic Video Dubbing (AVD)** has been tackled from multiple perspectives. In [8], authors developed the AVD pipeline where, after TTS output, lip movements are modified to match the speech. VideoDubber [9] approached AVD as a machine translation problem and proposed a translation system based on reference speech duration. Chronopoulou et al. [10] proposed a model that jointly translates and predicts the durations from speechSeamlessM4T [22] performs speech-to-speech translation while maintaining the prosody of input speech but does not consider alignment. Dubbing, alternatively defined as re-recording dialogues due to poor recording conditions, has been addressed by NeuralDubber [23] Text and the corresponding

videos are provided to synthesize speech aligning with the video. Other works like [24], [20, 25] and [17] also tackle similar problems. Such methods fail when a different text is given in the same or different language. On the contrary, our method works even for scenarios not trained on multilingual videos.

**Multimodal LLM:** Various methods have been suggested for integrating images and videos into text-based LLMs. LLaVA [26] suggested that simply by appending a linear layer after a visual representation layer (CLIP) and employing it as a prompt, the vision modality can be integrated into LLMs. On the other hand, Flamingo [27] presents a visual language model where the vision modality is incorporated into the prompt, and cross-attention layers within the LLM are introduced and trained accordingly. Other approaches focus on image and video comprehension and storytelling. However, our method introduces controllability to GPT-based TTS systems by leveraging video input. GPT based TTS primarily utilizes LLM based GPT decoder and we have used both LLM and GPT term interchangeably throughout this paper.

### 3. Proposed Method

The core of our model is a pre-trained, multilingual TTS based on GPT2 architecture. The key reasons for preferring GPT2 over its recent architectures are smaller model size and its wider adaptability in the SOTA TTS architectures [2–4]. We integrate lip-reading features to direct the TTS output Figure 1, aiming to attain synchronization between the generated speech and the observed lip movements in the reference video. DubWise models the speech  $S$  conditioned on lip-features  $F_{lip}$  and input text  $C_t$ .

$$p(S|C_t, F_{lip}) = \prod_{l=1}^L p(S_l|S_{<l}, C_t, F_{lip}) \quad (1)$$

where  $C_t = \text{translate}(C_s)$  is translated text for source text ( $C_s$ ).

We use an LLM-based TTS, XTTS [5] as the backbone. The system utilizes byte-pair tokenizer to encode input text. Reference speaker audio  $S_r$  uttering arbitrary sentence is used to clone the voice. We employ a speaker encoder network to capture the speaker’s unique traits. This network extracts speaker embeddings, which a perceiver [28] subsequently processes to derive a fixed-dimensional representation. To capture the overall scene and mood of the utterance, videoCLIP [29] features of a reference video is used. This is given to the perceiver along with the speaker embedding to capture overall speaker style, while also training the perceiver. A one-hot vector-based language ID is used to specify the language of the input text. These three elements – language ID  $L_t$ , speaker embedding  $S_r$ , and the input text  $C_t$  – are then concatenated to form the prompt, which is subsequently fed into the GPT model. This entire system is pre-trained on 16 languages. During training, the GPT model learns to predict the next token autoregressively. A pre-trained DVAE is utilized to tokenize the audio data for GPT model training.

**Video-encoder:** Given a reference video of a person speaking in a source language  $L_s$  that you want to dub to a target language  $L_t$ , the cropped lip-region  $V_{lip}$  from the silent source video is retrieved. The lip-representation features  $F_{lip}$  are extracted using lip-reading model  $M_{lip}$ . This work employs a lip encoder model trained on masked prediction-based self-supervised learning. We leverage the SOTA pre-trained AV-HuBERT [30] model as the lip feature extractor. During training, this model takes both video (cropped lip region) and audio as inputs. A selection of audio and video frames are masked,

and the transformer component within the model is trained to predict the iteratively refined cluster IDs. The initial cluster IDs are derived through k-means clustering of Mel-Frequency Cepstral Coefficients (MFCC) features extracted from the audio. Finally, the model undergoes fine-tuning with only video features as input, using a transformer decoder designed for lip-reading task to predict text from observed lip movements.

We utilize these features extracted by the AV-HuBERT model before the Language Model (LM) Transformer head as the lip features.

$$F_{lip} = M_{lip}(V_{lip}) \quad (2)$$

**Integration:** This work explores two approaches to incorporating video features into a TTS system for introducing video-based controllability. The first approach involves concatenating video features with speaker embeddings and text into a single prompt fed to the GPT model. This results in an extended prompt length, and the entire GPT model undergoes fine-tuning based on this modified prompt structure.

$$\text{prompt} = [L_t, S_{ref}, F_{lip}, C_t] \quad (3)$$

The second approach leverages cross-attention mechanisms within the LLM architecture. We introduced cross-attention layers between the GPT model, which directly attend to the video features. To ensure the generated speech duration aligns with the reference video, we extend the video features to match the expected speech duration. This simplifies training by enabling linear attention towards the video features without requiring explicit duration specification. We achieve this by applying a combination of convolution and transposed convolution to the video features, effectively learning the necessary expansion without manual intervention. An end token is then concatenated to mark the sequence termination. During training, all other model parameters remain fixed, while only the cross-attention layer and the newly introduced transposed convolution layers are fine-tuned for the language modeling task.

The losses in the model include cross-entropy loss on audio-tokens ( $CE_{audio}$ ), scaled ( $\alpha$ ) text tokens loss ( $CE_{text}$ ) and scaled ( $\beta$ ) duration loss ( $duration\_loss$ ). As the embedding layers remain pretrained and unaltered, they can be used to infer predicted token IDs. We identify the end of a sequence by detecting the occurrence of a pre-defined stop token within the predicted sequence and ground-truth audio sequence. The proposed duration loss is calculated as the difference between the index of the end token in ground-truth audio and the predicted end token.

$$\text{loss} = CE_{audio} + \alpha * CE_{text} + \beta * \text{duration\_loss} \quad (4)$$

The HiFi-GAN [31] vocoder-based decoder was utilized to extract audio from the GPT output’s latent representation.

## 4. Experiments

This section provides details of the datasets and evaluation metrics employed for evaluation.

**Datasets:** The experiments are performed on Lip2Wav-chemistry dataset [23], [32]. It is a single-speaker, 9.2 hours long, YouTube lecture video based dataset. We also perform experiments on multi-speaker LRS2 dataset [33], which is a 29 hours long YouTube clips dataset. To extract the lip-reading features, the video sequences are resampled to a frame rate of 25 frames per second (fps). We utilize the  $S^3FD$  [34] face detector to detect facial key points, and crop a mouth-centered region-of-interest (ROI). This is used to extract the lip-reading

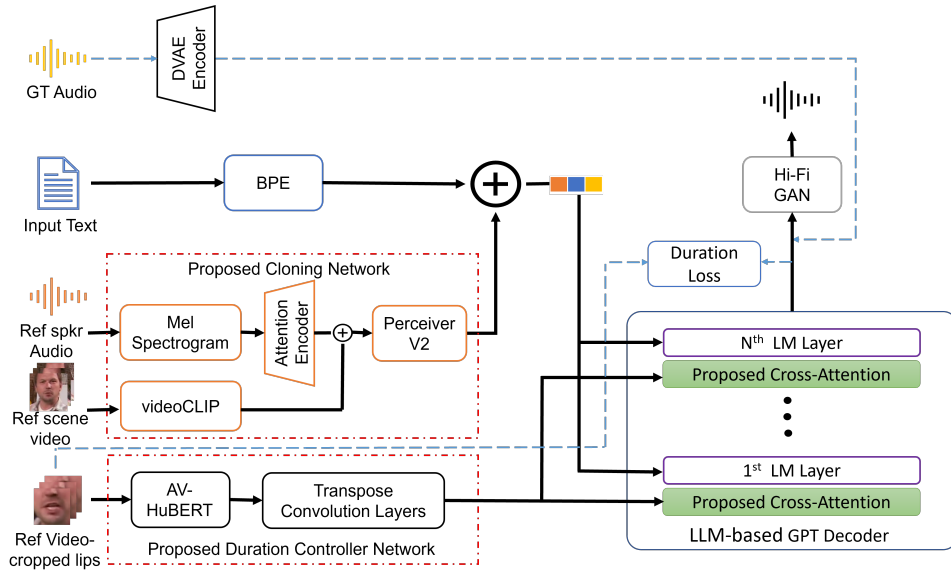


Figure 1: *Proposed Method: Tokenized reference-speaker audio and text form the model’s prompt (ground truth audio included during training only). Lip region video is fed through cross-attention. HiFi-GAN generates speech from the output.*

features. The corresponding text is tokenized using Byte-pair-encoding. For evaluations using translated text, we generate the translation using IndicTrans2 [35], an off-the-shelf machine translation model.

**Baseline Methods:** We compare our results with following SOTA methods. FastSpeech2 [16] proposed phoneme based global duration controllability in the generated audio by speed up or slowed down by a factor. YourTTS [15] leverages a stochastic duration predictor while training. XTTS [5] is GPT2-based TTS model which doesn’t explicitly have duration controllability. To adjust the length of XTTS-generated speech, we employ the Wave Similarity Overlap and Add (WSOLA) algorithm [11] within the FFmpeg toolkit. This approach is referred to as *XTTS+WSOLA*. HPMDubbing [17] is a multimodal dubbing pipeline for same-text scenarios, where corresponding text and video are inputs. We evaluate it for same-text and different-text scenarios.

**Evaluation Metrics:** We do both subjective and objective evaluations to demonstrate the efficacy of the model. In objective evaluation, duration modelling capability, is measured using the ratio of durations of synthesized speech to reference video (DR, ideal value is 1) and difference between durations of synthesized speech and reference video (DD, ideal value is 0). To evaluate the intelligibility and performance of TTS, we calculate the word error rate (WER) and character error rate (CER). Alignment of speech to lip-movements - lip-sync - is evaluated using Lip Sync Error - Confidence (LSE-C) and Lip Sync Error - Distance (LSE-D) [36]. AV-Offset is the time offset between audio and video. The measurements are taken using a pretrained SyncNet model [37]. In subjective evaluation, we compute the Mean Opinion Score (MOS) on overall quality and intelligibility for the speech samples generated by different methods. 20 subjects (aged between 25 to 35 years and no known hearing impairment) took part in the subjective evaluation.

## 5. Results

We present results in three different scenarios, namely, same-text, different-text, and cross-lingual, as shown in Table 1.

Demo samples are available at <sup>1</sup>. In the same-text condition, the same text as in the reference video is input, to synthesize audio samples. On the other hand, in different-text condition, audio samples are synthesized using text different from that in the videos. Here, we changed the sentence structure by  $\pm 50\%$  in terms of word-count from the video’s original text. We also evaluated the proposed approach in a cross-lingual setting in which Hindi audio is generated by using an English speaker as a reference for voice cloning.

For the same text condition, Dubwise achieves the best DR and DD close to ground truth except for FastSpeech2 and XTTS+WSOLA. We can see that overall global duration controllability is achieved by speeding up or slowing down the FastSpeech2 and XTTS+WSOLA TTS output. However, it introduces distortion and affects the intelligibility of speech, which is evident from the higher WER and CER scores compared to Dubwise. Regarding intelligibility evaluation, Dubwise outperforms other baselines in terms of WER and CER. On the LSE evaluation side, Dubwise achieves the best LSE-D and LSE-C scores except for the HPMDubbing. HPMDubbing is trained to synthesize audio based on the audio-video pair using the same text in the video. It tries to lip-sync with audio for the same text, but it fails to maintain intelligibility and constrains the duration of generated audio, which is evident from the high WER, CER, DR, and DD. The proposed Dubwise overcomes this by maintaining intelligibility and speech quality while minimizing duration differences. The proposed approach also tries to achieve lip sync at the same time.

Different-text scenario (Table 1), also displays similar trend as above. Our method has a lower WER and CER with a decent duration ratio. In terms of Lip-Sync-Error, Dubwise achieves the best score for LSE-D and LSE-C, apart from the HPMDubbing approach. HPMDubbing’s WER and CER scores are way higher as they are highly unintelligible and distorted. This is due to HPMDubbing being trained to synthesize audio using audio-video pairing with the same text as in the video. Another drawback of HPMDubbing is it does not have cross-lingual support.

Similar trend is found for cross-lingual scenario, where au-

<sup>1</sup><https://nirmesh-sony.github.io/DubWise/>

dio is synthesized using Hindi text and English reference video to control the generated audio duration. The DubWise approach preserves intelligibility and speech quality while minimizing variations in duration. Additionally, the proposed approach aims to achieve lip synchronization simultaneously. As HP-Mdubbing and FastSpeech2 doesnot support cross-lingual support, we could not compare DubWise approach with them.

Table 1: Objective evaluation on Lip2Wav-Chem dataset

Method	Duration Control		Intelligibility		Lip Sync Error		
	DR	DD	WER ↓	CER ↓	LSE-D ↓	LSE-C ↑	AV offset ↓
Same Text							
Ground Truth	1	0	4.05	2.02	7.32	7.19	0.6
FastSpeech2	0.94	0.29	7.87	2.79	12.76	2.37	7.43
YourTTS	1.12	0.53	27.0	14.85	13.01	1.48	7.27
Baseline_XTTS2	1.26	1.335	8.1	4.68	13.01	2.25	8.15
XTTS2+WSOLA	<b>0.99</b>	<b>0.019</b>	9.19	4.38	12.96	2.12	8.16
HPMDubbing	0.837	0.724	6.29	3.75	<b>6.97</b>	<b>7.17</b>	<b>3.97</b>
DubWise (Ours)	1.077	0.432	<b>4.13</b>	<b>2.48</b>	12.1	2.82	5.98
Different Text							
FastSpeech2	0.947	0.23	22.48	13.28	13.65	1.61	9.53
YourTTS	1.014	0.597	31.92	18.49	13.36	1.61	8.7
XTTS	1.29	1.92	16.72	10.19	13.51	1.78	9.67
XTTS+WSOLA	<b>0.995</b>	<b>0.019</b>	21.17	13.08	13.39	1.64	9.08
HPMDubbing	0.784	1.026	119.01	92.82	<b>8.94</b>	<b>6.40</b>	<b>4.18</b>
DubWise (Ours)	0.94	1.03	<b>11.94</b>	<b>8.13</b>	13.31	1.91	9.54
Different Language Text							
YourTTS	1.12	0.69	56.23	43.85	13.43	1.55	8.21
XTTS	1.32	1.46	34.34	25.17	13.44	1.67	10.14
XTTS+WSOLA	<b>0.99</b>	<b>0.018</b>	38.38	29.06	13.39	1.65	9.35
DubWise (Ours)	1.19	0.932	<b>22.76</b>	<b>13.59</b>	<b>13.33</b>	<b>1.58</b>	<b>8.91</b>

Table 2: Objective evaluation on LRS2 dataset

Method	Duration Control		Intelligibility		Lip Sync Error		
	DR	DD	WER ↓	CER ↓	LSE-D ↓	LSE-C ↑	AV offset ↓
Same Text							
Ground truth	1	0	33.04	27.16	6.23	8.41	0.06
XTTS	1.96	1.41	<b>22.10</b>	<b>15.45</b>	11.55	3.02	8.27
DubWise (Ours)	<b>1.18</b>	<b>0.28</b>	38.88	29.35	<b>10.16</b>	<b>4.35</b>	<b>3.56</b>
Different Language Text							
XTTS	3.15	3.13	49.24	40.22	12.23	2.38	10.33
DubWise (Ours)	<b>1.29</b>	<b>0.46</b>	<b>46.77</b>	<b>36.53</b>	<b>11.60</b>	<b>2.86</b>	<b>7.61</b>




Figure 2: The subjective evaluation of proposed DubWise approach and baseline methods. We obtained  $p$ -value  $< 0.004$

To demonstrate multispeaker capabilities, we also evaluate the LRS2 dataset for same-text and translated Hindi text scenarios and compare the duration controllability aspect with respect to the baseline XTTS model. We observe that DubWise is able to achieve good values for duration difference and duration ratio. The WER is high due to noisy data and small samples used for speaker-reference voice. We also see that DubWise achieves almost 100%  $SLC_p$  in both same-text and translated-text scenarios. From subjective evaluation, we can see that it also supports observations in the objective evaluation. In particular, the proposed DubWise approach achieves the highest score for audio quality, as shown in Figure 2.

### 5.1. Ablation Experiments

We investigated different architectures for the DubWise model through ablation analysis (refer to Table 3). In Ablation 1, we integrated the video features by concatenating them to the input prompt of GPT-TTS. But this resulted in a higher Word Error Rate (WER) compared to other methods. Ablation 2 using

cross-attention to integrate video information with GPT-TTS, achieved the lowest WER and lip-sync error. In Ablation 3, videoCLIP features were incorporated for global style control, but this led to performance degradation compared to Ablation 2. Ablation 4 incorporated duration loss with cross-attention, which improved Duration Ratio compared to Ablation2. By incorporating duration loss and videoCLIP with cross-attention, DubWise achieved the best balance between speech quality (measured by WER and lip-sync error) and control over speech duration (DR). The percentage of sentences matching the desired duration was found to be maximum in this case.

Table 3: Ablation Study

Method	Duration Control		Intelligibility		Lip Sync Error		
	DR (Ideal=1)	DD (Ideal=0)	WER ↓	CER ↓	LSE-D ↓	LSE-C ↑	AV offset ↓
Ablation1	1.07	<b>0.43</b>	8.94	5.92	11.93	3.01	5.71
Ablation2	1.06	<b>0.43</b>	<b>4.39</b>	<b>2.20</b>	<b>11.84</b>	<b>3.07</b>	<b>5.52</b>
Ablation3	1.08	0.46	4.63	2.50	12.11	2.79	5.71
Ablation4	1.06	0.45	4.63	2.50	12.15	2.83	5.69
DubWise	<b>1.04</b>	0.45	4.86	2.73	12.07	2.84	5.77

Figure 3 illustrates an example of a mel-spectrogram, demonstrating how the proposed model captures pauses (high-lighted within the blue box) and lip movements in the synthesized speech in the considered cross-lingual scenario compared to the SOTA algorithm.




Figure 3: Spectrographic analysis of the reference English speech, and corresponding translated Hindi synthesized speech using DubWise and different baselines. Here, reference English sentence is "I have to break six moles of carbon hydrogen bonds, because there's two carbon hydrogen bonds and each acetylene."

## 6. Summary and Conclusion

In this paper, we propose a novel method, DubWise: Multimodal Large Language Model (LLM)-based Text-to-Speech (TTS), which can control the speech duration of synthesized speech in such a way that it aligns well with the speaker's lip movements given in the reference video even when the spoken text is different or in a different language. To accomplish this, we proposed to utilize cross-modal attention techniques in a pre-trained GPT-based TTS, while combining linguistic tokens from text, speaker identity tokens via voice cloning network and video tokens via proposed duration controller network. We demonstrate the effectiveness of our system on Lip2Wav-Chemistry and LRS2 datasets. Also, the proposed method achieves improved lip sync and naturalness compared to the SOTAs for the same language but different text (i.e., non-parallel) and the different language different text (i.e., cross-lingual) scenarios. In future, one should work on improving the global duration controllability as well as word level and phrase-level duration controllability.

## 7. References

- [1] T. Saeki, S. Takamichi, and H. Saruwatari, "Incremental text-to-speech synthesis using pseudo lookahead with large pretrained language model," *IEEE Signal Processing Letters*, vol. 28, pp. 857–861, 2021.
- [2] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, I. Sutskever *et al.*, "Language models are unsupervised multitask learners," *OpenAI journal*, vol. 1, no. 8, p. 9, 2019.
- [3] J. Betker, "Better speech synthesis through scaling," *arXiv preprint arXiv:2305.07243*, 2023.
- [4] M. Łajszczak *et al.*, "BASE TTS: lessons from building a billion-parameter text-to-speech model on 100k hours of data," *arXiv preprint arXiv:2402.08093*, 2024.
- [5] G. Eren and The Coqui TTS Team, "Coqui TTS," Jan. 2021. [Online]. Available: <https://github.com/coqui-ai/TTS>
- [6] E. Casanova, K. Davis, E. Gölge, G. Gökner, I. Gulea, L. Hart, A. Aljafari, J. Meyer, R. Morais, S. Olayemi, and J. Weber, "Xtts: a massively multilingual zero-shot text-to-speech model," in *INTERSPEECH*, 2024, pp. 1–5.
- [7] B. Caramiaux, F. Lotte, J. Geurts, G. Amato, M. Behrmann, *et al.*, "Ai in the media and creative industries," *Research Report, New European Media (NEM)*. 2019, pp.1-35., 2019.
- [8] Y. Yang, B. Shillingford, Y. Assael, M. Wang, W. Liu, Y. Chen, Y. Zhang, E. Sezener, L. C. Cobo, M. Denil *et al.*, "Large-scale multilingual audio visual dubbing," *arXiv preprint arXiv:2011.03530*, 2020.
- [9] Y. Wu, J. Guo, X. Tan, C. Zhang, B. Li, R. Song, L. He, S. Zhao, A. Menezes, and J. Bian, "Videodubber: Machine translation with speech-aware length control for video dubbing," in *Proceedings of the AAAI*, vol. 37, no. 11, 2023, pp. 13 772–13 779.
- [10] A. Chronopoulou, B. Thompson, P. Mathur, Y. Virkar, S. M. Lakew, and M. Federico, "Jointly optimizing translations and speech timing to improve isochrony in automatic dubbing," *arXiv preprint arXiv:2302.12979*, 2023.
- [11] W. Verhelst and M. Roelands, "An overlap-add technique based on waveform similarity (wsola) for high quality time-scale modification of speech," in *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, 1993, pp. 554–557.
- [12] J. Driedger and M. Müller, "A review of time-scale modification of music signals," *Applied Sciences*, vol. 6, no. 2, p. 57, 2016.
- [13] Y. Stylianou, "Voice transformation: a survey," in *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2009, pp. 3585–3588.
- [14] S. Tomar, "Converting video formats with ffmpeg," *Linux Journal*, vol. 2006, no. 146, p. 10, 2006.
- [15] E. Casanova, J. Weber, C. D. Shulby, A. C. Junior, E. Gölge, and M. A. Ponti, "Yourtts: Towards zero-shot multi-speaker tts and zero-shot voice conversion for everyone," in *International Conference on Machine Learning*. PMLR, 2022, pp. 2709–2720.
- [16] Y. Ren, C. Hu, X. Tan, T. Qin, S. Zhao, Z. Zhao, and T.-Y. Liu, "Fastspeech 2: Fast and high-quality end-to-end text to speech," in *ICLR*, 2021.
- [17] G. Cong, L. Li, Y. Qi, Z.-J. Zha, Q. Wu, W. Wang, B. Jiang, M.-H. Yang, and Q. Huang, "Learning to dub movies via hierarchical prosody models," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 14 687–14 697.
- [18] Q. Chen, M. Tan, Y. Qi, J. Zhou, Y. Li, and Q. Wu, "V2c: visual voice cloning," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 21 242–21 251.
- [19] J. Schroeter, J. Ostermann, H. P. Graf, M. Beutnagel, E. Cosatto, A. Syrdal, A. Conkie, and Y. Stylianou, "Multimodal speech synthesis," in *IEEE International Conference on Multimedia and Expo. (ICME)*, vol. 1, 2000, pp. 571–574.
- [20] R. Liu, B. Sisman, G. Gao, and H. Li, "Expressive tts training with frame and style reconstruction loss," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 1806–1818, 2021.
- [21] J. Parker, R. Maia, Y. Stylianou, and R. Cipolla, "Expressive visual text to speech and expression adaptation using deep neural networks," in *ICASSP*, 2017, pp. 4920–4924.
- [22] L. Barrault, Y.-A. Chung, M. C. Meglioli, D. Dale, N. Dong, P.-A. Duquenne, H. Elsahar, H. Gong, K. Heffernan, J. Hoffman *et al.*, "Seamless4t-massively multilingual & multimodal machine translation," *arXiv preprint arXiv:2308.11596*, 2023.
- [23] C. Hu, Q. Tian, T. Li, W. Yuping, Y. Wang, and H. Zhao, "Neural dubber: Dubbing for videos according to scripts," *Advances in neural information processing systems*, vol. 34, pp. 16 582–16 595, 2021.
- [24] M. Hassid, M. T. Ramanovich, B. Shillingford, M. Wang, Y. Jia, and T. Remez, "More than words: In-the-wild visually-driven prosody for text-to-speech," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 10 587–10 597.
- [25] J. Lu, B. Sisman, R. Liu, M. Zhang, and H. Li, "Visualtts: Tts with accurate lip-speech synchronization for automatic voice over," in *ICASSP*, 2022, pp. 8032–8036.
- [26] H. Liu, C. Li, Q. Wu, and Y. J. Lee, "Visual instruction tuning," *Advances in neural information processing systems*, vol. 36, 2024.
- [27] J.-B. Alayrac, J. Donahue, P. Luc, A. Miech, I. Barr, Y. Hasson, K. Lenc, A. Mensch, K. Millican, M. Reynolds *et al.*, "Flamingo: a visual language model for few-shot learning," *Advances in Neural Information Processing Systems*, vol. 35, pp. 23 716–23 736, 2022.
- [28] A. Jaegle, F. Gimeno, A. Brock, O. Vinyals, A. Zisserman, and J. Carreira, "Perceiver: General perception with iterative attention," in *International conference on machine learning*. PMLR, 2021, pp. 4651–4664.
- [29] H. Xu, G. Ghosh, P.-Y. Huang, D. Okhonko, A. Aghajanyan, F. Metze, L. Zettlemoyer, and C. Feichtenhofer, "Videoclip: Contrastive pre-training for zero-shot video-text understanding," in *EMNLP*, Punta Cana, Dominican Republic, 2021.
- [30] B. Shi, W.-N. Hsu, K. Lakhotia, and A. Mohamed, "Learning audio-visual speech representation by masked multimodal cluster prediction," in *International Conference on Learning Representations*, 2022.
- [31] J. Kong, J. Kim, and J. Bae, "Hifi-gan: Generative adversarial networks for efficient and high fidelity speech synthesis," *Advances in neural information processing systems*, vol. 33, pp. 17 022–17 033, 2020.
- [32] K. Prajwal, R. Mukhopadhyay, V. P. Nambodiri, and C. Jawahar, "Learning individual speaking styles for accurate lip to speech synthesis," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 13 796–13 805.
- [33] T. Afouras, J. S. Chung, A. Senior, O. Vinyals, and A. Zisserman, "Deep audio-visual speech recognition," *IEEE transactions on pattern analysis and machine intelligence*, vol. 44, no. 12, pp. 8717–8727, 2018.
- [34] S. Zhang, X. Zhu, Z. Lei, H. Shi, X. Wang, and S. Z. Li, "S3fd: Single shot scale-invariant face detector," in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 192–201.
- [35] J. Gala *et al.*, "Indictrans2: Towards high-quality and accessible machine translation models for all 22 scheduled indian languages," *Transactions on Machine Learning Research*, 2023. [Online]. Available: <https://openreview.net/forum?id=vfT4YuzAYA>
- [36] K. Prajwal, R. Mukhopadhyay, V. P. Nambodiri, and C. Jawahar, "A lip sync expert is all you need for speech to lip generation in the wild," in *Proceedings of the 28th ACM international conference on multimedia*, 2020, pp. 484–492.
- [37] J. S. Chung and A. Zisserman, "Out of time: automated lip sync in the wild," in *Computer Vision-ACCV 2016 Workshops: ACCV 2016 International Workshops, Taipei, Taiwan, November 20-24, 2016, Revised Selected Papers, Part II 13*. Springer, 2017, pp. 251–263.