

V2C: Visual Voice Cloning

Qi Chen^{1*} Mingkui Tan² Yuankai Qi¹ Jiaqiu Zhou^{2,3} Yuanqing Li^{2,3†} Qi Wu^{1†}

¹University of Adelaide ²South China University of Technology ³Pazhou Lab

{qi.chen04, qi.wu01}@adelaide.edu.au, qykshr@gmail.com

{auyqli, mingkuitan}@scut.edu.cn, mszjq@mail.scut.edu.cn

Abstract

Existing Voice Cloning (VC) tasks aim to convert a paragraph of text to a speech with desired voice specified by a reference audio. This has significantly boosted the development of artificial speech applications. However, there also exist many scenarios that cannot be well reflected by these VC tasks, such as movie dubbing, which requires the speech to be with emotions consistent with the movie plots. To fill this gap, in this work we propose a new task named Visual Voice Cloning (V2C), which seeks to convert a paragraph of text to a speech with both desired voice specified by a reference audio and desired emotion specified by a reference video. To facilitate research in this field, we construct a dataset, V2C-Animation, and propose a strong baseline based on existing state-of-the-art (SoTA) VC techniques. Our dataset contains 10,217 animated movie clips covering a large variety of genres (e.g., Comedy, Fantasy) and emotions (e.g., happy, sad). We further design a set of evaluation metrics, named MCD-DTW-SL, which help evaluate the similarity between ground-truth speeches and the synthesised ones. Extensive experimental results show that even SoTA VC methods cannot generate satisfying speeches for our V2C task. We hope the proposed new task together with the constructed dataset and evaluation metric will facilitate the research in the field of voice cloning and broader vision-and-language community. Source code and dataset will be released in <https://github.com/chenqi008/V2C>.

1. Introduction

Voice Cloning (VC) [2, 9, 20, 30] aims to convert a paragraph of text to speech with the desired voice from a reference audio. However, there exist many applications in the real world that require the generated speeches not only using a template voice but also being with rich emotions (e.g., angry, happy, and sad), such as movie dubbing. This is be-

*Partial of the work was performed when Qi Chen was an intern at Pazhou Lab.

†Corresponding authors.

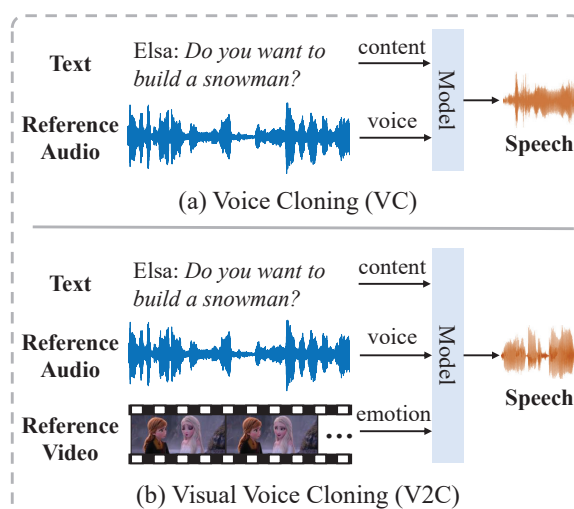


Figure 1. (a) Voice Cloning (VC) vs. (b) Visual Voice Cloning (V2C). Given an input triplet (i.e., subtitle/text, reference audio, and target video), our V2C task seeks to convert the text into a speech, which should be with the voice of reference audio and the emotion derived from reference video. Note that the reference audio only provides an expected voice while its content is irrelevant.

yond the scope of conventional VC tasks (Figure 1(a)), as no extra guiding information can be used to generate desired tones and rhythms. Considering that we humans accomplish the movie dubbing task with the most reference from visual observations (e.g., watching the movie to grasp the emotion of characters), we propose an extension task of VC, namely Visual Voice Cloning (V2C).

An example of the proposed V2C task is shown in Figure 1(b). Unlike the conventional VC task, which converts text to speech only aided by a reference audio, our V2C task takes a triplet (text/subtitle, reference audio, reference video) as input and expects a resulting speech with the same voice but varying emotions derived from the reference video. The text/subtitle is the content that the generated speech needs to cover. The reference audio includes a pre-recorded voice of the target speaker from a different

clip. And we aim to generate a speech with the voice in the reference audio and the character’s visual emotion from the reference video, speaking the content in the given text.

The new task poses several novel challenges. First, the conventional Voice Cloning (VC) methods [2, 6, 9, 20, 30] cannot well solve the V2C task as they focus only on how to convert the input text to speech with the voice/tone exhibited in the reference audio, without considering the emotion and context of the new speech. However, in our V2C task (e.g., movie dubbing) the voice emotion is crucial for generating human-like speech. Second, in our V2C task, the voice emotion should be derived from the reference video rather than the reference audio from an irrelevant clip. Taking movie dubbing as an example, it requires humans to grasp the emotions of characters by watching the corresponding movie clips and observing their performances (e.g., facial expressions or actions). Although several improved VC methods [37,47] also try to inject the voice emotion into their generated speech, they capture both emotion and voice from the reference audio, which cannot satisfy the requirements of V2C task. In our V2C task, an ideal method should be able to disentangle voice and emotion from the reference audios and the reference videos, respectively.

As there is no off-the-shelf dataset suitable for the V2C task, we collect the first V2C-Animation dataset to facilitate the research in this field. It comprises 10,217 video clips with audios and subtitles, covering 26 animated movies with 153 characters (i.e., speakers) in total. Our V2C dataset covers three modalities (i.e., text, audio and video) unlike the existing text-to-speech datasets [19,31,49,51] or movie description datasets [36,41], which only focus on text and audio, or text and video. Besides, we also provide emotion annotation (e.g., happy or sad) for each audio and video clip like [14]. To alleviate the impact from background music, we only extract the sound channel of the centre speaker, which mainly focuses on the sound of the speaking character. In this way, we ensure that all the audio clips only contain the sound from speaking characters.

To address the above challenges of V2C task, based on the widely used Text-to-Speech (TTS) framework (i.e., FastSpeech2 [34]), we propose a new method called Visual Voice Cloning Network (V2C-Net), considering the emotion information derived from the reference video frames. Moreover, based on MCD [24], we design an evaluation metric, called MCD-DTW weighted by Speech Length (MCD-DTW-SL), seeking to evaluate the generated speech effectively and automatically.

In summary, our contribution include:

- We propose a new task, namely Visual Voice Cloning (V2C). Given a triplet (i.e., text/subtitle, reference audio and reference video), the task seeks to convert the text into a speech with voice and emotion derived from reference audio and reference video, respectively.

- We collect the first V2C-Animation dataset, consisting of 26 animated movies, 153 characters, 10,217 video clips with aligned audios and subtitles, covering three modalities (text, audio, video) and speakers’ emotion.
- We design a new method, called Visual Voice Cloning Network (V2C-Net). Besides, to evaluate the generated speech automatically, we provide an advanced automatic evaluation metric, named MCD-DTW-SL.

2. Related Work

As the V2C is a new task, here we briefly review several closely related works in the fields of Text to Speech, Voice Cloning, and Prosody Transfer.

Text to Speech. Many text-to-speech (TTS) synthesis methods [3, 8, 21, 26, 46, 50] have been proposed to generate natural speech from text. Then, based on WaveNet, Deep Voice [3] divides a TTS model into several modules, which are optimised independently. Wang *et al.* [46] propose a new framework Tacotron, which integrates all the necessary stages in text-to-speech synthesis and enables that the speech synthesis model can be optimised in an end-to-end manner. Recently, TransformerTTS [26] introduces the structure of transformer [44] into TTS task while Ren *et al.* [35] propose a more efficient transformer (i.e., FastSpeech) by using non auto-regressive generation method. Based on FastSpeech, they further design an improved FastSpeech2 [34], which seeks to control the generated speech via the adjustment of pitch and energy. However, the TTS task mainly focuses on how to convert natural language text to speech in a correct pronounce. Instead, our V2C task requires the generated speech to be additionally with a suitable voice emotion and tone.

Voice Cloning. Unlike the TTS method, which synthesises speech only with a single voice, voice cloning (VC) task [11, 32, 40] seeks to generate multiple speeches with different voices. Based on Deep Voice [3] and Tacotron [46], Deep Voice 2 [11] map the voices from different speakers into a common space and use the low-dimensional embedding from this common space as a condition to aid the generation process. Jia *et al.* [20] propose a multi-speaker TTS framework, which consists of three sub-modules (i.e., encoder, synthesizer and vocoder), which is able to synthesise a high-quality speech from given text. More recent extensions [2, 6, 9, 30] focus on synthesising the voice of an unseen person using a few samples only. Specifically, to synthesise a person’s voice from only a few audio samples, Arik *et al.* [2] study two approaches: speaker adaption and speaker encoding. The speaker adaption seeks to fine-tune a trained multi-speaker model for an unseen speaker using a few audio-text pairs while the speak encoding tries to directly estimate the speaker embedding from audios of an unseen speaker. Chen *et al.* [9] propose an adaptive TTS system by using the meta-learning approach. Different from

VC task, our V2C additionally requires the prosody/tone of generated speeches to match with the reference video.

Prosody Transfer. To produce realistic speech, prosody transfer (PT) [5, 17, 37, 38, 43, 47, 48] seeks to grasp prosody from reference audios. Specifically, extending from Tacotron [46], Skerry-Ryan *et al.* [37] propose an encoder architecture to learn a representation of prosody from reference spectrogram slices, which are derived from the reference audio. Global Style Tokens (GSTs) [47] models the styles of different speakers using an interpretable embedding, which can be used as a condition when transferring the different speaking styles. Based on the variational autoencoder (VAE) framework [22], Hsu *et al.* [17] design a neural sequence-to-sequence TTS model, which categorises the speaking styles into several latent attributes and hence controls the speaking style via adapting these attributes. To transfer speaking style which is underrepresented in the dataset, Whitehill *et al.* [48] propose an adversarial cycle-consistent training procedure for a multi-reference neural TTS system. Overall, the goal of prosody transfer is to capture both the voice and emotion from the reference audio, and therefore the task is defined without the use of the information from visual side. By contrast, V2C task is proposed to infer the voice emotion from a reference video, which has many real-world applications such as movie dubbing.

3. V2C Task and V2C-Animation Dataset

3.1. Problem Definition for V2C Task

Given a triplet $\mathcal{Z} = \{Z_{text}, Z_{audio}, Z_{video}\}$ (*i.e.*, text, reference audio and reference video), our Visual Voice Cloning (V2C) task aims to generate a speech \mathcal{Y}_{wave} (*i.e.*, a waveform in time-domain) from text Z_{text} , which should use the voice of reference audio Z_{audio} and be with the emotion derived from reference video Z_{video} . In Figure 1, we take movie dubbing as an example. Given a movie clip (*i.e.*, reference video), a subtitle (*i.e.*, text) and a reference audio, we seek to synthesise a speech from subtitle according to both the character’s emotion derived from movie and the voice from reference audio.

3.2. Dataset Construction for V2C Task

The dataset for V2C task should cover all three modalities and the samples from different modalities need to be aligned with each other. As there is no an off-the-shelf dataset suitable for this new task, we collect the first V2C dataset, called V2C-Animation.

Data Collection. We search for Blu-ray animated movies with the corresponding subtitles and then select a set of 26 movies of diverse genres. Specifically, we first cut the movies into a series of video clips according to the subtitle files. Here, we use an SRT type subtitle file. In addition to subtitles/texts, the SRT file also contains starting



Figure 2. An example of how to cut a movie into a series of video clips according to SRT subtitle files. Note that the SRT files contain both starting and ending time-stamps for each video clip. No. 1340 refers to the sequential number of the current utterance.

and ending time-stamps to ensure the subtitles match with video and audio, and sequential number of subtitles (*e.g.*, No. 1340 in Figure 2), which indicates the index of each video clip. Based on the SRT file, we cut the movie into a series of video clips using the FFmpeg toolkit [42] (an automatic audio and video processing toolkit) and then extract the audio from each video clip by FFmpeg as well. Note that, to alleviate the impact from the background music, we only extract the sound channel of the centre speaker, which mainly focuses on the sound of the speaking character.

Data Annotation and Organisation. Inspired by the organisation of LibriSpeech dataset [31], we categorise the obtained video clips, audios and subtitles into their corresponding characters (*i.e.*, speakers) via a crowd-sourced service. To ensure that the characters appeared in the video clips are the same as the speaking ones, we manually remove the data example that does not satisfy the requirement. Then, following the categories of FER-2013 [14] (a dataset for human facial expression recognition), we divide the collected video/audio clips into 8 types including angry, happy, sad, *etc.* In this way, we totally collect a dataset with 10,217 video clips with paired audios and subtitles. All of the annotations, time-stamps of the mined movie clips and a tool to extract the triplet data will be released. We randomly split 60% samples as training data, 10% samples as validation data and 30% samples as testing data.

3.3. V2C-Animation Dataset vs. Related Datasets

We compare our V2C-Animation dataset against ones of the VC, TTS, and PT tasks. In addition, we consider some Movie Description (MD) datasets and Lip Reading Sentence (LRS) datasets, which contain both video and text like ours. Specifically, the VC/TTS/PT datasets include LJ Speech [19], LibriSpeech [31], VCTK [49] and LibriTTS [51], the MD datasets involve MPII-MD [36] and MovieQA [41], while LRS datasets contain LRS2 [10]¹. From Table 1, the statistic results demonstrate that our V2C-

¹The LRS2 dataset has two subsets: LRS2-main and LRS2-pretrain. On LRS2-pretrain, the utterance of each video may contain multiple sentences. But each video only corresponds to a single sentence on LRS2-main. There is some overlap between LRS2-pretrain and LRS2-main sets.

Dataset	Text	Audio	Video	Identity	Emotion	#Movies	#Video Clips	#Audio Clips	#Speakers	Avg. S	Avg. A/V (s)
LJ Speech [19]	✓	✓		✓		-	-	13100	1	17.23	6.57
LibriSpeech [31]	✓	✓		✓		-	-	250698	2484	32.55	14.10
VCTK [49]	✓	✓		✓		-	-	44070	108	7.41	3.59
LibriTTS [51]	✓	✓		✓		-	-	375086	2456	16.86	5.62
MPII-MD [36]	✓		✓			94	68337	-	-	-	3.88
MovieQA [41]	✓		✓			140	6771	-	-	6.20	202.67
LRS2-main [10]	✓	✓	✓			-	48164	48164	-	7.13	1~2
LRS2-pretrain [10]	✓	✓	✓			-	96318	96318	-	21.43	~10
V2C-Animation	✓	✓	✓	✓	✓	26	10217	10217	153	6.51	2.40

Table 1. We compare our V2C-Animation dataset with several existing multi-modal datasets. “Identity” and “Emotion” indicate whether the datasets contain annotations corresponding to the speaker’s identity and emotion. The notations “#Movies”, “#Video Clips”, “#Audio Clips” and “#Speakers” refer to the number of movies, videos, audios and speakers/characters, respectively. “Avg. S” indicates the average length of subtitle while “Avg. A/V” is the average duration of audio/video.

Dataset	LJ Speech	LibriSpeech
Avg. P (Hz)	127.27 ± 11800.96	88.15 ± 7313.39
Dataset	LibriTTS	V2C-Animation (Ours)
Avg. P (Hz)	93.97 ± 9295.67	117.99 ± 16910.77

Table 2. We compare the average and variance of pitch from our V2C-Animation dataset and the related datasets. “Avg. P” refers to the average values of pitches with the corresponding variances.

Animation dataset is unique, covering all the three modalities (*i.e.*, text, audio and video) with both identity and emotion annotations, while most others only focus on the two modalities and all of them are without emotion annotations.

To further compare the differences between our V2C-Animation dataset and related datasets, following [37], we visualise the pitch tracks of the samples from our dataset and others. Specifically, we randomly select an audio sample from LJ Speech, LibriSpeech and LibriTTS, respectively. Due to the varying lengths of audios, for a fair comparison, we crop two seconds of audio from each compared sample. As shown in Figure 3, the audio pitches from the existing datasets are more smooth and the ranges of frequency (Hz) are narrower than ours. Moreover, we provide average and variance values of the pitch tricks. Table 2 shows that the variance of our V2C-Animation is largest, which further demonstrates that our proposed dataset has a wider range of frequency (Hz). Both the visual and statistical results demonstrate that our V2C-Animation dataset is more challenging due to the varied prosody.

4. Visual Voice Cloning Network (V2C-Net)

For the V2C task, we propose a baseline model, called Visual Voice Cloning Network (V2C-Net), which is based on a widely used TTS framework FastSpeech2 [34]. As shown in Figure 4, our model contains three main components: a multi-modal encoder, a synthesizer and a vocoder.

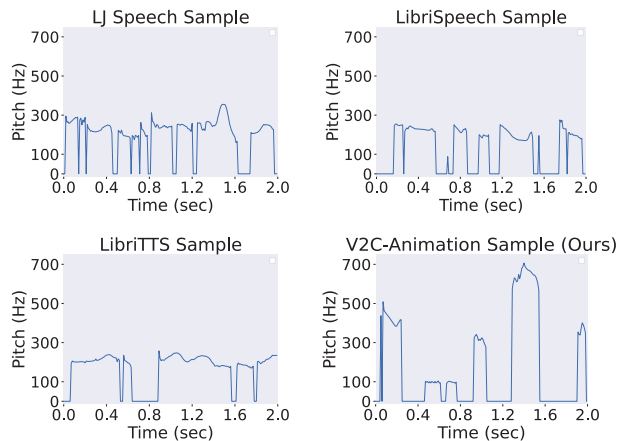


Figure 3. Examples from the existing Text-to-Speech (TTS) datasets (*i.e.*, LJ speech, LibriTTS and LibriSpeech) and our V2C-Animation dataset. A pitch of 0 Hz refers to an unvoiced segment.

We input a triplet (*i.e.*, text, reference audio, and reference video) into the encoder and output three types of features (*i.e.*, phoneme, speaker, and emotion). Based on these features, we use the synthesizer to generate a mel-spectrogram (see Figure 4 right side), which is a time-frequency representation of the audio signal. Last, we convert the generated mel-spectrogram into waveform (*i.e.*, speech, see Figure 4 right-bottom) by the vocoder.

4.1. Multi-modal Encoder for Feature Extraction

Given a triplet $\mathcal{Z} = \{Z_{text}, Z_{audio}, Z_{video}\}$, the output feature/embedding \mathcal{X} from the multi-modal encoder is

$$\mathcal{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_L\} = f(\mathcal{Z}), \quad (1)$$

where L is the length of input sentence (*i.e.*, the number of phoneme). $f(\cdot)$ is the multi-modal encoder, mainly containing three sub-modules: a text encoder f_{txt} , a speaker encoder f_{spk} , and an emotion encoder f_{emo} . Here, Z_{text} ,

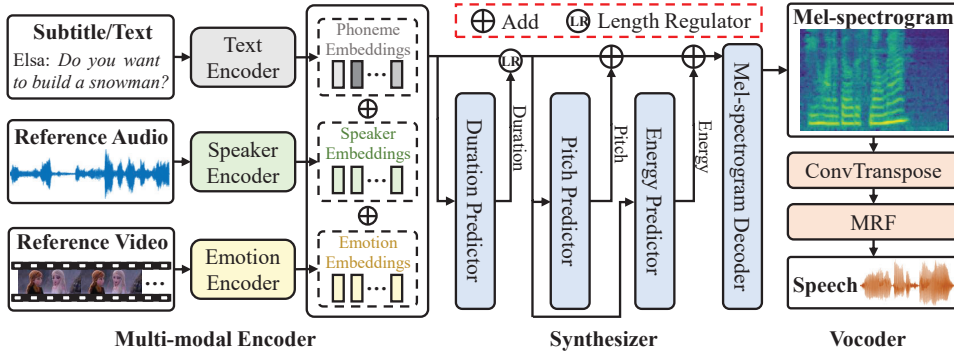


Figure 4. Overview of V2C-Net. It consists of three main components: a multi-modal encoder, a synthesizer and a vocoder. A triplet (*i.e.*, text, reference audio and reference video) is fed into the multi-modal encoder (Sec. 4.1) and it outputs three types of embeddings. Based on these embeddings, the synthesizer (Sec. 4.2) generates a mel-spectrogram. Finally, the mel-spectrogram is converted into the waveform (*i.e.*, speech) by the vocoder (Sec. 4.3).

Z_{audio} and Z_{video} indicate the input text, reference audio and reference video, respectively. We obtain the i -th output feature $\mathbf{x}_i = \mathbf{o}_i \oplus \mathbf{u} \oplus \mathbf{v}$, where \mathbf{o}_i is the embedding of the i -th phoneme derived from text Z_{text} . The embeddings \mathbf{u} and \mathbf{v} are the outputs of f_{spk} and f_{emo} , respectively. The notation \oplus indicates the operation of element-wise add.

Text encoder. Following the structure of FastSpeech2, we take 4 Feed-Forward Transformer (FFT) blocks [35] as our text encoder. Based on such a text encoder, we produce a series of phoneme embeddings $\mathcal{O} = \{\mathbf{o}_1, \dots, \mathbf{o}_L\}$ from an input text Z_{text} . Mathematically, the process can be defined as $\mathcal{O} = f_{txt}(Z_{text})$.

Speaker encoder. To explore the voice characteristics of different speakers, we adopt a speaker encoder f_{spk} , which has the same architecture as [45], comprising 3 LSTM layers and a linear layer. The speaker encoder first converts a sequence of mel-spectrogram frames, derived from the reference audio, to a series of hidden embeddings by LSTM, and then maps the last hidden embedding to a fixed-dimensional vector via the linear layer. For convenience, we define the process as $\mathbf{u} = f_{spk}(\sigma(Z_{audio}))$, where \mathbf{u} refers to the speaker embedding while σ is a mapping function, converting the audio from waveform to mel-spectrogram.

Emotion encoder. To exploit the emotion from video, we design an emotion encoder E_E , which captures the embedding of the whole video clip Z_{video} . To be specific, we use the I3D model [7] as our emotion encoder f_{emo} and calculate the emotion embedding \mathbf{v} by $\mathbf{v} = f_{emo}(Z_{video})$.

4.2. Synthesizer for Mel-Spectrogram Generation

To generate the mel-spectrogram from the conditional phoneme embedding $\mathcal{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_L\}$, we introduce a synthesizer inspired by FastSpeech2 [34], and obtain the predicted mel-spectrogram frames $\mathcal{Y} = \{\mathbf{y}_1, \dots, \mathbf{y}_T\}$, where T is the number of mel-spectrogram frames. Here, the synthesizer mainly contains four parts: a duration pre-

dictor, a pitch predictor, an energy predictor, and a mel-spectrogram decoder. Loss function of the synthesizer is

$$\mathcal{L}_S = \mathcal{L}_{mel} + \lambda_1 \mathcal{L}_{dur} + \lambda_2 \mathcal{L}_{pitch} + \lambda_3 \mathcal{L}_{energy}, \quad (2)$$

where \mathcal{L}_{mel} , \mathcal{L}_{dur} , \mathcal{L}_{pitch} and \mathcal{L}_{energy} refer to the losses of mel-spectrogram, duration predictor, pitch predictor and energy predictor, respectively. The λ_1 , λ_2 and λ_3 are hyper-parameters, and we set $\lambda_1 = \lambda_2 = \lambda_3 = 1$ in practice. The details are depicted in the following.

Duration Predictor. To alleviate the problem of length mismatch between the input embedding \mathcal{X} and mel-spectrogram frames \mathcal{Y} (*i.e.*, $L \neq T$), we introduce a duration predictor S_d , which takes the embedding \mathcal{X} as input and predicts the duration $\mathcal{D} = \{d_1, \dots, d_i, \dots, d_L\} = \{S_d(\mathbf{x}_1), \dots, S_d(\mathbf{x}_i), \dots, S_d(\mathbf{x}_L)\}$ of each phoneme embedding. The i -th phoneme duration d_i indicates the number of copies for i -th phoneme embedding \mathbf{x}_i . Then, we use a Length Regulator (\mathcal{LR}):

$$\mathcal{X}_{mel} = \mathcal{LR}(\mathcal{X}, \mathcal{D}) = \mathcal{LR}(\mathcal{X}, S_d(\mathcal{X})), \quad (3)$$

where \mathcal{X}_{mel} is extended phoneme embedding² with length T . To optimise the duration predictor, we use Montreal forced alignment (MFA) [28] tools to obtain the ground-truth phoneme duration sequence, and then calculate a mean square error (MSE) loss between ground-truth $\hat{\mathcal{D}}$ and predicted \mathcal{D} . Formally, the loss can be defined as

$$\mathcal{L}_{dur} = \frac{1}{L} \sum_{i=1}^L (\hat{d}_i - d_i)^2. \quad (4)$$

Pitch and Energy Predictors. To affect the prosody and volume of speech, following [34], we employ a pitch predictor S_p and an energy predictor S_e , respectively. Specifically, to predict the pitch contour, we use continuous

²For example, if $d_i = 2$, \mathcal{X}_{mel} would be $\mathcal{X}_{mel} = \{\dots, \mathbf{x}_i, \mathbf{x}_i, \dots\}$. For simplicity, we redefine \mathcal{X}_{mel} as $\mathcal{X}_{mel} = \{\mathbf{x}_1, \dots, \mathbf{x}_t, \dots, \mathbf{x}_T\}$

wavelet transform (CWT) to convert the continuous pitch into pitch spectrogram [16, 39], and take it as ground-truth to optimise the pitch predictor S_p by MSE loss:

$$\mathcal{L}_{pitch} = \frac{1}{T} \sum_{t=1}^T (\hat{p}_t - p_t)^2, \quad (5)$$

where \hat{p}_t and $p_t = S_p(\mathbf{x}_t)$ denote the t -th ground-truth and predicted pitch value, respectively. For energy, we follow the operation in [34] that calculate an L2-norm of the amplitude of each short-time Fourier transform (STFT) frame and take it as energy. The corresponding loss function is

$$\mathcal{L}_{energy} = \frac{1}{T} \sum_{t=1}^T (\hat{e}_t - e_t)^2, \quad (6)$$

where \hat{e}_t and $e_t = S_e(\mathbf{x}_t)$ are the t -th ground-truth and predicted energy value, respectively. Last, we encode each pitch and energy value into the corresponding embedding by the embedding layers ϕ and φ separately, and then add the pitch and energy embeddings into the extended phoneme embedding \mathcal{X}_{mel} . Mathematically, the mel-spectrograms \mathcal{Y} can be generated by

$$\mathcal{Y} = g(\mathcal{X}_{mel} \oplus \phi(S_p(\mathcal{X}_{mel})) \oplus \varphi(S_e(\mathcal{X}_{mel}))), \quad (7)$$

where $g(\cdot)$ refers to a mel-spectrogram decoder, consisting of 6 FFT blocks [35]. To optimise the predicted mel-spectrogram, we use the loss function:

$$\mathcal{L}_{mel} = \frac{1}{T} \sum_{t=1}^T \|\hat{\mathbf{y}}_t - \mathbf{y}_t\|, \quad (8)$$

where $\hat{\mathbf{y}}_t$ denotes the t -th frame of the ground-truth mel-spectrogram while $\mathbf{y}_t \in \mathcal{Y}$ is the predicted one.

4.3. Vocoder for Speech Synthesis

In Figure 4, to convert the generated mel-spectrogram \mathcal{Y} into time-domain waveform \mathcal{Y}_{wave} , we use HiFi-GAN [23] as our vocoder, which mainly focuses on the raw waveform generation from mel-spectrogram by GANs [13]. The generator of HiFi-GAN can be divided into two major modules: a transposed convolution (ConvTranspose) network and a multi-receptive field fusion (MRF) module. Specifically, we first upsample the mel-spectrogram \mathcal{Y} by ConvTranspose, which seeks to take an alignment between the length of the output features and the temporal resolution of raw waveforms. Then, we feed the upsampled features into the MRF module, which consists of multiple residual blocks [15], and take the sum of outputs from these blocks as our predicted waveform. Here, we follow the settings of [23] that use the residual blocks with different kernel sizes and dilation rates to ensure different receptive fields. We optimise the vocoder via the objective function that contains an LSGAN-based loss [27], a mel-spectrogram loss [18], and a feature matching loss [25]. In practice, we use the ‘‘universal’’ checkpoint provided in the official code of HiFi-GAN without fine-tuning on any of the synthesised mel-spectrograms.

5. Experiments

We evaluate the quality of generated speech in terms of three aspects: 1) objective evaluation, 2) subjective evaluation, and 3) identity and emotion accuracy. The objective and subjective evaluation metrics aim to assess the quality of generated speeches by comparing with ground-truth ones. By contrast, the identity accuracy and emotion accuracy focus on whether the generated speeches involve the desired voice (*i.e.*, identity) and emotion, respectively. We provide both quantitative and qualitative results on the V2C-Animation dataset. More details are in the following.

5.1. Evaluation Metrics

Objective Evaluation Metric. To assess the quality of generated speech, we use Mel Cepstral Distortion (MCD) [24] metric, which compares Mel Frequency Cepstral Coefficient (MFCC) vectors $\mathcal{C} = \{\mathbf{c}_1, \mathbf{c}_2, \dots, \mathbf{c}_i, \dots, \mathbf{c}_M\}$ and $\mathcal{C}' = \{\mathbf{c}'_1, \mathbf{c}'_2, \dots, \mathbf{c}'_j, \dots, \mathbf{c}'_N\}$ derived from the generated speech and ground truth, respectively. We sum the Euclidean distance over the first K ($K = 13$) MFCC values:

$$\text{MCD}(\mathcal{C}, \mathcal{C}') = \frac{1}{T} \sum_{t=1}^T d(\mathbf{c}_t, \mathbf{c}'_t) = \frac{1}{T} \sum_{t=1}^T \sqrt{\sum_{k=1}^K (c_{t,k} - c'_{t,k})^2}, \quad (9)$$

where $M = N = T$ refers to the number of speech/audio frames. The $c_{t,k}$ and $c'_{t,k}$ denote the k -th MFCC value of the t -th speech frame from generated and ground-truth speeches, respectively, while $\mathbf{c}_t = (c_{t,1}, c_{t,2}, \dots, c_{t,K})$ and $\mathbf{c}'_t = (c'_{t,1}, c'_{t,2}, \dots, c'_{t,K})$.

Note that the MCD metric requires the lengths of two input speeches to be the same (*i.e.*, $M = N$). When $M \neq N$, the existing voice cloning methods like [37] simply extend the shorter speech to the length of longer one by padding 0 for the time-domain waveform. In this way, the value of MCD may become extremely large if the mismatching occurs at the beginning of two speeches. To avoid this issue, Battenberg *et al.* [4] use an improved MCD metric, called MCD-DTW, which adopts the Dynamic Time Warping (DTW) [29] algorithm to find the minimum MCD between two speeches. However, MCD-DTW would achieve a better value as long as there is a match between two speeches. This is not reasonable as a better generated speech should have a similar length with the ground truth.

To alleviate the above issues, we propose a MCD-DTW weighted by Speech Length (MCD-DTW-SL), which evaluates both the length and the quality of alignment between two speeches. In MCD-DTW-SL, to evaluate whether the two speeches (*i.e.*, \mathcal{C} and \mathcal{C}') are aligned, we still use DTW algorithm to calculate the minimum distance among them. Specifically, we compute the cumulative distance $\gamma_{i,j} = d(\mathbf{c}_i, \mathbf{c}'_j) + \min(\gamma_{i-1,j-1}, \gamma_{i-1,j}, \gamma_{i,j-1})$, where $\gamma_{i,j}$ is the minimum cumulative distance from index $(1, 1)$

Methods	MCD ↓	MCD-DTW ↓	MCD-DTW-SL ↓	Id. Acc. ↑	Emo. Acc. ↑	MOS-naturalness ↑	MOS-similarity ↑
Ground Truth	00.00	00.00	00.00	90.62	84.38	4.61 ± 0.15	4.74 ± 0.12
SV2TTS [20]	21.08	12.87	49.56	33.62	37.19	2.03 ± 0.22	1.92 ± 0.15
SV2TTS* [20]	17.41	11.16	15.92	38.21	41.24	3.20 ± 0.20	3.09 ± 0.33
FastSpeech2 [34]	12.08	10.29	10.31	59.38	53.13	3.86 ± 0.07	3.75 ± 0.06
V2C-Net (Ours)	11.79	10.09	10.05	62.50	56.25	3.97 ± 0.06	3.90 ± 0.06

Table 3. Comparison with the state-of-the-art methods. We provide the results of both objective (*i.e.*, MCD, MCD-DTW and MCD-DTW-SL) and subjective evaluation metrics (*i.e.*, MOS-naturalness and MOS-similarity). “Id. Acc.” and “Emo. Acc.” are the identity and emotion accuracy of the generated speech, respectively. The method with “*” refers to a variant taking video (emotion) embedding as an additional input. “Ground Truth” denotes the results on ground-truth samples. ↑ (↓) means that the higher (lower) value is better.

to (i, j) . Then, we obtain the objective minimum distance $\gamma_{M,N}$ by accumulating R distances in total, where $\max(M, N) \leq R < M + N - 1$. Besides, considering the influence of lengths, we design a simple but effective coefficient $\eta = \frac{\max(M,N)}{\min(M,N)}$. Formally, we calculate the metric

$$\text{MCD-DTW-SL}(\mathcal{C}, \mathcal{C}') = \frac{\eta}{R} \cdot \gamma_{M,N}. \quad (10)$$

Subjective Evaluation Metric. To further evaluate the quality of generated speech, we conduct a human study by using a subjective evaluation metric. Specifically, following the settings in [20], we use a Mean Opinion Score (MOS) evaluation approach based on subjective listening tests. In this approach, we use the Absolute Category Rating (ACR) scale [33] with rating scores from 1 to 5 (*i.e.*, from “Bad” to “Excellent”) in 0.5 point increments. Based on such an approach, we mainly evaluate the generated speeches with respect to naturalness and similarity. 1) *MOS-naturalness*: to assess the naturalness of the generated speech, we randomly sample 100 generated audios from the testing set and divide them into 4 groups. Each group is rated by a single rater. 2) *MOS-similarity*: to evaluate whether the generated speech is well aligned with the desired voice and prosody, we compare each generated speech with the ground-truth one from the same speaker. We use the same samples as when evaluating MOS-naturalness above. Each pair is rated by the rater according to the similarity between two speeches.

Identity and Emotion Accuracy. To evaluate whether the generated speech carries proper speaker identity and emotion, we propose an identity accuracy and an emotion accuracy, respectively. The identity accuracy aims to verify whether the generated speech can be recognised as the same speaker as the input reference audio. Similarly, the emotion accuracy reflects whether the generated speech contains the same emotion like the reference video. To this end, we first use GE2E [45] model as our audio encoder to obtain a fixed-dimensional audio embedding for each speech. Then, based on the audio embeddings $\{\mathbf{a}_{i1}, \dots, \mathbf{a}_{iH}\}$ ³ of the i -th speaker, we obtain the centroid \mathbf{a}'_i of the i -th speaker by $\mathbf{a}'_i = \mathbb{E}_d[\mathbf{a}_{ih}] = \frac{1}{H} \sum_{d=1}^H \mathbf{a}_{ih}$, where H refers to the number of audio belonging to the i -th speaker. Last, we compare

³The embeddings have been normalised using L2 norm.

the cosine similarity between the embedding of generated speech and each centroid, and then classify it into the category the most similar centroid belongs to. Similarly, we can calculate the emotion accuracy in the same way as well.

5.2. Quantitative Evaluation

To evaluate the performance of our method, we compare V2C-Net with several state-of-the-art methods. In Table 3, our V2C-Net consistently outperform the existing VC models (*e.g.*, the MCD-DTW-SL of our V2C-Net is 10.05 while the SV2TTS model only achieves 49.56). Besides, we propose a variant of SV2TTS model, called SV2TTS*, which takes the video embedding derived from our method as an additional input. From Table 3, SV2TTS* achieves better performance compared with SV2TTS, which further demonstrates the effectiveness of our video component. Note that the Id. Acc. and Emo. Acc. on the ground truth are not 100%. This is because our pre-trained identity and emotion classification models are not perfect.

5.3. Qualitative Evaluation

To further assess the quality of generated speech, we show the visualised results of the proposed methods, baseline method and ground-truth, respectively. In Figure 5, compared with FastSpeech2, both the energy (volume) curve and the fundamental frequency curve (*i.e.*, F_0 curve) of the mel-spectrogram generated by our V2C-Net is more similar to the ground-truth (GT) ones. Notably, as the duration of audio should be predicted as well, the lengths of the generated audio and the GT one may be different.

5.4. Effect of Reference Audio and Video

To investigate the effect of reference audio and video, we compare the generated speeches by removing them alternately and show the quantitative results (*i.e.*, identity and emotion accuracy) in Table 4. The results show that with the control of reference audio, our V2C-Net achieves higher identity accuracy obviously than the counterpart without reference audio (*i.e.*, from 25.00% to 59.38%). After further incorporating the information of reference video, the model obtains the best performance on both two metrics.




Figure 5. Mel-spectrogram of generated and ground-truth audios. Orange curves are F_0 contours, where F_0 is fundamental frequency of audio. Purple curves refer to energy (volume) of audio. Horizontal axis is duration of audio. We highlight main difference via red circle.

	ref. A	ref. V	Id. Acc. \uparrow	Emo. Acc. \uparrow
V2C-Net	\checkmark	\checkmark	25.00	47.61
	\checkmark	\checkmark	59.38	53.13
	\checkmark	\checkmark	62.50	56.25

Table 4. Effect of reference audio and video. “ref. A” and “ref. V” denote the reference audio and reference video separately. “Id. Acc.” and “Emo. Acc.” are identity and emotion accuracy of the generated speech, respectively. \uparrow means higher value is better.

Task	Dataset	MCD \downarrow	MOS-naturalness \uparrow
VC	VCTK [49]	-	4.07 ± 0.06
	LibriSpeech [31]	-	3.98 ± 0.06
	Multi-speaker [37]	12.37	-
V2C	V2C-Animation (Ours)	21.08	2.03 ± 0.22

Table 5. Comparisons on the difficulties of the conventional Voice Cloning (VC) and Visual Voice Cloning (V2C) tasks. We show the performance of SV2TTS [20] trained on different datasets. VCTK, LibriSpeech and Multi-speaker are widely used in VC task. \uparrow (\downarrow) means higher (lower) value refers to better performance.

5.5. Comparing Difficulties of V2C and VC Tasks

To investigate whether our V2C task is more challenging than the conventional VC task, we compare results of the SV2TTS method [20] on both our V2C-Animation dataset and the existing VC datasets (*e.g.*, VCTK [49], LibriSpeech [31] and Multi-speaker [37]). In Table 5, SV2TTS obtains 4.07 ± 0.06 and 3.89 ± 0.06 MOS-naturalness on VCTK and LibriSpeech, respectively, which are higher than that on our V2C-Animation dataset (2.03 ± 0.22 in Table 5). Besides, the SV2TTS model achieves 12.37 MCD value on Multi-speaker dataset [37], which is better than the results of the same model on our V2C-Animation dataset (*i.e.*, 21.08 in Table 5). It demonstrates that the proposed V2C task is more non-trivial as the same VC model obtains worse results on our V2C-Animation dataset than others.

5.6. Future Work and Discussion on Social Impacts

Limitations and Future Work. In the future, we may extend V2C-Net in two aspects. First, to grasp the emotion from video, we simply use I3D model as emotion encoder. However, it may not well disentangle emotion from char-

acter identity (see results in Table 4). To alleviate this issue, we may design an emotion-aware loss to capture more discriminative emotion features. Second, we integrate the multi-modal features (*i.e.*, text, audio, video) by a simple operation (*i.e.*, element-wise add), which may result in sub-optimal performance. Thus, a more promising model for feature fusion is necessary, *e.g.*, COOT [12] or VATT [1].

Discussion on Social Impacts. We provide a new V2C task, which takes advantage of many real-world applications, *e.g.*, movie dubbing or restoring the ability to communicate naturally for users who have lost their voice. However, the technology has a risk to be used maliciously for fake voice generation, which may be misused to financial scams by combining with video deepfakes. To avoid this issue, in this paper, we only focus on the voice generation based on animated movies without any personally identifiable information, *e.g.*, face of the real person.

6. Conclusion

In this paper, we propose a novel task, Visual Voice Cloning (V2C), extended from conventional Voice Cloning. It seeks to convert a paragraph of text to speech with the desired voice and emotion from reference audio and video, respectively. To facilitate the research of this new task, we collect the first V2C-Animation dataset. We also design a V2C baseline method, namely Visual Voice Cloning Network (V2C-Net), based on FastSpeech2 (a widely used TTS framework). Moreover, to assess the quality of the generated speech, we propose a variant of MCD-DTW, called MCD-DTW-SL, which is weighted by speech length. The experimental results demonstrate the effectiveness of our V2C-Net, but it is still far from saturation.

Acknowledgments

Jiaqiu Zhou and Mingkui Tan were supported by Ministry of Science and Technology Foundation Project (2020AAA0106900), Key Realm R&D Program of Guangzhou 202007030007. Qi Wu was funded by ARC DE190100539.

References

- [1] Hassan Akbari, Linagzhe Yuan, Rui Qian, Wei-Hong Chuang, Shih-Fu Chang, Yin Cui, and Boqing Gong. Vatt: Transformers for multimodal self-supervised learning from raw video, audio and text. *arXiv preprint arXiv:2104.11178*, 2021. 8
- [2] Sercan Ö. Arik, Jitong Chen, Kainan Peng, Wei Ping, and Yanqi Zhou. Neural voice cloning with a few samples. In *NeurIPS*, 2018. 1, 2
- [3] Sercan Ö Arik, Mike Chrzanowski, Adam Coates, Gregory Diamos, Andrew Gibiansky, Yongguo Kang, Xian Li, John Miller, Andrew Ng, Jonathan Raiman, et al. Deep voice: Real-time neural text-to-speech. In *ICML*, pages 195–204, 2017. 2
- [4] Eric Battenberg, RJ Skerry-Ryan, Soroosh Mariooryad, Daisy Stanton, David Kao, Matt Shannon, and Tom Bagby. Location-relative attention mechanisms for robust long-form speech synthesis. In *ICASSP*, pages 6194–6198, 2020. 6
- [5] Yanyao Bian, Changbin Chen, Yongguo Kang, and Zhenglin Pan. Multi-reference tacotron by intercross training for style disentangling, transfer and control in speech synthesis. In *International Speech Communication Association*, 2019. 3
- [6] Merlijn Blaauw, Jordi Bonada, and Ryunosuke Daido. Data efficient voice cloning for neural singing synthesis. In *ICASSP*, pages 6840–6844, 2019. 2
- [7] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *CVPR*, pages 6299–6308, 2017. 5
- [8] Mingjian Chen, Xu Tan, Bohan Li, Yanqing Liu, Tao Qin, Sheng Zhao, and Tie-Yan Liu. Adaspeech: Adaptive text to speech for custom voice. *ICLR*, 2021. 2
- [9] Yutian Chen, Yannis M. Assael, Brendan Shillingford, D. Budden, Scott E. Reed, H. Zen, Q. Wang, Luis C. Cobo, Andrew Trask, Ben Laurie, Çağlar Gülçehre, Aäron van den Oord, Oriol Vinyals, and N. D. Freitas. Sample efficient adaptive text-to-speech. In *ICLR*, 2019. 1, 2
- [10] Joon Son Chung, Andrew Senior, Oriol Vinyals, and Andrew Zisserman. Lip reading sentences in the wild. In *CVPR*, pages 3444–3453, 2017. 3, 4
- [11] Andrew Gibiansky, Sercan Ö. Arik, G. Diamos, J. Miller, Kainan Peng, Wei Ping, Jonathan Raiman, and Yanqi Zhou. Deep voice 2: Multi-speaker neural text-to-speech. In *NeurIPS*, 2017. 2
- [12] Simon Ging, Mohammadreza Zolfaghari, Hamed Pirsiavash, and Thomas Brox. Coot: Cooperative hierarchical transformer for video-text representation learning. *NeurIPS*, 2020. 8
- [13] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *NeurIPS*, 2014. 6
- [14] Ian J Goodfellow, Dumitru Erhan, Pierre Luc Carrier, Aaron Courville, Mehdi Mirza, Ben Hamner, Will Cukierski, Yichuan Tang, David Thaler, Dong-Hyun Lee, et al. Challenges in representation learning: A report on three machine learning contests. In *International Conference on Neural Information Processing*, pages 117–124, 2013. 2, 3
- [15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 6
- [16] Keikichi Hirose and Jianhua Tao. *Speech Prosody in Speech Synthesis: Modeling and generation of prosody for high quality and flexible speech synthesis*. Springer, 2015. 6
- [17] Wei-Ning Hsu, Y. Zhang, Ron J. Weiss, H. Zen, Y. Wu, Yuxuan Wang, Yuan Cao, Y. Jia, Z. Chen, Jonathan Shen, P. Nguyen, and Ruoming Pang. Hierarchical generative modeling for controllable speech synthesis. In *ICLR*, 2019. 3
- [18] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *CVPR*, pages 1125–1134, 2017. 6
- [19] Keith Ito and Linda Johnson. The lj speech dataset. <https://keithito.com/LJ-Speech-Dataset/>, 2017. 2, 3, 4
- [20] Ye Jia, Y. Zhang, Ron J. Weiss, Q. Wang, Jonathan Shen, Fei Ren, Z. Chen, P. Nguyen, Ruoming Pang, I. Lopez-Moreno, and Y. Wu. Transfer learning from speaker verification to multispeaker text-to-speech synthesis. In *NeurIPS*, 2018. 1, 2, 7, 8
- [21] Nal Kalchbrenner, Erich Elsen, Karen Simonyan, Seb Noury, Norman Casagrande, Edward Lockhart, Florian Stimberg, Aaron Oord, Sander Dieleman, and Koray Kavukcuoglu. Efficient neural audio synthesis. In *ICML*, pages 2410–2419, 2018. 2
- [22] Diederik P. Kingma and M. Welling. Auto-encoding variational bayes. *ICLR*, 2014. 3
- [23] Jungil Kong, Jaehyeon Kim, and Jaekyoung Bae. Hifi-gan: Generative adversarial networks for efficient and high fidelity speech synthesis. *NeurIPS*, 2020. 6
- [24] Robert Kubichek. Mel-cepstral distance measure for objective speech quality assessment. In *Proceedings of IEEE Pacific Rim Conference on Communications Computers and Signal Processing*, 1993. 2, 6
- [25] Kundan Kumar, Rithesh Kumar, Thibault de Boissiere, Lucas Gestein, Wei Zhen Teoh, Jose Sotelo, Alexandre de Brébisson, Yoshua Bengio, and Aaron Courville. Melgan: Generative adversarial networks for conditional waveform synthesis. *NeurIPS*, 2019. 6
- [26] Naihan Li, Shujie Liu, Yanqing Liu, Sheng Zhao, and Ming Liu. Neural speech synthesis with transformer network. In *AAAI*, pages 6706–6713, 2019. 2
- [27] Xudong Mao, Qing Li, Haoran Xie, Raymond YK Lau, Zhen Wang, and Stephen Paul Smolley. Least squares generative adversarial networks. In *ICCV*, pages 2794–2802, 2017. 6
- [28] Michael McAuliffe, Michaela Socolof, Sarah Mihuc, Michael Wagner, and Morgan Sonderegger. Montreal forced aligner: Trainable text-speech alignment using kaldi. In *Interspeech*, pages 498–502, 2017. 5
- [29] Meinard Müller. Dynamic time warping. *Information Retrieval for Music and Motion*, pages 69–84, 2007. 6
- [30] Eliya Nachmani, Adam Polyak, Yaniv Taigman, and Lior Wolf. Fitting new speakers based on a short untranscribed sample. In *ICML*, pages 3683–3691, 2018. 1, 2
- [31] Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur. Librispeech: an asr corpus based on public domain audio books. In *ICASSP*, 2015. 2, 3, 4, 8

- [32] Wei Ping, Kainan Peng, Andrew Gibiansky, Serkan Ö. Arik, A. Kannan, Sharan Narang, Jonathan Raiman, and J. Miller. Deep voice 3: Scaling text-to-speech with convolutional sequence learning. In *ICLR*, 2018. 2
- [33] ITUT Rec. P. 800: Methods for subjective determination of transmission quality. *International Telecommunication Union, Geneva*, 1996. 7
- [34] Yi Ren, Chenxu Hu, Xu Tan, Tao Qin, Sheng Zhao, Zhou Zhao, and Tie-Yan Liu. Fastspeech 2: Fast and high-quality end-to-end text to speech. In *ICLR*, 2021. 2, 4, 5, 6, 7
- [35] Yi Ren, Yangjun Ruan, X. Tan, Tao Qin, Sheng Zhao, Zhou Zhao, and T. Liu. Fastspeech: Fast, robust and controllable text to speech. In *NeurIPS*, 2019. 2, 5, 6
- [36] Anna Rohrbach, Marcus Rohrbach, Niket Tandon, and Bernt Schiele. A dataset for movie description. In *CVPR*, pages 3202–3212, 2015. 2, 3, 4
- [37] RJ Skerry-Ryan, Eric Battenberg, Ying Xiao, Yuxuan Wang, Daisy Stanton, Joel Shor, Ron Weiss, Rob Clark, and Rif A Saurous. Towards end-to-end prosody transfer for expressive speech synthesis with tacotron. In *ICML*, pages 4693–4702, 2018. 2, 3, 4, 6, 8
- [38] Daisy Stanton, Yuxuan Wang, and RJ Skerry-Ryan. Predicting expressive speaking style from text in end-to-end speech synthesis. In *Spoken Language Technology Workshop*, pages 595–602, 2018. 3
- [39] Antti Santeri Suni, Daniel Aalto, Tuomo Raitio, Paavo Alku, Martti Vainio, et al. Wavelets for intonation modeling in hmm speech synthesis. In *ISCA Workshop on Speech Synthesis*, 2013. 6
- [40] Yaniv Taigman, Lior Wolf, A. Polyak, and Eliya Nachmani. Voiceloop: Voice fitting and synthesis via a phonological loop. In *ICLR*, 2018. 2
- [41] Makarand Tapaswi, Yukun Zhu, Rainer Stiefelhagen, Antonio Torralba, Raquel Urtasun, and Sanja Fidler. Movieqa: Understanding stories in movies through question-answering. In *CVPR*, pages 4631–4640, 2016. 2, 3, 4
- [42] Suramya Tomar. Converting video formats with ffmpeg. *Linux Journal*, page 10, 2006. 3
- [43] Rafael Valle, Jason Li, R. Prenger, and Bryan Catanzaro. Mellotron: Multispeaker expressive voice synthesis by conditioning on rhythm, pitch and global style tokens. In *ICASSP*, pages 6189–6193, 2020. 3
- [44] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, pages 5998–6008, 2017. 2
- [45] Li Wan, Quan Wang, Alan Papir, and Ignacio Lopez Moreno. Generalized end-to-end loss for speaker verification. In *ICASSP*, pages 4879–4883, 2018. 5, 7
- [46] Yuxuan Wang, R. Skerry-Ryan, Daisy Stanton, Y. Wu, Ron J. Weiss, Navdeep Jaitly, Z. Yang, Y. Xiao, Z. Chen, S. Bengio, Quoc V. Le, Yannis Agiomyrgiannakis, R. Clark, and R. Saurous. Tacotron: Towards end-to-end speech synthesis. In *International Speech Communication Association*, 2017. 2, 3
- [47] Yuxuan Wang, Daisy Stanton, Yu Zhang, RJ-Skerry Ryan, Eric Battenberg, Joel Shor, Ying Xiao, Ye Jia, Fei Ren, and Rif A Saurous. Style tokens: Unsupervised style modeling, control and transfer in end-to-end speech synthesis. In *ICML*, pages 5180–5189, 2018. 2, 3
- [48] M. Whitehill, S. Ma, Daniel McDuff, and Yale Song. Multi-reference neural tts stylization with adversarial cycle consistency. In *International Speech Communication Association*, 2020. 3
- [49] Junichi Yamagishi, Christophe Veaux, Kirsten MacDonald, et al. Cstr vctk corpus: English multi-speaker corpus for cstr voice cloning toolkit. 2019. 2, 3, 4, 8
- [50] Yuzi Yan, Xu Tan, Bohan Li, Tao Qin, Sheng Zhao, Yuan Shen, and Tie-Yan Liu. Adaspeech 2: Adaptive text to speech with untranscribed data. In *ICASSP*, pages 6613–6617, 2021. 2
- [51] Heiga Zen, Viet Dang, Rob Clark, Yu Zhang, Ron J Weiss, Ye Jia, Zhifeng Chen, and Yonghui Wu. Libritts: A corpus derived from librispeech for text-to-speech. *arXiv preprint arXiv:1904.02882*, 2019. 2, 3, 4