

VideoDubber: Machine Translation with Speech-Aware Length Control for Video Dubbing

Yihan Wu^{1*}, Junliang Guo², Xu Tan², Chen Zhang³, Bohan Li³, Ruihua Song^{1†},
Lei He³, Sheng Zhao³, Arul Menezes⁴, Jiang Bian²

¹ Gaoling School of Artificial Intelligence, Renmin University of China

² Microsoft Research Asia,

³ Microsoft Azure Speech,

⁴ Microsoft Azure Translation

yihanwu@ruc.edu.cn, {junliangguo, xuta, zhangche, bohli, helei, szhao, arulm, jiang.bian}@microsoft.com

Abstract

Video dubbing aims to translate the original speech in a film or television program into the speech in a target language, which can be achieved with a cascaded system consisting of speech recognition, machine translation and speech synthesis. To ensure the translated speech to be well aligned with the corresponding video, the length/duration of the translated speech should be as close as possible to that of the original speech, which requires strict length control. Previous works usually control the number of words or characters generated by the machine translation model to be similar to the source sentence, without considering the isochronicity of speech as the speech duration of words/characters in different languages varies. In this paper, we propose VideoDubber, a machine translation system tailored for the task of video dubbing, which directly considers the speech duration of each token in translation, to match the length of source and target speech. Specifically, we control the speech length of generated sentence by guiding the prediction of each word with the duration information, including the speech duration of itself as well as how much duration is left for the remaining words. We design experiments on four language directions (German \rightarrow English, Spanish \rightarrow English, Chinese \leftrightarrow English), and the results show that VideoDubber achieves better length control ability on the generated speech than baseline methods. To make up the lack of real-world datasets, we also construct a real-world test set collected from films to provide comprehensive evaluations on the video dubbing task.

1 Introduction

Video dubbing¹ aims to translate the original speech in a movie or TV from a source language to the speech in a target language. Generally, automatic video dubbing system consists of three cascaded sub-tasks, i.e., Automatic

*Work done during the first author is interning at Microsoft Research Asia.

†Corresponding author

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹There are usually two scenarios for video dubbing: 1) translating the speech of a video from one language to another; 2) generating speech in the same language according to the text transcripts of a video. In this paper, we focus on the first video dubbing for translation scenario.

Speech Recognition (ASR) (Yu and Deng 2016), Neural Machine Translation (NMT) (Vaswani et al. 2017a) and Text-to-Speech (TTS) (Tan et al. 2021, 2022). Specifically, ASR transcribes the original speech into text in the source language, which can be omitted when subtitles are available. NMT aims to translate text in the source language to the target language, and TTS finally synthesizes the translated text into target speech. Different from speech-to-speech translation (Wahlster 2013), video dubbing requires strict isochronous constraints between the original source speech and the synthesized target speech to ensure that the speech matches the original video footage in terms of length and thus provide immersive watching feeling, which brings additional challenges.

In the cascaded video dubbing system, both NMT and TTS can determine the length of translated speech, i.e., the number of translated words is determined by the NMT model, while the length of the synthesized speech can be controlled by TTS via adjusting the pause and duration of words. Previous automatic video dubbing systems control the speech isochrony separately, by first controlling the number of generated target words/characters similar to that of the source words, and then adjusting the speaking rate to ensure the duration of synthesized speech similar to that of the source speech. However, as the speech duration of words/characters in different languages varies, the same number of words/characters in the source and target sentences does not guarantee the same length of speech. Therefore, TTS has to adjust the speaking rate of each word in a wide range to match the total speech length, which will affect the fluency and naturalness of synthesized speech (e.g., the speaking rates may be inconsistent with those of adjacent sentences).

In this paper, we tackle the length control problem in a systematic way. Our idea starts from a simple motivation: matching the length of the synthesized speech with the original one while avoiding speaking rate inconsistency of adjacent sentences. To preserve high hearing quality, we should assign the responsibility for length control more to NMT and less to TTS, which requires NMT can be controlled by original speech length while maintaining translation quality. To achieve this goal, we propose the following techniques. Firstly, to control the speech duration when translating the

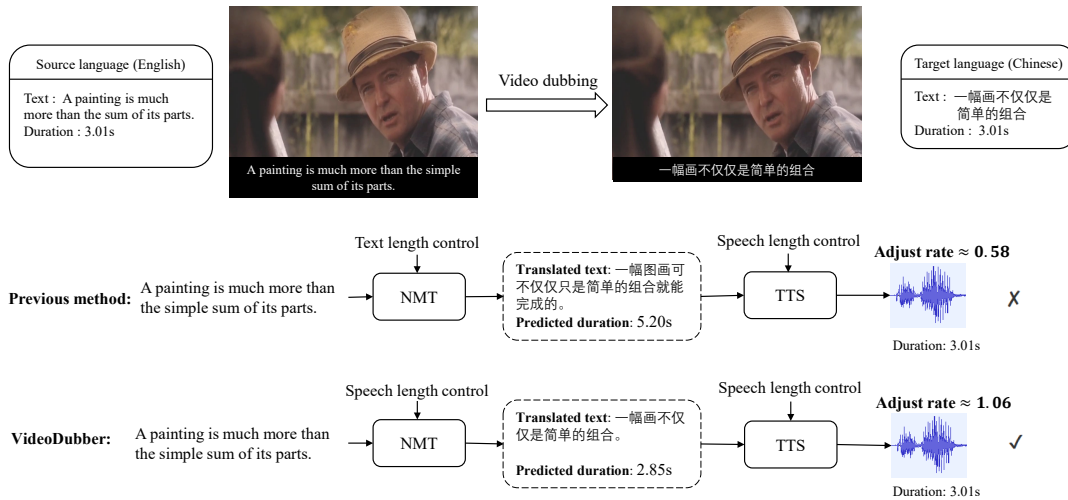


Figure 1: An example of video dubbing from English to Chinese. When translating the source sentence “A painting is much more than the simple sum of its parts.” in NMT stage, its corresponding speech duration is 3.01s in original speech. Previous method control the number of words of translated sentence, but the length of corresponding speech is very different from the original speech without adjustment. The translation result from our method considering speech length directly, thus generate translated speech that are very close to the original one. Then the TTS model generates the desired translation result to speech with a little duration adjustment.

target sentence, we make the model aware of how much duration is left at each time-step as well as the speech duration of each word, by representing it as two kinds of positional embeddings which are added to the original one and taken as input of the decoder. Secondly, we introduce a special pause word [P] which is inserted between each target word, in order to control the speech length more smoothly by considering the prosody (language-related tempo, rhythm and pause) of speech through adjusting the duration of [P]. Finally, while training, we calculate the alignment between the target speech and text to obtain the speech duration of each target word instead of using only the number of words/characters. Also we introduce a neural network component named duration predictor to explicitly model the speech duration of each word. In inference, given the source speech length, the duration predictor works together with the decoder to generate high quality translations with similar speech length to the source. Moreover, considering the scarcity of real-world video dubbing dataset (i.e., motion pictures with golden cross-lingual source and target speech available), we construct a test set collected from dubbed films to provide comprehensive evaluations of video dubbing systems.

The main contributions of this work are summarized as follows:

- We propose VideoDubber, a machine translation system tailored for the task of video dubbing, which directly controls the speech length of generated sentence by calculating and encoding the speech duration of each word into the model.
- We conduct experiments on four language directions, with both objective and subjective evaluation metrics. Experiment results shows that VideoDubber achieves

better length control compared with baseline method in all languages.

- Due to the scarcity of the video dubbing test dataset, we construct a real-world test set collected from films to provide comprehensive evaluation on the video dubbing task.

2 Background

We first provide the problem definition of the video dubbing task.

Problem Definition Given the source speech s with duration T_s and its transcription $x = (x_1, \dots, x_{T_x})$, we aim to generate the translated text sequence $y = (y_1, \dots, y_{T_y})$ as well as the synthesized speech t with duration T_t , where T_t should be as close to T_s as possible, while maintaining the high translation quality of y as well as the rhythmic and fluency of t .

Existing video dubbing works (Öktem, Farrús, and Bonafonte 2019; Federico et al. 2020b; Lakew et al. 2021; Virkar et al. 2021; Sharma et al. 2021; Effendi et al. 2022; Lakew et al. 2022; Virkar et al. 2022; Tam et al. 2022) are usually based on a cascaded speech-to-speech translation system (Federico et al. 2020a) with ad-hoc designs, mainly concentrating on the Neural Machine Translation (NMT) and Text-To-Speech (TTS) stages. In the NMT stage, related works achieve the length control by assuming that similar number of words/characters should have similar speech length, and therefore encourage a model to generate target sequence with similar number of words/characters to the source sequence (Federico et al. 2020a). They compute the length ratio between source and target text sequences as the

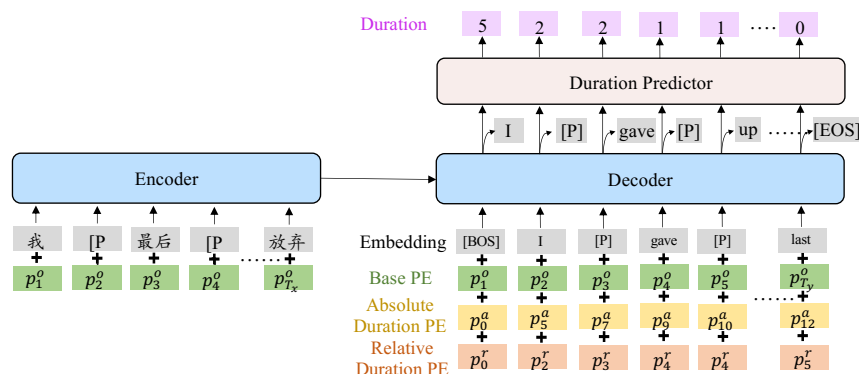


Figure 2: The overall architecture of VideoDubber with speech-aware length control for video dubbing. PE stands for positional embedding and [P] indicates the special pause token. We set $N = 5$ and follow Equation (2) to calculate the relative duration PE.

verbosity information during training, and control translated speech length by setting length ratio to 1 during inference. Lakew et al. (2022) introduce a character-level length constraint to the NMT model, by augmenting the training set with length controlled back translation. Moreover, following previous works (Takase and Okazaki 2019; Lakew, Gangi, and Federico 2019), Tam et al. (2022) encode the number of characters that are going to be generated into the position embedding to incorporate the length control regularization. These methods only control the number of generated words/characters to be similar to the source sequence, without considering the discrepancy of the speech duration of words/characters in source and target languages.

The TTS stage synthesizes the translated text sequence to speech with isochronous constraints, by adjusting the speaking rate to ensure the duration of the generated speech fits the source speech (Sharma et al. 2021; Effendi et al. 2022). Further, Effendi et al. (2022) present a duration model which can predict and adjust the phoneme duration by uniform normalization or non-isoelectric normalization (Effendi et al. 2022). To relax the duration adjustment effort and match the pauses and phrases in the target to the source speech, several works (Federico et al. 2020a,b; Virkar et al. 2021) introduce an additional prosodic alignment module before the TTS stage. It inserts pause tokens into the translated text considering the speaking rate match between corresponding source-target phrases and the linguistic plausibility of the chosen split points. However, this two-stage rule-based approach cannot take advantage of natural pauses and rhythm from speech directly. Moreover, it only inserts pauses between phrases, ignoring that the pauses in different granularities usually correspond to different durations. As the audience is sensitive to the speaking rate change of speech, the inconsistent speaking rates of adjacent sentences, e.g., faster at the previous sentence while slower at the current one, will destroy the rhythm and fluency of the speech. Therefore, in this paper, we try to alleviate the burdens of length control on TTS and concentrate on the NMT stage, by proposing a speech-aware length control model to directly control the speech duration instead of the number of tokens of the gen-

erate target sequence.

3 Proposed Method

We introduce VideoDubber in this section. Specifically, we control the speech length of generated sentence by guiding the prediction of each word with duration information. An illustration of the model architecture is shown in Figure 2.

3.1 Overview

To match the length of synthesized speech with the original one while maintaining fluency and naturalness, we should assign the responsibility for length control more to NMT and less to TTS. Thus, the goal becomes how to achieve speech-aware length control on NMT while maintaining high translation quality. To achieve this goal, we first obtain the speech duration $d = (d_1, \dots, d_{T_y})$ of each target word and then incorporate it into the NMT model.

Specifically, we integrate the duration information by designing two kinds of duration-aware positional embeddings. One is *absolute* duration positional embedding which indicates the accumulated duration information at the current time-step, the other is *relative* duration positional embedding calculated as the ratio between the absolute duration and the total duration, indicating how much duration is left for future tokens. In this way, the model is trained to jointly consider the semantic and duration information when making predictions. We also take pauses into consideration to control the speech length with more flexibility by introducing a special pause token [P].

To obtain the speech duration for each target word, while training, we first calculate the alignment between speech and text in phoneme-level through Montreal forced alignment (MFA) (McAuliffe et al. 2017)². Then, we introduce a duration predictor, which is a neural network component consisting of convolutional layers and being inserted on the top of the decoder, to predict the duration of each word conditioned on the hidden output of decoder. Next in inference, given a

²MFA is an open-source tool for speech-text alignment with good performance.

total speech length (i.e., the length of source speech), the decoder will determine the appropriate translation conditioned on both the semantic representation and the duration information at each step in an autoregressive manner as shown in Figure 2. We dive into details in the following section.

3.2 Model Architecture

Duration-aware Position Embedding To take account of the speech duration in translation, we integrate the duration information via designing duration-aware positional embeddings for each token. Here, we design two kinds of positional embedding to provide the absolute and relative speech-aware duration information respectively.

We denote the original positional embedding in Transformer (Vaswani et al. 2017b) as p^o . Then given the i -th target token y_i and its duration d_i , we define the absolute duration positional embedding p_i^a by the accumulated speech duration up-to-now, which can be written as:

$$p_i^a = PE\left(\sum_{j=1}^i d_j\right), \quad (1)$$

where $PE(\cdot)$ indicates the lookup function over the sinusoidal positional embedding matrix.

The introduced p^a tells the model the up-to-now speech length of generated words. Moreover, to make the model aware of how much duration is left and therefore plan the generation by choosing words with the appropriate duration, we also introduce a relative duration positional embedding p^r to incorporate the global duration information. It is calculated as the ratio between the accumulated duration and the total duration:

$$p_i^r = PE\left(q_N\left(\frac{\sum_{j=0}^i d_j}{\sum_{j=0}^{T_y} d_j}\right)\right), \quad (2)$$

where $q_N(x) = \lfloor x \times N \rfloor$ quantizes the float duration ratio from $[0, 1]$ into integers within $[0, N]$. Finally, given the word embedding of the i -th target token w_i , we add the three positional embeddings with it to construct the input to decoder:

$$h_i = w_i + p_i^o + p_i^a + p_i^r. \quad (3)$$

The proposed two position embeddings are complementary to each other, and we evaluate their effectiveness in experiments.

Pause Token We further consider the pauses between words and adjust their duration to ensure the flexibility of length control, i.e., we can increase or decrease the duration of pauses automatically to share some responsibility for length control, which can avoid uneven speaking rate between adjacent sentences and is crucial for the fluency of speech. As a byproduct, pauses can also be utilized to model the prosody of the speech. Specifically, we explicitly model the pauses in speech by introducing a special token [P], which is inserted between each word (instead of sub-words) in both the source and target sentences. For example, the phrase “A painting” is tokenized as “A pain@@

ting” after conducting Byte-Pair Encoding (BPE) (Sennrich, Haddow, and Birch 2015), which is then transformed to “A [P] pain@@ ting” after inserting the pause token. The duration of the pause token can also be obtained through MFA, and it varies in different context. In this way, the model learns to predict a pause token after each word with appropriate duration in context, providing a more flexible way for speech length control.

Duration Predictor Following FastSpeech 2 (Ren et al. 2021), we use MFA tool to compute ground truth word duration for training. To make the model able to predict the speech duration of each token when the ground truth duration is absent in inference, we introduce an extra neural network component called duration predictor, which is built on top of the decoder. While training, it takes the hidden representation of each token as input, and predicts the golden duration of each token, which represents how many mel-frames correspond to this token, and is converted into the logarithmic domain for ease of prediction. The duration predictor is optimized with the mean square error (MSE) loss, taking the ground truth duration as training target. Denoting the duration sequence as $d = (d_1, \dots, d_{T_y})$, the prediction can be written as:

$$p(d|x, y) = \prod_{i=1}^{T_y} p(d_i|x, y_{<i}; \theta^{du}), \quad (4)$$

where θ^{du} indicates the parameters of duration predictor.

3.3 Discussion

Overall Length Control Pipeline After text translation with speech-aware length control as we introduced in Section 3.2, we further achieve more precise length control through duration adjustment in TTS stage. We employ AdaSpeech 4 (Wu et al. 2022), a zero-shot TTS model which can adjust duration directly. Instead of uniform duration adjustment which scales the duration of each phoneme by a same factor, we only adjust the duration of vowels and keep the length of consonants unchanged, since the duration of consonants is not significantly changed when people slow down or speed up their speech in natural (Kruspe 2015). With this systematic approach on NMT and TTS, we can achieve precise and flexible speech length control, so as to ensure the translated speech to be well aligned with the corresponding video in video dubbing scenario.

About the Pause Token Some related works (Federico et al. 2020a; Virkar et al. 2021) also introduce the pause token into the video dubbing framework, by inserting the pause token between phrases in the translated text with a dynamic programming algorithm. VideoDubber differs from theirs from two perspectives. Firstly, we provide more fine-grained control of the pause by inserting the pause token into each word instead of phrase. Secondly, the speech duration of each pause token can be different in different context and end-to-end learned by the model, which provides better flexibility on length control and prosody modeling of synthesized speech.

About Video Dubbing Dataset Due to the scarcity of video dubbing dataset, we usually conduct experiments on speech-to-speech translation datasets, where the lengths of the source and target speech in the training and test data are not guaranteed to be similar. Thus, this is actually different from the real-world video dubbing, where the length of source and target speech exactly match. It is not accurate to evaluate the semantic translation quality of the video dubbing system using the reference translation that is not in same length of the source speech (i.e., the controlled speech length in NMT model). Therefore, we try to alleviate this issue in two aspects: 1) In evaluation, we use the length of source and target speech respectively to control the NMT model, where the length of source speech is consistent with what we practically use in video dubbing, and the length of target speech is used to evaluate how our model performs when the length of reference sentence is the same as the control length. 2) We additionally construct a real-world video dubbing test set from dubbed films, where the source and target speech have exactly the same length. In this way, the constructed test set can be used to evaluate the video dubbing system in practical scenarios.

Another potential issue is that the NMT model is trained to generate text with the target speech length (which is different from the source speech length), but in inference, it is used to generate text with the source speech length, which causes mismatch. However, it is really hard to mitigate this issue since large-scale training data with the same source and target length is very difficult to collect. Nevertheless, we plan to construct a training dataset with the same speech length in the source and target utterances by knowledge distillation, where the distilled target sentences are controlled by our NMT model to have the same speech length with the source sentences. We leave it for future work.

4 Real-World Video Dubbing Test Set

Considering the scarcity of real-world video dubbing dataset (i.e., motion pictures with golden cross-lingual source and target speech), we construct a test set collected from dubbed films to provide comprehensive evaluations of video dubbing systems.

Specifically, we select nine popular films translated from English to Chinese, which are of high manual translation and dubbing quality, and contain rich genres such as love, action, scientific fiction, etc. We cut 42 conversation clips from them with the following criteria: 1) The clip duration is around 1 ~ 3 minutes. 2) More than 10 sentences are involved in each clip, which contains both long and short sentences. 3) The face of speaker is visible mostly during his or her talks, especially visible lips at the end of speech.

After that, we extract audios from video clips and remove background sounds to get clean speech. Then, we perform speech recognition and manual correction to obtain text transcripts in source and target languages. In addition, we split clips into sentences based on semantics and speakers. We discard silence frames between sentences to make sure there is no more than 0.5s of silence at the beginning and end of each split.

Finally, we obtain the test dataset with a total duration of 1 hour, including the original movie, source speech with transcripts, and human-dubbed speech with transcripts. It contains 892 sentences for sentence-level evaluation³.

5 Experimental Setup

5.1 Datasets

Training Data We train VideoDubber on four language directions: Chinese → English (Zh-En), English → Chinese (En-Zh), German → English (De-En), Spanish → English (Es-En). Since the absence of real-world video dubbing dataset, we train and test VideoDubber in the speech translation dataset. For language direction from Others → English, we use public speech-to-speech translation dataset CVSS (Jia et al. 2022), which contains multilingual-to-English speech-to-speech translation corpora derived from the CoVoST2 dataset (Wang et al. 2021). For the direction from English to Chinese, we use the En-Zh subset of MuST-C (Cattoni et al. 2021), an English → Others speech translation corpus built from English TED Talks. Since the Must-C dataset does not have corresponding speech in the target language, we use a well trained Chinese TTS model, FastSpeech 2 (Ren et al. 2021), to generate Chinese speech from the translated text.

Evaluation Data We consider two datasets in evaluation. One is the standard test set from the speech translation dataset that follows the settings of previous works, and the other is the constructed real-world test set for the video dubbing scenario, to test the overall performance in the real-world setting.

5.2 Evaluation Metrics

We adopt both subjective and objective evaluation metrics to evaluate the compared video dubbing systems.

Objective Metrics We use BLEU and speech length compliant to evaluate the semantic quality of the translated text and the length control ability of the model.

- **BLEU.** BLEU (Papineni et al. 2002) score is a widely used automatic evaluation metric of machine translation task. We report the tokenized BLEU score to keep consistent with previous works.
- **Speech Length Compliant.** To evaluate the speech length control performance, inspired by Lakew et al. (2022), we design the Speech Length Compliant (SLC_p) score. SLC_p refers to the percentage of sentences whose $ratio \in [1 - p, 1 + p]$, where the ratio is defined as:

$$ratio = \frac{\sum_{i=0}^{T_{y'}} d_i}{\sum_{j=0}^{T_x} d_j}, \quad (5)$$

where d_i, d_j represent the duration of token i, j from source speech and translated speech respectively. SLC_p refers to the percentage of sentences that $ratio \in [1 - p, 1 + p]$. Federico et al. (2020b) set $p = 0.4$ to

³The test set can be built following <https://speechresearch.github.io/videodubbing/>

Settings	Es-En			De-En			Zh-En			En-Zh		
	BLEU	SLC_p		BLEU	SLC_p		BLEU	SLC_p		BLEU	SLC_p	
		0.4	0.2		0.4	0.2		0.4	0.2		0.4	0.2
Transformer	41.70	74.14	41.89	39.92	71.46	42.19	15.46	80.33	41.80	15.31	38.24	13.82
Baseline	32.61	75.09	41.03	31.32	70.65	41.19	13.60	82.64	54.37	14.12	57.76	24.09
VideoDubber (Target)	41.67	80.30	52.87	39.10	78.52	48.82	14.66	91.43	68.82	14.32	72.19	60.05
VideoDubber (Source)	39.60	79.39	45.31	36.64	75.76	44.83	13.77	90.10	57.03	13.95	74.55	41.96

Table 1: BLEU scores and SLC_p score (\uparrow) on CVSS Es-En, CVSS De-En, CVSS Zh-En, MuST-C En-Zh. Here, we set $p = 0.4$ and $p = 0.2$ respectively. We compare our method with Transformer (Vaswani et al. 2017b) (without any length control strategy) and baseline (Federico et al. 2020a) (with text length control). And “VideoDubber (Source)” and “VideoDubber (Target)” represents controlling the translation with the length of the source speech and golden target speech respectively.

evaluate the fluency of synthesized audio, while we find that $p = 0.2$ is a more reasonable interval for users from the user study (i.e., speaking rate speedup by $1.4x$ or slowdown by $0.6x$ will cause large inconsistency that affects user watching experience). Therefore, we use both $p = 0.2$ and $p = 0.4$ and compare their corresponding SLC_p in the experiments respectively.

Subjective Metrics In addition to objective metrics, we also introduce some subjective metrics to evaluate VideoDubber on the real-world video dubbing scenario.

- **Translation quality.** Translation quality measures the quality of translated text of the dubbed video through human evaluation. It should be mentioned that judges will rate the translation quality according to the original video, which is more suitable for video dubbing scenario.
- **Isochronism.** Isochronism score measures whether the original video matches the translated speech in video dubbing scenario.
- **Naturalness.** Naturalness score measures the overall performance of the video with dubbed speech. The judges will consider both the translation quality and the isochronicity.

5.3 Model Configurations

We follow the Transformer (Vaswani et al. 2017b) based encoder-decoder framework as the backbone of our translation model. For all settings, we set the hidden dimension as 512 for the model, 2048 for the feed-forward layers. For both encoder and decoder, we use 6 layers and 8 heads for the multi-head attention. The duration-aware positional embeddings are added only at the input to the decoder stacks, and the encoder architecture is consistent with the original Transformer.

The duration predictor consists of a two-layer convolutional network with ReLU activation, each followed by the layer normalization and the dropout layer. The extra linear layer outputs a scalar, which is the predicted token duration. Note that this module is jointly trained with the Transformer model to predict the duration of each token with the mean square error (MSE) loss. We predict the duration in the logarithmic domain, which makes them easier to train. It should

be mentioned that duration predictor is only used in inference stage, since the ground truth token duration can be directly used in training.

In TTS stage, we employ a pre-trained zero-shot TTS model, AdaSpeech 4 (Wu et al. 2022). Given a reference speech, it can synthesize voice for unseen speakers. Here, we regard source speech as reference, generating speech with duration adjustment.

5.4 Baseline

We compare VideoDubber with well applied baseline (Federico et al. 2020a), which control the length of translated text in NMT by using the number of words. According to the target/source token number ratio, it splits the sentences from training dataset into three groups (short, normal, long). At training phrase, it pretends the corresponding length ratio token to source sequence. At inference phrase, the desired length token (i.e. normal) is pretended to source sequence to control the length of translated text.

6 Results

6.1 Automatic Evaluation

We show the machine translation quality and the length control performance of related models on four language directions in Table 1. As discussed in Section 3.3, we also list the results when controlling the translation with the length of the golden target speech, to show the upper-bound performance of our model. When considering the length control ability measured by SLC_p , VideoDubber consistently outperforms considered baselines with a large margin, illustrating that the proposed speech-aware length control achieves better speech duration isochronism than controlling the numbers of words/characters (Federico et al. 2020a). Especially, on the En-Zh direction, we find that VideoDubber brings significant improvements over the token number controlling baseline. The reason may be the large inconsistency between the token number and speech duration in Chinese, i.e., the token in Chinese may contain several Chinese characters instead of only a sub-word in English. As a result, even small mismatch on the token number may cause large difference on the speech duration in Chinese. VideoDubber alleviates this problem by directly controlling the speech length.

Settings	Es-En			De-En			Zh-En			En-Zh		
	BLEU	SLC _p		BLEU	SLC _p		BLEU	SLC _p		BLEU	SLC _p	
		0.4	0.2		0.4	0.2		0.4	0.2		0.4	0.2
VideoDubber	39.60	79.39	45.31	36.64	75.76	44.83	13.77	90.10	57.03	13.95	74.55	41.96
<i>w/o</i> abs PE	41.58	75.70	42.88	39.92	71.46	42.19	15.46	80.33	41.80	15.31	42.88	28.43
<i>w/o</i> rel PE	40.53	74.57	42.30	34.83	71.60	39.46	14.73	82.56	52.34	15.21	54.55	30.50
<i>w/o</i> original PE	39.62	73.98	41.55	37.68	73.46	43.42	13.02	82.38	52.29	14.05	53.95	29.65

Table 2: Ablation study on CVSS Es-En, CVSS De-En, CVSS Zh-En, MuST-C En-Zh. “abs PE” indicates absolute positional embedding and “rel PE” indicates relative positional embedding. Note that unless otherwise stated, “VideoDubber” refers to our proposed method which controls the translation with the length of source speech.

Settings	Objective			Subjective		
	BLEU	SLC _{0.4}	SLC _{0.2}	Translation Quality	Isochronous	Naturalness
Transformer (Vaswani et al. 2017b)	7.21	41.75%	27.25%	3.45	2.97	3.00
Baseline (Federico et al. 2020a)	6.56	52.17%	38.16%	3.64	3.30	3.40
VideoDubber	7.31	79.17%	66.67%	3.98	3.41	3.60

Table 3: Objective and subjective evaluation on real-world video dubbing test set. We compare our method with Transformer (Vaswani et al. 2017b) (without any length control strategy) and baseline (Federico et al. 2020a) (with text length control).

Regarding the BLEU scores, although with slightly drop compared with the vanilla Transformer model, VideoDubber outperforms the length control baseline, especially in Es-En and De-En. The baseline method performs coarse length control, with the same length control embedding for all sentences. It may lead to unstable translation ability in some datasets. In addition, it is worth noting that the source and target sentences in the test set do not have the same speech length. It is unfair to judge the translation quality with BLEU scores. Therefore we propose a new real-world test set and conduct fair comparisons on it, as illustrated in Section 6.3.

6.2 Ablation Study

To verify the effectiveness of the duration-aware positional embedding, we conduct ablation studies on three kinds of PEs on four language directions, as shown in Table 2. We can find that the absolute and relative duration PE are both crucial to achieve better speech-aware length control results.

6.3 Results on Real-world Test Set

To compare the performance of related methods on the real-world video dubbing scenario, we conduct experiments on the real-world test set constructed by us. Results are shown in Table 3.

With objective evaluation, VideoDubber achieves the highest BLEU and SLC_p score compared with Transformer and the token number control baseline. It proves that in the real test set when the speech isochronous is considered, the proposed NMT model with speech-aware length control can achieve better isochronous control ability as well as the translation quality.

Also we conduct subjective evaluation to evaluate the translation quality, the synchronization with the original

film footage, and the overall quality of the synthesized speech. We hired 8 judges who are good at both Chinese and English, asking them to rate the samples generated from different method on the 5 scale. From Table 3, aligned with the objective evaluation, we observe that VideoDubber achieves the highest translation quality and the isochronicity of speech. Moreover, our method achieves significantly better performance regarding the Naturalness score, which reflects the overall quality of the automatic dubbed video. Demo samples are available at <https://speechresearch.github.io/videodubbing/>.

7 Conclusion

In this paper, we propose VideoDubber, a machine translation model with speech-aware length control for the task of video dubbing. To ensure the translated speech to be well aligned with the original video, we directly consider the speech duration of each token in translation. Firstly, we guide the prediction of each word with the duration information, by representing it as two kinds of positional embeddings. Secondly, we introduce a special pause word [P], which is inserted between each word, in order to control the speech length more smoothly by considering the prosody. Thirdly, we construct a real-world test set collected from dubbed films to provide more actual evaluations of video dubbing systems. Experimental results demonstrate that VideoDubber achieves better translation quality and isochronous control ability than related works. For future work, we plan to construct a training dataset with the same speech length in the source and target utterances by knowledge distillation, where the distilled target sentences are controlled by our NMT model to have the same speech length with the source sentence.

References

- Cattoni, R.; Di Gangi, M. A.; Bentivogli, L.; Negri, M.; and Turchi, M. 2021. MuST-C: A multilingual corpus for end-to-end speech translation. *Computer Speech & Language*, 66: 101155.
- Effendi, J.; Virkar, Y.; Barra-Chicote, R.; and Federico, M. 2022. Duration Modeling of Neural TTS for Automatic Dubbing. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 8037–8041. IEEE.
- Federico, M.; Enyedi, R.; Barra-Chicote, R.; Giri, R.; Isik, U.; Krishnaswamy, A.; and Sawaf, H. 2020a. From Speech-to-Speech Translation to Automatic Dubbing. In *Proceedings of the 17th International Conference on Spoken Language Translation*, 257–264. Online: Association for Computational Linguistics.
- Federico, M.; Virkar, Y.; Enyedi, R.; and Barra-Chicote, R. 2020b. Evaluating and Optimizing Prosodic Alignment for Automatic Dubbing. In *INTERSPEECH*, 1481–1485.
- Jia, Y.; Ramanovich, M. T.; Wang, Q.; and Zen, H. 2022. CVSS Corpus and Massively Multilingual Speech-to-Speech Translation. *CoRR*, abs/2201.03713.
- Kruspe, A. M. 2015. Training Phoneme Models for Singing with "Songified" Speech Data. In *ISMIR*, 336–342.
- Lakew, S. M.; Federico, M.; Wang, Y.; Hoang, C.; Virkar, Y.; Barra-Chicote, R.; and Enyedi, R. 2021. Machine translation verbosity control for automatic dubbing. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 7538–7542. IEEE.
- Lakew, S. M.; Gangi, M. A. D.; and Federico, M. 2019. Controlling the Output Length of Neural Machine Translation. In *IWSLT*. Association for Computational Linguistics.
- Lakew, S. M.; Virkar, Y.; Mathur, P.; and Federico, M. 2022. Isometric mt: Neural machine translation for automatic dubbing. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 6242–6246. IEEE.
- McAuliffe, M.; Socolof, M.; Mihuc, S.; Wagner, M.; and Sonderegger, M. 2017. Montreal Forced Aligner: Trainable Text-Speech Alignment Using Kaldi. In *INTERSPEECH*, 498–502. ISCA.
- Öktem, A.; Farrús, M.; and Bonafonte, A. 2019. Prosodic Phrase Alignment for Machine Dubbing. In Kubin, G.; and Kacic, Z., eds., *Interspeech 2019, 20th Annual Conference of the International Speech Communication Association, Graz, Austria, 15-19 September 2019*, 4215–4219. ISCA.
- Papineni, K.; Roukos, S.; Ward, T.; and Zhu, W. 2002. Bleu: a Method for Automatic Evaluation of Machine Translation. In *ACL*, 311–318. ACL.
- Ren, Y.; Hu, C.; Tan, X.; Qin, T.; Zhao, S.; Zhao, Z.; and Liu, T. 2021. FastSpeech 2: Fast and High-Quality End-to-End Text to Speech. In *ICLR*. OpenReview.net.
- Sennrich, R.; Haddow, B.; and Birch, A. 2015. Neural machine translation of rare words with subword units. *arXiv preprint arXiv:1508.07909*.
- Sharma, M.; Virkar, Y.; Federico, M.; Barra-Chicote, R.; and Enyedi, R. 2021. Intra-Sentential Speaking Rate Control in Neural Text-To-Speech for Automatic Dubbing. In *Interspeech*, 3151–3155.
- Takase, S.; and Okazaki, N. 2019. Positional Encoding to Control Output Sequence Length. In *NAACL-HLT (1)*, 3999–4004. Association for Computational Linguistics.
- Tam, D.; Lakew, S. M.; Virkar, Y.; Mathur, P.; and Federico, M. 2022. Isochrony-Aware Neural Machine Translation for Automatic Dubbing. In Ko, H.; and Hansen, J. H. L., eds., *Interspeech 2022, 23rd Annual Conference of the International Speech Communication Association, Incheon, Korea, 18-22 September 2022*, 1776–1780. ISCA.
- Tan, X.; Chen, J.; Liu, H.; Cong, J.; Zhang, C.; Liu, Y.; Wang, X.; Leng, Y.; Yi, Y.; He, L.; et al. 2022. NaturalSpeech: End-to-End Text to Speech Synthesis with Human-Level Quality. *arXiv preprint arXiv:2205.04421*.
- Tan, X.; Qin, T.; Soong, F.; and Liu, T.-Y. 2021. A survey on neural speech synthesis. *arXiv preprint arXiv:2106.15561*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 2017a. Attention is All you Need. In *NIPS*, 5998–6008.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 2017b. Attention is All you Need. In Guyon, I.; von Luxburg, U.; Bengio, S.; Wallach, H. M.; Fergus, R.; Vishwanathan, S. V. N.; and Garnett, R., eds., *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, 5998–6008.
- Virkar, Y.; Federico, M.; Enyedi, R.; and Barra-Chicote, R. 2021. Improvements to prosodic alignment for automatic dubbing. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 7543–7574. IEEE.
- Virkar, Y.; Federico, M.; Enyedi, R.; and Barra-Chicote, R. 2022. Prosodic Alignment for off-screen automatic dubbing. *arXiv preprint arXiv:2204.02530*.
- Wahlster, W. 2013. *Verbmobil: foundations of speech-to-speech translation*. Springer Science & Business Media.
- Wang, C.; Wu, A.; Gu, J.; and Pino, J. 2021. CoVoST 2 and Massively Multilingual Speech Translation. In Hermansky, H.; Cernocký, H.; Burget, L.; Lamel, L.; Scharenborg, O.; and Motlíček, P., eds., *Interspeech 2021, 22nd Annual Conference of the International Speech Communication Association, Brno, Czechia, 30 August - 3 September 2021*, 2247–2251. ISCA.
- Wu, Y.; Tan, X.; Li, B.; He, L.; Zhao, S.; Song, R.; Qin, T.; and Liu, T. 2022. AdaSpeech 4: Adaptive Text to Speech in Zero-Shot Scenarios. In Ko, H.; and Hansen, J. H. L., eds., *Interspeech 2022, 23rd Annual Conference of the International Speech Communication Association, Incheon, Korea, 18-22 September 2022*, 2568–2572. ISCA.
- Yu, D.; and Deng, L. 2016. *Automatic speech recognition*, volume 1. Springer.