

FT2TF: First-Person Statement Text-To-Talking Face Generation

Xingjian Diao, Ming Cheng, Wayner Barrios, SouYoung Jin
 Dartmouth College

{xingjian.diao.gr, ming.cheng.gr, wayner.j.barrios.quiroga.gr, souyoung.jin}@dartmouth.edu

Abstract

Talking face generation has gained immense popularity in the computer vision community, with various applications including AR, VR, teleconferencing, digital assistants, and avatars. Traditional methods are mainly audio-driven, which have to deal with the inevitable resource-intensive nature of audio storage and processing. To address such a challenge, we propose **FT2TF** – First-Person Statement Text-To-Talking Face Generation, a novel one-stage end-to-end pipeline for talking face generation driven by first-person statement text. Different from previous work, our model only leverages visual and textual information without any other sources (e.g., audio/landmark/pose) during inference. Extensive experiments are conducted on LRS2 and LRS3 datasets, and results on multi-dimensional evaluation metrics are reported. Both quantitative and qualitative results showcase that FT2TF outperforms existing relevant methods and reaches the state-of-the-art. This achievement highlights our model’s capability to bridge first-person statements and dynamic face generation, providing insightful guidance for future work.

1. Introduction

Talking face generation, with the goal of creating a sequence of faces to simulate the appearance of a person talking, has surged in significance amid the metaverse boom and technological advancements. This pursuit has gained substantial popularity and importance within the computer vision community, influencing applications including augmented reality [17], virtual reality [9], computer games [32, 33], teleconferencing [3, 7], video dubbing [11, 19], special effects for movies [46], and the development of digital assistants [36] and avatars [9, 52].

Traditional methods for talking face generation mainly rely on audio input, which leverages visual and audio cues. These methods, exemplified by works like [18, 31, 49, 50, 52], show impressive performance. However, they come with a significant drawback: the need for high-quality audio. Compared to audio-driven methods, text-driven ap-

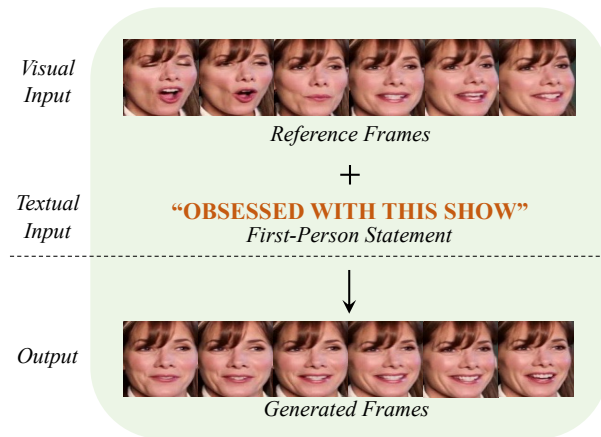


Figure 1. **High-quality text-driven talking face generation.** FT2TF aims to generate realistic talking faces using two inputs: (i) reference talking face frames and (ii) first-person statement text.

proaches offer unparalleled advantages, notably in terms of:

- **Data Storage:** Text data storage requires far less space than audio data. For instance, storing the text of a 300-page article requires a mere 128 KB, while storing audio for the same content demands several gigabytes.
- **Data Transmission:** Current audio-driven talking face generation methods necessitate high-quality audio, imposing stringent requirements on data transmission conditions and recording environments [37]. In contrast, text-driven talking face generation methods operate effectively with basic and reliable network transmission conditions.

Addressing such challenges holds significant value not only for the scientific community but also for the industry, aiming to mitigate the environmental impact of AI technologies [35]. Consequently, one crucial question has arisen: *Is it feasible to substitute audio with text inputs while ensuring detailed facial expressions of generated frames?*

There exist two primary types of text-driven face generation methods. The first involves using *third-person* text descriptions where a facial expression is described from an external perspective. These descriptions are leveraged to

generate realistic faces using end-to-end GAN architectures [13, 28, 44]. The second path employs *first-person* statements which are the speaker’s spoken captions captured in text form. While this approach aligns more closely with talking face generation, it has not been thoroughly explored. Previous methods mainly conduct a two-stage framework including text-to-speech and speech-to-face procedures [30, 48]. These approaches generate audio waveforms from text statements and further generate faces based on synthesized audio [34]. Therefore, methods that follow a one-stage end-to-end pattern to generate talking faces driven by first-person statements remain unestablished. Based on this, our objective is to revolutionize the synthesis of coherent and natural talking face videos through first-person statements, encompassing well-synchronized lip movements, accurate cheek actions, detailed textures, and overall facial expressions, as illustrated in Figure 1. To overcome the challenges above, we propose **First-Person Statement Text-To-Talking Face Generation (FT2TF)** – a one-stage end-to-end pipeline that effectively generates high-quality talking faces from visual-textual cross-modality. Different from previous work, our pipeline encapsulates the entire face generation process without any other sources (*e.g.*, audio, landmarks, poses) being used during inference.

In summary, our contribution is threefold:

- **One-stage End-to-end Continuous-Frame Talking Face Generation.** We propose FT2TF, a one-stage end-to-end pipeline that generates realistic continuous-frame talking faces by integrating visual and textual input.
- **Effective and Efficient Text-Driven Talking Face Generation.** FT2TF utilizes only visual and textual information during inference. It enhances face naturalness, synchronizes lip movements, and improves detailed facial expressions while using fewer trainable parameters.
- **State-of-the-Art Performance.** We conduct extensive experiments on the LRS2 [1] and LRS3 [2] datasets with both qualitative and quantitative results to demonstrate the effectiveness of FT2TF. This strongly proves that our model can consistently generate high-quality talking faces across different conditions.

2. Related Work

2.1. Audio-Driven Talking Face Generation

The majority of the work on talking face generation is audio-driven which needs both audio and video ground truth as the input, focusing on lip region/whole face generation.

Lip Generation. As a groundwork, Chen *et al.* proposes ATVGnet [5] that uses audio signal and reference image to generate the region around the lip, while the remaining parts of the face remain static from frame to frame. Based on that,

Wav2Lip [31] generates dynamically changed faces through lip masking and reconstruction strategy. Following the previous lip masking strategy, IP_LAP [49] works as state-of-the-art by using a Transformer-based [39] encoder to predict the landmarks for further lip generation. Although these methods generate high-quality dynamic faces, all of them *only* focus on lip generation, failing to capture facial expression details of the entire face.

Whole Face Generation. Zhou *et al.* [52] disentangles audio sources into content and speaker embeddings, followed by an LSTM-based module for landmark prediction to generate the talking head. Based on that, models that use input sources other than audio have been proposed. PC-AVS [50] utilizes the pose source to control the head pose of the synthesized faces, while EAMM [18] uses pose sequence, emotion source, and audio ground truth to generate the target video manipulated by different emotion guidances. In contrast to existing methods that involve multiple input resources for face generation, our model *only* uses text and reference face frames as the input source. This ensures the efficient generation of high-quality facial representations, free from input requirements such as constraints of audio quality.

2.2. Text-Driven Talking Face Generation

The exploration of talking face generation based on first-person statements has not been well-established. Initially, models using Variational Autoencoder (VAE) [21, 26, 40, 43] and CRNN [4, 38] are designed to generate video based on text input. However, the generated frames by these models are unrealistic and over-smoothing, especially in long-sequence video generation tasks [29, 43]. Text-driven talking face generation has been divided mainly into two directions: using first-person statements or third-person descriptions. Methods driven by third-person descriptions [13, 28, 44] aim to learn the mapping between objective text descriptions and synthesized video. The most relevant approaches using first-person statements are constructed in a two-stage architecture with two separate steps: converting text to speech and then generating faces from the speech [22, 23, 30, 45, 48]. The latest work, TTFS [16], integrates text-to-speech and face-generation modules to synthesize talking face and audio from input text jointly. Consequently, one-stage methods that generate talking faces directly from first-person text captions remain unexplored [25, 27]. Addressing this research gap, our model is constructed as a one-stage end-to-end framework that directly leverages reference talking face frames and first-person text caption to generate high-quality talking faces efficiently. Since relevant exploration has not been widely established, our solution aims to bridge first-person text and dynamic face synthesis to provide insightful guidance for future work.

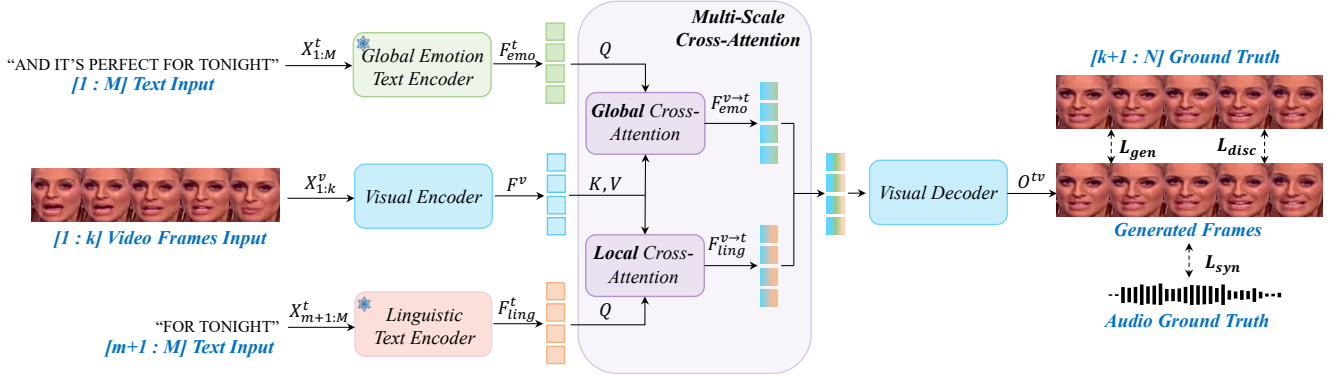


Figure 2. **Overview of FT2TF Pipeline.** The FT2TF pipeline employs two specialized Text Encoders, the Global Emotion Text Encoder, and the Linguistic Text Encoder, for extracting emotional and linguistic text features, respectively. Additionally, a Visual Encoder is utilized to extract visual features. Afterward, it leverages a Multi-Scale Cross-Attention Module for visual-textual fusion. The resulting visual-textual representations are fed to a Visual Decoder to synthesize talking face frames.

3. FT2TF Pipeline

We address the challenging problem of generating realistic talking face videos from *first-person* text descriptions. Formally, given M text caption tokens $X_{1:M}^t$ as the complete sentence and the initial k talking face frames $X_{1:k}^v$ as input, our objective is to synthesize the subsequent talking face frames $X_{k+1:N}^v$. In this context, t represents the textual content, and v represents the visual content. The text captions $X_{1:M}^t$ and the talking face frames $X_{1:N}^v$ are aligned to correspond to each other.

To tackle the problem, we present a novel pipeline – FT2TF: First-Person Statement Text-to-Talking Face Generation, as illustrated in Figure 2. We use two specialized text encoders – the Global Emotion Text Encoder captures overall emotional tone from the complete text caption $X_{1:M}^t$, and the Linguistic Text Encoder encodes semantic compression from the reference text caption $X_{m+1:M}^t$ which corresponds to frames $X_{k+1:N}^v$. Alongside these, a Visual Encoder is employed to extract visual features from reference frames $X_{1:k}^v$. Inspired by [24], we design and utilize a Multi-Scale Cross-Attention Module that integrates Global and Local Cross-Attention modules to align features from two modalities into the same distribution [24, 39], aiming to ensure a comprehensive visual-textual fusion across varying scales and feature dimensions. The resulting visual-textual representations are fed to a Visual Decoder to synthesize talking face frames $X_{k+1:N}^v$.

3.1. Multimodal Encoders

In this Section, we explore the multimodal encoders utilized in FT2TF. Our investigation concentrates on the design insights and functionalities of the Visual Encoder, Global Emotion Text Encoder, and Linguistic Text Encoder. As shown in Figure 2, the features output from these three

encoders are denoted as F^v , F_{emo}^t , and F_{ling}^t , respectively.

Visual Encoder. The CNN-based Visual Encoder plays a pivotal role in extracting image features and facial details from talking face videos. Considering the properties of ResBlocks to handle the temporal dynamics and variations in visual input without massive data or strong data augmentations [15], the Visual Encoder is built with ResBlocks as the backbone. The extracted visual features, F^v , will then be integrated with the textual ones through the Multi-Scale Cross-Attention module. Furthermore, Table 5 illustrates the results of various backbones of the Visual Encoder.

Global Emotion Text Encoder. The Global Emotion Text Encoder harnesses the capabilities of the pretrained Emoberta [20] model. Emoberta serves as the cornerstone for encoding the overarching emotional context embedded within textual descriptions. In light of this, it presents itself as a prime contender for global attention. Using the entire face frames, we can direct our exclusive focus toward text emotions and their face representations.

Linguistic Text Encoder. We adopt to use the pretrained GPT-Neo [10] as the Linguistic Text Encoder. In our choice to employ this model, we find motivation in the distinctive characteristics of the GPT-Neo, as opposed to the Emoberta [20]. GPT-Neo has undergone training using a vast language dataset, which empowers it with enhanced linguistic compression for every statement conveyed by the talking face. Consequently, we intend to narrow our attention solely to the region of the lips, aligned with the phonetic articulation of each word, as can be seen in Figure 3.

3.2. Multi-Scale Cross-Attention and Visual Decoder

Given the encoded textual (F_{emo}^t , F_{ling}^t) and visual features (F^v), we design a Multi-Scale Cross-Attention module that integrates *Global* and *Local* Cross-Attention mod-

ules. Specifically, features extracted from the Visual Encoder, F^v , are used in conjunction with the textual features from the Global Emotion Text Encoder, F_{emo}^t , as input to the Global Cross-Attention. In this context, F^v serves as both the key (K) and the value (V), while F_{emo}^t functions as the query (Q). A similar process occurs with Local Cross-Attention where F^v serves as K and V and the output of the Linguistic Text Encoder, F_{ling}^t , acts as Q .

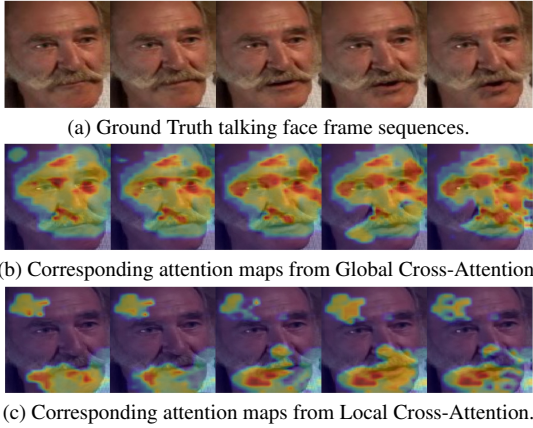


Figure 3. **Visualization of Global and Local Cross-Attention.** The red areas indicate high attention weights, while the blue areas indicate the opposite.

Global Cross-Attention. This module serves as a pivotal component for infusing global emotional cues from the textual description into the video context. It combines the encoded video features F^v with the emotional text description ones F_{emo}^t and generates the fusion features $F_{emo}^{v \rightarrow t}$. As shown in Figure 3b, this integration empowers the model to effectively incorporate emotional expression into facial components such as cheeks, eyes, and eyebrows, enabling more expressive and emotionally accurate face synthesis.

Local Cross-Attention. This module enriches the model’s understanding of local language features in the text. Specifically, by fusing the features $F_{ling}^{v \rightarrow t}$ from F^v and F_{ling}^t , it facilitates the generation of localized and context-aware facial expressions and details. Since the linguistic information is mainly presented by mouth movements, the Local Cross-Attention results in the highest weight in the area around the mouth, as shown in Figure 3c. This highly enhances the model’s ability to capture fine-grained linguistic nuances and improves the overall coherence of the generated face.

Multi-Scale Cross-Attention. The collective use of Global and Local Cross-Attention ensures a balanced fusion of global emotional context and local linguistic features. As corroborated in Section 4, this multimodal fusion mechanism enables FT2TF to generate talking face videos with natural lip synchronization and emotionally expressive facial details, resulting in more realistic and engaging outputs.

Alignment. In the realm of text-video talking face synthesis, a primary concern revolves around the alignment between textual inputs and video frames. Given the variability in speech rates and temporal intervals between words, especially when considering text-only inputs, we address this challenge by circumventing the alignment problem between text and video frames. Instead, we leverage extensive audiovisual data alongside a highly efficient lip synthesis expert module to ensure the synthesis of text-video talking faces. Operating on an end-to-end basis, our model necessitates correspondence between the video input and the textual content, thereby facilitating the generation of lip images that closely approximate the ground truth during synthesis. As illustrated in Figure 3, the emotional content and latent linguistic information of the text also influence different facial regions through Multi-Scale Cross-Attention. This phenomenon indicates that our model has successfully aligned features from different distributions.

Visual Decoder. The Visual Decoder is constructed with transpose convolution blocks, aligning each convolutional kernel with its counterpart in the Visual Encoder and matching the channel dimensions to those of $F_{emo}^{v \rightarrow t}$ and $F_{ling}^{v \rightarrow t}$. In this way, the output resolution of the Visual Decoder aligns with the model input. The Visual Decoder conducts channel-wise convolution operations on the two text features and visual features, aiming to reconstruct the visual information embedded in the text features.

The output of the Global and Local Cross-Attention ($F_{emo}^{v \rightarrow t}, F_{ling}^{v \rightarrow t}$) are concatenated and subsequently input into the Visual Decoder. The Visual Decoder generates faces from the $k + 1$ frame to the N frame $O_{k+1}^{tv}, \dots, O_N^{tv}$.

3.3. Pipeline Training

Three key loss functions are constructed for pipeline training.

Generation Loss. The generation loss, L_{gen} , measures the pixel-wise discrepancy between the generated frames and the ground truth ones, with the length represented by N_{gen} . This loss primarily focuses on supervising the accurate generation of pixel-level details to ensure the fidelity of the synthesized frames, as shown below:

$$L_{gen} = \frac{1}{N_{gen}} \sum_{i=1}^{N_{gen}} \|X_{k+i}^v - O_{k+i}^{tv}\|_1. \quad (1)$$

Synthesis Loss. To improve lip synchronization and the naturalness of lip movements [31], the synthesis loss, L_{syn} , leverages the audio component by integrating the Mel-spectrum generated by the Face Synthesizer. While it is not explicitly expressed in mathematical terms, L_{syn} involves the alignment of the synthesized talking face frames with the corresponding audio data A_{k+i} :

$$L_{syn} = \frac{1}{N_{gen}} \sum_{i=1}^{N_{gen}} SyncNet(O_{k+i}^{tv}, A_{k+i}). \quad (2)$$

Discriminator Loss. The discriminator loss [12], L_{disc} , aims to guide the encoder-decoder process to produce more realistic and coherent images. It operates by utilizing a discriminator that provides binary predictions O_{pred} based on whether the input is the generated one or the ground truth. Specifically, the corresponding labels Y_{disc} are assigned the label 1 for the generated data and 0 for the ground truth, producing a binary cross-entropy loss:

$$L_{disc} = -\frac{1}{N_{gen}} \sum_{i=1}^{N_{gen}} (Y_{disc} \cdot \log(O_{pred}) + (1 - Y_{disc}) \cdot \log(1 - O_{pred})). \quad (3)$$

Eventually, the total loss L_{total} is expressed as follows:

$$L_{total} = \lambda_1 L_{gen} + \lambda_2 L_{syn} + \lambda_3 L_{disc}. \quad (4)$$

4. Experiment

4.1. Datasets

We conduct experiments on two widely used and substantial large-scale text-audio-visual datasets: LRS2 [1] and LRS3 [2].

- **LRS2 [1]:** The dataset comprises 48,164 video clips sourced from BBC outdoor broadcasts, divided into training, validation, and test sets with 45,839, 1,082, and 1,243 clips, respectively. Each clip is paired with a text caption of up to 100 characters. To ensure uniform input lengths, we retain videos with frame counts between 30 and 35.
- **LRS3 [2]:** The dataset comprises over 400 hours of content from 5,594 indoor TED/TEDx talks, totaling 151,819 videos with corresponding text captions. Specifically, it includes 32,000 training/validation utterances and 1,452 test utterances. Consistent with the LRS2 dataset, we retain videos with frame counts between 30 and 35.

Although the two datasets share similar names, it is important to emphasize that they have **no overlap**. The LRS2 and LRS3 videos differ significantly in terms of scenes, data collection methods, character actions, environmental conditions, and other attributes.

4.2. Comparison Methods

We conduct a comprehensive evaluation of our method against state-of-the-art approaches for talking face generation. To demonstrate the effectiveness and efficiency of our model, we compare it with state-of-the-art audio-driven and text-driven models on the LRS2 and LRS3 datasets. These models are categorized based on two criteria:

- **Generation Task:** Methods are categorized into either *Lip Region* or *Whole Face*, with the former generating only the lip region and the latter generating the entire face.
- **Modality:** Methods are classified based on their input modality, including *Audio+Video (A+V)*, *Text+Video (T+V)*, *Audio+Video+Pose/Landmarks (A+V+P/L)*.

4.3. Evaluation Metrics

We employ multiple evaluation metrics across various criteria, ranging from evaluating video synthesis quality to evaluating facial expression representation. To evaluate the similarity between generated results and ground truth ones, we apply Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity (SSIM) [41]. In addition, we utilize Learned Perceptual Image Patch Similarity (LPIPS) [47] and Frechet Inception Distance (FID) [14] for assessing feature-level similarity. Considering identity preservation, we leverage cosine similarity (CSIM) between identity vectors extracted by the ArcFace [8] network. Moreover, aligning with the approach outlined in [42], we measure the normalized lip landmarks distance (LipLMD) between generated and ground truth frames to evaluate lip synchronization quality. The scores for all methods are calculated over the **entire face** for a fair comparison.

4.4. Comparison with State-of-the-art

4.4.1 Quantitative Results

As shown in Table 1, FT2TF achieves state-of-the-art performance across multiple evaluation metrics for talking face generation, as evidenced by the quantitative results on both LRS2 and LRS3 datasets.

Video Quality Enhancement. FT2TF achieves state-of-the-art results across video quality evaluation metrics (PSNR, SSIM, and LPIPS). For example, on the LRS2 dataset, FT2TF outperforms TTFS [16], the current state-of-the-art T+V method, with a 9.6% improvement in PSNR, a 4.9% improvement in SSIM, and a significantly lower LPIPS value, reduced by 66.5%. Moreover, FT2TF consistently outperforms lip-specific methods, even those using ground truth for the remaining parts of the face, further validating its superior performance. Similar results are observed on the LRS3 dataset, highlighting FT2TF’s ability to generate talking faces with high-quality and visual realism, closely resembling real-world data.

Method	Dataset	Generation	Input	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	FID \downarrow	LipLMD \downarrow	CSIM \uparrow
ATVGnet [6]		Lip Region	A+V	11.55	0.3944	0.5575	223.26	0.27401	0.1020
Wav2Lip [31]		Lip Region	A+V	27.92	0.8962	0.0741	43.46	0.02003	0.5925
IP_LAP [49]		Lip Region	A+V+L	32.91	0.9399	0.0303	27.87	0.01293	0.6523
EAMM [18]	LRS2	Whole Face	A+V+P	15.17	0.4623	0.3398	91.95	0.15191	0.2318
PC-AVS [51]		Whole Face	A+V+P	15.75	0.4867	0.2802	110.60	0.07569	0.3927
MakeItTalk [52]		Whole Face	A+V	17.25	0.5562	0.2237	76.57	0.05024	0.5799
TTFS [16]		Whole Face	T+V	30.28	0.9436	0.0173	19.62	0.03184	0.8149
FT2TF (Ours)		Whole Face	T+V	33.20	0.9901	0.0058	22.52	0.02356	0.9642
ATVGnet [6]		Lip Region	A+V	10.90	0.3791	0.5667	190.49	0.30564	0.1176
Wav2Lip [31]		Lip Region	A+V	28.45	0.8852	0.0683	49.60	0.02001	0.5909
IP_LAP [49]		Lip Region	A+V+L	32.97	0.9222	0.0310	29.96	0.01353	0.6385
EAMM [18]	LRS3	Whole Face	A+V+P	15.37	0.4679	0.3812	108.83	0.17818	0.2689
PC-AVS [51]		Whole Face	A+V+P	15.60	0.4732	0.3321	115.25	0.10611	0.3537
MakeItTalk [52]		Whole Face	A+V	17.78	0.5607	0.2788	97.99	0.08432	0.5465
TTFS [16]		Whole Face	T+V	29.75	0.9317	0.0206	13.98	0.03061	0.8243
FT2TF (Ours)		Whole Face	T+V	31.63	0.9895	0.0085	15.97	0.02545	0.9665

Table 1. Comparison with state-of-the-art talking face generation methods on different evaluation metrics on LRS2 and LRS3 Datasets. Bold results indicate the best among *all* methods, while the results in the gray background indicate the best among *whole face generation* methods. “V” - Visual Modality. “A” - Audio Modality. “T” - Textual Modality. “L” - Landmark Source. “P” - Pose Source.




Figure 4. Efficiency comparison. Our analysis includes a comparison of SSIM scores against the total trainable parameters for various models, demonstrating our model’s efficiency.

Data Distribution Alignment. In terms of FID, which evaluates the alignment of data distributions, FT2TF achieves scores of 22.52 on LRS2 and 15.97 on LRS3, demonstrating competitive performance compared to state-of-the-art methods [16]. The low FID values indicate that FT2TF generates frames with distributions statistically close to real ones, underscoring the model’s capacity to produce natural-looking talking faces.

Lip Synchronization and Identity Preservation. Although FT2TF is not specifically designed for lip generation, it achieves comparable lip synchronization performance (LipLMD) to the state-of-the-art lip-specific method, IP_LAP [49], and surpasses all whole-face generation meth-

ods. Furthermore, FT2TF excels in identity preservation, as evidenced by the highest CSIM scores. This underscores its capability to maintain accurate facial features and expressions across frames, ensuring the realistic and coherent generation of talking faces.

Robustness Across Datasets. FT2TF demonstrates consistent and robust performance across distinct and non-overlapping datasets. Despite the variations between the two datasets described in Section 4.1, FT2TF maintains stable, high-quality generation capabilities. This underscores the model’s adaptability and its strong generalization ability across diverse video sources.

Efficiency Analysis. Figure 4 compares the total trainable parameters across various models, including both lip-specific and face-generation approaches. Specifically, compared to the state-of-the-art methods that only learn lip-around features instead of the whole face (IP_LAP [49], Wav2Lip [31]), our model achieves a higher SSIM score while using a competitive number of trainable parameters. Moreover, our model outperforms other cutting-edge face generation models significantly. Quantitatively, we achieve a 78.01% improvement in SSIM with only 2.47% more trainable parameters than MakeItTalk [52]. In comparison to other face-generation models listed in the figure, we present substantial advantages in terms of both trainable parameters and performance, with an average performance improvement of 122.88% and an average parameter decrease of 68.94%.



Figure 5. **Qualitative comparison with state-of-the-art methods on LRS2 and LRS3.** The three models on the left specialize in lip generation, whereas the others are designed to generate entire faces. TTFS [16] and our model are text-driven, whereas the remaining methods are audio-driven. Our model consistently generates the most detailed and accurate talking faces across diverse roles, genders, and ages under different lighting and head-poses conditions.



Figure 6. **Qualitative comparison demonstration in continuous talking face frames Generation.** Our model demonstrates enhanced capability in generating talking faces with temporally coherent and accurately synchronized lip movements that align with the spoken text.

4.4.2 Qualitative Results

We process all talking face frames using the same face detector [31], enabling a more detailed comparison of facial expressions. Figure 5 shows the comparison of generation quality between our model and the relevant methods on LRS2 and LRS3. Our results stand out as the most similar to ground truth in diverse conditions, including different character roles and identities, various lighting, and face angles. Benefiting from the Global Emotion Text Encoder,

Linguistic Text Encoder, and Multi-Scale Cross-Attention modules, our model accurately preserves facial expression details, such as teeth, face shape, head pose, skin color, and texture details. Furthermore, we evaluate multiple consecutive video frames, as shown in Figure 6. Specifically, a temporal comparison with the state-of-the-art face generation methods shows that our model constructs the most accurate and consistent video frames in the time domain. In addition to the visual demonstration, we also report quantitative results of temporal generation in Section 4.5.

Method	Temp Change	Lip Sync	Facial Details	Video Quality
MakeItTalk [52]	1.83	1.26	1.70	1.87
IP_LAP [49]	3.13	2.90	2.99	2.82
FT2TF (Ours)	4.50	4.58	4.77	4.43

Table 2. **User study on the state-of-the-art lip/face generation methods.** Mean Opinion Scores (MOS) are reported based on four aspects: **Temp Change**: temporal transition smoothness. **Lip Sync**: lip synchronization accuracy. **Facial Details**: facial detail naturalness. **Video Quality**: overall video quality. We conducted this experiment over three different methods.

4.5. User Study

We conduct a user study to evaluate videos synthesized by state-of-the-art lip and face generation methods. 30 participants with diverse roles, genders, and ages are invited to watch 4 videos from each method and evaluate on a scale of 1 to 5 for four different aspects: temporal transition smoothness, lip synchronization accuracy, facial detail naturalness, and overall video quality. According to the Mean Opinion Scores (MOS) reported in Table 2, our model outperforms the state-of-the-art lip/face generation methods from all evaluation domains.

4.6. Ablation Analysis

We conduct a comprehensive analysis of the ablation studies on the FT2TF to investigate the impact of key components, including the Cross-Attention mechanisms, loss terms, and different visual encoder backbones. The primary goal is to analyze the model’s performance and demonstrate the contribution of each component in generating high-quality face frames from first-person text statements.

Method	PSNR↑	SSIM↑	FID↓	CSIM↑
w/o Global CA	23.03	0.9101	40.59	0.8965
w/o Local CA	21.49	0.8975	40.56	0.8925
w/o Both CA	22.73	0.7317	127.38	0.5682
FT2TF (Ours)	33.20	0.9901	22.52	0.9642

Table 3. **Ablation study on Cross-Attention mechanisms.** We conduct experiments on the Cross-Attention (CA) mechanisms on LRS2.

Cross-Attention Mechanisms. The ablation study on the Cross-Attention module is shown in Table 3. By removing either Global or Local Cross-Attention, a significant drop in all the metrics can be observed. Since emotional coherence and linguistic information are enhanced from the Cross-Attention module as discussed in Section 3.2, the model without it fails to generate high-quality and distribution-aligned faces with the person’s identity being preserved.

Role of Loss Terms. We delve into the role of specific loss

Method	PSNR↑	SSIM↑	FID↓	CSIM↑
w/o L_{syn}	28.92	0.9389	40.59	0.8965
w/o L_{disc}	19.90	0.8736	25.66	0.9305
FT2TF (Ours)	33.20	0.9901	22.52	0.9642

Table 4. **Ablation study on loss terms.** Specifically, we conduct ablation study on L_{syn} and L_{disc} on LRS2.

terms, L_{syn} and L_{disc} , as shown in Table 4. L_{syn} focuses on improving the naturalness of lip movements and facial expression details, which leads to a significant gain in all quantitative metrics. L_{disc} guides the model to generate more realistic frames through binary predictions by the discriminator, contributing to accurate and expressive faces.

Backbone	PSNR↑	SSIM↑	FID↓	CSIM↑
ViT-B	31.30	0.9690	27.61	0.9378
ViT-L	32.89	0.9690	20.60	0.9434
ResBlocks	33.20	0.9901	22.52	0.9642

Table 5. **Ablation study on Visual Encoder backbones.** We construct the model with ViT/ResBlocks as the encoder backbone on LRS2.

Encoder Backbones. We investigate the effectiveness of different backbone structures of the Visual Encoder, including ViT with different scales (ViT-B/L) and ResBlocks. As shown in Table 5, our model achieves stable and effective results with various encoder backbones, demonstrating the flexibility of our pipeline structure. Moreover, since ResBlocks retains information at different scales and inherently captures spatial hierarchies in the visual input, it results in a lead in PSNR, SSIM, and CSIM compared to ViT. Considering the parameter and data efficiency, we utilize ResBlocks as the Visual Encoder backbone.

5. Conclusion

This paper addresses the challenge of generating talking faces through the FT2TF pipeline, which uses first-person text captions to replace traditional audio input. Our approach ensures the generation of faces with natural and detailed expressions. By incorporating two large language models, we extract emotional and linguistic features from the text. Moreover, the Multi-Scale Cross-Attention module effectively fuses visual and textual features, contributing to the synthesis of emotionally expressed talking faces. Extensive experiments on the LRS2 and LRS3 datasets highlight FT2TF’s effectiveness in generating realistic, natural, and emotionally expressive talking faces from first-person text.

References

- [1] Triantafyllos Afouras, Joon Son Chung, Andrew Senior, Oriol Vinyals, and Andrew Zisserman. Deep audio-visual speech recognition. *Transactions on Pattern Analysis and Machine Intelligence*, 2022. 2, 5
- [2] Triantafyllos Afouras, Joon Son Chung, and Andrew Zisserman. Lrs3-ted: a large-scale dataset for visual speech recognition. *arXiv preprint arXiv:1809.00496*, 2018. 2, 5
- [3] Mandy M Archibald, Rachel C Ambagtsheer, Mavourneen G Casey, and Michael Lawless. Using zoom videoconferencing for qualitative data collection: perceptions and experiences of researchers and participants. *International Journal of Qualitative Methods*, 2019. 1
- [4] Yogesh Balaji, Martin Renqiang Min, Bing Bai, Rama Chelappa, and Hans Peter Graf. Conditional gan with discriminative filter generation for text-to-video synthesis. In *International Joint Conference on Artificial Intelligence*, 2019. 2
- [5] Lele Chen, Ross K Maddox, Zhiyao Duan, and Chenliang Xu. Hierarchical cross-modal talking face generation with dynamic pixel-wise loss. In *Conference on Computer Vision and Pattern Recognition*, 2019. 2
- [6] Lele Chen, Ross K Maddox, Zhiyao Duan, and Chenliang Xu. Hierarchical cross-modal talking face generation with dynamic pixel-wise loss. In *Conference on Computer Vision and Pattern Recognition*, 2019. 6
- [7] Ana-Paula Correia, Chenxi Liu, and Fan Xu. Evaluating videoconferencing systems for the quality of the educational experience. *Distance Education*, 2020. 1
- [8] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *Conference on Computer Vision and Pattern Recognition*, 2019. 5
- [9] Matthew C Fysh, Iliyana V Trifonova, John Allen, Cade McCall, A Mike Burton, and Markus Bindemann. Avatars with faces of real people: A construction method for scientific experiments in virtual reality. *Behavior Research Methods*, 2022. 1
- [10] Leo Gao, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason Phang, Horace He, Anish Thite, Noa Nabeshima, et al. The pile: An 800gb dataset of diverse text for language modeling. *arXiv preprint arXiv:2101.00027*, 2020. 3
- [11] Pablo Garrido, Levi Valgaerts, Hamid Sarmadi, Ingmar Steiner, Kiran Varanasi, Patrick Perez, and Christian Theobalt. Vdub: Modifying face video of actors for plausible visual alignment to a dubbed audio track. *Computer Graphics Forum*, 2015. 1
- [12] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, 2014. 5
- [13] Ligong Han, Jian Ren, Hsin-Ying Lee, Francesco Barbieri, Kyle Olszewski, Shervin Minaee, Dimitris Metaxas, and Sergey Tulyakov. Show me what and tell me how: Video synthesis via multimodal conditioning. In *Conference on Computer Vision and Pattern Recognition*, 2022. 2
- [14] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in Neural Information Processing Systems*, 2017. 5
- [15] Khawar Islam. Recent advances in vision transformer: A survey and outlook of recent work. *arXiv preprint arXiv:2203.01536*, 2022. 3
- [16] Youngjoon Jang, Ji-Hoon Kim, Junseok Ahn, Doyeop Kwak, Hong-Sun Yang, Yoon-Cheol Ju, Il-Hwan Kim, Byeong-Yeol Kim, and Joon Son Chung. Faces that speak: Jointly synthesising talking face and speech from text. In *Conference on Computer Vision and Pattern Recognition*, 2024. 2, 5, 6, 7
- [17] Ana Javornik, Ben Marder, Jennifer Brannon Barhorst, Graeme McLean, Yvonne Rogers, Paul Marshall, and Luk Warlop. ‘what lies behind the filter?’ uncovering the motivations for using augmented reality (ar) face filters on social media and their effect on well-being. *Computers in Human Behavior*, 2022. 1
- [18] Xinya Ji, Hang Zhou, Kaisiyuan Wang, Qianyi Wu, Wayne Wu, Feng Xu, and Xun Cao. Eamm: One-shot emotional talking face via audio-based emotion-aware motion model. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference*, 2022. 1, 2, 6
- [19] Hyeonwoo Kim, Mohamed Elgharib, Michael Zollhöfer, Hans-Peter Seidel, Thabo Beeler, Christian Richardt, and Christian Theobalt. Neural style-preserving visual dubbing. *Transactions on Graphics*, 2019. 1
- [20] Taewoon Kim and Piek Vossen. Emoberta: Speaker-aware emotion recognition in conversation with roberta. *arXiv preprint arXiv:2108.12009*, 2021. 3
- [21] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. In *International Conference on Learning Representations*, 2014. 2
- [22] Prajwal KR, Rudrabha Mukhopadhyay, Jerin Philip, Abhishek Jha, Vinay Namboodiri, and CV Jawahar. Towards automatic face-to-face translation. In *International Conference on Multimedia*, 2019. 2
- [23] Rithesh Kumar, Jose Sotelo, Kundan Kumar, Alexandre De Brebisson, and Yoshua Bengio. Obamanet: Photo-realistic lip-sync from text. *arXiv preprint arXiv:1801.01442*, 2017. 2
- [24] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International Conference on Machine Learning*, 2023. 3
- [25] Lincheng Li, Suzhen Wang, Zhimeng Zhang, Yu Ding, Yixing Zheng, Xin Yu, and Changjie Fan. Write-a-speaker: Text-based emotional and rhythmic talking-head generation. In *AAAI Conference on Artificial Intelligence*, 2021. 2
- [26] Yitong Li, Martin Min, Dinghan Shen, David Carlson, and Lawrence Carin. Video generation from text. In *AAAI Conference on Artificial Intelligence*, 2018. 2
- [27] Jinglin Liu, Zhiying Zhu, Yi Ren, Wencan Huang, Baoxing Huai, Nicholas Yuan, and Zhou Zhao. Parallel and high-fidelity text-to-lip generation. In *AAAI Conference on Artificial Intelligence*, 2022. 2

- [28] Yifeng Ma, Suzhen Wang, Yu Ding, Bowen Ma, Tangjie Lv, Changjie Fan, Zhipeng Hu, Zhidong Deng, and Xin Yu. Talkclip: Talking head generation with text-guided expressive speaking styles. *arXiv preprint arXiv:2304.00334*, 2023. [2](#)
- [29] Shentong Mo and Miao Xin. Long-term head pose forecasting conditioned on the gaze-guiding prior. In *Conference on Computer Vision and Pattern Recognition*, 2021. [2](#)
- [30] Dan Oneatǎ, Beáta Lőrincz, Adriana Stan, and Horia Cucu. Flexlip: A controllable text-to-lip system. *Sensors*, 2022. [2](#)
- [31] KR Prajwal, Rudrabha Mukhopadhyay, Vinay P Namboodiri, and CV Jawahar. A lip sync expert is all you need for speech to lip generation in the wild. In *International Conference on Multimedia*, 2020. [1](#), [2](#), [4](#), [6](#), [7](#)
- [32] Tianyang Shi, Yi Yuan, Changjie Fan, Zhengxia Zou, Zhenwei Shi, and Yong Liu. Face-to-parameter translation for game character auto-creation. In *International Conference on Computer Vision*, 2019. [1](#)
- [33] Tianyang Shi, Zhengxia Zuo, Yi Yuan, and Changjie Fan. Fast and robust face-to-parameter translation for game character auto-creation. In *AAAI Conference on Artificial Intelligence*, 2020. [1](#)
- [34] Hyoung-Kyu Song, Sang Hoon Woo, Junhyeok Lee, Seungmin Yang, Hyunjae Cho, Youseong Lee, Dongho Choi, and Kang-wook Kim. Talking face generation with multilingual tts. In *Conference on Computer Vision and Pattern Recognition*, 2022. [2](#)
- [35] Guglielmo Tamburrini. The ai carbon footprint and responsibilities of ai scientists. *Philosophies*, 2022. [1](#)
- [36] Justus Thies, Mohamed Elgharib, Ayush Tewari, Christian Theobalt, and Matthias Nießner. Neural voice puppetry: Audio-driven facial reenactment. In *European Conference on Computer Vision*, 2020. [1](#)
- [37] Mukhiddin Toshpulatov, Wookey Lee, and Suan Lee. Talking human face generation: A survey. *Expert Systems with Applications*, 2023. [1](#)
- [38] Sergey Tulyakov, Ming-Yu Liu, Xiaodong Yang, and Jan Kautz. Mocogan: Decomposing motion and content for video generation. In *Conference on Computer Vision and Pattern Recognition*, 2018. [2](#)
- [39] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, 2017. [2](#), [3](#)
- [40] Jacob Walker, Carl Doersch, Abhinav Gupta, and Martial Hebert. An uncertain future: Forecasting from static images using variational autoencoders. In *European Conference on Computer Vision*, 2016. [2](#)
- [41] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *Transactions on Image Processing*, 2004. [5](#)
- [42] Tianyi Xie, Liucheng Liao, Cheng Bi, Benlai Tang, Xiang Yin, Jianfei Yang, Mingjie Wang, Jiali Yao, Yang Zhang, and Zejun Ma. Towards realistic visual dubbing with heterogeneous sources. In *International Conference on Multimedia*, 2021. [5](#)
- [43] Tianfan Xue, Jiajun Wu, Katherine L Bouman, and William T Freeman. Visual dynamics: Stochastic future generation via layered cross convolutional networks. *Transactions on Pattern Analysis and Machine Intelligence*, 2018. [2](#)
- [44] Jianhui Yu, Hao Zhu, Liming Jiang, Chen Change Loy, Weidong Cai, and Wayne Wu. Celebv-text: A large-scale facial text-video dataset. In *Conference on Computer Vision and Pattern Recognition*, 2023. [2](#)
- [45] Lingyun Yu, Jun Yu, and Qiang Ling. Mining audio, text and visual information for talking face generation. In *International Conference on Data Mining*, 2019. [2](#)
- [46] Longwen Zhang, Chuxiao Zeng, Qixuan Zhang, Hongyang Lin, Ruixiang Cao, Wei Yang, Lan Xu, and Jingyi Yu. Video-driven neural physically-based facial asset for production. *Transactions on Graphics*, 2022. [1](#)
- [47] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Conference on Computer Vision and Pattern Recognition*, 2018. [5](#)
- [48] Sibozhang, Jiahong Yuan, Miao Liao, and Liangjun Zhang. Text2video: Text-driven talking-head video synthesis with personalized phoneme-pose dictionary. In *International Conference on Acoustics, Speech and Signal Processing*, 2022. [2](#)
- [49] Weizhi Zhong, Chaowei Fang, Yinqi Cai, Pengxu Wei, Gangming Zhao, Liang Lin, and Guanbin Li. Identity-preserving talking face generation with landmark and appearance priors. In *Conference on Computer Vision and Pattern Recognition*, 2023. [1](#), [2](#), [6](#), [8](#)
- [50] Hang Zhou, Yasheng Sun, Wayne Wu, Chen Change Loy, Xiaogang Wang, and Ziwei Liu. Pose-controllable talking face generation by implicitly modularized audio-visual representation. In *Conference on Computer Vision and Pattern Recognition*, 2021. [1](#), [2](#)
- [51] Hang Zhou, Yasheng Sun, Wayne Wu, Chen Change Loy, Xiaogang Wang, and Ziwei Liu. Pose-controllable talking face generation by implicitly modularized audio-visual representation. In *Conference on Computer Vision and Pattern Recognition*, 2021. [6](#)
- [52] Yang Zhou, Xintong Han, Eli Shechtman, Jose Echevarria, Evangelos Kalogerakis, and Dingzeyu Li. Makeltalk: speaker-aware talking-head animation. *Transactions On Graphics*, 2020. [1](#), [2](#), [6](#), [8](#)