YINGRUO FAN, The University of Hong Kong, Hong Kong ZHAOJIANG LIN*, The Hong Kong University of Science and Technology, Hong Kong JUN SAITO, Adobe Research, USA WENPING WANG, Texas A&M University, USA TAKU KOMURA[†], The University of Hong Kong, Hong Kong



Fig. 1. Given a speech signal (audio and text) and a 3D face mesh, our method produces the expressive 577502 lip-synchronized 3D facial animation with realistic facial expressions. Please see our supplementary video.

Speech-driven 3D facial animation with accurate lip synchronization has been widely studied. However, synthesizing realistic motions for the entire face during speech has rarely been explored. In this work, we present a joint audio-text model to capture the contextual information for expressive speech-driven 3D facial animation. The existing datasets are collected to cover as many different phonemes as possible instead of sentences, thus limiting the capability of the audio-based model to learn more diverse contexts. To address this, we propose to leverage the contextual text embeddings extracted from the powerful pre-trained language model that has learned rich contextual representations from large-scale text data. Our hypothesis is that the text features can disambiguate the variations in upper face expressions, which are not strongly correlated with

*Work done at HKUST

[†]Corresponding author

Authors' addresses: Yingruo Fan, The University of Hong Kong, Hong Kong, Hong Kong, yingruo@hku.hk; Zhaojiang Lin, The Hong Kong University of Science and Technology, Hong Kong, Hong Kong, zlinao@ust.hk; Jun Saito, Adobe Research, USA, jsaito@adobe.com; Wenping Wang, Texas A&M University, USA, wenping@cs.hku.hk; Taku Komura, The University of Hong Kong, Hong Kong, Hong Kong, taku@cs.hku.hk.

© 2022 Association for Computing Machinery.

This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in Proceedings of the ACM on Computer Graphics and Interactive Techniques, https: //doi.org/10.1145/3522615.

the audio. In contrast to prior approaches which learn phoneme-level features from the text, we investigate the high-level contextual text features for speech-driven 3D facial animation. We show that the combined acoustic and textual modalities can synthesize realistic facial expressions while maintaining audio-lip synchronization. We conduct the quantitative and qualitative evaluations as well as the perceptual user study. The results demonstrate the superior performance of our model against existing state-of-the-art approaches.

CCS Concepts: • **Computing methodologies** → **Animation**.

Additional Key Words and Phrases: joint audio-text model, expressive speech-driven 3D facial animation, contextual text embeddings, pre-trained language model

ACM Reference Format:

Yingruo Fan, Zhaojiang Lin, Jun Saito, Wenping Wang, and Taku Komura. 2022. Joint Audio-Text Model for Expressive Speech-Driven 3D Facial Animation. *Proc. ACM Comput. Graph. Interact. Tech.* 5, 1, Article 5775021 (May 2022), 15 pages. https://doi.org/10.1145/3522615

1 INTRODUCTION

Speech-driven 3D facial animation has garnered tremendous interest in computer graphics and vision. It can be used in a wide array of potential applications, including computer games, filmmaking, 3D telepresence system and other human-computer interactive interfaces. Generating expressive lip-synchronized 3D facial animation from speech remains challenging. There appear to be several reasons for this. First, the subtle changes in upper face expressions are weakly correlated with the audio [Cudeiro et al. 2019; Karras et al. 2017]. The same audio signal can be associated with different upper facial motion sequences and vice versa. Second, synthesizing natural facial muscle movements is difficult due to the complicated geometric structure of human faces [Edwards et al. 2016]. Third, the acoustic speech signal can vary from person to person due to the identity and physical differences between speakers [Sjerps et al. 2019].

Most existing studies on audio-driven 3D facial animation [Edwards et al. 2016; Suwajanakorn et al. 2017; Taylor et al. 2017; Zhou et al. 2018] focus on the issue of lip-synchronization, which might lack realism in the upper face such as eyelids and brows. However, it is important to animate a 3D talking avatar that produces speech utterances with vivid facial expressions. A small collection of studies [Karras et al. 2017; Richard et al. 2021; Wang et al. 2021] emphasize on the synthesized facial motions of the entire face as well as the lip-sync. The latest state-of-the-art approach [Richard et al. 2021] adopts a cross-modality loss to disentangle the audio-correlated and audio-uncorrelated face motions, thus ensuring the plausible animation of the upper face. Similar to [Karras et al. 2017; Richard et al. 2021], our focus is animating a 3D face mesh with realistic muscle movements in both the upper and lower face parts. On the other hand, we aim to predict the geometry of a 3D talking head instead of the pixel values of a 2D talking head image [Chen et al. 2020, 2018, 2019; Chung et al. 2017; Das et al. 2020; Ji et al. 2021; Mittal and Wang 2020; Prajwal et al. 2020; Vougioukas et al. 2020; Wiles et al. 2018; Zeng et al. 2020; Zhou et al. 2019, 2021]. In this paper, we present a novel approach that combines acoustic and textual modalities to animate a 3D talking head with vivid facial expressions, as illustrated in Figure 1. Given the audio and text signal, our method generates expressive 3D face animation sequences which have synchronized lip movements to the audio and realistic facial muscle movements (e.g., eye and eyebrow motions) for the entire face.

The human speech signal inherently involves both acoustic and textual features. Emotions can manifest acoustically in measures such as the speaking rate and the fundamental frequency. On the other hand, text contents can also convey the emotional state of the speaker. For example, if the speaker frowns and says "It is the most dreadful news" in a neutral voice, then it would be ambiguous when inferring the facial expression from the acoustic features. In this case, the emotional state is mostly reflected through the expression of language, e.g., the word "dreadful".

Inspired by this, we investigate the hypothesis that integrating acoustic and textual context could improve speech-driven 3D facial animation.

The pre-trained transformer-based representations [Devlin et al. 2018; Radford et al. 2019] have proven to be successful in various natural language processing tasks. Our motivation is that the pre-trained language model has learned rich contextual information, since it has been trained on the large scale text corpora. While for the existing 3D audio-visual datasets of facial animation [Fanelli et al. 2010; Karras et al. 2017; Richard et al. 2021], they only have a limited number of sentences. Training a model with audio alone might limit its ability to learn more diverse contexts. Intuitively, the high-level textual features could aid in understanding the emotional context of the speech. Previous studies on speech-driven gesture generation [Kucherenko et al. 2020; Yoon et al. 2020] have demonstrated the effect of the text modality and verified that using text and audio modalities together can further improve their results. Therefore, we exploit the text embeddings from GPT-2 [Radford et al. 2019] and incorporate them as part of input features to our model. This is different from previous linguistics-based methods which usually extract the phoneme-level features from the transcript. To the best of our knowledge, there has been no previous attempt at exploring the language model to resolve the ambiguity of facial expression variations for speech-driven 3D facial animation.

In our method, the language and acoustic modalities are modeled through two subnetworks separately. Furthermore, it is desirable to take into account the bimodal interactions between audio and text cues. To this end, we build a fusion layer, named Tensor Fusion [Zadeh et al. 2017], in our model to disentangle unimodal and bimodal dynamics. The Tensor Fusion layer explicitly models the unimodal and bimodal interactions based on the Cartesian product of audio and text embeddings. Ideally, considering audio and text modalities and their interactions should capture a wide range of variations in the speech. This allows the model to learn more expressive joint representations for facial expressions.

In summary, our main contributions include:

- Exploiting the contextual text embeddings for speech-driven 3D facial animation. Considering the potential use of a strong pre-trained language model, this work has implications for producing more expressive facial motions using audio and text signals as inputs.
- Providing meaningful insights into the effect of the text modality on the upper part of the predicted face mesh by performing visualization analysis.
- Extensive evaluation of the proposed approach on a 3D audio-visual corpus [Fanelli et al. 2010] and a perceptual user study of facial animation quality in terms of realism. The results show that our approach can synthesize more realistic 3D facial animations, as compared to existing state-of-the-art approaches.

2 RELATED WORK

Speech-driven facial animation has been a popular topic in computer graphics and vision. In computer vision, it is also known as talking face generation. Here we mainly review the previous speech-driven 3D facial animation approaches, where the output is the 3D mesh animation. The related works can be categorized into two broad groups: linguistics-based methods and audio-based methods.

2.1 Linguistics-based Methods

The concept of phonemes is well developed in speech animation. Viseme [Fisher 1968] is derived from a set of phonemes that have similar visual appearances on the lips. The popular dominance model [Massaro et al. 2012] uses the dominance functions of consecutive phonemes and determines

the influence of each phoneme on the respective facial animation parameters. Taylor et al. [2012] propose the one-to-many phoneme to viseme mappings, a.k.a. dynamic visemes, using a data-driven method. Xu et al. [2013] define the animation curves for a canonical set of visemes to achieve the coarticulation effects. Alternatively, Edwards et al. [2016] introduce the linguistic rules to animate the phonemes of a JALI (Jaw and Lip) rig, which is built upon the Facial Action Coding System (FACS). The follow-up work VisemeNet [Zhou et al. 2018] predicts a stream of phoneme-groups and facial landmarks from audio, and produces the JALI parameters used for animating the 3D face.

Previous linguistics-based approaches focus on synthesizing the synced mouth movements and explore diverse forms of complex linguistic rules, including phoneme-to-viseme mapping and coarticulation. Unlike them, we aim to discover high-level semantic features that might correspond to the emotional state, as such text representations would be useful for facial expression synthesis.

2.2 Audio-based Methods

Recently, a variety of approaches for animating the 3D talking avatar from audio have been proposed. The earlier method [Cao et al. 2005] implements a support vector machine classifier to detect the emotional state of the input audio and then synthesizes the 3D animations. Pre-labeling emotion categories is required, such as happy, sad and angry, for all the sentences. Karras et al. [2017] disambiguate the variations in facial expression by introducing a latent emotional state vector. They collect 3-5 minutes of animation data for two subjects. However, their model is person-specific and not applicable to unseen speakers. Taylor et al. [2017] propose a sliding window approach on phoneme subsequences and generalize their model to other avatars via re-targeting. Pham, Cheung, and Pavlovic [2017] obtain the blendshape coefficients from the 2D videos first. Subsequently, the predicted coefficients are regarded as ground truth for training an audio-to-face model. Despite the generalization ability, both methods [Pham et al. 2017; Taylor et al. 2017] rely on 2D videos rather than high-resolution 3D face scans, which may influence the quality of the resulting animation. The CNN-based model VOCA [Cudeiro et al. 2019] extracts the audio features from a pre-trained DeepSpeech model [Hannun et al. 2014] and then learns the mapping of audio to the 3D vertex offsets of a face model. However, as can be noted, VOCA only learns the facial motions that are mostly present in the lower part of the face. In a recent work, MeshTalk [Richard et al. 2021], the cross-modality loss is employed to disentangle the information about upper and lower face motion, thus generating full facial 3D animation from speech.

The approaches closely related to ours are VOCA [Cudeiro et al. 2019], MeshTalk [Richard et al. 2021] and the Audio2Face model [Karras et al. 2017] in the sense that their models are trained on high-resolution 3D face scans and are designed to decode the learned representations to 3D vertex space. Different from their methods using audio features alone, we propose to incorporate the contextual embeddings from Transformer-based GPT-2 to aid in understanding the emotional context, in order to produce a more diverse range of facial expressions.

3 METHOD

3.1 Overview of the Architecture

The pipeline of our proposed method is illustrated in Figure 2. Overall, the network is composed of four main components: the audio encoder, the text encoder, the tensor fusion module and the decoder. The speaker identity is represented as a one-hot vector. The audio-text embeddings are mapped into a high-dimensional space of 3D vertex offsets, i.e., 70110 outputs. In the end, the predicted vertex offsets are added to the vertices of the neutral template mesh, generating the expressive 3D face animation sequence.

3.2 Problem Formulation

Let us assume a video sequence of T face meshes $Y = [y_t]_{t=1:T}$, each of which has V vertices in 3D coordinates. Given a sequence of acoustic features $X^a = [x^a_t]_{t=1:T}$ as well as the textual features $X^l = [x^l_t]_{t=1:T}$, our goal is to predict the corresponding 3D facial animation sequence $\tilde{Y} = [\tilde{y}_t]_{t=1:T}$. Since the speech signal and the facial movements are synchronized temporally, we explicitly align the audio and text features to the corresponding video frame.

3.3 Audio Encoding

Inspired by [Oord et al. 2016; Richard et al. 2021], we adopt a stack of four dilated temporal convolutional layers as the first part of the audio encoder. Dilated convolutions [Yu and Koltun 2015] were first proposed in computer vision for context aggregation. Here we use the time-dilated convolution so that the feature maps and output can keep the high-resolution information of the input in the temporal domain. Formally the dilated discrete convolution operation can be described as follows:

$$(\mathbf{F} *_{l} \mathbf{k})(\mathbf{p}) = \sum_{s+lt=\mathbf{p}} \mathbf{F}(s)\mathbf{k}(t), t \in [-r, r],$$
(1)

where $*_l$ denotes the convolution with dilation factor l, **F** is a discrete function, **p** is the feature map, and **k** is the filter with size $(2r + 1)^2$. Particularly, $*_l$ supports exponentially increasing the receptive field without losing resolution. Hence, we empirically use a dilation factor that is exponentially increased with respect to the network depth. Each convolution layer has the same number of filters and the same kernel size, followed by the LeakyReLU activation. Besides, the residual connections are used to maintain the temporal information about input acoustic features.

Similar to [Cudeiro et al. 2019], we indicate the speaker identity with a one-hot vector to model the speaking style. This one-hot vector is concatenated to the output of the final convolution layer of the audio encoder, forming the input of the following fully-connected layer. During inference, altering the one-hot vector can manipulate the output animation in different speaking styles.

3.4 Text Encoding

The introduction of transformers [Vaswani et al. 2017] has shown promise in the field of language and vision. Transformer-based models have been applied to capture the contextual information of different modalities such as text and image [Ramesh et al. 2021]. To incorporate the contextualized semantic features, we adopt the pre-trained Transformer-based language model, GPT-2 [Radford et al. 2019], to process the input text modality. Specifically, each sentence is encoded into 768 features per word. Due to different sequence lengths of text and audio, we upsample the GPT-2 features to match the length of the audio feature sequence according to the duration time of each word. The text encoder consists of two fully-connected layers, a LeakyReLU layer and an LSTM layer. The outputs of the text encoder are taken as the features for the text modality.

3.5 Tensor Fusion

At the fusion stage, we apply a tensor fusion layer [Zadeh et al. 2017] to disentangle the unimodal dynamics and bimodal interactions. Prior studies on multimodal fusion [Poria et al. 2016; Zadeh et al. 2017] have demonstrated that this method can better capture multimodal interactions than the simple concatenation operation. The tensor fusion layer first obtains the unimodal feature vectors, namely the audio embedding H^a and the language embedding H^l , by passing the unimodal inputs X^a and X^l into the audio encoder and the text encoder, respectively. Subsequently, the tensor



Fig. 2. Overview of our method. First, the raw audio is converted into mel spectrogram whereas the text representations are extracted from the pre-trained GPT-2. Second, these extracted features are fed to the respective encoders to produce the audio and text embeddings. The one-hot speaker embedding is concatenated to the output of the final convolution layer of the audio encoder. Third, we employ a tensor fusion layer to fuse the encoded contextual features of audio and text modalities. Finally, the decoder produces a sequence of animated 3D face meshes, each of which is represented as a 23370-dimensional vector in 3D vertex coordinates.

fusion layer produces the multimodal output representation H^m by performing a differentiable outer product operation, which is computed by:

$$\mathbf{H}^{m} = \begin{bmatrix} \mathbf{H}^{a} \\ 1 \end{bmatrix} \otimes \begin{bmatrix} \mathbf{H}^{l} \\ 1 \end{bmatrix}, \tag{2}$$

where \otimes denotes the outer product. The interpretation of tensor fusion is illustrated in Figure 2. The multimodal dynamics are generated by adding an extra constant dimension with value 1. This results in a 2-D Cartesian space, where the horizontal and vertical axes are defined by $[\mathbf{H}^a \ 1]^T$ and $[\mathbf{H}^l \ 1]^T$, respectively.

3.6 3D Facial Animation Generation

In the decoder module, the sequence of **T** fused features $\mathbf{H}^{\mathbf{m}} = [\mathbf{h}^{\mathbf{m}}_{t}]_{t=1:\mathbf{T}}$ obtained in Equation 2 is fed to a fully-connected layer to project it to a lower-dimensional output vector space. After that, two bidirectional LSTMs are applied to model temporal dependencies. The bidirectional LSTM makes use of both the preceding and succeeding context information by processing the sequence both forward and backward. Finally, the output sequence of two bidirectional LSTMs is used as the input sequence of a fully-connected layer, which maps the learned representation to 3D vertex space. Therefore, the decoding flow can be described as:

$$[\tilde{\mathbf{y}}_t]_{t=1:\mathrm{T}} = \mathrm{FC}(\mathrm{BLSTM}(\mathrm{FC}([\mathbf{h}^m_t]_{t=1:\mathrm{T}}))). \tag{3}$$

During the training phase, we minimize the mean squared difference between the ground-truth vertices $\mathbf{Y} = [\mathbf{y}_t]_{t=1:T}$ and the corresponding outputs $\tilde{\mathbf{Y}} = [\tilde{\mathbf{y}}_t]_{t=1:T}$ produced by the decoder, which is computed as follows:

$$\mathcal{L} = \sum_{t=1}^{\mathrm{T}} \sum_{v=1}^{\mathrm{V}} \left\| \tilde{\mathbf{y}}_{t,v} - \mathbf{y}_{t,v} \right\|^2.$$
(4)

4 EXPERIMENTS

4.1 Dataset

We train and evaluate our models using the publicly-available audio-visual dataset BIWI [Fanelli et al. 2010]. This corpus involves 14 native speakers (8 females and 6 males), each of which utters 40 short English sentences. The corresponding text transcriptions are also provided. The 3D geometry data of the performance of the speaker is acquired at 25 fps. All face meshes have the same topology and the same number of vertices (23370). For each speaker, the static face mesh with a neutral expression is provided as the template. Each sentence lasts 4.67s long on average and is recorded twice: with and without emotion. We use the sentences recorded in the emotional context for our experiments and remove the sequences that are not complete in terms of the tracked 3D face scans or the audio files.

We divide the dataset into training, validation and testing parts. The training set is composed of 192 sentences, pronounced by six subjects (each subject utters 32 sentences). The validation part contains 24 sentences, pronounced by the same six subjects (each subject utters 4 sentences). The testing part includes two sets of test sentences: the first are 4 sentences of the same six subjects in the training set, i.e., 24 sentences (Test set A); the second is a set of 4 sentences of eight subjects not seen in the training set, i.e., 32 sentences (Test set B). No overlap of sentences exists in the training, validation, and test sets, and there is no overlap of subjects presents in the training set and Test set B. Here Test set A is used for quantitive evaluation whereas Test set B represents a more challenging set where the subjects are unseen.

As the 3D motion capture process is very expensive, a smaller training set is deemed desirable. The pioneer work [Karras et al. 2017] demonstrates a small amount of 3D data (3-5 minutes) are enough for training a deep network. Hence, we consider using 192 sentences (about 18 minutes in total) to train our network is feasible. Also, our model benefits from the language model that has been trained on a large scale text corpora. This makes our system easier to generalize well to novel sentences.

4.2 Implementation Details

4.2.1 Data pre-processing. To extract the audio features, we utilize the librosa [McFee et al. 2015] library to transform the raw audio to 128-channel mel-frequency power spectrograms and convert the power spectrograms to decibel (dB) units. Besides, we synchronize the input audio sequence to the corresponding 3D face animation sequence at 25 fps. Hence, the input audio feature vector for every visual frame has 128 dimensions. For the text features, we extract the word representations from the pre-trained GPT-2 small model [Radford et al. 2019] (12 layers, 768-dimensional hidden state). To ensure that the input audio and text feature sequences share the same time steps, we employ the Gentle forced aligner [Ochshorn and Hawkins 2017] to find the start and end timestamps of each word within a sentence. Given the known duration time of each word, we apply zero padding to the frames that do not contain a word as there is usually a pause between words. For each word that contains semantic information, we repeat the extracted word representation according to the word's duration time. The text feature vector for every visual frame has 768 dimensions. Then, we set each text feature vector to the average of its surrounding text feature vectors (previous 8 vectors and future 7 vectors). After pre-processing, both the input audio and text feature sequences have the same length as the 3D face animation sequence.

Table 1.	Quantitative evaluation of different methods. We report the mean absolute error (MAE) between			
the ground-truth data and the generated results for AU coding consistency evaluation.				

Methods	Upper Face AUs	Lower Face AUs
VOCA	0.278 ± 0.133	0.229 ± 0.093
MeshTalk	0.236 ± 0.124	0.230 ± 0.096
Ours	0.221 ± 0.117	0.212 ± 0.085

Table 2. Perceptual user study results. The number reports the percentage (%) of A/B tests where A is chosen over B.

Model Pairs	Full Face	Upper Face	Lips
Ours/VOCA	84.20 ± 6.21	84.20 ± 6.17	84.90 ± 6.81
Ours/MeshTalk	71.35 ± 3.13	68.92 ± 5.00	71.88 ± 3.65
Ours/GT	19.79 ± 5.73	24.48 ± 9.39	19.27 ± 6.14

4.2.2 Network architectures. The overall architecture of the model is shown in Figure 2. The audio encoder first applies 4 dilated convolutional layers each equipped with 128 filters of size 3 × 3. The dilation factors of these convolutions are 1, 2, 4 and 8, respectively. Each dilated convolutional layer is followed by the LeakyReLU activation. Additionally, the encoder uses residual connections between layers. The output of the dilated convolutions is then concatenated with a 6-dimensional one-hot vector. Finally, the concatenated feature vector goes through a 128-unit FC layer. The text encoder consists of a 128-unit FC layer, a LeakyReLU layer, a 64-unit FC layer and a 64-unit LSTM layer. The tensor fusion layer takes a 128-dimensional audio embedding vector and a 64-dimensional text embedding vector as inputs and produces a high-dimensional feature vector. The subsequent decoder consists of a 128-unit FC layer, two 128-unit bidirectional LSTM layers and a 70110-unit FC layer.

4.2.3 Training setup. All the experiments are implemented using PyTorch. In the training stage, the optimization function is Adam, with a constant learning rate of 1e-4 and a batch size of 1. We train our model for 100 epochs and apply it to the testing data directly.

4.3 Comparisons Against State-of-the-Arts

We compare our method with the state-of-the-art approaches in speech-driven 3D facial animation, VOCA [Cudeiro et al. 2019] and MeshTalk [Richard et al. 2021]. We use the publicly available implementation for VOCA and train the model on BIWI. Since the implementation of MeshTalk is not publicly available yet, we reproduce their method to the best of our understanding.

4.3.1 Quantitative Evaluation. Given the many-to-many mappings between the speech signal and the facial expressions, it is not adequate to use the norm on the prediction error between the original and generated outputs to measure the quality of speech-driven 3D facial animation, as suggested by the previous study [Cudeiro et al. 2019]. The related works [Cudeiro et al. 2019; Karras et al. 2017] only include the perceptual and qualitative evaluations. Since we focus on the realism of facial expressions, we perform the quantitative evaluation based on the popular facial action unit representation FACS [Eckman and Friesen 1978]. FACS is the "gold standard" for quantifying facial muscle movements, a.k.a., Action Units (AUs). The combination of different AUs can represent all possible facial expressions. Here we first render both the ground-truth and generated 3D geometry

data with texture as 2D videos with a resolution of 800×800 . Then we utilize OpenFace [Baltrušaitis et al. 2016] to detect the AU intensities for the rendered videos. The mean absolute error (MAE) is used as the performance metric for evaluating the AU coding consistency between the ground-truth and generated results. We categorize all AUs into two groups: upper face AUs (AU1,2,4,5,6,7,9 and 45) and lower face AUs (AU10,12,14,15,17,20,23,25 and 26), and compute the average MAEs for these two groups, respectively. The quantitative evaluation for all the methods on Test set **A** of the BIWI dataset is summarized in Table 1.

Since the subtle facial movements are difficult to quantify, we recommend watching the clips in the supplementary video to assess the quality of the results.

4.3.2 Perceptual Evaluation. We carry out the perceptual user study on Amazon Mechanical Turk (AMT). We compare our approach to VOCA, MeshTalk and the captured ground-truth data using all sentences of Test set **B**, i.e., 32 sentences each spoken by a subject not seen in the training set. For VOCA and our method, we obtain the prediction results conditioned on all six training identities. Therefore, 192 (32 sentences \times 6 identities) pairs are evaluated for each row in Table 2. In total, 576 A vs. B pairs are created. In our study, a HIT consists of four pairs, one of which is used for qualification test (ground-truth vs. result produced by a not well trained model). Each HIT is evaluated by 3 judges and a total of 576 HITs are collected. Turkers are required to make the right choice for the qualification test before they are allowed to submit HITs. For each pair, two animations with the same subject and audio are presented side-by-side and three questions are designed in terms of the full face, the upper face and the lips. Turkers are instructed to watch the videos and asked to determine which animation looks more realistic for each given question. We shuffle the order of the videos and also randomize which method is A and which is B for each pair to eliminate any bias.

In Figure 3, we show the designed user interface on Amazon Mechanical Turk (AMT). The recruitment requirement is that the participants must have an HIT (human intelligence task) approval rate of at least 95% and have finished over 5000 HITs. 576 HITs were created and a total of 156 unique workers completed our HITs.

Table 2 shows the comparison of success rates of our method against VOCA, MeshTalk, and the ground truth. In comparison to VOCA, our method wins about 84% of the pairwise comparisons for the three face parts. This is not surprising as VOCA does not synthesize the upper face movements well. Turkers perceive the animation results of our method more realistic than those of MeshTalk, with the success rate of 71.35% for the full face, 68.92% for the upper face and 71.88% for the lips. As expected, the animations synthesized by our approach are generally perceived as less realistic than the ground truth, with the success rate of 19.79% for the full face, 24.48% for the upper face and 19.27% for the lips. The result is similar to that reported in [Cudeiro et al. 2019], as inferring from speech alone can not synthesize the subtle subject-specific details of the unseen subjects [Cudeiro et al. 2019].

4.3.3 Qualitative Evaluation. The supplementary video includes qualitative comparisons with the captured performance data, VOCA and MeshTalk. We also compare our results with the dynamic viseme method [Taylor et al. 2017] and the emotion-driven method [Karras et al. 2017] using the footage from their supplementary videos. To validate the ability to synthesize different speaking styles for the same speech, we provide the generated results by conditioning on all six training subjects. To assess the generalization capability of our model, we test our model using the speech clips extracted from TED videos on YouTube. As shown in the supplementary video, our method produces realistic and natural-looking facial animations, synthesizes various speaking styles and generalizes to unseen voices. More importantly, the generated facial motions appear to be expressive when the emotional state of the speaker is obvious. We can notice that the talking avatar shows

Instructions:

Please watch the four short videos (duration 4~7s) of two animated talking heads. You need to choose the talking heac (the left or the right) that moves more naturally in terms of the full face, the upper face and the lips. **Reminder 1:** Please **turn on the sound on your computer** while you are watching the videos.

Reminder 2: Some of the videos (one or two) are qualification testing videos. Your task might get rejected if you make the choices randomly



Comparing the two full faces (Left and Right), which one is more realistic?

 \bigcirc The Left one is more realistic

 \bigcirc The Right one is more realistic

Comparing the upper part of two faces, which one is more realistic?

 \bigcirc The Left one is more realistic

 \bigcirc The Right one is more realistic

Comparing the lips of two faces, which one is more realistic?

○ The Left one is more realistic

○ The Right one is more realistic

Fig. 3. Screen shot of the user interface on AMT.

more expressive facial movements when saying certain emotional words, e.g., "great", "thank you", etc. Please refer to the supplementary video for qualitative results.

4.4 Ablation Study

4.4.1 Ablation on the input modalities. To explore the importance of different input modalities, we carry out an ablation experiment to compare the performance among unimodal and bimodal models. From Table 3, the bimodal model "Audio+Text" performs best, outperforming the unimodal models "Audio Only" and "Text Only". It can be observed that using text alone has relatively large errors for the lower face AUs, whereas using audio alone degrade the performance for the upper face AUs. This suggests that the text information is beneficial in learning the variations in the upper face expressions and the audio information has a more strong correlation with the mouth movements. With both the audio and text modalities incorporated, the bimodal model leads to the best performance for the upper and lower face AUs. This indicates that audio and text modalities are complementary to each other. We also provide qualitative examples for modality ablation in our supplementary video. From the videos, we find that using the text alone generates expressive muscle movements in the upper face part, e.g., opening the eyes wide when pronouncing the word "thank".

I3D'22, May 03-05, 2022,



Fig. 4. t-SNE visualization of the text embedding. The data points represent the projected text features. The color of each point corresponds to the AU5 intensity predicted from the corresponding rendered video frame. Left: the projected text features from all testing sentences of a subject. Right: the projected text features from one testing sentence of the same subject.

4.4.2 Ablation on the fusion methods. Additionally, we conduct an experiment to compare the performance of the bimodal model using different fusion methods. As presented in Table 3, the tensor fusion strategy brings improvements compared to the simple concatenation method. This implies that the tensor fusion layer can better combine the audio and text cues. The 2-D Cartesian space might capture more variations in the speech, and hence allow a more expressive space for modeling the facial expression. Besides, we conduct qualitative comparisons between the two fusion strategies to show the visual differences. As shown in the supplementary video, we observe that the results generated by the model with the tensor fusion strategy have better temporal stability in the mouth region.

4.5 Analysis of the Text Embedding

Figure 4 shows a visualization of the extracted text features from the final output of the text encoder. We project the text features of all the testing sentences of a testing subject to 2-dimensional space using t-SNE. We can see that the text features of the words within the same sentence are usually clustered. To better understand the text embeddings, we zoom in on the data points of one sentence ("What an excellent room you have, sir. Such expensive furnishings!") of the same subject, as shown in the right side of Figure 4. Given the known duration time of each word within the sentence, we put the words close to the corresponding data point. We also visualize the corresponding generated facial expression to uncover how the facial expression changes when the word is changed. Colors from purple to yellow denote lower and higher AU5 (Upper Lid Raiser) intensity. In addition to the clusters, we observe that the AU5 intensity is higher for some words such as "excellent". According to the relationship between AUs and emotions [Friesen et al. 1983], AU5 is correlated with surprise, which also implicitly reflects the emotional state of the speaker. As can be noticed, there are some isolated data points in the lower part of the figure, meaning that those frames do not cover a word.

Methods	Upper Face AUs	Lower Face AUs
Audio Only	0.251	0.224
Text Only	0.247	0.253
Audio+Text (<i>C</i>)	0.225	0.218
Audio+Text (TF)	0.221	0.212

Table 3. Ablation experiments on different input modalities and different fusion methods. The evaluation metric is mean absolute error. C = Concatenate, TF = Tensor Fusion.



Fig. 5. Correlation between the encoded features (left: audio; right: text) and the predicted vertex offsets. The correlation is measured by the absolute value of the Pearson Correlation Coefficient. Colors from blue to red represent lower and higher strength of correlation.

Additionally, we analyze which regions of the generated face mesh are influenced by the audio and text modalities. Specifically, after extracting audio and text features from their encoders, we calculate the absolute value of the Pearson correlation coefficient between the specific modality features and each predicted vertex offset (in total 23370 vertices). The greater the absolute value is, the stronger the relationship is. As shown in Figure 5, color codes are for strength of correlation. The Pearson analysis suggests that the audio modality has a strong association with vertices in the mouth region, and is moderately correlated with vertices in the cheek region. In contrast, the text modality shows a strong relationship with vertices in the upper face region, and also influences vertices in the cheek and upper lip regions.

5 CONCLUSION, LIMITATIONS, AND FUTURE WORK

In this work, we present a novel approach for expressive speech-driven 3D facial animation based on audio and text. We pioneer the use of the pre-trained language models and feel that the contextual text embeddings offer great promise to this field for producing more expressive facial motions. Our method exhibits better performance against previous approaches, which we show through quantitative and qualitative analysis and perceptual user studies. However, the proposed model also presents limitations and sometimes fails to model the lip closure well when encountering certain plosives. On the other hand, due to the limitations of the current data, our model lacks some additional motions such as eye gaze and head movement. Future work may include studying the correlations between the speech signal and other subtle motions to achieve a higher level of conversational realism. Also, the use of one-hot encoding potentially limits the scaling ability of our system to a larger number of subjects. This can be addressed by replacing the one-hot vector with an identity encoder. In the future, we will collect 3D visual-audio data of many subjects and

take into account subject-specific details by using an identity encoder, in order to improve the generalization ability of our model to arbitrary subjects.

ACKNOWLEDGMENTS

A part of this research is based on results obtained from a project, JPNP21004, commissioned by the New Energy and Industrial Technology Development Organization (NEDO).

REFERENCES

- Tadas Baltrušaitis, Peter Robinson, and Louis-Philippe Morency. 2016. Openface: an open source facial behavior analysis toolkit. In 2016 IEEE Winter Conference on Applications of Computer Vision (WACV). IEEE, 1–10.
- Yong Cao, Wen C Tien, Petros Faloutsos, and Frédéric Pighin. 2005. Expressive speech-driven facial animation. ACM Transactions on Graphics (TOG) 24, 4 (2005), 1283–1302.
- Lele Chen, Guofeng Cui, Celong Liu, Zhong Li, Ziyi Kou, Yi Xu, and Chenliang Xu. 2020. Talking-head generation with rhythmic head motion. In *Proceedings of the European Conference on Computer Vision*. Springer, 35–51.
- Lele Chen, Zhiheng Li, Ross K Maddox, Zhiyao Duan, and Chenliang Xu. 2018. Lip movements generation at a glance. In *Proceedings of the European Conference on Computer Vision*. 520–535.

Lele Chen, Ross K Maddox, Zhiyao Duan, and Chenliang Xu. 2019. Hierarchical cross-modal talking face generation with dynamic pixel-wise loss. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 7832–7841.

Joon Son Chung, Amir Jamaludin, and Andrew Zisserman. 2017. You said that? *arXiv preprint arXiv:1705.02966* (2017). Daniel Cudeiro, Timo Bolkart, Cassidy Laidlaw, Anurag Ranjan, and Michael J Black. 2019. Capture, learning, and synthesis

- of 3D speaking styles. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 10101–10111. Dipanjan Das, Sandika Biswas, Sanjana Sinha, and Brojeshwar Bhowmick. 2020. Speech-driven facial animation using
- cascaded gans for learning of motion and texture. In *European Conference on Computer Vision*. Springer, 408–424. Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional trans-
- formers for language understanding. arXiv preprint arXiv:1810.04805 (2018).
- P Eckman and W Friesen. 1978. Facial action coding system: A technique for the measurement of facial movement. *Consulting Psychologists Press* (1978).
- Pif Edwards, Chris Landreth, Eugene Fiume, and Karan Singh. 2016. JALI: an animator-centric viseme model for expressive lip synchronization. ACM Transactions on graphics (TOG) 35, 4 (2016), 1–11.
- Gabriele Fanelli, Juergen Gall, Harald Romsdorfer, Thibaut Weise, and Luc Van Gool. 2010. A 3-d audio-visual corpus of affective communication. *IEEE Transactions on Multimedia* 12, 6 (2010), 591–598.
- Cletus G Fisher. 1968. Confusions among visually perceived consonants. *Journal of speech and hearing research* 11, 4 (1968), 796–804.
- Wallace V Friesen, Paul Ekman, et al. 1983. EMFACS-7: Emotional facial action coding system. Unpublished manuscript, University of California at San Francisco 2, 36 (1983), 1.
- Awni Hannun, Carl Case, Jared Casper, Bryan Catanzaro, Greg Diamos, Erich Elsen, Ryan Prenger, Sanjeev Satheesh, Shubho Sengupta, Adam Coates, et al. 2014. Deep speech: Scaling up end-to-end speech recognition. *arXiv preprint arXiv:1412.5567* (2014).
- Xinya Ji, Hang Zhou, Kaisiyuan Wang, Wayne Wu, Chen Change Loy, Xun Cao, and Feng Xu. 2021. Audio-driven emotional video portraits. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 14080–14089.
- Tero Karras, Timo Aila, Samuli Laine, Antti Herva, and Jaakko Lehtinen. 2017. Audio-driven facial animation by joint end-to-end learning of pose and emotion. *ACM Transactions on Graphics (TOG)* 36, 4 (2017), 1–12.
- Taras Kucherenko, Patrik Jonell, Sanne van Waveren, Gustav Eje Henter, Simon Alexandersson, Iolanda Leite, and Hedvig Kjellström. 2020. Gesticulator: A framework for semantically-aware speech-driven gesture generation. In Proceedings of the 2020 International Conference on Multimodal Interaction. 242–250.
- DW Massaro, MM Cohen, M Tabain, J Beskow, and R Clark. 2012. Animated speech: research progress and applications. *Audiovisual Speech Processing* (2012), 309–345.
- Brian McFee, Colin Raffel, Dawen Liang, Daniel PW Ellis, Matt McVicar, Eric Battenberg, and Oriol Nieto. 2015. librosa: Audio and music signal analysis in python. In *Proceedings of the 14th python in science conference*, Vol. 8. Citeseer, 18–25.
- Gaurav Mittal and Baoyuan Wang. 2020. Animating face using disentangled audio representations. In Proceedings of the IEEE Winter Conference on Applications of Computer Vision. 3290–3298.

RM Ochshorn and Max Hawkins. 2017. Gentle forced aligner. github. com/lowerquality/gentle (2017).

Aaron van den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew Senior, and Koray Kavukcuoglu. 2016. Wavenet: A generative model for raw audio. arXiv preprint arXiv:1609.03499 (2016).

- Hai X Pham, Samuel Cheung, and Vladimir Pavlovic. 2017. Speech-driven 3D facial animation with implicit emotional awareness: a deep learning approach. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops. 80–88.
- Soujanya Poria, Iti Chaturvedi, Erik Cambria, and Amir Hussain. 2016. Convolutional MKL based multimodal emotion recognition and sentiment analysis. In 2016 IEEE 16th international conference on data mining (ICDM). IEEE, 439–448.
- KR Prajwal, Rudrabha Mukhopadhyay, Vinay P Namboodiri, and CV Jawahar. 2020. A lip sync expert is all you need for speech to lip generation in the wild. In Proceedings of the 28th ACM International Conference on Multimedia. 484–492.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog* 1, 8 (2019), 9.
- Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. 2021. Zero-shot text-to-image generation. *arXiv preprint arXiv:2102.12092* (2021).
- Alexander Richard, Michael Zollhoefer, Yandong Wen, Fernando De la Torre, and Yaser Sheikh. 2021. MeshTalk: 3D Face Animation from Speech using Cross-Modality Disentanglement. *arXiv preprint arXiv:2104.08223* (2021).
- Matthias J Sjerps, Neal P Fox, Keith Johnson, and Edward F Chang. 2019. Speaker-normalized sound representations in the human auditory cortex. *Nature communications* 10, 1 (2019), 1–9.
- Supasorn Suwajanakorn, Steven M Seitz, and Ira Kemelmacher-Shlizerman. 2017. Synthesizing obama: learning lip sync from audio. ACM Transactions on Graphics (ToG) 36, 4 (2017), 1–13.
- Sarah Taylor, Taehwan Kim, Yisong Yue, Moshe Mahler, James Krahe, Anastasio Garcia Rodriguez, Jessica Hodgins, and Iain Matthews. 2017. A deep learning approach for generalized speech animation. ACM Transactions on Graphics (TOG) 36, 4 (2017), 1–11.
- Sarah L Taylor, Moshe Mahler, Barry-John Theobald, and Iain Matthews. 2012. Dynamic units of visual speech. In Proceedings of the 11th ACM SIGGRAPH/Eurographics conference on Computer Animation. 275–284.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*. 5998–6008.
- Konstantinos Vougioukas, Stavros Petridis, and Maja Pantic. 2020. Realistic speech-driven facial animation with gans. International Journal of Computer Vision 128, 5 (2020), 1398–1413.
- Qianyun Wang, Zhenfeng Fan, and Shihong Xia. 2021. 3D-TalkEmo: Learning to Synthesize 3D Emotional Talking Head. arXiv preprint arXiv:2104.12051 (2021).
- Olivia Wiles, A Koepke, and Andrew Zisserman. 2018. X2face: A network for controlling face generation using images, audio, and pose codes. In *Proceedings of the European Conference on Computer Vision*. 670–686.
- Yuyu Xu, Andrew W Feng, Stacy Marsella, and Ari Shapiro. 2013. A practical and configurable lip sync method for games. In *Proceedings of Motion on Games*. 131–140.
- Youngwoo Yoon, Bok Cha, Joo-Haeng Lee, Minsu Jang, Jaeyeon Lee, Jaehong Kim, and Geehyuk Lee. 2020. Speech gesture generation from the trimodal context of text, audio, and speaker identity. *ACM Transactions on Graphics (TOG)* 39, 6 (2020), 1–16.
- Fisher Yu and Vladlen Koltun. 2015. Multi-scale context aggregation by dilated convolutions. *arXiv preprint arXiv:1511.07122* (2015).
- Amir Zadeh, Minghai Chen, Soujanya Poria, Erik Cambria, and Louis-Philippe Morency. 2017. Tensor fusion network for multimodal sentiment analysis. arXiv preprint arXiv:1707.07250 (2017).
- Dan Zeng, Han Liu, Hui Lin, and Shiming Ge. 2020. Talking Face Generation with Expression-Tailored Generative Adversarial Network. In *Proceedings of the 28th ACM International Conference on Multimedia*. 1716–1724.
- Hang Zhou, Yu Liu, Ziwei Liu, Ping Luo, and Xiaogang Wang. 2019. Talking face generation by adversarially disentangled audio-visual representation. In *Proceedings of the AAAI Conference on Artificial Intelligence*. 9299–9306.
- Hang Zhou, Yasheng Sun, Wayne Wu, Chen Change Loy, Xiaogang Wang, and Ziwei Liu. 2021. Pose-controllable talking face generation by implicitly modularized audio-visual representation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 4176–4186.
- Yang Zhou, Zhan Xu, Chris Landreth, Evangelos Kalogerakis, Subhransu Maji, and Karan Singh. 2018. Visemenet: Audiodriven animator-centric speech animation. ACM Transactions on Graphics (TOG) 37, 4 (2018), 1–10.

A QUALIFICATION TEST FOR USER STUDY

In the user study, the qualification test is designed to avoid Turkers to select the answers randomly. If they failed to pass the test, a warning message would pop up when they clicked the submit button, as shown in Figure 6.

I3D'22, May 03-05, 2022,

Please answer the following questions, after you watch the video.				
Comparing the two full faces (Left and Right), which on The Left one is more realistic	You failed to pass the hidden test. Please check the video			
The Right one is more realistic Comparing the upper part of two faces which one is more	multiple times, your assignments will be rejected without further notice.			
The Left one is more realistic	Close			
The Right one is more realistic Comparing the lips of two faces, which one is more realistic?				
 The Left one is more realistic The Right one is more realistic 				
SUBMIT				

Fig. 6. Warning message of the hidden test.