Enhancing Speech-Driven 3D Facial Animation with Audio-Visual Guidance from Lip Reading Expert

Han EunGi^{1*}, Oh Hyun-Bin^{2*}, Kim Sung-Bin², Corentin Nivelet Etcheberry², Suekyeong Nam³, Janghoon Joo³, Tae-Hyun Oh^{1,2,4}

¹Grad. School of Artificial Intelligence and ²Dept. of Electrical Engineering, POSTECH, Korea ³KRAFTON, Korea

⁴Inst. for Convergence Research and Education in Advanced Technology, Yonsei University, Korea

Abstract

Speech-driven 3D facial animation has recently garnered attention due to its cost-effective usability in multimedia production. However, most current advances overlook the intelligibility of lip movements, limiting the realism of facial expressions. In this paper, we introduce a method for speech-driven 3D facial animation to generate accurate lip movements, proposing an audiovisual multimodal perceptual loss. This loss provides guidance to train the speech-driven 3D facial animators to generate plausible lip motions aligned with the spoken transcripts. Furthermore, to incorporate proposed audio-visual perceptual loss, we devise an audio-visual lip reading expert leveraging its prior knowledge about correlations between speech and lip motions. We validate the effectiveness of our approach through broad experiments, showing noticeable improvements in lip synchronization and lip readability performance.

Index Terms: Speech-driven 3D Facial Animation, Audio-Visual Speech Recognition, Multimodal Perceptual Loss

1. Introduction

The field of speech-driven 3D facial animation is growing rapidly, with a focus on generating realistic facial expressions from speech signals. Animating 3D faces in practical applications often requires retouching or post-correction through manual intervention by skilled animators, which demands substantial human resources and costs. In that sense, speech-driven 3D facial animation is widely receiving attention in industries such as entertainment, gaming, and virtual communication [8, 9, 10, 11], enhancing user experience and immersion.

There have been significant advancements towards adopting learning-based approaches in speech-driven 3D facial animation [7, 12, 13, 14, 15, 16, 17, 18]; *e.g.*, 3D facial movement generation is modeled by 1D convolutions over the temporal dimension (VOCA [7]) or Transformer layers (FaceFormer [13] and CodeTalker [14]). These methods focus on model architectures and exhibit appealing performance. Nonetheless, they primarily focus on minimizing Euclidean distance between ground truth and predicted mesh vertices, overlooking the importance of generating perceptually natural and intelligible lip movements, which is crucial for human visual comprehension.

To mitigate the unsatisfactory results, particularly in the lip region, we propose a method to guide the speech-driven 3D facial animation to better understand how the lip moves according to the spoken words, thus generating more plausible lip shapes. Specifically, we introduce an audio-visual multimodal perceptual loss by leveraging a lip reading expert [1], which incorporates both visual and speech inputs, to facilitate the speech-driven 3D facial animator to learn more speechrelated information and generate plausible lip movements. Furthermore, we implement the perceptual loss part in a two-stage training scheme. In the initial stage, we integrate the lip reading expert, which has been trained on extensive 2D talking face datasets [2, 3, 4, 5]. Subsequently, in the second stage, we finetune the lip reading expert to optimize it on 3D facial datasets [6, 7], concurrently with training a 3D facial animator. This not only allows us to leverage prior knowledge of large 2D datasets but to reduce the domain gap from 3D face rendering.

While our proposed method is novel, there have been related attempts [19, 20, 15] of using lip reading networks to enhance the intelligibility of lip shapes in their respective tasks. SelfTalk [15] introduces a training framework that jointly trains a 3D face reconstruction module with a lip reading module, so that reconstructed 3D faces to be guided generate accurate transcripts synchronized with lip movements. In contrast to our method, they rely on visual input to train the lip reading module, utilizing limited 3D datasets [6, 7], without leveraging prior knowledge from large-scale 2D talking face datasets [2, 3, 4, 5].

SPECTRE [20], which is closely related to our approach, demonstrates that leveraging a lip reading expert [21] with prior knowledge can improve 3D facial reconstruction performance by minimizing the feature distance of lip movements between original and rendered videos. They also rely on a visual-only expert, while our method leverages speech and visual modalities together in our perceptual loss.

We conduct extensive experiments on BIWI [6] and VO-CASET [7] and show that our method is effective on different speech-driven 3D facial animator baselines (*i.e.*, Face-Former [13] and CodeTalker [14]) compared to the model without 2D prior knowledge or speech cues of the lip reading expert. We further measure the quality of lip movements from lip readability's perspectives, such as Viseme Error Rate (VER) and Character Error Rate (CER) following SPECTRE [20].

Our main contributions can be summarized as follows:

- We propose an audio-visual perceptual loss, which guides the speech-driven 3D facial animator to learn more speechrelated information and generate plausible lip movements.
- We devise an audio-visual lip reading expert tailored for the audio-visual perceptual loss, achieved via a two-stage training strategy: incorporating prior knowledge from extensive

^{*}These authors contributed equally.

Project Page: https://3d-talking-head-avguide. github.io/



Figure 1: Overview of our proposed framework. We adopt the audio-visual lip reading expert [1] trained on the large-scale 2D datasets [2, 3, 4, 5] and finetune it on 3D datasets [6, 7] concurrently with training the 3D facial animator. Given an input speech signal, a 3D facial animator regresses a sequence of 3D facial meshes and the following lip reading expert predicts the spoken transcript considering both the input speech signal and the sequence of lip regions of output faces.

2D talking face datasets in the initial stage, followed by finetuning of the lip reading expert on 3D talking face datasets.

2. Methods

In this section, we describe the proposed method utilized in the 3D facial animator baselines (*i.e.*, FaceFormer [13] and CodeTalker [14]) in detail. The proposed framework consists of two components: a 3D facial animator and a speech-informed lip reading expert. The 3D facial animator regresses a sequence of 3D face vertices from input speech signal, while the lipreading expert maps the lip shape sequence, which is rendered with a differentiable face renderer, to textual representations. The key idea of our method is to leverage the prior knowledge of the lip reading expert and to incorporate the audio modality to the expert for guiding the 3D face animation model to generate more plausible lip shapes. Figure. 1 illustrates the whole pipeline of our proposed framework.

2.1. Speech-Driven 3D Facial Animator

The 3D facial animator learns to regress a sequence of 3D facial movements from an input speech signal. The regression process can be formulated as follows: Let $\mathbf{Y}_{1:\mathbf{T}} = (\mathbf{y}_1,...,\mathbf{y}_{\mathbf{T}})$ denotes a sequence of ground truth 3D face vertices, where \mathbf{T} is the length of facial scan sequences and $\mathbf{y}_t \in \mathbb{R}^{\mathbf{V} \times 3}$ represents the face mesh of each frame which consists of \mathbf{V} vertices. In addition, let $\mathbf{A}_{1:\mathbf{T}'} = (\mathbf{a}_1,...,\mathbf{a}_{\mathbf{T}'})$ be a sequence of speech representation, where \mathbf{T}' is the length of the input speech signal. The 3D facial animator predicts a sequence of 3D face vertices $\hat{\mathbf{Y}}_{1:\mathbf{T}} = (\hat{\mathbf{y}}_1,...,\hat{\mathbf{y}}_{\mathbf{T}})$ given a speech signal $\mathbf{A}_{1:\mathbf{T}'}$ as:

$$\hat{\mathbf{Y}}_{1:\mathbf{T}} = \text{FacialAnimator}_{\theta_1}(\mathbf{A}_{1:\mathbf{T}'}), \quad (1)$$

where θ_1 denotes the weight of the 3D facial animator. After generating the complete 3D facial motion sequence, the 3D facial animator is trained to predict accurate facial shapes by minimizing the Mean Squared Error (MSE) between the outputs of the 3D facial animator $\hat{\mathbf{Y}}_{1:\mathbf{T}}$ and the ground truth $\mathbf{Y}_{1:\mathbf{T}}$:

$$\mathcal{L}_{mse} = \sum_{\mathbf{t}=1}^{\mathbf{T}} \sum_{\mathbf{v}=1}^{\mathbf{V}} \|\hat{\mathbf{y}}_{\mathbf{t},\mathbf{v}} - \mathbf{y}_{\mathbf{t},\mathbf{v}}\|^{2}.$$
 (2)

2.2. Speech-Informed Lip Reading Expert

The speech-driven 3D facial animators exhibit impressive performance in lip synchronization ability. However, solely minimizing Euclidean distance between the ground truth and predicted face vertices is not sufficient to generate intelligible lip movements. To generate realistic lip movements , we incorporate a powerful lip reading expert [1] that has the prior knowledge of the correlation between lip motions and their corresponding text content, which is learned from the extensive 2D talking face datasets [2, 3, 5, 4]. Specifically, we render the sequence of 3D face vertices from the 3D face animator into 2D video frames $\hat{\mathbf{I}}_{1:T}$ with a differentiable face renderer as:

$$\mathbf{\tilde{I}}_{1:\mathbf{T}} = \text{Renderer}(\mathbf{\tilde{Y}}_{1:\mathbf{T}}).$$
(3)

We crop the rendered gray-scale video frames $\hat{\mathbf{I}}_{1:\mathbf{T}}$ around the lip regions, resulting in the sequence of lip-cropped video frames $\hat{\mathbf{L}}_{1:\mathbf{T}}$. Then, the sequence of lip-cropped video frames $\hat{\mathbf{L}}_{1:\mathbf{T}}$ is fed to the lip reading expert.

Apart from SelfTalk [15], our speech-informed lip reading expert not only exploits prior knowledge from the 2D largescale datasets but also incorporates both visual and speech information. This approach produces facial mesh deformation that better corresponds to the speech, compared to models with a lip reading expert that only considers lip shapes (refer to Sec. 3.4). The lip reading expert predicts the transcript given both the sequence of lip-cropped video frames $\hat{\mathbf{L}}_{1:\mathbf{T}}$ and a sequence of speech representation $\mathbf{A}_{1:\mathbf{T}'}$ as:

$$\hat{\mathbf{s}}_{1:\mathbf{T}} = \text{LipExpert}_{\theta_2}(\hat{\mathbf{L}}_{1:\mathbf{T}}, \mathbf{A}_{1:\mathbf{T}'}), \tag{4}$$

where θ_2 is the weight of the lip reading expert. Following [1], we incorporate the joint CTC/attention loss [22] into our objective function, which penalizes the error between the predicted transcript $\hat{s}_{1:T}$ and ground truth $s_{1:T}$. The CTC loss \mathcal{L}_{ctc} and the attention loss \mathcal{L}_{ce} are for learning the alignment between the predicted and actual transcripts, respectively. Thus, the lip expert is finetuned with Audio-Visual (AV) perceptual loss \mathcal{L}_{av} which can be represented as a weighted sum of two losses:

$$\mathcal{L}_{av} = \lambda_{ctc} \mathcal{L}_{ctc} + \lambda_{ce} \mathcal{L}_{ce}, \qquad (5)$$

where λ_* denotes the loss weight, respectively. We utilize the AV perceptual loss \mathcal{L}_{av} to guide the 3D facial animator to generate the output lip movements that are comprehensible enough to guess the spoken words.

2.3. Training Details

The objective function. As mentioned in SPECTRE [20], naïvely imposing the CTC loss to improve lip movement quality

invokes face distortion since the model may prioritize achieving perfect lip reading recognition, a common phenomenon observed in adversarial attacks [23, 24]. To address this issue, we add relative lip vertex loss \mathcal{L}_{rlv} as a regularizer which retains the spatial structure of the lip regions. The relative lip vertex loss \mathcal{L}_{rlv} is calculated as the Mean Squared Error (MSE) between the output lip vertices and the ground truth lip vertices:

$$\mathcal{L}_{\text{rlv}} = \sum_{\mathbf{t}=1}^{\mathbf{T}} \sum_{\mathbf{v}=1}^{\mathbf{V}_{L}} \|\hat{\mathbf{y}}_{\mathbf{t},\mathbf{v}} - \mathbf{y}_{\mathbf{t},\mathbf{v}}\|^{2}, \tag{6}$$

where V_L denotes the number of lip region vertices.

To sum up, we train the facial animator and finetune the speech-informed lip reading expert with the objective function:

$$\mathcal{L} = \lambda_{\rm mse} \mathcal{L}_{\rm mse} + \lambda_{\rm av} \mathcal{L}_{\rm av} + \lambda_{\rm rlv} \mathcal{L}_{\rm rlv},\tag{7}$$

where λ_* denotes the loss weight, respectively.

Implementation details. We conduct experiments on existing speech-driven 3D facial animation models, i.e., Face-Former [13] and CodeTalker [14], using them as a 3D facial animator in our framework. Since CodeTalker separately trains the discrete motion prior and the auto-regressive facial animator, we train only the latter using pre-trained motion prior part. We use the Adam optimizer and set the learning rate to 1e-4 without weight decay for 100 epochs. Note that we reproduce all the 3D face animator baselines using publicly accessible codes and configurations. We feed raw audio inputs into the Wav2Vec2.0 [25] encoder and extract the audio features from the last hidden state, following [13, 14, 15]. We adopt the same architecture of Auto-AVSR [1] as our lipreading expert, which consists of a visual encoder, an audio encoder, a multi-layer perceptron (MLP), a projection layer, and a transformer decoder. We train the lip reading expert with joint CTC/attention loss [22] on the LRS2 [2], LRS3 [3], AVSpeech [4], and Voxceleb2 [5] datasets, following the same training procedure of Auto-AVSR. All experiments are conducted on a single NVIDIA A6000 GPU.

3. Experiments

3.1. Experimental setup

3.1.1. Datasets

We use two publicly available 3D datasets, VOCASET [7] and BIWI [6], for training and testing. Both of them include the audio-3D scan pairs of utterances spoken in English. We adopt the same data splits, *i.e.*, training, validation, and test splits, following the FaceFormer [13] and CodeTalker [14].

VOCASET. VOCASET [7] provides 480 audio-facial motion sequences for 12 subjects, captured at 60 Frames Per Second (FPS). The dataset comprises 255 sentences, some of which are spoken by multiple speakers. The facial meshes are aligned with the FLAME [26] topology, containing 5023 vertices.

BIWI. BIWI [6] is a 3D audio-visual dataset with 40 unique sentences, which are all shared across 14 subjects and captured at a 25 FPS. The dataset provides dynamic 3D face geometry aligned with 23,390 vertices. Each utterance is repeated twice with and without emotion, and we used the emotional subset. There are 190 training sentences, 24 validation sentences, and two test datasets: BIWI-Test-A, containing 24 sentences from 6 subjects seen during training, and BIWI-Test-B, including 32 sentences from 8 unseen subjects.

Table 1: Quantitative evaluation results on BIWI-Test-A.

Methods	$\begin{array}{c} \text{LVE} \downarrow \\ (\times 10^{-4} \text{ mm}) \end{array}$	$\operatorname{CER} \downarrow$	$\text{VER}\downarrow$
FaceFormer [13]	6.0449	72.588%	68.777%
+ AV Guidance	5.5061	68.423%	62.422%
CodeTalker [14]	5.3711	72.592%	65.593%
+ AV Guidance	4.8403	70.711%	63.299%

Table 2: Quantitative evaluation results on VOCASET test split.

Methods	$\begin{array}{c} \text{LVE} \downarrow \\ (\times 10^{-5} \text{ mm}) \end{array}$	$\operatorname{CER} \downarrow$	VER \downarrow
FaceFormer [13]	3.2496	76.244%	66.932%
+ AV Guidance	3.0987	71.589%	60.250%
CodeTalker [14]	4.0557	75.988%	65.105%
+ AV Guidance	3.9884	75.971%	64.912%

3.1.2. Evaluation metrics

Lip synchronization. To measure lip synchronization performance, we calculate Lip Vertex Error (LVE), which is a widely used metric for speech-driven 3D facial animation evaluation. It computes the maximal L2 error by comparing all lip vertices of each predicted frame to the ground truth and takes the average over all frames in the test set.

Lip readability. LVE alone may not be enough to evaluate the lip movements, especially in the aspect of lip readability. As a complement, we measured Character Error Rate (CER) and Viseme Error Rate (VER) between the actual and the predicted text from the lip reading expert. VER is calculated by converting the predicted and actual text to visemes using the Amazon Polly phoneme-to-viseme mapping [27], following [20].

3.2. Quantitative Results

We incorporate our method into the 3D facial animator baselines (*i.e.*, FaceFormer and CodeTalker) and calculate the Lip Vertex Error (LVE), Character Eror Rate (CER) and Viseme Error Rate (VER) for all sequences in the BIWI-TEST-A and VOCASET-Test datasets. As shown in Table. 1 and 2, Audio-Visual Guidance (AV Guidance), which includes 1) prior knowledge from the extensive 2D talking face dataset, 2) the relative lip vertex loss, and 3) speech information into the lip reading expert, improves all the evaluation metrics compared to the 3D facial animator baselines on both datasets. This indicates that our method helps to generate intelligible lip movements. In particular, the LVE on the BIWI-Test-A is 10% lower than the baselines, which shows the effectiveness of our proposed audiovisual perceptual loss.

3.3. Qualitative Results

We also conduct qualitative evaluation to assess the effectiveness of our method. Figure. 2 illustrates visual comparisons between the ground truth, FaceFormer [13], CodeTalker [14], FaceFormer with AV Guidance, and CodeTalker with AV Guidance on both the VOCASET test split and BIWI-Test-A. To ensure fair comparisons, we provide all models with the same speaking style as conditional input. To evaluate the lip synchronization quality, we show several frames of output facial animations generated from the same audio input for each method. We observe that our proposed methods produce more accurate lip closure movements compared to the baselines, accurately representing fully closed lip motions, particularly for syllables such



Figure 2: Qualitative comparisons of output facial movements on VOCASET and BIWI. Compared to the 3D facial animator baselines [13, 14], the outputs of our method show better lip synchronization quality for both lip closure and opening words.



Figure 3: t-SNE visualization for features of audiovisual/visual-only lip reading expert. Distinct separation of features for words "just" (red) and "must" (blue) is observed in the audio-visual lip reading expert. However, with visual-only input, features become entangled, hindering distinction.

as "/m/" or "/b/". Additionally, our method exhibits improved lip synchronization on mouth-opening as well, such as the syllable "/ Λ /". These comparison results underscore the effectiveness of employing an AV Guidance in achieving intelligible lip motions and accurately discerning various pronunciations.

3.4. Ablation Study

We conduct ablation studies to demonstrate the effectiveness of the different components of our method. Specifically, we investigate the impact of AV Guidance, which comprises of prior knowledge of lip reading expert, the relative lip vertex loss, and the audio-visual perceptual loss.

Impact of prior knowledge of lip reading expert. To investigate the effectiveness of the prior knowledge from the lip reading expert, we jointly train the 3D facial animator and the lip reading expert from scratch on the BIWI-Test-A dataset. Table. 3 shows noticeable degradation in lip vertex error on baseline models [13, 14] without prior knowledge. These results indicate that the prior knowledge of lip reading expert plays an important role in generating more accurate lip motions, guided by the correlation between lip motion and spoken text in 2D

Table 3: Ablation study for our components on BIWI-Test-A.

Methods	$LVE \downarrow (\times 10^{-4} \text{ mm})$
FaceFormer [13] + AV Guidance	5.5061
w/o prior knowledge	5.9344
w/o relative lip vertex loss \mathcal{L}_{rlv}	5.9023
w/o speech information	6.0352
CodeTalker [14] + AV Guidance	4.8403
w/o prior knowledge	5.4271
w/o relative lip vertex loss \mathcal{L}_{rlv}	5.6524
w/o speech information	5.3155

talking face datasets [2, 3, 4, 5].

Impact of relative lip vertex loss. We investigate the impact of removing the relative lip vertex loss \mathcal{L}_{rlv} by optimizing the baseline with AV Guidance but excluding the \mathcal{L}_{rlv} loss. Table. 3 show a deterioration in lip shape generation on both baseline models [13, 14], when removing the relative lip vertex loss. This underscores the crucial role of this loss as a regularizer for retaining the spatial structure of the lip regions. To assess the effectiveness of the relative lip vertex loss in AV Guidance regarding the improvement, we conduct an additional experiment on BIWI-Test-A. In this experiment, we optimize the baseline model with its original loss and the relative lip vertex loss. FaceFormer with \mathcal{L}_{rlv} and CodeTalker with \mathcal{L}_{rlv} reveals 6.0976×10^{-4} mm and 5.2639×10^{-4} mm of LVE, respectively, showing subtle improvement or even decrease of performance compared to the baselines. This suggests that imposing more regularization on the spatial structure of the lip regions, without the other components, is ineffective for accurate lip movements.

Impact of audio-visual perceptual loss. We investigate how combining speech signals with lip reading enhances multimodal perceptual loss. As shown in Table. 3, leveraging speech information for predicting the spoken transcript leads to more effective learning signals related to speech information being transmitted to the 3D facial animator. Since a single lip motion can correspond to multiple spoken texts, predicting the transcript from both visual and speech information yields better quality transcripts compared to predictions solely from visual information, i.e., lip motions. Consequently, with improved transcript prediction, the 3D facial animator is trained to generate more intelligible lip movements. We visualize the features of the lip reading expert using t-SNE [28] in Figure. 3, illustrating cases of utilizing audio-visual information versus visual-only information. Visual-only input entangles the features, hindering distinction of "just" (red) and "must" (blue). The audio-visual lip reading expert is shown to be able to guide lip movement.

4. Conclusion

In this paper, we introduce a method to guide the speech-driven 3D facial animation in comprehending lip movements corresponding to spoken words, thereby enhancing the realism of lip shapes. Our method proposes an audio-visual perceptual loss, which aids the speech-driven 3D facial animator in acquiring additional speech-related knowledge to produce plausible lip movements. We develop an audio-visual lip reading expert using a two-stage training approach: initial integration of prior knowledge from extensive 2D talking face datasets, followed by fine-tuning on 3D datasets. Extensive experiments show the effectiveness of our method in improving both lip synchronization and the intelligibility of generated lip motion, crucial aspects for human visual understanding.

5. Acknowledgements

This research was supported by a grant from KRAFTON AI. This work was partially supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT) (No.RS-2023-00225630, Development of Artificial Intelligence for Text-based 3D Movie Generation; No.RS-2022-II220290, Visual Intelligence for Space-Time Understanding and Generation based on Multi-layered Visual Common Sense; No.RS-2019-II191906, Artificial Intelligence Graduate School Program (POSTECH))

6. References

- P. Ma, A. Haliassos, A. Fernandez-Lopez, H. Chen, S. Petridis, and M. Pantic, "Auto-avsr: Audio-visual speech recognition with automatic labels," in *IEEE International Conference on Acoustics*, *Speech, and Signal Processing (ICASSP)*, 2023.
- [2] J. Son Chung, A. Senior, O. Vinyals, and A. Zisserman, "Lip reading sentences in the wild," in *IEEE Conference on Computer Vi*sion and Pattern Recognition (CVPR), 2017.
- [3] T. Afouras, J. S. Chung, and A. Zisserman, "Lrs3-ted: a large-scale dataset for visual speech recognition," *arXiv preprint* arXiv:1809.00496, 2018.
- [4] A. Ephrat, I. Mosseri, O. Lang, T. Dekel, K. Wilson, A. Hassidim, W. T. Freeman, and M. Rubinstein, "Looking to listen at the cocktail party: A speaker-independent audio-visual model for speech separation," arXiv preprint arXiv:1804.03619, 2018.
- [5] J. S. Chung, A. Nagrani, and A. Zisserman, "Voxceleb2: Deep speaker recognition," arXiv preprint arXiv:1806.05622, 2018.
- [6] G. Fanelli, J. Gall, H. Romsdorfer, T. Weise, and L. Van Gool, "A 3-d audio-visual corpus of affective communication," *IEEE Transactions on Multimedia*, 2010.
- [7] D. Cudeiro, T. Bolkart, C. Laidlaw, A. Ranjan, and M. J. Black, "Capture, learning, and synthesis of 3d speaking styles," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [8] P. Edwards, C. Landreth, M. Popławski, R. Malinowski, S. Watling, E. Fiume, and K. Singh, "Jali-driven expressive facial animation and multilingual speech in cyberpunk 2077," in *Special Interest Group on Computer Graphics and Interactive Techniques Conference Talks*, 2020.
- [9] S. Taylor, T. Kim, Y. Yue, M. Mahler, J. Krahe, A. G. Rodriguez, J. Hodgins, and I. Matthews, "A deep learning approach for generalized speech animation," *ACM Transactions on Graphics (SIG-GRAPH)*, 2017.
- [10] I. Wohlgenannt, A. Simons, and S. Stieglitz, "Virtual reality," Business & Information Systems Engineering, 2020.
- [11] E. A. Boyle, E. W. MacArthur, T. M. Connolly, T. Hainey, M. Manea, A. Kärki, and P. Van Rosmalen, "A narrative literature review of games, animations and simulations to teach research methods and statistics," *Computers & Education*, 2014.
- [12] A. Richard, M. Zollhöfer, Y. Wen, F. De la Torre, and Y. Sheikh, "Meshtalk: 3d face animation from speech using cross-modality disentanglement," in *IEEE International Conference on Computer Vision (ICCV)*, 2021.
- [13] Y. Fan, Z. Lin, J. Saito, W. Wang, and T. Komura, "Faceformer: Speech-driven 3d facial animation with transformers," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [14] J. Xing, M. Xia, Y. Zhang, X. Cun, J. Wang, and T.-T. Wong, "Codetalker: Speech-driven 3d facial animation with discrete motion prior," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023.

- [15] Z. Peng, Y. Luo, Y. Shi, H. Xu, X. Zhu, H. Liu, J. He, and Z. Fan, "Selftalk: A self-supervised commutative training diagram to comprehend 3d talking faces," in ACM International Conference on Multimedia (MM), 2023.
- [16] R. Daněček, K. Chhatre, S. Tripathi, Y. Wen, M. Black, and T. Bolkart, "Emotional speech-driven animation with contentemotion disentanglement," in ACM Transactions on Graphics (SIGGRAPH Asia), 2023.
- [17] T. Karras, T. Aila, S. Laine, A. Herva, and J. Lehtinen, "Audiodriven facial animation by joint end-to-end learning of pose and emotion," ACM Transactions on Graphics (SIGGRAPH), 2017.
- [18] K. Sung-Bin, L. Hyun, D. H. Hong, S. Nam, J. Ju, and T.-H. Oh, "Laughtalk: Expressive 3d talking head generation with laughter," in *IEEE Winter Conf. on Applications of Computer Vision* (WACV), 2024, pp. 6404–6413.
- [19] J. Wang, X. Qian, M. Zhang, R. T. Tan, and H. Li, "Seeing what you said: Talking face generation guided by a lip reading expert," in *IEEE Conference on Computer Vision and Pattern Recognition* (CVPR), 2023, pp. 14 653–14 662.
- [20] P. P. Filntisis, G. Retsinas, F. Paraperas-Papantoniou, A. Katsamanis, A. Roussos, and P. Maragos, "Spectre: Visual speechinformed perceptual 3d facial expression reconstruction from videos," in *IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2023.
- [21] P. Ma, S. Petridis, and M. Pantic, "Visual speech recognition for multiple languages in the wild," *Nature Machine Intelligence*, 2022.
- [22] S. Watanabe, T. Hori, S. Kim, J. R. Hershey, and T. Hayashi, "Hybrid ctc/attention architecture for end-to-end speech recognition," *IEEE Journal of Selected Topics in Signal Processing*, 2017.
- [23] N. Akhtar and A. Mian, "Threat of adversarial attacks on deep learning in computer vision: A survey," *Ieee Access*, 2018.
- [24] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," Advances in Neural Information Processing Systems (NeurIPS), 2014.
- [25] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," *Advances in neural information processing systems*, vol. 33, pp. 12449–12460, 2020.
- [26] T. Li, T. Bolkart, M. J. Black, H. Li, and J. Romero, "Learning a model of facial shape and expression from 4D scans," ACM Transactions on Graphics (SIGGRAPH Asia), 2017.
- [27] Amazon Polly, "Developer guide," 2015, https://docs.aws. amazon.com/polly/latest/dg/viseme.html.
- [28] L. Van der Maaten and G. Hinton, "Visualizing data using t-sne." Journal of Machine Learning Research (JMLR), 2008.