

Emotional Listener Portrait: Realistic Listener Motion Simulation in Conversation

Luchuan Song¹ Guojun Yin² Zhenchao Jin³ Xiaoyi Dong⁴ Chenliang Xu¹

¹University of Rochester ²University of Science and Technology of China

³University of Hong Kong ⁴Shanghai AI Laboratory

{lsong11@ur., chenliang.xu@}rochester.edu, gjyin@mail.ustc.edu.cn,

blwx96@connect.hku.hk, dongxiaoyi@pjlab.org.cn

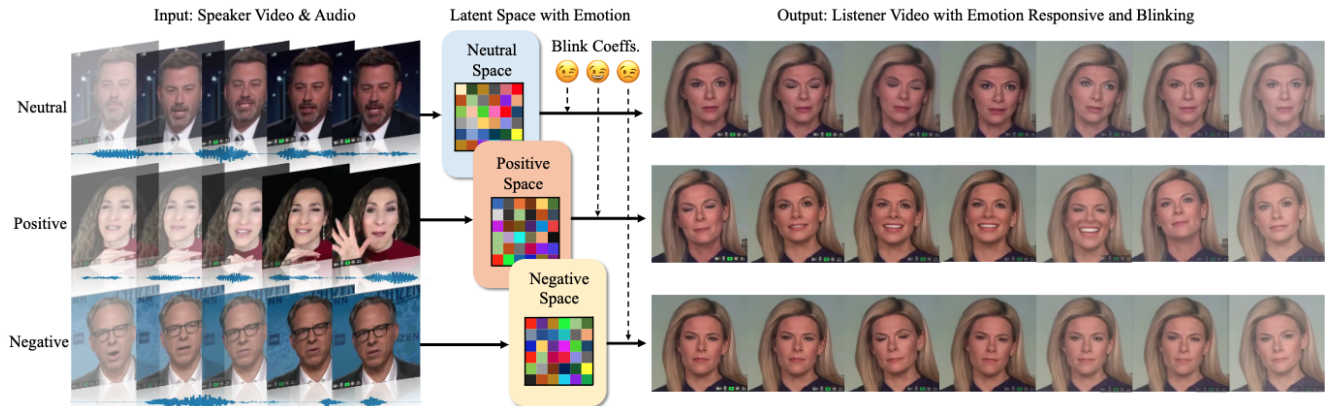


Figure 1: Illustration of our method on listener motion synthesis with the ternary emotional value as an example. Given different emotional speakers (e.g. neutral, positive, and negative), our approach generates corresponding listeners under the different emotional latent spaces.

Abstract

Listener head generation centers on generating non-verbal behaviors (e.g., smile) of a listener in reference to the information delivered by a speaker. A significant challenge when generating such responses is the non-deterministic nature of fine-grained facial expressions during a conversation, which varies depending on the emotions and attitudes of both the speaker and the listener. To tackle this problem, we propose the Emotional Listener Portrait (ELP), which treats each fine-grained facial motion as a composition of several discrete motion-codewords and explicitly models the probability distribution of the motions under different emotion in conversation. Benefiting from the “explicit” and “discrete” design, our ELP model can not only automatically generate natural and diverse responses toward a given speaker via sampling from the learned distribution but also generate controllable responses with a predetermined attitude. Under several quantitative metrics, our ELP exhibits significant improvements compared to previous methods.

1. Introduction

Listener Head Generation (LHG) technology aims to synthesize the motion of the listener in response to the speaker.

In contrast to speaker head generation (SHG) [10, 44, 21, 13, 35, 47, 9, 53, 45], which focuses on generating lip-speech synchronized portrait videos, LHG analyzes the talking semantics of the speaker automatically, without explicit guidance, to synthesize corresponding interactive motions of the listener. As shown in Figure 1, the listener reacts positively when the speaker shares happy, and vice versa. LHG can be employed in many applications, e.g. human-computer interaction [57, 25, 60], virtual reality [22, 24], metaverse [8, 7, 46] and media forensics [42, 37, 43, 17] etc.

The distinct nature of LHG, which necessitates a comprehensive modeling of the speaker’s motion [32], presents a significant hurdle in yielding realistic listener head. In the absence of audio-to-mouth matching evaluation, audiences are more inclined to discern subtle changes in facial expressions and head movements. However, the existing methods Responsive Listening Head Generation (RLHG) [61] and Learning2Listen [31] have ignored these key components. Specifically, RLHG [61] has replicated the regression experience from SHG [10], which weakens the non-deterministic properties and smoothes the listener motion. Meanwhile, although the motion categories in codebook proposed by Learning2Listen [31] alleviate this problem, the one-dimensional codebook from VQ-VAE [29] limits the

diversity of motion and emotional representation. Consequently, different emotional states are intricately intertwined within a single codebook, and the generated listener emotion tends to be the biased emotion in the training set. Moreover, neither method can simulate fine-grained facial motion under different emotions, such as minute alterations in the motion surrounding the eyes and movements of the mouth. To explore the solution toward superior LHG results, we focus on the two unresolved hurdles: (1) how to simulate finer-grained listener movements, including the head motion and expression details of the face, and (2) how to explicitly model emotions in the discrete space.

In this paper, we propose a novel method called **Emotional Listener Portrait (ELP)** for vivid listener head video generation. The visual and audio information from the speaker is combined together for the listener motion synthesis (the right part of Figure 1). (1) To overcome the limitation imposed by one codeword search, we have expanded the classification dimensions to facilitate the mapping of listener motion onto a higher-dimensional discrete space. The fine-grained listener movements correspond in high-dimensional discrete space, which offers greater capacity for the precise depiction of the listener’s facial expression and head pose than a single codeword. (2) Despite the expanded latent space on the codeword, explicit emotional representation remains unattainable. As such, building upon the increased space, we leverage emotion priors to split and rearrange the discrete space. More specifically, different emotions are rearranged into corresponding spaces, with the distance between these spaces being determined by the value range of codeword, shown in the middle of Figure 1 (it takes the ternary emotion as an example). The listener features (blink, facial and head motion) from the spaces with emotion are decoded into different emotional listeners, shown in the right part of Figure 1. It indicates that the distance between listeners with different emotions is widened, such as listeners in positive response by smiling, while listeners in negative tend to frown.

There are two modules in the **ELP**, the Adaptive Space Encoder (**ASE**) and Mesh-to-Video Renderer. In **ASE**, the discrete latent space obtained by the one-hot vectors argument maximum is concatenated and weighted according to the position with the prior emotion. Employing this approach leads to a further enlarge in the probability distribution distance, and the listener motion coefficients are learned from that. The Mesh-to-Video Renderer renders the photorealistic face from the mesh corresponding to the predicted coefficients with only a single portrait image of the listener.

We demonstrate the ability of our method through quantitative and qualitative experiments on two popular conversation portraits datasets, the ViCo [61] and the large-scale in the wild conversation videos collected by the Learn2Listen [31]. Our contributions are summarized:

- We propose a novel framework called **ELP** for emo-

tional listener head generation in dynamic conversation, which can improve the fidelity of the fine-grained generated listener with facial expression, head pose and blink *etc.*

- We introduce the Adaptive Space Encoder (**ASE**) to rearrange the latent space based on emotional priors to obtain more explicit emotional representation.
- Extensive experiments demonstrate that our method outperforms most existing methods in quantitative and qualitative results.

2. Related Work

Blink-Controlled Facial Animation. There are some existing methods [58, 41] for blink control on facial animation. FACIAL [58] first employs an eye attention map to locate the eye region, then uses constraints between the rendered pixels and the synthetic result to realize the realistic eye blinks. Sinha *et al.* [41] adopt the random one-hot encoding for blink control. They use blink one-hot encoding of input into landmark generator. The landmark then guides the generation of the face images. However, these are specifically designed for SHG. Regarding LHG, the methods used to predict blinks through speech are unsuitable as they rely on regression techniques that ultimately lead to a low blink frequency. That can lead to artifacts if the listener keeps their blink frequency extremely strange during a conversation.

Audio-Visual Emotion Representation. There are numerous approaches for extracting emotions from speaker videos. These include Automatic Speech Emotion Recognition [29], Advanced LSTM [48], and Cross-Reconstructed Emotion Disentanglement [21], which aim to decode emotions from speech alone. However, relying solely on speech feature extraction may overlook the visual modality, resulting in insufficient accuracy of the extracted features [33]. Schonveld *et al.* [38] employed a joint model of facial and speech features to extract more representative emotional features.

Face motion simulation via speech. It include two parts of work: speaker head generation (SHG) and listener head generation (LHG). Over the past few years, SHG has emerged as a burgeoning field. Chung *et al.* [11] and Shi *et al.* [39] extract speech features to simulate the corresponding face or mouth movement. Wu *et al.* [56] imitates arbitrary talking style with speaker speech and head motion features, this method of feature design is also introduced into our work. The LHG is different from the SHG, which pays more attention to the feedback of the listener motion to the speaker. Ahuja *et al.* [3] focuses on the non-verbal behaviours generation in human body, Bohus *et al.* [4] and Gratch *et al.* [15] study synchronized conversations agent motion in dyadic conversation adapt speech. More recently, Ng *et al.* [31] regress the discrete listener head motion with VQVAE [51] and Geng *et al.* [14] retrieve possible videos of listener face with large language model. In our work, we also learn the

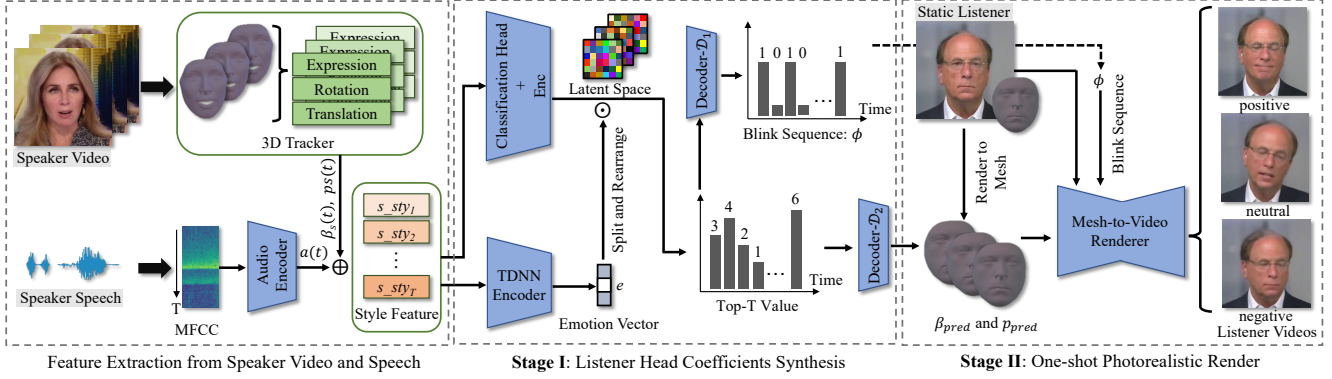


Figure 2: The overview of ELP. We take the static listener face, speaker video and the corresponding speaker speech as input. In the **stage I**, the Adaptive Space Encoder maps the discretized features combined with emotion to the listener motion parameters, then Mesh-to-Video Renderer generates the emotional listener videos from the obtained parameters in the **stage II**.

non-verbal listener motion in the dynamic communication.

3. Method

3.1. Overview

We propose a two-stage listener portrait synthesis framework (ELP) which synthesizes emotional listener videos with three inputs: one static portrait of the listener (Static Listener), speaker video (Speaker Video) and the corresponding speech (Speaker Speech). The overview pipeline is shown in Figure 2. We formalize the two stages as follows.

Stage I: Listener Coefficients Synthesis. In this stage, given the speaker video and corresponding speech, the facial and head pose motion ($\beta_s(t), p_s(t)$) are tracked from the video. It aims to generate the response listener facial head pose movement and blink coefficient sequence from the speech and $\beta_s(t), p_s(t)$. The ground truth facial and head parameters are captured from the monocular video with Gauss-Newton optimization [50], they are resulting in $\beta \in \mathbb{R}^{100T}$ (facial expression), $p \in \mathbb{R}^{6T}$ (head rotation). And the blink coefficients sequence is a binary list on time series ($\phi \in \mathbb{R}^{1T}$). T is the length of video.

Stage II: One-shot Photorealistic Render. In this second stage, our objective is to generate dynamic and photorealistic videos of the listener by utilizing the predicted results from the first stage and a single static portrait image of the listener.

To streamline the problem, we will introduce parameter definitions within our pipeline in Section 3.2, and go through the details of the two stages in Section 3.3 and Section 3.4, respectively.

3.2. Parameter Definitions

To begin this section, we define the initialization input cross-modal data: the MFCC feature [26] from speech and the facial/head motion parameters $\beta_s(t), p_s(t)$.

Speaker Style Features. Given the interplay between the

speaker and listener, it is necessary to study the stylistic features inherent in the speaker. We expect to employ these features in order to retrieve the emotional underpinnings of the conversation, as well as the listener’s response movements. Therefore, we define a computational rule, which extracts speaker style-related features:

$$s_sty(t) = a(t) \oplus \sigma(\beta(t)) \oplus \sigma\left(\frac{\partial\beta(t)}{\partial t}\right) \oplus \sigma\left(\frac{\partial p(t)}{\partial t}\right), \quad (1)$$

the $a(t)$ is the feature of MFCC via Audio Encoder, the \oplus is the concatenation, the $\sigma(\cdot)$ is the standard deviation. The $\beta(t), \sigma\left(\frac{\partial\beta(t)}{\partial t}\right), \sigma\left(\frac{\partial p(t)}{\partial t}\right)$ represent the fluctuation of the speaker’s movements relative to time, which can reflect the speaker’s visual movement style [56]. For speech features, some methods [12, 49] attempt to extract semantic features from pretrained speech-to-text models such as DeepSpeech [16]. However, it is not applicable to our method, as speech emotion and intonation within is important for listener, but are not included in purely semantic feature.

Emotion Vector. The emotion vector e is a one-hot embedding, which is used to determine which registered emotion codebooks are selected.

Blink Coefficients. We leverage the geometric ratio of eye landmarks to represent eyelid movement. When the ratio [6] exceeds a predefined threshold value, the current frame is classified as exhibiting blinking behavior, which is represented by the binary value 1, otherwise 0.

3.3. Listener Head Coefficients Synthesis.

In this stage, we present the Adaptive Space Encoder (ASE) for generating emotional listener coefficients. Specifically, the ASE takes the $s_sty_{1:T}$ as input. We consider s_sty to be a cross-modal feature from the speaker, containing various information required by the listener, including emotional value, utterance semantics, and response guidance. The latent speech feature $a(t)$ are obtained from the

backbone of ResNet-50 and Dropout from the input MFCC feature, shown in the right part of Figure 2.

The ASE is shown in the middle of Figure 2 and Figure 3(a), which reveals that ASE is composed of two encoders and decoders. We assume that in a clip of input speaker video (2 seconds), the emotion is constant. For the emotion classification branch, the TDNN [34] Encoder takes a series of $s_sty_{1:T} \in \mathbb{R}^{25T \times D_s}$ (T is length of video with 25 framerate, D_s is the dimension of each $s_sty(t)$) as input and encodes them into the predicted emotion vector e_{pred} . For the motion prediction branch, the ASE model leverages a multi-classification head mechanism to encode the continuous-valued stylized features to a discrete latent space, which is a classification probability distribution and then the predicted motions are sampled from this distribution. Through this mechanism, we obtain a motion space with discrete manifold. We achieve discretization of this categorical representation with gumbel-softmax [20, 35], which can be formalized by:

$$v_{t,h;1} = [\text{Gumbel-Softmax}(\mathbf{enc}(s_sty)_{t,h,1:V})]_{1:H;1}, \quad (2)$$

where the $\mathbf{enc}(s_sty)_{t,h,1:V}$ means the multi-modality feature s_sty encoded to $T \times H \times V$ -dimensional latent, which is still in continuous-value space. Then, the operation of Gumbel-Softmax [20] takes the $\mathbf{enc}(s_sty)_{t,h,1:V}$ into the maximum probability in V dimension, which represents the corresponding discrete motion code in the space be mapped by Equation 2, shown in Figure 3(b). The $v_{t,h;1}$ means each codeword value in the discrete space, the H is the number of latent classification heads and V is the number of categories. For the sake of simplicity, we refer to this space as the *Base Space* throughout the remainder of this paper, the *Base Space* is composed of one-hot vectors. However, emotions are still implicit in the *Base Space*. To capture these emotional responses explicitly, it is necessary to distinguish different emotion spaces.

Specifically, we split and rearrange the *Base Space* based on the emotion prior. The emotion vector e perform dot production with *Base Space, and the results are concatenated. The $e \in \mathbb{R}^{1 \times N}$ is a one-hot vector, N is the number of emotion types. Then the *Base Space* is expanded to $T \times H \times NV$ from $T \times H \times V$. We call the expanded result the *Transformed Space*. Finally, by computing the argument maximum on *Transformed Space*, we obtain the final discrete latent space \mathbb{U} . The codeword value v' in the \mathbb{U} is in the range of $\{v'_{1:T,1:H} | v'_{i,j} \in [1, 2, \dots, NV]\}$. It should be noted that the values within different emotional intervals are unique and non-overlapping. As an example, let us consider the two different emotion values for e . The \mathbb{U} corresponding to the first emotions will only take values in the range of $[1, \dots, V]$, and the other in $[V + 1, \dots, 2V]$. And it can be found that our method obtains \mathbb{U} space without additional dimensional overhead after splitting and rearrangement. Then,*

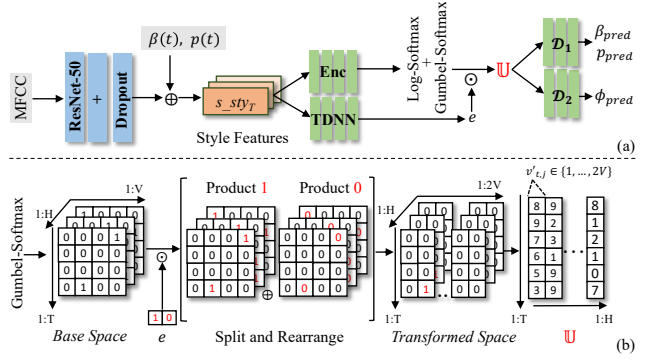


Figure 3: The structure and details of Adaptive Space Encoder. (a) The Adaptive Space Encoder takes speaker speech MFCC [26] feature, speaker facial ($\beta(t)$) and head motion ($p(t)$) coefficients as input, and outputs the listener’s facial and head motion. (b) The details of split and rearrangement for latent space \mathbb{U} ($N=2$ as an example), the *Base Space* is weighted and concatenated (\oplus) by the element values of e .

two different decoders \mathcal{D}_1 , \mathcal{D}_2 act on the \mathbb{U} to obtain the facial/head movements coefficients and blink sequences.

For the training of ASE model, we adopt the parameter series $\beta'(t)$, $p'(t)$ reconstructed from the monocular listener video as ground truth, and calculate $s_sty(t)$ frame by frame. Afterwards, we feed $s_sty(t)$ to the ASE model, yielding the predicted listener parameter $\beta_{pred}(t)$, $p_{pred}(t)$ from \mathcal{D}_1 and blink sequence $\phi_{pred}(t)$ from \mathcal{D}_2 . Based on the predicted $\beta_{pred}(t)$ and $p_{pred}(t)$, we apply the L_2 loss as follows:

$$\mathcal{L}_{L_2} = \sum_{t=1}^T \|\beta_{pred}(t) - \beta'(t)\|_2 + \|p_{pred}(t) - p'(t)\|_2. \quad (3)$$

We take the blink sequence as a decision at each time instance t , and employ the cross-entropy loss for binary classification, the $\phi(t)$ is the ground truth blink state:

$$\mathcal{L}_{CE_1} = - \sum_{t=1}^T \phi(t) \log \phi_{pred}(t) + [1 - \phi(t)] \log [1 - \phi_{pred}(t)], \quad (4)$$

we introduce regularization loss to suppress noises and encourage a more concentrated density distribution (a blinking action comprises several consecutive 1 value):

$$\mathcal{L}_{reg} = \sum_{t=2}^T \|\phi_{pred}(t) - \phi_{pred}(t-1)\|_1. \quad (5)$$

Additionally, for emotional constraints, TDNN encodes s_sty into vector e' , we also apply cross entropy on it, the e is the groundtruth one-hot emotion vector:

$$\mathcal{L}_{CE_2} = - \sum_{i=1}^N e_i \log e'_i + [1 - e_i] \log [1 - e'_i], \quad (6)$$

which is used for N classification of emotions, N is the type of emotion. The final loss function is then defined as:

$$\mathcal{L} = \mathcal{L}_{L_2} + \lambda_1 \mathcal{L}_{CE_1} + \lambda_2 \mathcal{L}_{CE_2} + \lambda_3 \mathcal{L}_{reg}, \quad (7)$$

the λ_1, λ_2 and λ_3 are three weight to balance these terms.

3.4. One-shot Photorealistic Render

To improve the visual generalization in our method, we also train a one-shot Mesh-to-Video translation network. Usually, it is very difficult to collect long-term person-specific portrait videos to train high-quality person-specific video, so we tend to use one-shot rendering method, which is resource free. Inspired by Siarohin *et al.* [40], we first train the motion capture module [40] on face mesh videos, and then use the face mesh videos to drive the motion of one-shot portrait image to synthesis the photorealistic video, we call it as Mesh-to-Video Renderer, shown in the right part in Figure 2. To control the blinking, we obtain the length L of each eye blink motion group on ϕ (each consecutive 1 is counted as one group), and the expression blendshape β_1, β_2 at the beginning and end of the motion group (the eye closure blendshape is marked as β_c). We interpolate $\beta_1 \rightarrow \beta_c$ and $\beta_c \rightarrow \beta_2$ respectively on the interval of $L/2$. The blendshape for eye blink is linearly weighted, simulating the eyelid position at each timestamp for physical blink or emotional events around eyes.

4. Experiments

4.1. Experimental Settings

Datasets. We evaluate our method on two of the most popular conversation portrait datasets, the ViCo [61] dataset and the dataset proposed by Learning2Listen [31]. The ViCo dataset contains rich samples of 483 video clips with 50 different listener identity and three emotion annotations (we set $N = 3$ for emotion for this dataset). The Learning2Listen [31] is a 72 hours versus 95 minutes dataset collected in the wild, which comes from Youtube with six identities but no emotion annotations. We adopt the pretrained speech emotion analyzer model [2] to extract emotions. We set the [“happy”] to positive and [“calm”, “fearful”, “sad”, “angry”] to unpositive (we set $N = 2$). To maintain the balance in this dataset, we assign more labels to unpositive emotions (the positive videos more than other emotions in this dataset). The one-shot photorealistic renderer is trained on the TalkingHead-1KH [55] datasets with 256×256 resolution.

Network Architectures. We apply the first 16 layers of ResNet-50 as the backbone to encode the input MFCC from 29T-dimension to 128T-dimension. The TDNN Encoder is composed of five hidden layer TDNN and three-layer MLP, which is used for the classification of N -dimension emotional one-hot vector from the 334T-dimension (128T

for $a(t)$, 100T for $\sigma(\beta(t))$ and $\sigma(\frac{\partial \beta(t)}{\partial t})$, 6T for $\sigma(\frac{\partial p(t)}{\partial t})$) speaker style feature. The Classification Head Encoder consists of 8 layers conv1d and 3 layers of LSTM with MLP, it takes the style feature as input, the dimension of output is $T \times H \times 3C$. The decoder \mathcal{D}_1 has two conv2d layers and the decoder \mathcal{D}_2 has four layers conv2d with three layers LSTM, their role is to recover the motion from latent space. The structure of one-shot photorealistic renderer (Mesh-to-Video Render) is from the First Order Motion Model [40].

Implementation details. For the convenience of training, we randomly clip the input to 2 second video (50 frames), and clip the output corresponding to 50 frames. The λ_1, λ_2 and λ_3 are 5, 5, 0.01. We set V and H to 64, 128 for the latent space. When optimizing, we adopt the AdamW optimizer [27] to train the ASE model with the initial learning rate of 1×10^{-3} . We train 10000 iterations with a batch size of 32 samples for ViCo [61] dataset and 50000 iterations with 128 batch size for Learning2Listen dataset [31].

Baselines. We compare with the state-of-the-art methods and the hand-craft methods. For the listener’s head generation, the hand-craft baseline is usually strong enough.

- Learning2Listen (L2L) [31]: It maps the motion patterns to realistic movements through the VQ-VAE [51] without the consideration of emotions and blinks.
- Responsive Listening Head Generation (RLHG) [61]: It regresses the listener motion from the speaker and audio. We use the official code for fairness.
- Perceptual Conversational Head (PCH) [19]: It wins the 1st in the ViCo challenge. Since only part of the code has been released, we reproduce it based on RLHG [61].
- NN-motion/NN-audio: For arbitrary input speaker motion or audio input, we find its nearest neighbor from the training set and use its corresponding listener motion as output. We follow it from the L2L [31].
- Random: We random select the facial and head motion parameters in the training data, and inject random small perturbations into the normal distribution.
- Discrete-Latent-Space Random (DLS-Random): Randomly generate $v'_{t;1:H}$ value, and use the pretrained $\mathcal{D}_1, \mathcal{D}_2$ to generate motion sequences.

Metrics. We choose the following metrics to evaluate the generated facial/head motion coefficients.

- Frechet Distance (FD) [18]: The L_1 distance to measure the difference between the generated facial/head motion and the ground truth.
- Variation for Diversity (V-D): Proposed by the L2L [31], the variance of the facial/head motion on time series.
- SI for Diversity (SI-D) [59]: From the L2L [31], it measures diverseness of predictions with k-means to facial/head motion, we report the average entropy (Shannon index).

Methods	Facial <i>Coeff.</i> (β)						Head <i>Coeff.</i> (p)						Blink (ϕ)
	FD \downarrow ($\times 10^3$)	V-D	SI-D	RPCC \downarrow	WTLCC (\rightarrow GT)	STS \downarrow ($\times 10^3$)	FD \downarrow ($\times 10^3$)	V-D	SI-D	RPCC \downarrow	WTLCC (\rightarrow GT)	STS \downarrow ($\times 10^3$)	WTLCC \uparrow
NN-motion	12.17	2.88	2.70	0.16	0.092	8.89	6.65	1.79	1.16	0.20	0.083	5.59	0.12
NN-audio	19.05	3.70	2.99	0.18	0.077	11.36	14.05	2.51	2.54	0.15	0.123	7.10	0.13
Random	56.21	4.57	3.92	0.28	0.047	40.26	35.10	2.54	2.29	0.15	0.043	25.08	0.11
DLS-Random	38.10	4.01	3.47	0.21	0.075	22.51	26.17	3.99	1.33	0.12	0.105	15.50	0.11
RLHG [61]	150.0*	3.72	0.41	0.20	0.011	29.00	143.1*	0.28	0.24	0.16	0.010	25.64	0.04
PCH** [19]	18.10	1.14	0.37	0.18	0.011	27.22	20.35	0.24	0.22	0.10	0.003	17.44	$1e^{-2}$
L2L [31]	4.20	3.02	2.79	0.11	0.094	5.44	1.93	0.97	2.45	0.04	0.017	3.38	$1e^{-2}$
GT	–	4.89	4.27	–	0.177	–	–	1.90	1.96	–	0.192	–	–
ELP	1.14	5.26	4.00	0.02	0.166	2.93	0.60	2.04	1.55	0.01	0.199	1.34	0.39
NN-motion	15.80	2.14	2.37	0.12	0.155	9.33	7.97	1.55	1.92	0.13	0.061	7.94	0.12
NN-audio	22.03	1.87	1.61	0.09	0.114	8.90	13.37	1.39	1.20	0.11	0.077	7.15	0.09
Random	47.02	2.26	2.51	0.58	0.044	20.41	24.70	1.84	1.91	0.19	0.075	13.80	0.11
DLS-Random	33.61	2.23	2.14	0.12	0.071	8.04	19.97	1.37	1.54	0.09	0.088	6.29	0.13
RLHG [61]	21.28	0.59	1.06	0.39	0.070	14.11	18.54	0.14	1.10	0.292	0.026	12.31	0.02
L2L [31]	3.55*	2.01*	2.48*	0.02*	0.130	7.39	0.81*	0.62*	1.82*	0.00*	0.004	6.01	$1e^{-2}$
GT	–	2.47	2.20	–	0.202	–	–	0.55	1.54	–	0.127	–	–
ELP	1.37	2.70	2.15	0.014	0.182	4.49	0.36	0.59	1.60	0.077	0.130	2.51	0.42

Table 1: Quantitative results on two different datasets. The above is the performance evaluation on the ViCo dataset [61], whereas the bottom is on the L2L dataset [31]. The \downarrow indicates lower is better, and \rightarrow GT means closer to GT is better. The * indicates that we directly follow the office report results and ** means we reproduce the PCH [19] on our system (no source code is provided). The best performances are highlighted in bold.

- Residual Pearson Correlation Coefficient (RPCC): The Pearson Correlation Coefficient (PCC) [28, 36] measure facial/head motion frame-by-frame, it is for listener covaries with the input speakers. We calculate the L_1 distance between generated PCC and PCC of ground truth.
- Windowed Time Lagged Cross Correlation (WTLCC) [5]: It is the correlation between the generated motion and the input speaker’s motion in the set time window. We calculate the L_1 distance between generated TLCC and ground truth with the window size as 4 seconds.
- Short Time Series distance (STS) [30]: The STS can measure the similarity on sampled data. We calculate the STS between generated listener motion and ground truth listener motion.

We report head and facial metrics separately. For blink sequence evaluation, we adopt the WTLCC, which has a good performance on handling the noise by the time shifting, especially for the binary sequences. It is worth noting that for the blink WTLCC we compute the correlation between the generated blink sequence and the real blink sequence. It is different from the WTLCC calculated on the facial and head motion. Please refer to the Appendix for the details about each metric.

4.2. Comparison Results

Quantitative Results. We retrain PCH [19] and L2L [31] on the ViCo datasets [61], retrain PCH [19] and RLHG [61] on L2L [31] datasets. Table 1 shows the quantitative comparison of listener head generation in the first stage. Based on

the metrics presented in Table 1, it is clear that our proposed **ELP** method outperforms other existing methods by a significant margin. Our method is about 400 times better than the current state-of-the-art method L2L [31] in terms of facial blink indicators (WTLCC in Blink ϕ), the L2L [31] struggles to synthesize realistic eye movements, our proposed **ELP** is capable of generating blinking movements that accurately reflect the conversation. In the realm of measuring motion diversity (V-D, SI-D), L2L [31] rely on the distance from ground truth as a performance metric, but we opt for a more comprehensive perspective. Our expectation is that listeners’ movements will showcase a greater range of diversity when the FD and STS metric (feature distance) are kept to a minimum. As demonstrated in Table 1, it is evident that even when FD and STS are kept to a minimum, our **ELP** approach remains highly effective in both the V-D and SI-D, which benefits from our latent space split and rearrange method. The RPCC and WTLCC (in Facial and Head *Coeff.*) measures the motion synchronization between the speaker and the listener, and it can be found that our method achieves an effect close to GT (the RPCC of GT is 0).

The weaker performance of regression-based methods [61, 19] demonstrates the regression methods are not suitable for the listener head synthesis, due to the inherent randomness involved in listener. This fact is made especially clear when evaluating key indicators such as FD, STS, V-D and SI-D *etc.* Despite L2L’s [31] impressive performance, it still falls behind our proposed method. This is primarily attributed to the fact that the codebook space of VQ-VAE is equivalent to an encoding head with $H=1$, while our approach not only considers emotional rearrangement but also

Methods	Coeff. ($\beta + p$)						Blink (ϕ)
	FD↓ ($\times 10^3$)	V-D	SI-D	RPCC↓	WTLCC (\rightarrow GT)	STS↓ ($\times 10^3$)	WTLCC↑
NN-motion+ ∇	18.30	3.52	3.77	0.59	0.104	19.91	0.12
NN-motion+ \cup	<u>12.82</u>	4.67	3.86	<u>0.36</u>	<u>0.175</u>	14.48	0.12
NN-audio+ ∇	34.49	4.13	4.24	0.51	0.128	16.67	0.05
NN-audio+ \cup	33.10	6.21	5.53	<u>0.33</u>	0.200	18.46	<u>0.09</u>
DLS-Rand.+ ∇	68.02	6.94	5.99	<u>0.31</u>	0.120	39.07	<u>0.14</u>
DLS-Rand.+ \cup	<u>64.27</u>	8.00	4.80	0.33	0.180	<u>38.01</u>	0.11
GT	–	6.79	6.23	–	0.369	–	–
Our+ ∇	4.27	3.64	3.31	0.09	0.190	6.12	0.31
Our+ \cup	<u>1.74</u>	7.31	5.55	<u>0.03</u>	<u>0.366</u>	<u>4.27</u>	<u>0.39</u>
NN-motion+ ∇	26.43	3.77	3.91	0.26	0.207	<u>14.33</u>	<u>0.15</u>
NN-motion+ \cup	23.77	3.69	4.29	<u>0.25</u>	0.216	18.27	0.12
NN-audio+ ∇	49.18	5.12	4.77	0.34	0.292	27.43	0.05
NN-audio+ \cup	<u>35.40</u>	3.26	2.81	<u>0.20</u>	0.191	<u>16.05</u>	<u>0.09</u>
DLS-Rand.+ ∇	62.38	7.83	6.06	<u>0.35</u>	0.107	27.04	<u>0.07</u>
DLS-Rand.+ \cup	<u>53.58</u>	3.60	3.68	<u>0.21</u>	0.159	<u>14.33</u>	<u>0.11</u>
GT	–	3.02	3.74	–	0.329	–	–
Our+ ∇	5.12	2.11	1.79	0.122	0.297	8.38	0.36
Our+ \cup	<u>1.73</u>	3.29	3.75	<u>0.091</u>	<u>0.312</u>	<u>7.02</u>	<u>0.42</u>

Table 2: Ablation study results for latent space \cup (with emotional representation) and ∇ (without emotional representation) on two different datasets. The above is on the ViCo dataset [61], the bottom is on the L2L dataset [31]. The ‘‘DLS-Rand.’’ is the baseline method Discrete-Latent-Space Random. The best results in each group are underlined.

utilizes a larger discrete coding space ($H=128$). The NN-motion/audio and Random/DLS-Random are some hard-to-beat baselines in terms of variance (V-D, SI-D), but they fall short in simulating motion correlation. In other words, their diversity is primarily based on random generation. As mentioned above, our method outperforms the baseline method in both diversity and motion correlation simulation.

Qualitative Results. We first compare with the photorealistic results of PCH [19] and RLHG [61] in Figure 4, for fairness, we manually align and compare the results displayed on the office website [1]. The listener head synthesized by RLHG [61] and PCH [19] exhibit two significant shortcomings compared to ours, 1) the facial movements are limited and lack variation, 2) the videos generated do not incorporate any background information, the **ELP** overcomes these issues. The visualization results compared with L2L [31] are shown in Figure 5, we capture the speaker videos and generate the photorealistic results with Vid2Vid [54], which is a person-specific trained renderer. Since there is no guidance and source code on the settings of renderer in the L2L report [31], we follow the guidance of Vid2Vid [54] office code for training and generation. We also compare the facial expression and head pose details through mesh visualization in Figure 5, which excludes the background and facial texture interference. It can be found that the results synthesized by our method have more diversity than L2L [31] in terms of parameters (face mesh) and photorealistic generated listener.




Figure 4: Comparisons with RLHG [61] and PCH [19]. Our method synthesizes more diverse facial and head movements, the blinking frames are highlighted with red arrow. The input static listener image from the same one, the speaker videos come the test set of ViCo datasets [61].




Figure 5: Comparisons with L2L [31]. The L2L [31] works on the facial and head coefficients level, we visualize the coefficients to mesh for fair comparison. The details of mesh and face are highlighted on the right part.

4.3. Ablation Studies

Analysis on the Latent Space Rearrangemet. We conduct experiments to demonstrate the implications and importance of our latent space decomposition based on emotion priors. Specifically, as a comparison, we do not calculate the discrete latent space \cup through the emotional one-hot vector, but directly obtain the discrete latent space through the argument maximum on the *Base Space*, which is called as ∇ in the following part. The baseline methods NN-motion, NN-audio and DLS-Random are also repeated in \cup and ∇ spaces, respectively. For a more intuitive understanding, we eval the facial coefficient (β) and head coefficient (p) in together. The ablation study results are shown in Table 2 (we add the metrics value corresponding to β and p for comparison),

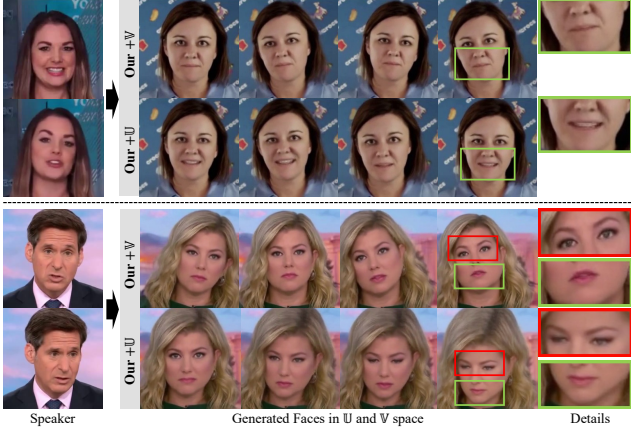


Figure 6: The visualization results of ablation study on the latent space rearrangement (\mathbb{U} and \mathbb{V}). The upper and lower parts correspond to positive and negative respectively. The details of generated face are on the right part.

	Coeff. ($\beta + p$)						Blink (ϕ)
	FD \downarrow ($\times 10^3$)	V-D	SI-D	RPCC \downarrow	WTLCC (\rightarrow GT)	STS \downarrow ($\times 10^3$)	WTLCC \uparrow
GT	-	3.02	3.74	-	0.329	-	-
L2L [31]	4.36	2.63	4.30	0.02	0.134	13.40	$1e^{-2}$
H=1	7.93	1.92	1.12	0.14	0.122	18.27	0.37
H=4	7.42	3.15	2.38	0.14	0.171	11.34	0.35
H=16	4.14	2.27	2.55	0.11	0.193	10.81	0.38
H=64	3.67	3.70	3.52	0.12	0.228	9.72	0.41
H=128	1.73	3.29	3.75	0.09	0.312	7.02	0.42
H=256	1.35	1.49	1.26	0.11	0.370	7.35	0.42

Table 3: Ablation study results on the size of latent space are revealed by gradually increasing the value of H from 1 to 256. The L2L [31] in this table is to compare with the results of H=1. We perform this experiment on the L2L [31] datasets.

we can find that \mathbb{U} space has a stable improvement in each metric compared to \mathbb{V} space, some outcomes are affected by randomness (RPCC on DLS-Rand.+ \mathbb{V} or \mathbb{U}). Then we visualize the results with \mathbb{V} space and \mathbb{U} space separately in Figure 6. From the Figure 6 above, we can observe that in the \mathbb{U} space, the positive listener exhibits a more pronounced grinning, while the listener in \mathbb{V} space only smiles slightly. On the other hand, from the below of Figure 6, the negative listener in \mathbb{U} space tends to show a more serious expression with more obvious frown. Furthermore, we adopt the t-SNE [52] to visualize the features decoded from the \mathbb{U} and \mathbb{V} spaces and take the corresponding emotion value as label. The emotional features in the \mathbb{V} space (in Figure 7(a)) appear to be poorly decomposed, with the feature corresponding to different emotions being coupled together. While the \mathbb{U} space (in Figure 7(b)) provides better feature distinction, as the value ranges of different emotions vary.




Figure 7: Visualization of the features from two different latent space. We use the ViCo dataset [61] for visualization. Three emotion are recorded in three different colors, there are 970 video clips of 2 seconds are collected.

Method	Facial Expression		Head Pose		Blink		All Aspect	
	Real	Mesh	Real	Mesh	Real	Mesh	Real	Mesh
RLHG [61]	1.4	1.2	1.5	1.2	1.6	1.0	1.5	1.1
PCH [19]	1.4	1.1	1.4	1.2	1.1	1.1	1.3	1.1
L2L [31]	3.7	2.3	3.1	2.3	1.0	1.0	3.3	2.9
GT	4.6	4.9	3.2	3.9	4.7	4.5	4.9	4.6
Our	3.9	4.2	4.7	4.9	4.3	4.5	3.7	4.1

Table 4: User study of our method with others. We calculate the average of the 5-scale scores for users. The “Real” means the photorealistic videos and “Mesh” means the mesh rendered videos. The facial and head movements can be evaluated more intuitively from the mesh videos.

Analysis on the size of Classification Head. We discuss the impact from the size of the discrete space \mathbb{U} . With the definition of \mathbb{U} , it is evident that the size of configurations is NV^H , which allows a vast motion space with a relatively small number of categories V . We set the default value of V to 64 and varied the size of H to evaluate the effect on performance. It is worth noting that when we set $H=1$, our ASE is equivalent to the VQ-VAE [29] used in L2L [31], with the codebook size of V . The results are shown in Table 3, when $H=1$, our method exhibits slightly lower performance compared to L2L [31], as L2L [31] utilizes a larger number of codewords than our default number of categories V . And the listener motion diversity and synchronicity increased (H from 1 to 128) with increasing latent space size. However, when H is excessively large (256), the performance of the model may be adversely affected, since the resulting latent space may be too large and hard to learn, ultimately leading to the difficulties in achieving an appropriate fit. Based on the aforementioned considerations, we set H to 128.

4.4. User Studies

We conduct user studies to compare the generated results from the human perspective evaluation. We collected questionnaires from 42 users through an online platform. The questionnaire includes the generated videos and ground truth under different emotion communication situations. To eliminate any facial texture interference caused by the renderer,

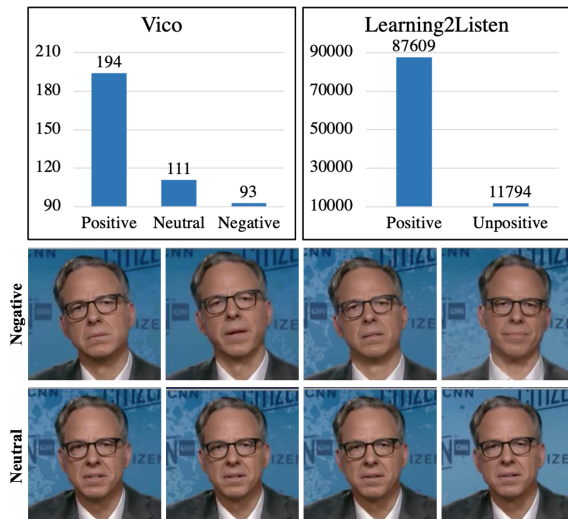


Figure 8: The visualization of unbalanced emotions and confusion annotation. The above shows the different emotion distributions in the RLHG [61] and Learning2Listen [31] datasets. For both datasets, the positive emotion is much higher than the other. We select the neutral and negative emotion from the ViCo [61] dataset¹, and visualize two different videos (each video represents one emotion). The emotions are not obvious.

we also used the corresponding mesh video to evaluate the user perception of the synthesized listeners. Each user was asked to answer a Likert-type scale with the following options for all videos [23], “Do you think the [×] in video is like a listener?” (1-disagree, 2-weakly disagree, 3-normal, 4-weakly agree, 5-agree), where the “[×]” is for four criteria evaluated: 1) facial expression, 2) head pose movements, 3) blink and 4) all aspects. We record the average score for each type. Usually, in the experiment of 5-scale score, we think that more than 4 points are close to the real performance. From the Table 4, it can be found that our method has made great progress compared to the state-of-the-art methods, especially for the evaluation of mesh video and blink criteria. However, there is still a gap between our synthesized videos and ground truth videos, which is mainly due to the limitations of the resource-free renderer.

5. Discussion and Reflections

In this work, we present a novel listener motion synthesis method that takes into account the emotion in the dynamic conversation videos. To our best knowledge, our approach is the first exploration of emotion space representation for the listener. Although our method has proven to be remarkably efficacious under the current conditions, there are still some

¹The neutral video id is @JohCZ6VIC70.000350.000358 in ViCo, the negative video id is @JohCZ6VIC70.000309.000324 in ViCo.

challenges. Two of the most tricky obstacles are (1) the emotional labels of listener in the dataset are unbalanced, and (2) the distance between emotions is unclear, that are visualized in Figure 8. In the Learning2Listen [31] dataset, the videos with positive emotion are about eight times more than the the negative ones, the source of the videos makes this unbalanced (the entertainment interview videos from Youtube). The distance between negative emotions and neutral emotions in the ViCo dataset [61] is not obvious, and it is even difficult to distinguish manually. We hope an extensive and sufficient listener dataset can be explored in the future.

References

- [1] Leaderboard. <https://vico.solutions/>.
- [2] Speech-emotion-analyzer. <https://github.com/MiteshPuthran/Speech-Emotion-Analyzer>.
- [3] Chaitanya Ahuja, Shugao Ma, Louis-Philippe Morency, and Yaser Sheikh. To react or not to react: End-to-end visual pose forecasting for personalized avatar during dyadic conversations. In *2019 International conference on multimodal interaction*, pages 74–84, 2019.
- [4] Dan Bohus and Eric Horvitz. Facilitating multiparty dialog with gaze, gesture, and speech. In *International Conference on Multimodal Interfaces and the Workshop on Machine Learning for Multimodal Interaction*, pages 1–8, 2010.
- [5] Steven M Boker, Jennifer L Rotondo, Minquan Xu, and Kadijah King. Windowed cross-correlation and peak picking for the analysis of variability in the association between behavioral time series. *Psychological methods*, 7(3):338, 2002.
- [6] Jan Cech and Tereza Soukupova. Real-time eye blink detection using facial landmarks. *Cent. Mach. Perception, Dep. Cybern. Fac. Electr. Eng. Czech Tech. Univ. Prague*, pages 1–8, 2016.
- [7] Aleksandra Cerekovic, Oya Aran, and Daniel Gatica-Perez. Rapport with virtual agents: What do human social cues and personality explain? *IEEE Transactions on Affective Computing*, 8(3):382–395, 2016.
- [8] Lele Chen, Chen Cao, Fernando De la Torre, Jason Saragih, Chenliang Xu, and Yaser Sheikh. High-fidelity face tracking for ar/vr via deep lighting adaptation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 13059–13069, 2021.
- [9] Lele Chen, Guofeng Cui, Celong Liu, Zhong Li, Ziyi Kou, Yi Xu, and Chenliang Xu. Talking-head generation with rhythmic head motion. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part IX*, pages 35–51. Springer, 2020.
- [10] Lele Chen, Ross K Maddox, Zhiyao Duan, and Chenliang Xu. Hierarchical cross-modal talking face generation with dynamic pixel-wise loss. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 7832–7841, 2019.
- [11] Soo-Whan Chung, Hong Goo Kang, and Joon Son Chung. Seeing voices and hearing voices: learning discriminative embeddings using cross-modal self-supervision. *arXiv preprint arXiv:2004.14326*, 2020.

- [12] Daniel Cudeiro, Timo Bolkart, Cassidy Laidlaw, Anurag Ranjan, and Michael J Black. Capture, learning, and synthesis of 3d speaking styles. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10101–10111, 2019.
- [13] Sefik Emre Eskimez, You Zhang, and Zhiyao Duan. Speech driven talking face generation from a single image and an emotion condition. *IEEE Transactions on Multimedia*, 2021.
- [14] Scott Geng, Revant Teotia, Purva Tendulkar, Sachit Menon, and Carl Vondrick. Affective faces for goal-driven dyadic communication. *arXiv preprint arXiv:2301.10939*, 2023.
- [15] David Greenwood, Stephen Laycock, and Iain Matthews. Predicting head pose in dyadic conversation. In *Intelligent Virtual Agents: 17th International Conference, IVA 2017, Stockholm, Sweden, August 27-30, 2017, Proceedings 17*, pages 160–169. Springer, 2017.
- [16] Awni Hannun, Carl Case, Jared Casper, Bryan Catanzaro, Greg Diamos, Erich Elsen, Ryan Prenger, Sanjeev Sathesh, Shubho Sengupta, Adam Coates, et al. Deep speech: Scaling up end-to-end speech recognition. *arXiv preprint arXiv:1412.5567*, 2014.
- [17] Yinan He, Bei Gan, Siyu Chen, Yichun Zhou, Guojun Yin, Luchuan Song, Lu Sheng, Jing Shao, and Ziwei Liu. Forgerynet: A versatile benchmark for comprehensive forgery analysis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4360–4369, 2021.
- [18] 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. *Advances in neural information processing systems*, 30, 2017.
- [19] Ailin Huang, Zhewei Huang, and Shuchang Zhou. Perceptual conversational head generation with regularized driver and enhanced renderer. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 7050–7054, 2022.
- [20] Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. *arXiv preprint arXiv:1611.01144*, 2016.
- [21] Xinya Ji, Hang Zhou, Kaisiyuan Wang, Wayne Wu, Chen Change Loy, Xun Cao, and Feng Xu. Audio-driven emotional video portraits. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14080–14089, 2021.
- [22] Patrik Jonell, Taras Kucherenko, Gustav Eje Henter, and Jonas Beskow. Let’s face it: Probabilistic multi-modal interlocutor-aware generation of facial gestures in dyadic settings. In *Proceedings of the 20th ACM International Conference on Intelligent Virtual Agents*, pages 1–8, 2020.
- [23] Hyeonwoo Kim, Pablo Garrido, Ayush Tewari, Weipeng Xu, Justus Thies, Matthias Niessner, Patrick Pérez, Christian Richardt, Michael Zollhöfer, and Christian Theobalt. Deep video portraits. *ACM Transactions on Graphics (TOG)*, 37(4):1–14, 2018.
- [24] Ruilong Li, Shan Yang, David A Ross, and Angjoo Kanazawa. Ai choreographer: Music conditioned 3d dance generation with aist++. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13401–13412, 2021.
- [25] Xian Liu, Qianyi Wu, Hang Zhou, Yuanqi Du, Wayne Wu, Dahua Lin, and Ziwei Liu. Audio-driven co-speech gesture video generation. *arXiv preprint arXiv:2212.02350*, 2022.
- [26] Beth Logan et al. Mel frequency cepstral coefficients for music modeling. In *Ismir*, volume 270, page 11. Plymouth, MA, 2000.
- [27] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.
- [28] Daniel S Messinger, Mohammad H Mahoor, Sy-Miin Chow, and Jeffrey F Cohn. Automated measurement of facial expression in infant–mother interaction: A pilot study. *Infancy*, 14(3):285–305, 2009.
- [29] Seyedmahdad Mirsamadi, Emad Barsoum, and Cha Zhang. Automatic speech emotion recognition using recurrent neural networks with local attention. In *2017 IEEE International conference on acoustics, speech and signal processing (ICASSP)*, pages 2227–2231. IEEE, 2017.
- [30] Carla S Möller-Levet, Frank Klawonn, Kwang-Hyun Cho, and Olaf Wolkenhauer. Fuzzy clustering of short time-series and unevenly distributed sampling points. In *Advances in Intelligent Data Analysis V: 5th International Symposium on Intelligent Data Analysis, IDA 2003, Berlin, Germany, August 28-30, 2003. Proceedings 5*, pages 330–340. Springer, 2003.
- [31] Evonne Ng, Hanbyul Joo, Liwen Hu, Hao Li, Trevor Darrell, Angjoo Kanazawa, and Shiry Ginosar. Learning to listen: Modeling non-deterministic dyadic facial motion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20395–20405, 2022.
- [32] Behnaz Nojavanasghari, Yuchi Huang, and Saad Khan. Interactive generative adversarial networks for facial expression generation in dyadic interactions. *arXiv preprint arXiv:1801.09092*, 2018.
- [33] Andrew Owens and Alexei A Efros. Audio-visual scene analysis with self-supervised multisensory features. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 631–648, 2018.
- [34] Vijayaditya Peddinti, Daniel Povey, and Sanjeev Khudanpur. A time delay neural network architecture for efficient modeling of long temporal contexts. In *Sixteenth annual conference of the international speech communication association*, 2015.
- [35] Alexander Richard, Michael Zollhöfer, Yandong Wen, Fernando De la Torre, and Yaser Sheikh. Meshtalk: 3d face animation from speech using cross-modality disentanglement. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1173–1182, 2021.
- [36] Marcel Riehle, Jürgen Kempkensteffen, and Tania M Lincoln. Quantifying facial expression synchrony in face-to-face dyadic interactions: Temporal dynamics of simultaneously recorded facial emg signals. *Journal of Nonverbal Behavior*, 41:85–102, 2017.
- [37] Andreas Rössler, Davide Cozzolino, Luisa Verdoliva, Christian Riess, Justus Thies, and Matthias Nießner. Faceforensics: A large-scale video dataset for forgery detection in human faces. *arXiv preprint arXiv:1803.09179*, 2018.
- [38] Liam Schoneveld, Alice Othmani, and Hazem Abdelkawy. Leveraging recent advances in deep learning for audio-visual emotion recognition. *Pattern Recognition Letters*, 146:1–7, 2021.

- [39] Bowen Shi, Wei-Ning Hsu, Kushal Lakhota, and Abdelrahman Mohamed. Learning audio-visual speech representation by masked multimodal cluster prediction. *arXiv preprint arXiv:2201.02184*, 2022.
- [40] Aliaksandr Siarohin, Stéphane Lathuilière, Sergey Tulyakov, Elisa Ricci, and Nicu Sebe. First order motion model for image animation. *Advances in Neural Information Processing Systems*, 32, 2019.
- [41] Sanjana Sinha, Sandika Biswas, Ravindra Yadav, and Brojeshwar Bhowmick. Emotion-controllable generalized talking face generation. *arXiv preprint arXiv:2205.01155*, 2022.
- [42] Luchuan Song, Zheng Fang, Xiaodan Li, Xiaoyi Dong, Zhenchao Jin, Yuefeng Chen, and Siwei Lyu. Adaptive face forgery detection in cross domain (supplementary material).
- [43] Luchuan Song, Xiaodan Li, Zheng Fang, Zhenchao Jin, Yuefeng Chen, and Chenliang Xu. Face forgery detection via symmetric transformer. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 4102–4111, 2022.
- [44] Luchuan Song, Bin Liu, Guojun Yin, Xiaoyi Dong, Yufei Zhang, and Jia-Xuan Bai. Tacr-net: Editing on deep video and voice portraits. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 478–486, 2021.
- [45] Luchuan Song, Bin Liu, and Nenghai Yu. Talking face video generation with editable expression. In *Image and Graphics: 11th International Conference, ICIG 2021, Haikou, China, August 6–8, 2021, Proceedings, Part III 11*, pages 753–764. Springer, 2021.
- [46] Luchuan Song, Guojun Yin, Bin Liu, Yuhui Zhang, and Nenghai Yu. Fsft-net: face transfer video generation with few-shot views. In *2021 IEEE International Conference on Image Processing (ICIP)*, pages 3582–3586. IEEE, 2021.
- [47] Supasorn Suwajanakorn, Steven M Seitz, and Ira Kemelmacher-Shlizerman. Synthesizing obama: learning lip sync from audio. *ACM Transactions on Graphics (ToG)*, 36(4):1–13, 2017.
- [48] Fei Tao and Gang Liu. Advanced lstm: A study about better time dependency modeling in emotion recognition. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2906–2910. IEEE, 2018.
- [49] Justus Thies, Mohamed Elgharib, Ayush Tewari, Christian Theobalt, and Matthias Nießner. Neural voice puppetry: Audio-driven facial reenactment. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XVI 16*, pages 716–731. Springer, 2020.
- [50] Justus Thies, Michael Zollhofer, Marc Stamminger, Christian Theobalt, and Matthias Nießner. Face2face: Real-time face capture and reenactment of rgb videos. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2387–2395, 2016.
- [51] Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Advances in neural information processing systems*, 30, 2017.
- [52] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008.
- [53] Konstantinos Vougioukas, Stavros Petridis, and Maja Pantic. Realistic speech-driven facial animation with gans. *International Journal of Computer Vision*, 128:1398–1413, 2020.
- [54] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Guilin Liu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. Video-to-video synthesis. *arXiv preprint arXiv:1808.06601*, 2018.
- [55] Ting-Chun Wang, Arun Mallya, and Ming-Yu Liu. One-shot free-view neural talking-head synthesis for video conferencing. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10039–10049, 2021.
- [56] Haozhe Wu, Jia Jia, Haoyu Wang, Yishun Dou, Chao Duan, and Qingshan Deng. Imitating arbitrary talking style for realistic audio-driven talking face synthesis. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 1478–1486, 2021.
- [57] Jun Yu and Chang Wen Chen. From talking head to singing head: a significant enhancement for more natural human computer interaction. In *2017 IEEE International Conference on Multimedia and Expo (ICME)*, pages 511–516. IEEE, 2017.
- [58] Chenxu Zhang, Yifan Zhao, Yifei Huang, Ming Zeng, Saifeng Ni, Madhukar Budagavi, and Xiaohu Guo. Facial: Synthesizing dynamic talking face with implicit attribute learning. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 3867–3876, 2021.
- [59] Yan Zhang, Mohamed Hassan, Heiko Neumann, Michael J Black, and Siyu Tang. Generating 3d people in scenes without people. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6194–6204, 2020.
- [60] Zhimeng Zhang, Lincheng Li, Yu Ding, and Changjie Fan. Flow-guided one-shot talking face generation with a high-resolution audio-visual dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3661–3670, 2021.
- [61] Mohan Zhou, Yalong Bai, Wei Zhang, Tiejun Zhao, and Tao Mei. Responsive listening head generation: A benchmark dataset and baseline. *arXiv preprint arXiv:2112.13548*, 2021.