

CustomListener: Text-guided Responsive Interaction for User-friendly Listening Head Generation

Xi Liu*, Ying Guo*, Cheng Zhen, Tong Li, Yingying Ao, Pengfei Yan†

Vision AI Department, Meituan

<https://customlistener.github.io/>

Abstract

Listening head generation aims to synthesize a non-verbal responsive listener head by modeling the correlation between the speaker and the listener in dynamic conversation. The applications of listener agent generation in virtual interaction have promoted many works achieving diverse and fine-grained motion generation. However, they can only manipulate motions through simple emotional labels, but cannot freely control the listener's motions. Since listener agents should have human-like attributes (e.g. identity, personality) which can be freely customized by users, this limits their realism. In this paper, we propose a user-friendly framework called CustomListener to realize the free-form text prior guided listener generation. To achieve speaker-listener coordination, we design a Static to Dynamic Portrait module (SDP), which interacts with speaker information to transform static text into dynamic portrait token with completion rhythm and amplitude information. To achieve coherence between segments, we design a Past Guided Generation module (PGG) to maintain the consistency of customized listener attributes through the motion prior, and utilize a diffusion-based structure conditioned on the portrait token and the motion prior to realize the controllable generation. To train and evaluate our model, we have constructed two text-annotated listening head datasets based on ViCo and RealTalk, which provide text-video paired labels. Extensive experiments have verified the effectiveness of our model.

1. Introduction

Listening Head Generation (LHG) is to generate listener motions (face expressions, head movements, etc.) that respond synchronously to the speaker during face-to-face communication. Good communication is a two-way process, and the roles of speakers and listeners are equally important [4]. Whether the listener responds promptly

*Equal Contribution

†Corresponding Author

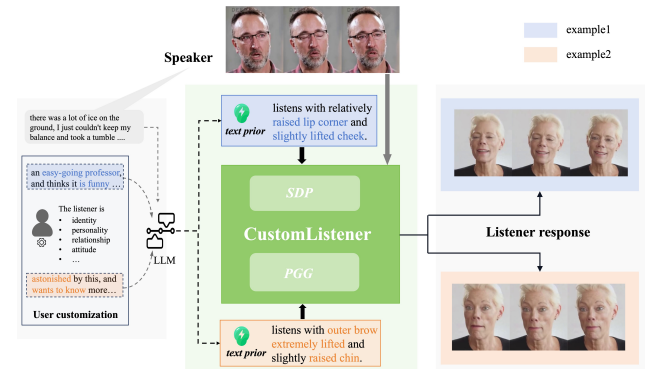


Figure 1. The process of text-guided listener generation in our CustomListener. The text prior provides the basic portrait style of the listener, which is input into CustomListener and combined with the speaker's information to obtain the listener's motions.

will affect whether the speaker can actively express the follow-up content [17, 19]. Unlike speaker head generation (SHG) [7, 16, 31, 38], listeners can only convey information through non-verbal feedback (nodding, frowning, etc.), and their responses are affected by both sides of the interaction (speaker's motion, tone, and listener's own personality and emotion state). Currently, many applications, such as digital avatar generation [8, 20] and human-computer interaction [43], use the listener agent to simulate more realistic conversations, which has promoted many studies in LHG.

Some early works manually integrate motions based on rules [1, 5, 35], or rely on 2D facial keypoints [11, 26] to manipulate the head, but they are limited in the variety of motions and the detailed control ability. Recent works [15, 21, 25, 34, 44] have introduced 3D face coefficients [9], combined with the post-rendering module to achieve more precise control. Specifically, RLHG [44] uses a sequence-sequence structure as the baseline method for decoding 3DMM coefficients. PCH [15] further improves the render module to inpaint image artifacts. L2L [25] quantizes listener motions into the discrete one-dimensional codebook by VQ-VAE [36], and ELP [34] expands the codebook to a composition of several discrete codewords to achieve more

fine-grained control. MFR-Net [21] designs a feature aggregation module to achieve better identity preservation and motion diversity.

Although the above methods can combine speaker-listener information to generate diverse and fine-grained listener motions, they have restrictions in the explicit control of the listener agent’s response. Specifically, L2L [25] focuses on the motion diversity but cannot control the motion under specific conditions. RLHG [44], PCH [15], MFR-Net [21], and ELP [34] achieves controllable listener motion conditioned on the input attitude. However, the attitude is limited to a few fixed labels (eg. positive, neutral, negative), posing limitations on the generation of realistic listeners: First, simple labels are not enough to give accurate responses. For example, for negative attitudes, the two emotions of doubt and anger correspond to different facial expressions. More importantly, a realistic listener agent should have its own identity (parent, friend, teacher, etc.), personality (lively, calm, etc.) and behavioral habits (e.g. frowning when thinking), which users can customize and set in advance. However, simple labels cannot achieve the control under these conditions.

Considering the above limitations of simple labels, in this paper, we propose a user-friendly framework called CustomListener to enable the text-guided generation, as shown in Figure 1. Users can pre-customize detailed attributes of the listener agent. Then using a large language model, we combine speech content and user-customized attributes to obtain the text prior of the listener’s basic portrait. In CustomListener, we seamlessly incorporate speaker information while being guided by the text prior, so as to generate realistic listener responses that are controllable and interactive.

For the guidance of text prior, due to the interactive nature of communication, listener motions cannot simply be regarded as text-conditioned motion generation, but faces the following challenges: 1) **Speaker-listener coordination**: the text provides the basic portrait style of the listener’s response. While to precisely express the listener’s emotional empathy towards the speaker, the response not only needs to complete the text-specified motion, but also needs to adjust its completion timing and rhythm, and fluctuate with the speaker’s semantics, intonation and movement amplitude. 2) **Listener motion coherence**: When motions conditioned on different text priors are combined into a long video, different segments should be coherent. Due to the customization, apart from ensuring smooth motion switching, we also need to guarantee that the user-customized listener’s behavioral habits reflected in past segments are maintained in the current segment. Maintaining the style of the past while still showing the motion of the current text can be challenging.

To solve the above challenges, in our CustomListener, we design the Static to Dynamic Portrait module (SDP) and the Past Guided Generation module (PGG) to achieve the *coordination* and the *coherence* respectively. In SDP,

we transform the static portrait provided by the text prior to be dynamic one, which interacts with audio to obtain the time-dependent information about the completion time and rhythm of the gradual motion changes, and refines the listener’s motions to make the speaker-listener fluctuations relevant. In PGG, we generate a motion prior based on the similarity of dynamic portrait tokens between adjacent segments, which provides reference information for features that need to be maintained in the current compared to the past. Then, the motion prior and the dynamic portrait token are input as conditions into the diffusion-based structure [14] with fixed noise at the segment connection, finally generating listener coefficients that enable coordination and coherence in interactions.

In summary, this paper has the following contributions:

- We argue that simple labels are not sufficient to achieve freely controllable listener generation, and thus propose the CustomListener framework that can freely control listener head motions guided by the text prior.
- We design a SDP module to achieve the speaker-listener coordination, so that the completion timing and rhythm of text-specified motions correspond to the speaker, and fluctuate with the speaker’s semantics, tone, and movements.
- For long-term generation, we design a PGG module to meet listener motion coherence between different text-prior-conditioned segments, so that the consistency of the listener’s customized behavioral habits between segments can be maintained and the motion switching is smooth.
- Extensive experiments on ViCo [44] and Realtalk [12] datasets confirm that our method can realize the state-of-the-art performance, and achieve the controllable and interactive realistic listener motion generation.

2. Related work

Listening head generation (LHG): Some early works [1, 5, 35] use rule-based methods to manually incorporate interactive motions. Later, data-driven methods [11, 26] generate 2D motions based on facial keypoints or movement frequencies, but 2D information cannot fully represent the details of facial movements.

Recently, motion generation based on 3D face coefficients, coupled with the post-processing renderer, has gradually become mainstream due to its stronger representation ability of facial movements. RLHG [44] proposes a representative ViCo dataset and constructs a baseline coefficient decoding method using a sequence-sequence structure. After that, PCH [15] assembles an enhanced renderer to make the generated results visually better and more stable by inpainting the image boundary and background. Learning2Listen (L2L) [25] proposes to use a sequence-encoding VQ-VAE [36] to learn the discrete codebook of listener motion. ELP [34] further utilizes a composition of multiple discrete motion-codewords to represent facial motion in a more

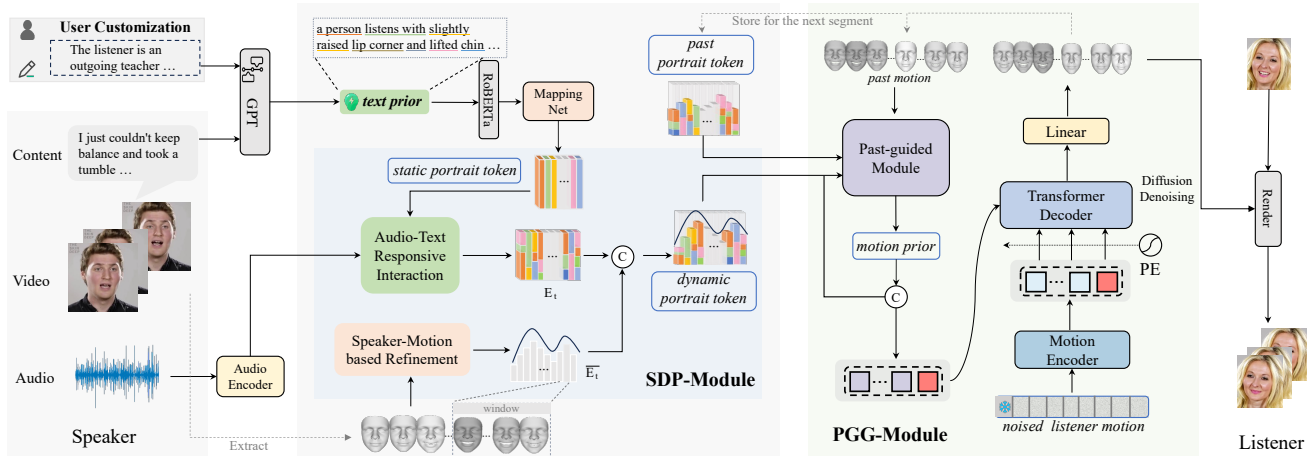


Figure 2. Overall framework of CustomListener. Given the text prior providing the listener’s static portrait style, SDP-Module transforms the static portrait into a dynamic one. Then in PGG-Module, the dynamic portrait token are combined with motion prior generated from Past-guided Module and are utilized as conditions of the diffusion-based structure to realize the controllable generation. The ‘C’ in the figure denotes concatenation and the small pink squares denotes diffusion time-step token.

fine-grained manner. MFR-Net [21] achieves the preservation of listener’s identity and the diversity of generation.

Regarding the listener’s motion status, L2L [25] focuses on the diversity but cannot explicitly control specific actions. RLHG [44], PCH [15], MFR-Net [21] achieve controllable generation conditioned on the input emotion, but the emotion is limited to a few labels and cannot specifically control the expression and posture of each emotion. Although ELP [34] represents the latent space under different emotions in a fine-grained manner through the codebook, what the user can control is still the input emotion rather than the fine-grained motions. In this way, the listener’s response to each emotion may depend on the data distribution of the training set.

Conditional motion generation: The motion generation task derives applications under a variety of conditions. For example, body motion conditioned on text descriptions [13, 18, 27, 28, 41], motion trajectory conditioned on scene image [3], and dance motion conditioned on music [39]. Recently, large language models promote the development of free-form text conditioned generation based on auto-encoder [28], hierarchical model [13], and diffusion models [18]. We also generate motions conditioned on free-form text. However, for our responsive interaction, in addition to focusing on the matching of motion and text like the above methods, we also need to make the motion and the speaker’s action, speech, tone and rhythm harmonious to express the listener’s empathy. In our paper, we adopt the SDP module to achieve this harmony. For long-term generation, to ensure the coherent fluency between clips under different text commands, some methods have adopted transition sampling [40], specifying noise [42], etc. However, the coherence for the custom listener, which has some motion habits, is more complex. It is not just a simple motion switch, but is related to past habits and styles, and our PGG module solves this challenge.

3. Method

3.1. Overview

In our CustomListener framework, under the guidance of the fine-grained text prior, we can synthesize natural and controllable listener motions, which can seamlessly integrate the speaker’s information (motion, speech content, audio) to give interactive non-verbal responses. Let $L_m \in \mathbb{R}^{T \times 70}$ denote the listener motions to be generated, and T is the number of frames, 70 denotes 64-dim expression and 6-dim pose. Given speaker’s motions $S_m \in \mathbb{R}^{T \times 70}$, audio S_a , speech content S_c and text prior R , our framework can be formulated as:

$$L_m = Custom(S_m, S_a, S_c; R) \quad (1)$$

The overview of our proposed framework is shown in Figure 2. Before generation, for each video segment, we use GPT to generate a text prior based on user-customized listener attributes and the speaker’s speech content, which provides the listener’s static portrait reference. Then, we extract the speaker’s audio features from the Audio Encoder and speaker’s motion coefficients from a pre-trained face reconstruction model [9]. Following this, we generate realistic listener motions through a two-stage process: (1) Static to Dynamic Portrait generation (SDP-Module): it can transform static portrait tokens into the dynamic through audio-text responsive interaction and speaker-motion based refinement. The dynamic portrait token includes the time-dependent information about the motion completion and rhythm, and makes the speaker-listener motion fluctuations correlated; (2) Past-guided Listener Motion Generation (PGG-Module): to produce coherent motions and maintain the listener’s previous behavioral habits in long-term generation, we use a past guided module to produce the motion prior, combined

with the diffusion denoising process with a fixed initial noise to generate the listener coefficients. Finally, listener videos are obtained through the post-rendering process.

3.2. Static to Dynamic Portrait generation

In this section, we first introduce the generation of the static portrait-token, and then present SDP-Module, which can convert static tokens into dynamic one by two stages: (1) Audio-text Responsive Interaction: static portrait tokens are converted into time-dependent portrait tokens, which reflects the completion time and rhythm of motions according to the semantics and intonation in the audio, ensuring gradual motion changes; (2) Speaker Motion-based Refinement: it refines the listener’s motions according to the fluctuations of speaker’s motions, producing the final dynamic portrait token combined with the time-dependent portrait token.

Static Portrait-token Generation To establish a better connection between the user-customized listener attributes (identity, personality, relationship, attitude, etc.) and the listener’s motions, we leverage the powerful GPT for semantic understanding of both user-customized contents and speaker’s speech contents, and generate the text prior that accurately describes the listener’s expressions and poses. Following this, we use RoBERTa [22], a robust and efficient model with high natural language understanding capability, to encode the text prior into text embeddings. To obtain portrait-related tokens, we further employ a mapping net to convert text embeddings into portrait tokens. We consider these portrait tokens to be static portrait tokens since the text prior describes the listener’s static portrait (eg. expressions and poses) in the entire video clip.

Audio-text Responsive Interaction Since static portrait tokens from text prior provide a comprehensive representation of listener head in a video segment, it lacks the time-dependent information. Thus, the model would be limited to generating constant motion and cannot effectively convey gradual variation in movement without a specific design. To address this critical issue, we make use of the semantic knowledge in speaker’s utterances, which serves as a guidance for motion shifts. Specifically, we choose the speaker audio to provide semantics instead of speech contents. The underlying intuition is the former contains not only semantic information but also the temporal information. This entire process can be split into two parts, part A and part B.

Part A In this part, we generate a weight matrix through responsive interaction between speaker audio features and the static portrait tokens obtained from text prior. It provides the correlation degree between the speaker’s semantics at each timestamp and the motion description in the text prior. For example, as shown in Figure 3, if the phrase “love you” in the speaker’s audio at time stamp k exhibits the strongest correlation with the phrase “raised lip corner” in the description texts, the corresponding weight value will reach its peak

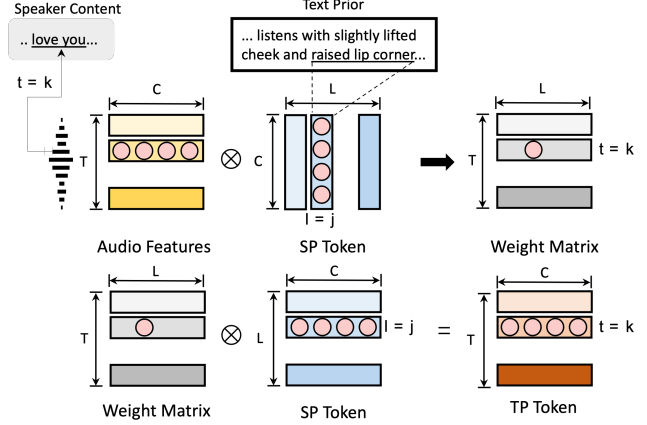


Figure 3. Illustration of Audio-text Responsive Interaction. We first generate a weight matrix by responsive interaction between audio features and static portrait-token (SP Token), and then generate time-dependent portrait token (TP Token) guided by weight matrix.

(indicated by a pink circle dot in the weight map). This weight matrix $M(A, H)$ can be formulated as Equation 2.

Part B After obtaining a weight matrix $M(A, H)$, we can generate time-dependent portrait tokens based on static portrait tokens with the guidance of $M(A, H)$. Specifically, as illustrated in Figure 3, the time-dependent portrait tokens (TP Token) at the time stamp k is generated with the highest weighting assigned to the static portrait tokens (SP Token) at $l = j$, which corresponds to the phrase “raised lip corner”. This part can be formulated as Equation 3.

$$M(A, H) = Sm\left(\frac{1}{\sqrt{\alpha}}(Linear(A) \times Linear(H)^T)\right), \quad (2)$$

$$E_t = \sum_{i=0}^{L-1} M(A, H)_{i,t} \cdot H_{i,j=0:C-1}, t \in \{0, \dots, T-1\} \quad (3)$$

where Sm denotes Softmax, $A \in \mathbb{R}^{T \times C}$ is the audio feature, $H \in \mathbb{R}^{L \times C}$ is the static portrait token, $M \in \mathbb{R}^{T \times L}$ is the weight map, α is a scaling factor, $E_t \in \mathbb{R}^{1 \times C}$ is the time-dependent portrait token at time stamp t . L is the length of text embeddings, and C is the feature dimension.

With the help of two parts above, we can obtain portrait tokens that contain time-dependent information about the rhythm of motion completion. This enables the model to learn the progressive changes in facial/head motions within a single video segment, as opposed to the constant motions.

Speaker Motion-based Refinement In human-to-human conversations, we believe that the amplitude of the speaker’s motion fluctuation also has an impact on the changes in the listener’s motion response[21, 32]. Therefore, we utilize the past speaker’s motion changes to refine the listener’s motion amplitude. This process can be well-designed formulated as:

$$\bar{E}_t = \sum_{i=t-1-w}^{t-1} \exp\left|\frac{\partial S_i}{\partial i}\right|, \quad (4)$$

$$\widetilde{E}_t = E_t \oplus \overline{E}_t, \quad (5)$$

where \oplus means concatenation, E_t , \overline{E}_t , \widetilde{E}_t is time-dependent portrait token, speaker motion-weighted portrait token and dynamic portrait token at frame t , respectively. S_i denotes speaker motion at frame i , w is the time-window length. For efficiency, we set w to 5, which means we only take the past five-frame speaker motions into account. For each time stamp t , the listener’s portrait token E_t is affected by the cumulative gradients of several previous speaker motions. This process can be illustrated in the lower part of the SDP-Module in Figure 2. Notably, when generating portrait token at frame t , we only focus on the speaker motions up to frame $t-1$. Due to the listener and the speaker communicating concurrently, the listener’s state at frame t should not be influenced by the speaker’s motion at the same frame t .

3.3. Past-guided Listener Motion Generation

The PGG-Module is to ensure that motions between segments are coherent, and the customized listener attributes in the previous segment can be maintained in the current segment. It contains Past-guided Module and diffusion-based motion generation module, which can generate motion prior based on past motions and then synthesize listener motions. Following this, a renderer is used to obtain listener videos.

Past-guided Module To produce long video with arbitrary length, we opt for a segment-by-segment generation approach during inference. However, we have observed that there are unsmooth pose transitions between segments. This issue can be attributed to the neglect of correlations between adjacent segments. To address this, we take past motions into consideration.

Since our final goal is to synthesize a customized listener, in addition to ensuring smooth transition, we also need to take the listener’s past behavioral habits into account (e.g. frowning when thinking). Hence, we sought assistance from dynamic portrait tokens of two adjacent segments to guide the generation of motion prior. Specifically, the weight we assign to past motions at different times is determined by the similarity between portrait tokens in the current segment and those in the last segment. Formally, we calculate the motion weight and obtain the motion prior $G_{t:t+L}$ by:

$$Sim_{t:t+L} = E_{t:t+L} \times E_{t-L:t}^T, \quad (6)$$

$$G_{t:t+L} = Sim_{t:t+L} \times P_{t-L:t}, \quad (7)$$

where \times denotes matrix multiplication, L is the segment length, t is the time stamp, $Sim_{t:t+L}$ denotes the similarity between portrait tokens in the current segment $E_{t:t+L}$ and those in the last segment $E_{t-L:t}$, which is the weight we assign to past motions $P_{t-L:t}$. The higher the similarity, the greater the weight we assign. This approach can not only encourage continuous transitions between two adjacent

segments, but also ensure the maintenance of customized listener’s behavioral habits between two adjacent sequences.

Diffusion Model To achieve diverse and natural head motion generation, our generation module is built on denoising diffusion probabilistic models (DDPM) [14].

During forward diffusion process q , we progressively inject noise into the listener motion x_0 until it becomes a Gaussian noise x_T , which can be defined as a Markov chain:

$$q(x_{1:T} | x_0) = \prod_{t=1}^T q(x_t | x_{t-1}), \quad (8)$$

$$q(x_t | x_{t-1}) = \mathcal{N}(x_t; \sqrt{1 - \beta_t}x_{t-1}, \beta_t \mathbf{I}), \quad (9)$$

where x_1, \dots, x_T are noised listener motions, β_1, \dots, β_T are the variance schedule. In the reverse process p_θ , we produce listener motion x_0 by removing noise from x_T step-by-step:

$$p_\theta(x_{t-1} | x_t, c) = \mathcal{N}(x_{t-1}; \mu_\theta(x_t, c, t), \Sigma_\theta(x_t, c, t)), \quad (10)$$

where c is a conditioning variable. We perform denoising process with a transformer decoder [37] conditioned on the motion prior and dynamic portrait tokens.

Render We utilize PIRenderer [30] for rendering, which can produce natural listener videos by incorporating a single listener reference image as well as the 3DMM coefficients.

4. Experiments

4.1. Datasets

Data Annotation Since the existing listener head datasets do not contain the full data required for training our model, we conduct additional data annotations on two datasets: (1) ViCo [44], a popular dataset in listener head generation; (2) RealTalk, proposed by [12], which serves as a database for retrieving listener videos in [12]. Firstly, to learn the relationships between the text prior and listener expressions, we follow the text annotation pipeline in TalkCLIP [24], and conduct text annotation on ViCo[44] and RealTalk[12]. Specifically, we employ [6] to acquire emotion labels, activated AUs as well as their intensities. Additionally, to learn realistic head movements such as nodding to convey agreement and shaking the head to signify disagreement, we use Hopenet [33] to detect head motions. Finally, we generate diverse text prior description for each video segment according to the format: [A person <EMOTION> and listens with <AU> (and <HEAD MOTION>)], where <EMOTION> denotes the emotional labels, <AU> denotes fine-grained expression-related texts based on AU activation intensities and Facial Action Coding System [10], and <HEAD MOTION> indicates head nodding/shaking labels. We use “and” to link when there is more than one activated AU. Detailed examples can be seen in Appendix A.

Methods	Testset	FD ↓			RTLCC ↓		RWTLCC ↓		FID $_{\Delta fm}$ ↓		SND ↓		V-D ↑	
		exp	angle	trans	exp	pose	exp	pose	exp	pose	exp	pose	exp	pose
RLHG* [44]	\mathcal{D}_{test}	15.03	7.90	6.55	0.113	0.160	0.108	0.160	11.65	0.90	2.90	0.09	0.69	0.10
	\mathcal{D}_{ood}	22.81	8.58	8.90	0.124	0.101	0.104	0.100	9.32	0.82	7.20	0.11	1.01	0.12
PCH* [15]	\mathcal{D}_{test}	19.02	13.26	7.74	0.124	0.163	0.109	0.198	11.59	0.96	3.10	0.13	0.31	0.12
	\mathcal{D}_{ood}	18.63	17.96	8.80	0.121	0.115	0.107	0.140	8.67	0.87	7.01	0.12	0.51	0.16
L2L* [25]	\mathcal{D}_{test}	14.89	7.35	6.34	0.102	0.153	0.101	0.113	8.64	0.87	2.62	0.09	0.57	0.20
	\mathcal{D}_{ood}	15.64	8.23	7.89	0.095	0.098	0.100	0.101	6.78	0.74	6.89	0.10	1.23	0.28
MFR-Net [†] [21]	\mathcal{D}_{test}	13.37	6.82	6.02	-	-	-	-	-	-	-	-	-	-
	\mathcal{D}_{ood}	14.70	8.12	6.37	-	-	-	-	-	-	-	-	-	-
Ours	\mathcal{D}_{test}	11.54	6.12	5.90	0.072	0.103	0.081	0.097	3.12	0.06	2.40	0.07	1.55	0.32
	\mathcal{D}_{ood}	12.67	7.49	6.01	0.081	0.073	0.084	0.075	3.59	0.07	5.94	0.05	1.25	0.30

Table 1. Comparisons of our model with other methods on \mathcal{D}_{test} and \mathcal{D}_{ood} of ViCo. **Bold** represents the best. The [†] means we directly refer to data in their paper and * denotes we retrain the model. The ↓ indicates lower is better and the ↑ indicates higher is better. The values of FD and FID $_{\Delta fm}$ are multiplied by 100. The quantitative results on RealTalk [12] are in Appendix B.

Data Construction For ViCo, which contains three parts: train set \mathcal{D}_{train} , test set \mathcal{D}_{test} , out-of-domain set \mathcal{D}_{ood} (all identities in \mathcal{D}_{ood} have not appeared in \mathcal{D}_{train}), we obtain the transcripts of the speaker’s audios by a speech recognition model [29], and get an extended version of ViCo. For RealTalk, to align the dimension of 3DMM coefficients with ViCo, we re-extract facial 3DMM coefficients using [9] and then obtain two subsets: \mathcal{D}_{train} and \mathcal{D}_{test} . These two extended datasets are important for future research on text-guided listener head generation task.

4.2. Experimental Settings

Loss Fuction During training, followed by [14], we perform denoising on noised listener motions x_t conditioned on motion prior and dynamic portrait token step-by-step, and our optimization goal is to predict noise term ϵ at each step. The objective function can be defined as a conditional diffusion loss as follows:

$$L_d = E_{x_0, t, \epsilon_t} \left[\|\epsilon_t - \epsilon_\theta(x_t, c, t)\|^2 \right], \quad (11)$$

where $\epsilon_t \sim \mathcal{N}(0, \mathbf{I})$ is the noise at step t , c denotes conditions, x_0 is the original listener motions.

Implementation Details The resolution of videos is 256×256 , and the FPS is 30. For training, we clip the original videos to several 60-frame video segments. We extract 3DMM coefficients from speaker videos and listener videos by [9], which serve as speaker motions and listener motions, respectively. Following MFR-Net [21], we extract 45-dim acoustic features from audio. We utilize AdamW optimizer [23] to train our model with a learning rate of $1e^{-4}$.

4.3. Quantitative Results

Metrics It’s non-trivial to measure the realism of generated listener motions. We choose the following metrics based on

Method	SSIM ↑	CPBD ↑	PSNR ↑	FID ↓
RLHG* [44]	0.56	0.12	17.39	27.70
PCH* [15]	0.58	0.15	18.48	21.29
L2L* [25]	0.58	0.16	17.67	22.13
MFR-Net [†] [21]	0.59	0.18	17.82	20.08
Ours	0.60	0.19	17.98	20.06

Table 2. Quantitative results with state-of-the-art methods on image quality. The best results are highlighted in bold.

that a good custom listener should have four characteristics: (1) natural-looking and diverse facial/head motions; (2) highly synchronous with the speaker motions; (3) highly coherent motions in a long video; (4) high image quality:

- **FD**: $L1$ distance between the predicted 70-dim listener motions and the ground-truth listener motions.
- **Variation for Diversity (V-D)**: Variance across the time series sequence of 3DMM coefficients.
- **Residual Time Lagged Cross Correlation (RTLCC)**: $L1$ distance between generated TLCC [2] and TLCC [2] of ground truth. RTLCC indicates the correlation between listener motions and lagged speaker motions.
- **Residual Windowed Time Lagged Cross Correlation (RWTLCC)**: RTLCC in a fixed window. We set the window length to 4 seconds following ELP [34].
- **FID $_{\Delta fm}$** : FID score of differences in 3DMM coefficients between adjacent frames, which indicates the temporal naturalness of facial motions, followed by [42].
- **SND**: Sequence Naturalness Distance proposed by [42], showing the distribution difference between generated and ground-truth motions.
- **SSIM, CPBD, PSNR and FID**: Common metrics for image quality evaluation.

Comparisons with Existing Methods We retrain PCH [15], RLHG [44] and L2L [25] on ViCo[44]. It should be




Figure 4. Visual results produced by CustomListener. All listener videos are generated conditioned on different pre-set text priors, the same speaker video (the 1st row) and the same reference listener image.

noted that, source codes of ELP [34] and MFR-Net [21] is unavailable. For MFR-Net [21], we utilize the data from the original paper. For ELP [34], the 3DMM coefficient extraction model used is different from ours, resulting in distinct dimensions of coefficients (e.g., $\beta \in \mathbb{R}^{100T}$ for ELP [34], $\beta \in \mathbb{R}^{64T}$ for ours), thus it is not reasonable to directly compare with evaluation data in ELP [34]. Therefore, we only provide its visual comparisons in Appendix C. Our quantitative comparisons consist of four parts: 1) precision and diversity of listener motions: as presented in Table 1, for the realism of listener motion generation, our proposed method achieves the lowest FD on both \mathcal{D}_{test} and \mathcal{D}_{ood} in ViCo. Meanwhile, the V-D of our model is comparable to other methods, which indicates that our model can maintain certain diversity when generating controllable listener motions. This may benefit from the probabilistic nature of the diffusion model. 2) speaker-listener synchronization: compared to all existing methods, CustomListener can decrease the RTLCC and the RWTLCC, which justifies the effectiveness of our dynamic portrait tokens. 3) temporal smoothness: in Table 1, the $FID_{\Delta fm}$ and SND of our model outperform other existing methods by a large margin, which validates that our model can mitigate un-smooth transitions. 4) image-level quality: although image quality is not our main focus, we still have a certain performance gain in related metrics, as shown in Table 2. These results indicate that our generated motion is closer to GT motion, thus achieving relatively high image-level quality without special design in rendering.



Figure 5. Qualitative comparisons with PCH [15], RLHG [44] conditioned on the same speaker and the same listener reference image.

4.4. Qualitative Results

Visual Results Since previous methods cannot generate listener head based on fine-grained texts, we first present visual results generated by CustomListener. As shown in Figure 4, given different texts, covering diverse listener’s personalities, identities and attitudes toward the speaker’s utterances, our framework can produce fine-grained non-verbal listener motions closely aligned with pre-customized texts. For example, if the listener is an impatient person who is not interested in the speaker’s words (the 3rd row in



Figure 6. Ablation study of PGM to verify that PGM can improve motion coherency. Each group includes adjacent four frames, and the frames enclosed by red box are the two frames at the transition point between two video clips.

Method	FD ↓ (x100)	RTLCC ↓	RWTLCC ↓	FID $_{\Delta fm}$ ↓ (x100)	SND ↓
w/o ATRI	20.67	0.206	0.203	5.01	4.53
w/o SMW	19.05	0.201	0.196	5.07	4.58
w/o SDP	21.97	0.213	0.210	5.90	4.60
w/o PGM	20.53	0.197	0.201	6.47	4.65
Ours	18.48	0.165	0.169	3.42	4.23

Table 3. Ablation study. Each cell represents the average results of 70-dim motions (64-dim exp + 6-dim pose) on \mathcal{D}_{test} and \mathcal{D}_{ood} .

Method	Similarity to GT	Overall Naturalness	Speaker-listener Coordination	Video Smooth
RLHG [44]	1.28	1.12	1.64	2.04
PCH [15]	2.08	2.40	2.48	2.28
L2L [25]	2.68	2.56	1.92	1.80
Ours	3.96	3.92	3.96	3.88

Table 4. User study results. The best results are highlighted in bold.

Figure 4), he may appear indifferent and show impatience. Additionally, the produced listener’s responses can fluctuate with the speaker’s semantics and movement, showing great effectiveness of our proposed dynamic portrait tokens.

Comparisons with Existing Methods In Figure 5, we compare our results with RLHG [44], PCH [15]. It can be seen that our method is capable of generating more realistic listener heads than other baselines. In the left four columns, both RLHG [44] and PCH [15] suffer from unnatural motions (e.g., disorderly opened mouth) and visible facial artifacts. In the right four columns, RLHG [44] still exhibits incorrect motions and PCH [15] can not display negative attitude accurately. While our method can produce natural listener head that highly aligns with the ground truth conditioned on the given texts. More detailed comparisons can be seen in Appendix C.

User Study We conduct user studies to evaluate from human perceptual perspective. We randomly choose 25 speaker videos from ViCo[44] and RealTalk[12] to generate listeners for each method. 28 people are asked to rate different methods(1-4, 4 is the best) for each video from 4 metrics: similarity to GT, overall naturalness, speaker-listener coordination and video smooth. The average scores of each method are shown in Table 4. Our method outperformed other methods in all aspects.

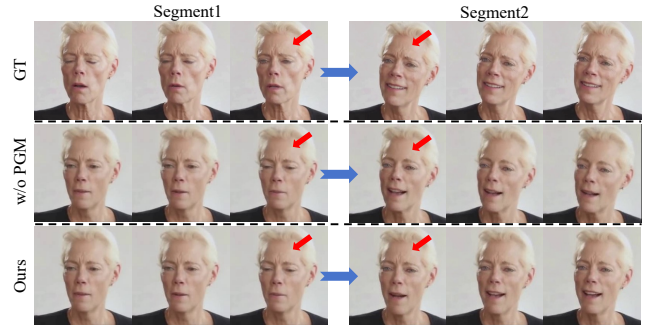


Figure 7. Ablation study of PGM. For different segments, the habit of frowning is maintained using PGM.

4.5. Ablation Study

SDP module In Section 3.2, we design a SDP module to generate dynamic portrait tokens for realizing progressive motion shifts that correspond well to the speaker. To validate its effectiveness, we conduct ablation study on Audio-text Responsive Interaction (ATRI) and Speaker Motion-based Weighting (SMW). The results are shown in Table 3. In addition to the improvement in FD, the improvement in the RTLCC and RWTLCC is more pronounced, which demonstrates the SDP module is crucial for enhancing speaker-listener synchronization, and each of its components matters.

Past-guided module To validate the effectiveness of our Past-guided module (PGM), we train networks without PGM. As presented in Table 3, without PGM, the metrics FID $_{\Delta fm}$ and SND, which are related to the temporal naturalness of motions, increase significantly, indicating a decrease in smoothness. This observation shows the necessity of PGM for mitigating unnatural transitions between video clips. The visual results are shown in Figure 6. Apart from this, our proposed PGM can maintain the consistency of listener’s customized behavioral habits between adjacent clips, as shown in Figure 7. The listener is customized a habit of frowning when thinking, and when segment 2 is guided by “lips mildly parted and lip corner raised slightly”, the frowning habit can be still maintained.

5. Conclusion

In this paper, we present a user-friendly framework called CustomListener, which can realize freely-controllable listener head generation conditioned on text priors. To achieve the speaker-listener coordination, we propose a SDP-module for dynamic portrait-token generation, which is capable of representing progressive motion changes. To ensure coherent listener responses in a long video, we design a PGG-module for smooth transition between adjacent video segments as well as maintaining customized behavioral habits. Comprehensive experiments demonstrate the superiority of our whole framework. In future work, it may be interesting to generalize our framework to listening body generation, yielding a complete and realistic listener agent.

References

- [1] 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. 1, 2
- [2] 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
- [3] Zhe Cao, Hang Gao, Karttikeya Mangalam, Qi-Zhi Cai, Minh Vo, and Jitendra Malik. Long-term human motion prediction with scene context. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part I 16*, pages 387–404. Springer, 2020. 3
- [4] Justine Cassell and Kristinn R Thorisson. The power of a nod and a glance: Envelope vs. emotional feedback in animated conversational agents. *Applied Artificial Intelligence*, 13(4-5): 519–538, 1999. 1
- [5] Justine Cassell, Catherine Pelachaud, Norman Badler, Mark Steedman, Brett Achorn, Tripp Becket, Brett Douville, Scott Prevost, and Matthew Stone. Animated conversation: rule-based generation of facial expression, gesture & spoken intonation for multiple conversational agents. In *Proceedings of the 21st annual conference on Computer graphics and interactive techniques*, pages 413–420, 1994. 1, 2
- [6] Di Chang, Yufeng Yin, Zongjian Li, Minh Tran, and Mohammad Soleymani. Libreface: An open-source toolkit for deep facial expression analysis. *arXiv preprint arXiv:2308.10713*, 2023. 5
- [7] 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. 1
- [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. 1
- [9] Yu Deng, Jiaolong Yang, Sicheng Xu, Dong Chen, Yunde Jia, and Xin Tong. Accurate 3d face reconstruction with weakly-supervised learning: From single image to image set. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pages 0–0, 2019. 1, 3, 6
- [10] Paul Ekman. Facial action coding system (facs). *A human face*, 2002. 5
- [11] Will Feng, Anitha Kannan, Georgia Gkioxari, and C Lawrence Zitnick. Learn2smile: Learning non-verbal interaction through observation. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4131–4138. IEEE, 2017. 1, 2
- [12] Scott Geng, Revant Teotia, Purva Tendulkar, Sachit Menon, and Carl Vondrick. Affective faces for goal-driven dyadic communication. *arXiv preprint arXiv:2301.10939*, 2023. 2, 5, 6, 8
- [13] Anindita Ghosh, Noshaba Cheema, Cennet Oguz, Christian Theobalt, and Philipp Slusallek. Synthesis of compositional animations from textual descriptions. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 1396–1406, 2021. 3
- [14] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. 2, 5, 6
- [15] 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 (MM’22)*, 2022. 1, 2, 3, 6, 7, 8
- [16] 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. 1
- [17] Adam Kendon. Movement coordination in social interaction: Some examples described. *Acta psychologica*, 32:101–125, 1970. 1
- [18] Jihoon Kim, Jiseob Kim, and Sungjoon Choi. Flame: Free-form language-based motion synthesis & editing. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 8255–8263, 2023. 3
- [19] Marianne LaFrance. Nonverbal synchrony and rapport: Analysis by the cross-lag panel technique. *Social Psychology Quarterly*, pages 66–70, 1979. 1
- [20] 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. 1
- [21] Jin Liu, Xi Wang, Xiaomeng Fu, Yesheng Chai, Cai Yu, Jiao Dai, and Jizhong Han. Mfr-net: Multi-faceted responsive listening head generation via denoising diffusion model. In *Proceedings of the 31th ACM International Conference on Multimedia (MM’23)*, 2023. 1, 2, 3, 4, 6, 7
- [22] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019. 4
- [23] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 6
- [24] Yifeng Ma, Suzhen Wang, Yu Ding, Bowen Ma, Tangjie Lv, Changjie Fan, Zhipeng Hu, Zhidong Deng, and Xin Yu. Talkclip: Talking head generation with text-guided expressive speaking styles. *arXiv preprint arXiv:2304.00334*, 2023. 5
- [25] 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 (CVPR)*, pages 20395–20405, 2022. 1, 2, 3, 6, 8
- [26] Behnaz Nojavanasghari, Yuchi Huang, and Saad Khan. Interactive generative adversarial networks for facial expres-

- sion generation in dyadic interactions. *arXiv preprint arXiv:1801.09092*, 2018. 1, 2
- [27] Mathis Petrovich, Michael J Black, and Gül Varol. Action-conditioned 3d human motion synthesis with transformer vae. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10985–10995, 2021. 3
- [28] Mathis Petrovich, Michael J Black, and Gül Varol. Temos: Generating diverse human motions from textual descriptions. In *European Conference on Computer Vision*, pages 480–497. Springer, 2022. 3
- [29] Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. Robust speech recognition via large-scale weak supervision. In *International Conference on Machine Learning*, pages 28492–28518. PMLR, 2023. 6
- [30] Yurui Ren, Ge Li, Yuanqi Chen, Thomas H Li, and Shan Liu. Pirenderer: Controllable portrait image generation via semantic neural rendering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13759–13768, 2021. 5
- [31] 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. 1
- [32] Daniel Richardson, Rick Dale, and Kevin Shockley. Synchrony and swing in conversation: Coordination, temporal dynamics, and communication. *Embodied communication in humans and machines*, pages 75–94, 2008. 4
- [33] Nataniel Ruiz, Eunji Chong, and James M Rehg. Fine-grained head pose estimation without keypoints. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 2074–2083, 2018. 5
- [34] Luchuan Song, Guojun Yin, Zhenchao Jin, Xiaoyi Dong, and Chenliang Xu. Emotional listener portrait: Neural listener head generation with emotion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 20839–20849, 2023. 1, 2, 3, 6, 7
- [35] Sinan Sonlu, Uğur Güdükbay, and Funda Durupinar. A conversational agent framework with multi-modal personality expression. *ACM Transactions on Graphics (TOG)*, 40(1): 1–16, 2021. 1, 2
- [36] Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Advances in neural information processing systems*, 30, 2017. 1, 2
- [37] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 5
- [38] Konstantinos Vougioukas, Stavros Petridis, and Maja Pantic. Realistic speech-driven facial animation with gans. *International Journal of Computer Vision*, 128:1398–1413, 2020. 1
- [39] Siqi Yang, Zejun Yang, and Zhisheng Wang. Longdancediff: Long-term dance generation with conditional diffusion model. *arXiv preprint arXiv:2308.11945*, 2023. 3
- [40] Zhao Yang, Bing Su, and Ji-Rong Wen. Synthesizing long-term human motions with diffusion models via coherent sampling. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 3954–3964, 2023. 3
- [41] Zhao Yang, Bing Su, and Ji-Rong Wen. Synthesizing long-term human motions with diffusion models via coherent sampling. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 3954–3964, 2023. 3
- [42] Zhentao Yu, Zixin Yin, Deyu Zhou, Duomin Wang, Finn Wong, and Baoyuan Wang. Talking head generation with probabilistic audio-to-visual diffusion priors. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7645–7655, 2023. 3, 6
- [43] 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. 1
- [44] Mohan Zhou, Yalong Bai, Wei Zhang, Ting Yao, Tiejun Zhao, and Tao Mei. Responsive listening head generation: A benchmark dataset and baseline. In *Proceedings of the European conference on computer vision (ECCV)*, 2022. 1, 2, 3, 5, 6, 7, 8