Responsive Listening Head Generation: A Benchmark Dataset and Baseline

Mohan Zhou^{1*†}[©], Yalong Bai^{2†}[©], Wei Zhang²[©], Ting Yao²[©], Tiejun Zhao^{1§}[©], and Tao Mei²[©]

¹Harbin Institute of Technology ²JD Explore Academy, Beijing, China {mhzhou99, ylbai}@outlook.com, {wzhang.cu, tingyao.ustc}@gmail.com, tjzhao@hit.edu.cn, tmei@jd.com

Abstract. We present a new listening head generation benchmark, for synthesizing responsive feedbacks of a listener (e.q., nod, smile) during a face-to-face conversation. As the indispensable complement to talking heads generation, listening head generation has seldomly been studied in literature. Automatically synthesizing listening behavior that actively responds to a talking head, is critical to applications such as digital human, virtual agents and social robots. In this work, we propose a novel dataset "ViCo", highlighting the listening head generation during a face-to-face conversation. A total number of 92 identities (67 speakers and 76 listeners) are involved in ViCo, featuring 483 clips in a paired "speaking-listening" pattern, where listeners show three listening styles based on their attitudes: positive, neutral, negative. Different from traditional speech-to-gesture or talking-head generation, listening head generation takes as input both the audio and visual signals from the speaker, and gives non-verbal feedbacks (e.g., head motions, facial expressions) in a real-time manner. Our dataset supports a wide range of applications such as human-to-human interaction, videoto-video translation, cross-modal understanding and generation. To encourage further research, we also release a listening head generation baseline, conditioning on different listening attitudes. Code & ViCo dataset: https://project.mhzhou.com/vico.

Keywords: Listening Head Generation · Video Synthesis

1 Introduction

Communication [5,24,34,42,52,54] is one of the most common activities that everybody engages in their daily lives. During a face-to-face communication [29], two persons shift their roles in turn between the speaker and listener, to effectively exchange information. The speaker verbally transmits information to the listener, while the listener provides real-time feedbacks to the speaker mostly through non-verbal behaviors such as *affirmative nod*, *smiling*, *head shake*.

^{*} This work was done at JD Explore Academy.

[†] Equal contribution. § Corresponding author.



Fig. 1: Illustrations of three related tasks and our proposed responsive listening head generation. (a) Speech-to-gesture translation: generates plausible gestures to go along with the given speech. (b) Speech to lip generation: produces lip-synchronization in talking-head video. (c) Talking head generation: synthesizes talking face video conditioned on the identity of the speaker, audio speech, and/or the speaker emotion. (d) Our proposed responsive listening head synthesizes videos in responding to the speaker video stream

Although static images, repeated frames, or pre-scripted animations are often used to synthesize listeners in practice, they are often rigid and not realistic enough to respond to the speaker appropriately. According to studies in social psychology and anthropology, listening is a function-specific [18] and conditioned behavior [3], where learnable patterns can be inferred from training data. First, common patterns of listeners are observed to express their viewpoints, symmetrical and cyclic motions were employed to signal 'yes', 'no' or equivalents; narrow linear movements occurred in phase with stressed syllables in the other's speech; wide, linear movements occurred during pauses in the other's speech. Even the duration of eye blinks of the listener is perceived as communicative signals in human face-to-face interaction [23]. Second, these patterns in listener motions are mainly affected by two signals: the attitude of listener [21], and signals from the speaker [10,16,35]. Different attitudes of the listener results in diverse facial expressions, e.g., attitude of agree is meant by a nod and accept, attitude of disbelieve is represented by the combination of head tilt and frown. Meanwhile, listening behavior is heavily affected by speaker motion and audio signals. For example, the flow of movement of listener may be rhythmically coordinated with the speech and motions by the speaker [28]. These psychological and ethological studies motivate us to propose a data-driven method for modeling listening behaviors for face-to-face communication.

There have been extensive research efforts on speaker-centric synthesis. As shown in Fig. 1, speech to gesture generation [17] learns a mapping between the audio signal and speaker's pose. Speech to lip generation [45] aims to refine the lip-synchronization of a given video input. Talking-head synthesis [11,56,62] tries to generate a vivid talking video of a specific speaker with facial animations from a still image and a clip of audio. However, these works only focus on the speaking role, while ignore the indispensable counterpart of listener. Notably during a face-to-face conversation, listening behavior is even more important, as proper feedbacks to the speaker (*e.g., nod, smile, eye contact, etc.*) are vital for a successful communication [37,49,50,51]. Through real-time feedbacks, listeners show how they are engaged (*e.g.*, *interested*, *understand*, *agree*, *etc.*) to the speech, such that conversation gets more accessible for both participates.

In this work, we propose a new task to highlight listener-centric generation. Specifically, listening-head generation aims to synthesize a video of listening head, conditioning on the corresponding talking-head video of the speaker and the identity information of the listener, as shown in Fig. 1d. Proper reactions of the listener are expected to coordinate with the input talking video. This task is critical to a wide range of applications including virtual anchors, digital influencers, customer representatives, digital avatar in Metaverse, wherever involves interactive communication.

To address this, we construct a high-quality speaker-listener dataset, named ViCo, by capturing the high-definition video data from public conversations between two persons containing frontal faces on the same screen. The data strictly follows the principle that a video clip contains only uniquely identified listener and speaker, and requires that the listener has responsive non-verbal feedback to the speaker. After data cleaning, we further annotate the listener with three different attitudes: positive, neutral and negative. In total, our ViCo dataset contains 483 video clips of 76 listeners responding to 67 speakers. Compared to speaker-centric datasets such as MEAD [56] and VoxCeleb2 [12], ViCo highlights the listener role, making an indispensable couterpart to those speaker-centric ones. Compared to SEMAINE [36] (human interacts with a limited artificial agent) and MAHNOB Laughter [44] (people watching movies), ViCo features real persons in real conversations, such that natural reactions between genuine humans during a conversation make a key difference.

Together with the dataset, we propose a listening-head generation baseline method. We are aware that previous speaker-centric tasks are usually modeled in an idiosyncratic way (different speakers are modeled independently). However, listening behavior patterns are typically well coordinated with the speaker video. Thus we decouple the identity features from the listener and focus on learning the general motion patterns of responsive listening behaviors. We model listening head generation as a video-to-video translation task, by designing a sequenceto-sequence architecture to sequentially decode the listener's head motion and expression. Through quantitative evaluation and user study, we show our baseline is able to automatically capture the salient moments of speaker video and responds properly with clear motions and expressions.

2 Related Works

Active listener Active listening is an effective communication skill that not only means focusing fully on the speaker but also actively showing the nonverbal signals of listening with attitude. Usually, the active listener would mirror some facial expressions used by the speaker or shows more eye contact with the speaker. Active listening have shown its positive effects in many areas, such as teaching [26], medical consultations [15], team management [39], *etc.* In this paper, we aim to generate an active listener that could provide responsive feed-

back, the listener would understand the speaker's verbal and non-verbal signals first and then give proper feedback to the speaker.

Speaker-centered video synthesis Given time-varying signals and a reference still image of the speaker, the talking head synthesis task aims to generate a vivid clip for the speaker with the time-varying signals matched. Based on the different types of time-varying signals, we can group these tasks into two groups: 1) audio-driven talking head synthesis [11,45,60], 2) video-driven talking head synthesis [2,58]. The goal of the former one is to generate a video of the speaker that matches the audio. And the latter one is to generate videos of speakers with expressions similar to those in the video. This differs from our task: the "listener" is forced to perceive the speaker's visual and audio signals and make an active response. Our task does not focus on only a single person or transfers face expression and slight head movements from another person. There are two roles in our task: listener and speaker, and the listener should actively respond to the speaker with non-verbal signals.

Listening behaviors modeling Many applications and research papers have focused on speaking, while the "listener modeling" is seldomly explored. Gillies et al. [16] first propose the data-driven method that can generate an animated character that can respond to speaker's voice. This lacks the supervision of speaker visual signals, which is incomplete for responsive listener modeling. And this method can not be applied to realistic head synthesis. Heylen et al. [20] further studied the relationship between listener and speaker audio/visual signals from a cognitive technologies view. SEMAINE [36] records the conversation between a human and a limited artificial listener. MAHNOB Laughter database [44] focuses on studying laughter's behaviors when watching funny video clips. Apart from these related work, ALICO [8] corpus about active listener analysis is the most relevant dataset with our proposed task. However, it has not been made public and also not constructed from the real scene conversations. Moreover, the main objective of ALICO is for psychology analysis, the data mode of that dataset is vastly different from the audio-video corpora in computer vision area. In the past few years, the social AI intelligence [27,40] has been introduced to model the nonverbal social signals in triadic or multi-party interactions. Joo et al. [27] concerned with the overall posture and head movement of a person, and Oertel et al. [40] aims to mine listening motion rules for robotics controlling.

Dataset	Public	Environment	Style	Interact with Real
Gillies et al. [16]	×	Lab	Simulated	×
SEMAINE [36]	1	Lab	Simulated	1
Heylen et al. [20]	×	Lab	Simulated	1
MAHNOB Laughter [44]	1	Lab	Realistic	×
ALICO [8]	×	Lab	Realistic	1
Ours	1	Wild	Realistic	1

Table 1: Comparison with other listener-related datasets

Both related works only deal with the speaking status and ignore the speaker's content. What's more, they rarely care about two-person interactions nor pay attention to model the face in detail, which is also different from our task. A detailed comparison to exising listener-related datasets is shown in Tab. 1.

As far as we know, this is the first time to introduce the learning-based listening head generation task in computer vision area. In this work, we propose a formulation of responsive listening head generation and construct a public ViCo dataset for this task. Meanwhile, a baseline method is proposed for listening head synthesis by perceiving both speaker's audio/visual signals and preset attitude.

3 Task Overview

We present Responsive Listening Head Generation, a new task that challenges vision systems to generate listening heads actively responding to the speaker's face or/and audio in real-time. In particular, we need to understand the head motion, facial expression, including eye blinks, mouth movements, *etc.*, of the input speaker video frame, and simultaneously understand the speaker's voice, then synchronously generate the active listening face video conditioned by the given attitude.

Given an input video sequence $\mathcal{V}_t^s = \{v_1^s, \cdots, v_t^s\}$ of a speaker head in time stamps ranging from $\{1, ..., t\}$, and an corresponding audio signal sequence $\mathcal{A}_t^s = \{a_1, \cdots, a_t\}$ of the speaker, listening head generation aims to generate a listener's head v_{t+1}^l of the next time stamp:

$$v_{t+1}^l = \mathbf{G}(\mathcal{V}_t^s, \mathcal{A}_t^s, v_1^l, e), \tag{1}$$

where v_1^l is the reference head of the listener, e denotes the attitude of the listener. The whole generated listener video \mathcal{V}_{t+1}^l can be denoted as the concatenation of $\{v_2^l, \cdots, v_{t+1}^l\}$.

Listening attitude definition During conversation, after perceiving the signals from the speaker, the listener usually reacts with an active, responsive *attitude*, including epistemic attitudes (*e.g.*, agree, disagree) and affective attitudes (*e.g.*, like, dislike). In this work, we group the attitudes into three categories: positive, negative and neutral. Positive attitude



Fig. 2: During a conversation, different attitudes of the listener could show different pose and expression patterns

consists of agree, like, interested. Conversely, negative attitude consists of disagree, dislike, disbelieve, not interested. In general, attitude potentially guides the listener's behavior and consequently affects the conversation. Also, different attitude results in different facial expressions and behaviors of the listener [21], e.g. a smile appears as the most appropriate signal for like, a combination of smile and raise eyebrows could be a possibility for interested, disagree can be

meant by a head shake, *dislike* is represented by a frown and tension of the lips, *etc.* A listener example with different attitudes is illustrated in Fig. 2.

Feature extraction In this work, we extract the energy feature, temporal domain feature, and frequency domain feature of the input audio; and model the facial expression and head poses using 3DMM [6] coefficients.

For the audio, we extract the Mel-frequency cepstral coefficients (MFCC) feature with the corresponding MFCC Delta and Delta-Delta feature. Besides, the energy, loudness and zero-crossing rate (ZCR) are also embedded into audio features s_i for each audio clip a_i . The audio feature extracted from \mathcal{A}_t^s can be denoted as $\mathcal{S}_t^s = \{s_1, \dots, s_t\}$.

We leverage the state-of-the-art deep learning-based 3D face reconstruction model [14] for the videos to get the 3DMM [6] coefficients. Specially, for each image, we can get the reconstruction coefficients { $\alpha, \beta, \delta, p, \gamma$ } which denote the identity, expression, texture [9,43], pose and lighting [46], respectively. Further, we distinguished the 3D reconstruction coefficients into two parts: $\mathcal{I} = (\alpha, \delta, \gamma)$ to represent relatively fixed, identity-dependent features, and $m = (\beta, p)$ to represent relatively dynamic, identity-independent features. The identity-independent feature extracted from speaker videos can be denoted as $\mathcal{M}_t^s = \{m_1^s, \cdots, m_t^s\}$, where $m_i^s \in \mathbb{R}^{1 \times C_v}$ is the expression and pose feature of 3D reconstruction coefficients for the *i*-th frame v_i^s , where $C_v = |\beta| + |p|$.

Task definition To ignore identity-dependent features and learn general listener patterns that can be adapted to multiple listener identities, we use only the head motion and facial expression feature m for responsive listening head generation model training, and then adapt the identity-dependent features \mathcal{I} of different listener identities for visualization and evaluation. Thus our listening head synthesis task can be formulated as:

$$m_{t+1}^{l} = \mathbf{G}_{m}(\mathcal{M}_{t}^{s}, \mathcal{S}_{t}^{s}, m_{1}^{l}, e),$$

$$v_{t+1}^{l} = \mathbf{G}_{v}(m_{t+1}^{l}, \mathcal{I}^{l}, v_{1}^{l}),$$
(2)

where m_{t+1}^l is the dynamic feature predicted for listener's head, and \mathcal{I}^l denotes the identity-dependent features of the given listener. In a real implementation, we use T frame of speaker's audio and video for responsive listening head generation model training.

The 3D face rendering technology \mathbf{G}_v has been well studied in many recent related works [30,60]. Moreover, the face rendering models are usually identityspecific, so one may need to train the rendering model separately for each identity for better performance. To highlight the properties of the interactive digital human synthesis task, and decouple the critical factor in this task, our proposed responsive listening head synthesis model primarily focuses on the motion-related and identity-independent 3D facial coefficients prediction task \mathbf{G}_m , and use the pretrained rendering model [47] for simplified visualization.



Fig. 3: In ViCo, valid clips are selected in accordance with the standards that 1) both the speaker and listener behaviors are clearly visible, and 2) listeners are responsively engaged to the conversation. The facial regions of listener-speaker pairs are further cropped for constructing our ViCo dataset (right)

Table 2: Statistics of ViCo dataset. **#**ID indicates the number of identities. The same person/identity can play different roles with multiple attitudes

Attitude	$\# {\rm Videos}$	$\# {\rm Speaker}$	$\# {\rm Listener}$	#ID	# Clips	Duration
Positive	42	53	62	81	226	49 min 18 s
Neutral	35	38	48	63	134	$27\mathrm{min}$ 7 s
Negative	11	11	9	18	123	$18\min57\mathrm{s}$
Total	50	67	76	92	483	$95 \min 22 s$

4 Dataset Construction

We construct a dataset for responsive listening-head generation by capturing conversational video clips from YouTube containing two people's frontal faces. A *valid* video clip is required to meet the following conditions:

- The screen contains only two people, and one of them is speaking while the other is listening carefully.
- The frontal faces of both people are clearly visible. The facial expression is natural and stable.
- The listener actively responds to the speaker in a dynamic and real-time manner.

The annotators were asked to accurately record the start and end time of each *valid* clip, label the position of the speaker (left or right of the screen) and identify the attitude of the listener in the video. Cross-validation was applied among at least three annotators for each candidate clip for quality control. For each valid clip, we use the MTCNN [61] to detect the face regions in each frame, and then crop and resize the detected face regions to 384×384 resolution image sequence for model training and evaluation, as shown in Fig. 3.

Table 2 shows the statistical information of our annotated responsive listening head generation dataset ViCo. The proposed dataset contains rich samples of 483 video clips. We normalize all videos to 30 FPS, forming more than 0.1 million frames in total. Moreover, our dataset has following properties:



Fig. 4: The overall pipeline of our responsive listening head generation baseline. The speaker encoder aims to encode the head motion, facial expression and audio features. Starting from the fused feature from reference listener image, the listener decoder receives signals from speaker encoder in temporal order, and predicts the head motion and facial expression features. These features are adapted to reconstruct the 3DMM coefficients with the reference listener's identity-dependent features, and then fed to a neural renderer to generate realistic listening video

High quality All raw videos are of high resolutions (1920×1080) , so that the subtle differences between different attitudes and changing moods are well preserved. And audios are in 44.1 kHz/16 bit such that the speech-related features can be well preserved, too.

High diversity Our dataset contains various scenarios, including news interviews, entertainment interviews, TED discussions, variety shows, *etc.* These diverse scenarios provide rich semantic information and various listener patterns in different situations. The video clips length ranges from 1 to 71 seconds.

Realtime and interactivity Different from the existing talking head video datasets [1,12,32,56,59] which aim to generate head or face in synchronization with the audio signals, our dataset focus on the face-to-face *response*. These responses are generated by jointly understanding the speaker's audio, facial, and head motion signals, then adapting to different listener heads. It matters about mutual interaction rather than a monologue.

5 Responsive Listening Head Generation

Based on ViCo dataset, we propose a responsive listening head generation baseline. The overview of our approach is illustrated in Fig. 4.

5.1 Model architecture

According to the psychological knowledge, an active listener tends to respond based on speaker's audio [28] and visual signals [10,16,35] comprehensively. And at a given moment, the listener receives information from the speaker of that moment as well as information from history and adopts a certain attitude to present actions in response to the speaker. Thus, the goal of our model is to estimate the conditioned probability $P(\mathcal{M}_{t+1}^{l}|\mathcal{M}_{t}^{s}, \mathcal{S}_{t}^{s}, m_{1}^{l}, e)$, where the \mathcal{M}_{t}^{s} and \mathcal{S}_{t}^{s} are time-varying signals that the listener should respond to, and the reference listener feature m_{1}^{l} and attitude e constrain the pattern of the entire generated sequence.

Inspired by the sequence-to-sequence model [53], a multi-layer sequential decoder module \mathbf{G}_m is applied for modeling the time-sequential information of conversation. Unlike talking-head generation [11,56,60,62], which accepts an entire input of audio and then processes it using a bidirectional LSTM or attention layer; in our scenario, the model \mathbf{G}_m receives the streaming input of the speaker where future information is not available.

For the speaker feature encoder, at each time step t, we first extract the audio feature s_t and the speaker's head and facial expression representation m_t^s , then apply non-linear feature transformations following a multi-modal feature fusion function f_{am} to get the encoded feature of speaker. The representation of reference listener m_1^l and attitude e can be embedded as the initial state h_1 for the sequential motion decoder. At each time step t, taking the speaker's fused feature $f_{am}(s_t, m_t^s)$ as input, \mathbf{G}_m in Eq. 2 is functioned as updating current state h_{t+1} and generating the listener motion m_{t+1}^l , which contains two feature vectors, *i.e.* β_{t+1}^l for the expression and p_{t+1}^l for the head rotation and translation. Our responsive listening head generator supports an arbitrary length of speaker input. The procedure can be formulated as:

$$\beta_{t+1}^l, p_{t+1}^l = \mathbf{G}_m(h_t, f_{am}(s_t, m_t^s)).$$
(3)

For optimization, with the ground truth listener patterns denoted as $\hat{\mathcal{M}}_T^l = [\hat{m}_2^l, \hat{m}_3^l, \cdots, \hat{m}_T^l]$, we drop the last prediction m_{T+1}^l due to the lack of supervision signals and use L_2 distance to optimize the training procedure:

$$\mathcal{L}_{gen} = \sum_{t=2}^{T} \|\beta_t^l - \hat{\beta}_t^l\|_2 + \|p_t^l - \hat{p}_t^l\|_2.$$
(4)

Moreover, a motion constraint loss \mathcal{L}_{mot} is applied to guarantee the interframe continuity across $\hat{\mathcal{M}}_T^l$ is similar to the predicted \mathcal{M}_T^l :

$$\mathcal{L}_{mot} = \sum_{t=2}^{T} w_1 \|\mu(\beta_t^l) - \mu(\hat{\beta}_t^l)\|_2 + w_2 \|\mu(p_t^l) - \mu(\hat{p}_t^l)\|_2,$$
(5)

where $\mu(\cdot)$ measures the inter-frame changes of current frame and its adjacent previous frame, i.e., $\mu(\beta_t^l) = \beta_t^l - \beta_{t-1}^l$, w_1 and w_2 is a weight to balance the motion constraint loss and generation loss. The final loss function of our proposed listening head motion generation baseline can be formulated as:

$$\mathcal{L}_{total} = \mathcal{L}_{gen} + \mathcal{L}_{mot}.$$
 (6)

By optimizing \mathcal{L}_{total} , our model can generate attitude conditioned responsive listening head for a given speaker video and audio.

5.2 Implementation details

To verify that our model can learn a generic listening pattern rather than conditioning on any particular individual, we divide the ViCo dataset (\mathcal{D}) into three parts: i) training set \mathcal{D}_{train} for learning listener patterns, ii) test set \mathcal{D}_{test} for validating our model on in-domain data, and iii) out-of-domain (OOD) test set \mathcal{D}_{ood} for evaluating the generalization and transferability. In this case, all identities in \mathcal{D}_{test} have appeared in \mathcal{D}_{train} , while identities in the \mathcal{D}_{ood} do not overlap with those in \mathcal{D}_{train} .

We extract 45-dimensional acoustic features for audios, including 14-dim MFCC, 28-dim MFCC-Delta, energy, ZCR and loudness. There are multiple choices to implement \mathbf{G}_m , such as standard sequential model like LSTM [22], GRU [13], or a Transformer [55] decoder with sliding window [4]. Here we adopt LSTM for our baseline, since it has been widely used in many similar applications such as motion generation [48], and achieve stable state-of-the-art performance when training on small corpus [38]. Our listening head generation model is trained with AdamW [33] optimizer with a learning rate of 1×10^{-3} (decayed exponentially by 0.8 every 30 epochs), $\beta_1 = 0.9$ and $\beta_2 = 0.999$, for 300 epochs. For all experiments, we set hyper-parameter w_1 to 0.1 and w_2 to 1×10^{-4} .

5.3 Experimental results

Quantitative results Since we use a detached renderer rather than an endto-end pipeline, we can divide the assessment into two sides: the performance of listener generator \mathbf{G}_m and the visual effects of renderer \mathbf{G}_v . For the former one, we use use L_1 distance between the generated features and the groundtruth features (FD) to ensure the predicted fine-grained head and expression coefficients similar to the ground-truth. And for the latter one, we select the Fréchet Inception Distance (FID) [19], Structural SIMilary (SSIM) [57], Peak Signal-to-Noise Ratio (PSNR) and Cumulative Probability of Blur Detection (CPBD) [7] to evaluate the visual effects of renderer. The high-level metric on \mathbf{G}_m can help us analyze the model, and the low-level metrics on \mathbf{G}_v provide a baseline for the successors.

In Table 3, we report the FD of the 3D facial coefficients across different listening head generation methods, including: 1) "Random": generate frames from reference image but injecting small perturbations in a normal distribution to mimic random head motion. 2) "Simulation": simulating natural listening behavior by repeating the motion patterns sampled from \mathcal{D}_{train} . 3) "Simulation*": repeating the natural listening motion patterns sampled from \mathcal{D}_{train} with the corresponding attitude. 4) "Ours": our proposed responsive listening head generation method. The random permutations is the worst and not able to present a listener. Our method can reach the best performance which demonstrates the superiority of our algorithm over traditional non-parametric listeners. Also, we provide the evaluation results of zero-shot 3D face rendering [47] in Table 4, as a basic criterion for evaluation of future realistic face rendering research work.

Attitude Motio	Motion	Random		Simulation		Simulation*		Ours	
	MOUIOII	\mathcal{D}_{test}	\mathcal{D}_{ood}	\mathcal{D}_{test}	\mathcal{D}_{ood}	\mathcal{D}_{test}	\mathcal{D}_{ood}	\mathcal{D}_{test}	\mathcal{D}_{ood}
Positive	angle	17.92	17.99	9.86	10.48	9.79	11.57	6.79	9.72
	\exp	44.71	44.86	27.08	27.81	30.00	30.27	15.37	24.89
	trans	19.74	20.14	16.25	12.06	9.07	13.82	6.48	9.51
Neutral	angle	17.85	17.78	10.94	9.47	14.18	8.94	8.79	6.33
	\exp	44.26	44.29	27.37	29.50	26.44	27.56	13.61	23.51
	trans	19.98	20.17	8.47	12.27	11.53	9.40	6.68	8.95
Negative	angle	19.53	18.70	17.86	9.66	13.24	18.75	12.45	8.54
	\exp	45.68	44.62	29.57	28.78	31.62	27.04	16.98	18.99
	trans	19.69	20.92	8.06	10.69	24.42	11.09	6.35	5.81
Average	angle	18.04	18.11	10.81	9.91	11.24	12.58	7.79	8.23
	\exp	44.67	44.60	27.37	28.66	29.20	28.46	15.04	22.83
	trans	19.80	20.36	13.52	11.76	11.00	11.55	6.52	8.32

Table 3: The Feature Distance $(\times 100)$ of different listening head generation methods. Each cell in the table represent the feature distance of angle / expression / translation coefficients respectively. Lower is better

Table 4: Quantitative Results of Renderer \mathbf{G}_v on \mathcal{D}_{test} and \mathcal{D}_{ood}

Attitude	FID↓		$\mathrm{SSIM}\uparrow$		$\mathrm{PSNR}\uparrow$		CPBD↑	
	\mathcal{D}_{test}	\mathcal{D}_{ood}	\mathcal{D}_{test}	\mathcal{D}_{ood}	\mathcal{D}_{test}	\mathcal{D}_{ood}	\mathcal{D}_{test}	\mathcal{D}_{ood}
Positive	29.736	27.865	0.565	0.496	17.075	15.421	0.122	0.121
Neutral	36.551	27.366	0.686	0.544	21.220	17.703	0.106	0.113
Negative	46.277	28.406	0.610	0.528	16.870	16.709	0.219	0.211
Average	30.529	24.962	0.601	0.521	18.149	16.558	0.126	0.142

Qualitative results Further, we visualize the results of those generations to analyze the differences between different configurations intuitively. Given \mathcal{D}_{test} and \mathcal{D}_{ood} , we generate listener videos with a given attitude. We randomly select two sequences from each set and then down-sample to six frames to qualitatively visualize the generated results and the ground truth video. The results are shown in Fig. 5, we can find that our model is generally able to capture listener patterns (*e.g., eye, mouth, and head motion, etc.*), which may differ from the ground-truth while still making sense.

From these two groups, we observe that the "Random" patterns behave very confusing and messy. The "Simulation" patterns depend heavily on whether we can randomize to a given attitude, and as long as this fails, the result is bad. The "Simulation*" performs a little bit better, while its motion is also limited by the size and diversity of dataset, and intuitively, it cannot respond to the speaker in a dynamic manner. Our results are more visually plausible than others with head motions and expression changes.



Fig. 5: Qualitative comparison of listening head generation methods conditioned by the same reference frame (the first column of each group) and the same attitude (left: positive listener in \mathcal{D}_{test} , right: neutral listener in \mathcal{D}_{ood}). Our method can generate various, vivid and responsive listening motions for the given speaker video stream

We also provide the comparison results of listening head generation under different attitudes in Fig. 6. Obviously, the facial expression and head motions under different attitudes are expressive and distinguishable.

Ablation Studies We conduct an ablation study on the impact of different speaker signals for listening head modeling on \mathcal{D}_{test} . As Tab. 5 shown, the listening head driven by audio-only inputs prefers expression modeling but performs badly in head motion (angle, trans), while the model with visual-only inputs would capture the motion or sightline changes of the speaker during conversation and provide reasonable listening head action response. Removing these two modalities and simply mirror the speaker causes the worst listener. And with these two joint input signals, our model exhibits the best performance. That is, only when we look and hear, can we act as better responsive listeners.

Besides, we also report the performance of listener generation by giving wrong audio inputs and correct visual inputs in Tab. 5. It shows that Audio Only < Wrong Audio < Visual Only for expression FD, while Visual Only \approx Wrong Audio < Audio Only for head motion FD (lower is better). This also reveals that expression modeling depends more on audio inputs and head motion modeling is more corresponding to visual inputs.

Table 5: The averaged Feature Distance $(\times 100)$ of listener generations across all attitudes on \mathcal{D}_{test} .

Method	angle	\exp	trans
Audio only	8.69	17.19	8.49
Visual only	8.06	18.85	7.54
Wrong Audio	8.15	18.02	7.60
Mirroring	9.99	26.48	11.39
Ours	7.79	15.04	6.52



Fig. 6: Comparative results of our listening heads generation method conditioned by the same reference image and speaker video but different attitudes



Fig. 7: Diverse visual patterns can be observed from the generated listeningheads. Left: diverse expressions can be generated corresponding to different attitudes. Right: motion patterns including *shaking head*, *nod*, *glare*, *looking askance*, *pressing lip* and *focusing*

Diversity of Generations Fig. 7 illustrates the diversity of visual patterns learned by our method. The single images on the left side show our model can generate different expressions while the six groups of images on the right demonstrate the head motions and eye contacts can be modeled by our method.

Runtime Complexity With a given speaker's streaming video, it takes 52.5 ms to fit the 3DMM coefficients, and generate the next step listener's motion in 0.0372 ms, then render the motion to RGB images in 29.3 ms. The per-frame delay of this process on one Tesla-V100 GPU is 81.84 ms without any optimization strategy such as ONNX or TensorRT. The FPS of generated listening head videos can reach 12, and can be further improved to 19 with pipeline parallelism. The bottleneck of our proposed baseline method is the efficiency of 3D face reconstruction (fitting). Some video post-processing methods, such as video frame interpolation [25,31,41] can be used for real-time interaction.

User study Since responsive listening head modeling is a user-oriented task, it is essential to conduct user studies for evaluation. For \mathcal{D}_{test} and \mathcal{D}_{ood} , we had 10 volunteers doing two double-blind tests: 1) **Preference Test (PT)**. Given the shuffled tuple of \langle ground-truth listening head, generated listening head \rangle , along with the ground-truth attitude, speaker's audio and speaker's video, the volunteers were asked to pick the best listening head. 2) Attitude Matching Test (AttMatch). Given the randomly shuffled listening heads gen-

Table 6: (a) The result (mean / variance) of preference test. "Equal": the generated results and the ground-truth are visually equivalent. "PD Better": the generated results are more in line with human perception. (b) Mean precision and precision variance of attitude matching test

(a) Preference Test			(b) Atti	(b) Attitude Matching Test				
PT	\mathcal{D}_{test}	\mathcal{D}_{ood}	AttMatch	\mathcal{D}_{test}	\mathcal{D}_{ood}			
Equal	31.3 / 17.1	13.9 / 6.2	Positive	66.7 / 15.7	69.0 / 7.1			
GT Better	56.2 / 14.9	65.3 / 6.1	Negative	63.9 / 16.4	50.0 / 12.8			
PD Better	12.5 / 9.3	20.8 / 4.6	Neutral	69.4 / 9.2	53.6 / 3.9			

erations conditioned by three different attitudes for each identity, the volunteers were asked to identify one attitude for each listening head.

For the first user-study, each volunteer was asked to check 30 samples in a preference test. As shown in Tab. 6a, fpr \mathcal{D}_{test} and \mathcal{D}_{ood} , volunteers voted that nearly 43.8% and 44.7% of the generated heads can get equal or even better rating than the ground-truth heads respectively, which verified that our model could generate responsive listeners consistent with subjective human perceptions, and even can be reached confused as real ones.

Our second user-study results are shown in Tab. 6b. Each volunteer was given 90 generated listening clips under three different attitudes. For each attitude, we calculate the mean precision and precision variance across all volunteers for analysis. The results show that both in \mathcal{D}_{test} and \mathcal{D}_{ood} , our model is capable of generating listening head yielding to the required attitude. The precision for negative and neutral attitude is slightly decreased in \mathcal{D}_{ood} , which might be caused by the unbalanced attitude distributions.

6 Conclusion

In this paper, we define the responsive listening head generation task. It aims to generate a responsive video clip for a listener with the understanding of the speaker's facial signals and voices. Further, the high-quality responsive listener dataset (ViCo) is contributed for addressing this problem. The responsive listener generation baseline can synthesis active listeners, which are more consistent with human perception. We expect that ViCo could benefit the face-to-face communication modeling in computer vision area and facilitate the applications in more scenarios, such as intelligence assistance, virtual human, *etc*.

Ethical Impact The ViCo dataset will be released only for research purposes under restricted licenses. The responsive listening patterns are identity-independent, which reduces the abuse of facial data. The only potential social harm is "fake content". However, different from talking head synthesis, responsive listening can hardly harm information fidelity.

Acknowledgement This work was supported by the National Key R&D Program of China under Grant No. 2020AA0108600.

Responsive Listening Head Generation: A Benchmark Dataset and Baseline

References

- 1. Afouras, T., Chung, J.S., Zisserman, A.: Lrs3-ted: a large-scale dataset for visual speech recognition. arXiv preprint arXiv:1809.00496 (2018)
- Bansal, A., Ma, S., Ramanan, D., Sheikh, Y.: Recycle-gan: Unsupervised video retargeting. In: Proceedings of the European conference on computer vision (ECCV). pp. 119–135 (2018)
- 3. Barker, L.L.: Listening behavior. (1971)
- 4. Beltagy, I., Peters, M.E., Cohan, A.: Longformer: The long-document transformer. arXiv preprint arXiv:2004.05150 (2020)
- 5. Berger, C.R.: Interpersonal communication: Theoretical perspectives, future prospects. Journal of communication (2005)
- Blanz, V., Vetter, T.: A morphable model for the synthesis of 3d faces. In: Proceedings of the 26th annual conference on Computer graphics and interactive techniques. pp. 187–194 (1999)
- Bohr, P., Gargote, R., Vhorkate, R., Yawle, R., Bairagi, V.: A no reference image blur detection using cumulative probability blur detection (cpbd) metric. International Journal of Science and Modern Engineering 1(5) (2013)
- Buschmeier, H., Malisz, Z., Skubisz, J., Wlodarczak, M., Wachsmuth, I., Kopp, S., Wagner, P.: Alico: A multimodal corpus for the study of active listening. In: LREC 2014, Ninth International Conference on Language Resources and Evaluation, 26-31 May, Reykjavik, Iceland. pp. 3638–3643 (2014)
- Cao, C., Weng, Y., Zhou, S., Tong, Y., Zhou, K.: Facewarehouse: A 3d facial expression database for visual computing. IEEE Transactions on Visualization and Computer Graphics 20(3), 413–425 (2013)
- 10. Cassel, N.N.W.W.: Elements of face-to-face conversation for embodied conversational agents, embodied conversational agents (2000)
- 11. Chung, J.S., Jamaludin, A., Zisserman, A.: You said that? arXiv preprint arXiv:1705.02966 (2017)
- Chung, J.S., Nagrani, A., Zisserman, A.: Voxceleb2: Deep speaker recognition. arXiv preprint arXiv:1806.05622 (2018)
- 13. Chung, J., Gulcehre, C., Cho, K., Bengio, Y.: Empirical evaluation of gated recurrent neural networks on sequence modeling. arXiv preprint arXiv:1412.3555 (2014)
- Deng, Y., Yang, J., Xu, S., Chen, D., Jia, Y., Tong, X.: Accurate 3d face reconstruction with weakly-supervised learning: From single image to image set. In: IEEE Computer Vision and Pattern Recognition Workshops (2019)
- Fassaert, T., van Dulmen, S., Schellevis, F., Bensing, J.: Active listening in medical consultations: Development of the active listening observation scale (alos-global). Patient education and counseling 68(3), 258–264 (2007)
- Gillies, M., Pan, X., Slater, M., Shawe-Taylor, J.: Responsive listening behavior. Computer animation and virtual worlds 19(5), 579–589 (2008)
- Ginosar, S., Bar, A., Kohavi, G., Chan, C., Owens, A., Malik, J.: Learning individual styles of conversational gesture. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3497–3506 (2019)
- Hadar, U., Steiner, T.J., Rose, F.C.: Head movement during listening turns in conversation. Journal of Nonverbal Behavior 9(4), 214–228 (1985)
- Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. Advances in neural information processing systems **30** (2017)

- 16 M. Zhou et al.
- Heylen, D., Bevacqua, E., Pelachaud, C., Poggi, I., Gratch, J., Schröder, M.: Generating listening behaviour. In: Emotion-oriented systems, pp. 321–347. Springer (2011)
- Heylen, D., Bevacqua, E., Tellier, M., Pelachaud, C.: Searching for prototypical facial feedback signals. In: International Workshop on Intelligent Virtual Agents. pp. 147–153. Springer (2007)
- Hochreiter, S., Schmidhuber, J.: Long short-term memory. Neural computation 9(8), 1735–1780 (1997)
- 23. Hömke, P., Holler, J., Levinson, S.C.: Eye blinks are perceived as communicative signals in human face-to-face interaction. PloS one **13**(12), e0208030 (2018)
- Honeycutt, J.M., Ford, S.G.: Mental imagery and intrapersonal communication: A review of research on imagined interactions (iis) and current developments. Annals of the International Communication Association 25(1), 315–345 (2001)
- 25. Huang, Z., Zhang, T., Heng, W., Shi, B., Zhou, S.: Real-time intermediate flow estimation for video frame interpolation. In: Proceedings of the European Conference on Computer Vision (ECCV) (2022)
- Jalongo, M.R.: Promoting active listening in the classroom. Childhood Education 72(1), 13–18 (1995)
- Joo, H., Simon, T., Cikara, M., Sheikh, Y.: Towards social artificial intelligence: Nonverbal social signal prediction in a triadic interaction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10873– 10883 (2019)
- Kendon, A.: Movement coordination in social interaction: Some examples described. Acta psychologica 32, 101–125 (1970)
- Kendon, A., Harris, R.M., Key, M.R.: Organization of behavior in face-to-face interaction. Walter de Gruyter (2011)
- 30. Kim, H., Garrido, P., Tewari, A., Xu, W., Thies, J., Niessner, M., Pérez, P., Richardt, C., Zollhöfer, M., Theobalt, C.: Deep video portraits. ACM Transactions on Graphics (TOG) 37(4), 1–14 (2018)
- Kong, L., Jiang, B., Luo, D., Chu, W., Huang, X., Tai, Y., Wang, C., Yang, J.: Ifrnet: Intermediate feature refine network for efficient frame interpolation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1969–1978 (2022)
- 32. Li, L., Wang, S., Zhang, Z., Ding, Y., Zheng, Y., Yu, X., Fan, C.: Write-a-speaker: Text-based emotional and rhythmic talking-head generation. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 35, pp. 1911–1920 (2021)
- Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. arXiv preprint arXiv:1711.05101 (2017)
- Luhmann, N.: What is communication? Communication theory 2(3), 251–259 (1992)
- 35. Maatman, R., Gratch, J., Marsella, S.: Natural behavior of a listening agent. In: International workshop on intelligent virtual agents. pp. 25–36. Springer (2005)
- McKeown, G., Valstar, M., Cowie, R., Pantic, M., Schroder, M.: The semaine database: Annotated multimodal records of emotionally colored conversations between a person and a limited agent. IEEE transactions on affective computing 3(1), 5–17 (2011)
- 37. McNaughton, D., Hamlin, D., McCarthy, J., Head-Reeves, D., Schreiner, M.: Learning to listen: Teaching an active listening strategy to preservice education professionals. Topics in Early Childhood Special Education 27(4), 223–231 (2008)
- Melis, G., Kočiský, T., Blunsom, P.: Mogrifier lstm. arXiv preprint arXiv:1909.01792 (2019)

Responsive Listening Head Generation: A Benchmark Dataset and Baseline

- Mineyama, S., Tsutsumi, A., Takao, S., Nishiuchi, K., Kawakami, N.: Supervisors' attitudes and skills for active listening with regard to working conditions and psychological stress reactions among subordinate workers. Journal of occupational health 49(2), 81–87 (2007)
- 40. Oertel, C., Jonell, P., Kontogiorgos, D., Mora, K.F., Odobez, J.M., Gustafson, J.: Towards an engagement-aware attentive artificial listener for multi-party interactions. Frontiers in Robotics and AI p. 189 (2021)
- Park, J., Lee, C., Kim, C.S.: Asymmetric bilateral motion estimation for video frame interpolation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 14539–14548 (2021)
- 42. Parker, J., Coiera, E.: Improving clinical communication: a view from psychology. Journal of the American Medical Informatics Association **7**(5), 453–461 (2000)
- Paysan, P., Knothe, R., Amberg, B., Romdhani, S., Vetter, T.: A 3d face model for pose and illumination invariant face recognition. In: 2009 sixth IEEE international conference on advanced video and signal based surveillance. pp. 296–301. Ieee (2009)
- Petridis, S., Martinez, B., Pantic, M.: The mahnob laughter database. Image and Vision Computing 31(2), 186–202 (2013)
- 45. Prajwal, K., Mukhopadhyay, R., Namboodiri, V.P., Jawahar, C.: A lip sync expert is all you need for speech to lip generation in the wild. In: Proceedings of the 28th ACM International Conference on Multimedia. pp. 484–492 (2020)
- 46. Ramamoorthi, R., Hanrahan, P.: An efficient representation for irradiance environment maps. In: Proceedings of the 28th annual conference on Computer graphics and interactive techniques. pp. 497–500 (2001)
- Ren, Y., Li, G., Chen, Y., Li, T.H., Liu, S.: Pirenderer: Controllable portrait image generation via semantic neural rendering. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 13759–13768 (2021)
- Richard, A., Zollhöfer, M., Wen, Y., De la Torre, F., Sheikh, Y.: Meshtalk: 3d face animation from speech using cross-modality disentanglement. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 1173–1182 (2021)
- 49. Robertson, K.: Active listening: more than just paying attention. Australian family physician **34**(12) (2005)
- 50. Rogers, C.R., Farson, R.E.: Active listening (1957)
- 51. Rost, M., Wilson, J.: Active listening. Routledge (2013)
- 52. Stacks, D.W., Salwen, M.B.: An integrated approach to communication theory and research. Routledge (2014)
- Sutskever, I., Vinyals, O., Le, Q.V.: Sequence to sequence learning with neural networks. In: Advances in neural information processing systems. pp. 3104–3112 (2014)
- 54. Tomasello, M.: Origins of human communication. MIT press (2010)
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. Advances in neural information processing systems **30** (2017)
- Wang, K., Wu, Q., Song, L., Yang, Z., Wu, W., Qian, C., He, R., Qiao, Y., Loy, C.C.: Mead: A large-scale audio-visual dataset for emotional talking-face generation. In: European Conference on Computer Vision. pp. 700–717. Springer (2020)
- Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. IEEE transactions on image processing 13(4), 600–612 (2004)

- 18 M. Zhou et al.
- 58. Wu, W., Zhang, Y., Li, C., Qian, C., Loy, C.C.: Reenactgan: Learning to reenact faces via boundary transfer. In: Proceedings of the European conference on computer vision (ECCV). pp. 603–619 (2018)
- Zhang, C., Ni, S., Fan, Z., Li, H., Zeng, M., Budagavi, M., Guo, X.: 3d talking face with personalized pose dynamics. IEEE Transactions on Visualization and Computer Graphics (2021)
- Zhang, C., Zhao, Y., Huang, Y., Zeng, M., Ni, S., Budagavi, M., Guo, X.: Facial: Synthesizing dynamic talking face with implicit attribute learning. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3867–3876 (2021)
- Zhang, K., Zhang, Z., Li, Z., Qiao, Y.: Joint face detection and alignment using multitask cascaded convolutional networks. IEEE Signal Processing Letters 23(10), 1499–1503 (2016)
- 62. Zhu, H., Luo, M.D., Wang, R., Zheng, A.H., He, R.: Deep audio-visual learning: A survey. International Journal of Automation and Computing pp. 1–26 (2021)