



# Speech-Driven Personalized Gesture Synthetics: Harnessing Automatic Fuzzy Feature Inference

Fan Zhang , Zhaohan Wang , Xin Lyu , Siyuan Zhao , Mengjian Li , Weidong Geng , Naye Ji , Hui Du ,  
Fuxing Gao , Hao Wu , and Shunman Li 

**Abstract**—Speech-driven gesture generation is an emerging field within virtual human creation. However, a significant challenge lies in accurately determining and processing the multitude of input features (such as acoustic, semantic, emotional, personality, and even subtle unknown features). Traditional approaches, reliant on various explicit feature inputs and complex multimodal processing, constrain the expressiveness of resulting gestures and limit their applicability. To address these challenges, we present *Persona-Gestor*, a novel end-to-end generative model designed to generate highly personalized 3D full-body gestures solely relying on raw speech audio. The model combines a fuzzy feature extractor and a non-autoregressive Adaptive Layer Normalization (AdaLN) transformer diffusion architecture (DiTs-based). The fuzzy feature extractor harnesses a fuzzy inference strategy that automatically infers implicit, continuous fuzzy features. These fuzzy features, represented as a unified latent feature, are fed into the AdaLN

transformer. The AdaLN transformer introduces a conditional mechanism that applies a uniform function across all tokens, thereby effectively modeling the correlation between the fuzzy features and the gesture sequence. This module ensures a high level of gesture-speech synchronization while preserving naturalness. Finally, we employ the diffusion model to train and infer various gestures. Extensive subjective and objective evaluations on the Trinity, ZEGGS, and BEAT datasets confirm our model's superior performance to the current state-of-the-art approaches. *Persona-Gestor* improves the system's usability and generalization capabilities, setting a new benchmark in speech-driven gesture synthesis and broadening the horizon for virtual human technology.

**Index Terms**—Speech-driven, gesture synthesis, fuzzy inference, AdaLN, diffusion, transformer, DiTs.

## I. INTRODUCTION

RECENT advancements have significantly expanded the use of 3D virtual human technology. Its growing appeal spans numerous applications, including animation, gaming, digital receptionists, and human-computer interaction. A major task in this research area is to create credible, personalized co-speech gestures. Speech-driven gesture generation through deep learning provides a cost-effective solution, eliminating the need for manual intervention associated with conventional motion capture systems.

However, the primary challenge in speech-driven gesture generation is precisely identifying the vast array of input conditions necessary for driving gesture synthesis. This complexity arises because an extensive range of factors, including acoustics, semantics, emotions, personality traits, and demographic variables like gender, age, etc., shape co-speech gestures.

Previous approaches [1], [2], [3], [4], [5], [6], [7] have explored the use of manual labels and diverse feature inputs to facilitate the synthesis of personalized gestures. Nonetheless, these methods depend heavily on various unstructured feature inputs and require complex multimodal processing. These approaches present a significant barrier to the practical application and broader adoption of virtual human technologies.

The fuzzy inference strategy, which pertains to the concept of fuzzy logic [8], is particularly useful in the field for dealing with uncertain or imprecise information. The fuzzy inference strategy is known for its effectiveness in speech-emotion recognition [9] and audio classification [10]. These methods do not necessarily require explicit classification outputs but instead provide fuzzy feature information, which broadens the explicit discrete space into an expansive implicit continuous fuzzy space.

Manuscript received 30 September 2023; revised 26 February 2024; accepted 21 April 2024. Date of publication 24 April 2024; date of current version 4 September 2024. This work was supported in part by the “Pioneer” and “Leading Goose” R&D Program of Zhejiang under Grant 2023 C01212, in part by the National Key Research and Development Program of China under Grant 2022YFF0902305, in part by the Public Welfare Technology Application Research Project of Zhejiang under Grant LGF21F020002 and Grant LGF22F020008, in part by the Key Program and development projects of Zhejiang Province of China under Grant 2021C03137, and in part by the Key Lab of Film and TV Media Technology of Zhejiang Province under Grant 2020E10015. Recommended for acceptance by S.-H. Lee. (*Corresponding author: Naye Ji.*)

Fan Zhang is with the Faculty of Humanities and Arts, Macau University of Science and Technology, Macau 999078, China, also with the College of Media Engineering, Communication University of Zhejiang, Hangzhou 310019, China, and also with the Research Center for Artificial Intelligence and Fine Arts, Zhejiang Lab, Hangzhou, Zhejiang 311121, China (e-mail: fanzhang@cuz.edu.cn).

Zhaohan Wang and Xin Lyu are with the School of Animation and Digital Arts, Communication University of China, Beijing 100024, China (e-mail: 2022201305j6018@cuc.edu.cn; lvxinlx@cuc.edu.cn).

Siyuan Zhao is with the Faculty of Humanities and Arts, Macau University of Science and Technology, Macau 999078, China (e-mail: 2109853jai30001@student.must.edu.mo).

Mengjian Li is with the Research Center for Artificial Intelligence and Fine Arts, Zhejiang Lab, Hangzhou, Zhejiang 311121, China (e-mail: limengjian@zhejianglab.com).

Weidong Geng is with the College of Computer Science and Technology, Zhejiang University, Hangzhou 310027, China, and also with the Research Center for Artificial Intelligence and Fine Arts, Zhejiang Lab, Hangzhou, Zhejiang 311121, China (e-mail: gengwd@zju.edu.cn).

Naye Ji, Hui Du, Fuxing Gao, and Hao Wu are with the College of Media Engineering, Communication University of Zhejiang, Hangzhou 310019, China (e-mail: jinaye@cuz.edu.cn; duhui@cuz.edu.cn; fuxing@cuz.edu.cn; 210207140@stu.cuz.edu.cn).

Shunman Li is with the Zhejiang Institute of Economics and Trade, Hangzhou, Zhejiang 310006, China (e-mail: 2017000018@zjiet.edu.cn).

Supplementary videos and code can be accessed at <https://zf223669.github.io/Diffmotion-v2-website/>.

This article has supplementary downloadable material available at <https://doi.org/10.1109/TVCG.2024.3393236>, provided by the authors.

Digital Object Identifier 10.1109/TVCG.2024.3393236

The information in fuzzy space better aligns with the actual scenario. Relevantly, research in psychology highlights the significance of various factors in speech [11], [12], [13], [14]. These factors, called fuzzy features, are intricately intertwined with co-speech gestures. These studies present novel opportunities for synthesizing personalized gestures based solely on speech audio, thereby simplifying the feature inputs and reducing the complexity of multimodal processing.

Another challenge in this field is ensuring a high level of gesture-speech synchronization while preserving naturalness. Recent developments have focused on the application of Transformer and Diffusion-based models. This methodological shift has led to substantial progress in the efficiency and flexibility of gesture-generation technologies. Key examples of such innovative efforts include Taming [15], Diffuse Style Gesture [2], Diffuse Style Gesture+[3], GestureDiffuClip [16], and LDA [4]. Yet, these approaches encounter challenges with either insufficient or excessive correlation between gesture and speech, reducing the naturalness of the generated gestures.

The success of the Diffusion Transformers (DiTs) in text2-image [17] and Stable Diffusion 3<sup>1</sup> and text2video generation tasks, such as Sora,<sup>2</sup> which incorporates AdaLN, marks a significant advancement. This framework introduces a conditional mechanism that applies a uniform function across all tokens, enhancing the model’s ability to represent conditional and output features. This conditional mechanism also holds promise for effectively enhancing the ability to model the intricate mapping between speech and gestures. While the original DiTs take discrete text prompts as conditional inputs, its adaptability for sequence-to-sequence tasks, such as speech-driven gesture generation, presents an area of exploration.

In this study, we propose *Persona-Gestor*, a novel approach aimed at synthesizing personalized gestures solely from raw speech audio. This model innovatively introduces a fuzzy feature inference strategy within its condition extractor and incorporates AdaLN in a diffusion-based transformer module. *Persona-Gestor* transitions from explicit conditions to a nuanced, continuous representation of fuzzy features by employing fuzzy inference, which captures a broad spectrum of stylistic nuances and specific audio details. These features are integrated into a unified latent representation, synthesizing intricate 3D full-body gestures. Adopting AdaLN significantly enhances the model’s capability to depict the nuanced relationship between speech and gestures. Leveraging a diffusion process, the framework can generate diverse gesture outputs, showcasing the potential for high fidelity in gesture synthesis as shown in Fig. 1.

For clarity, our contributions are summarized as follows:

- We pioneering introduce the fuzzy feature inference strategy that enables driving a wider range of personalized gesture synthesis from speech audio alone, removing the need for style labels or extra inputs. This fuzzy feature extractor improves the system’s usability and generalization capabilities. To the best of our knowledge, it is the first




Fig. 1. Each pose depicted is personalized gestures generated solely relying on raw speech audio. *Persona-Gestor* offers a versatile solution that bypasses complex multimodal processing and enhances user-friendliness.

approach that uses fuzzy features to generate co-speech personalized gestures.

- We combined *AdaLN transformer architecture within the diffusion model to enhance the Modeling of the gesture-speech interplay*. We demonstrate that this architecture can generate gestures that achieve an optimal balance of natural and speech synchronization.
- *Extensive subjective and objective evaluations reveal that our model is superior to the current state-of-the-art approaches*. These results show the remarkable capability of our method in generating credible, speech-appropriateness, and personalized gestures.

## II. RELATED WORK

The present discussion offers a succinct overview of the conditional extraction mechanism and generative models within speech-driven gesture generation.

### A. Condition Extraction Mechanism

Recent advancements in co-speech gesture generation systems have incorporated various unstructured conditional information as input.

Selecting optimal representations for conditional input is a crucial research challenge in creating virtual human motions [18], [19]. For accurate reflection of gestures that match the auditory perception, prevalent research [7], [20], [21], [22] utilizes preprocessed audio features, such as MFCCs, log amplitude spectrogram, etc. Li et al. [6] develop a model for direct audio-to-gesture mapping. Despite these methods capturing acoustic nuances, the quest for richer feature sets continues. This has prompted investigations into the WavLM model, a refined, pre-trained wav2vec framework, for enhanced speech

<sup>1</sup><https://stability.ai/news/stable-diffusion-3>

<sup>2</sup><https://openai.com/sora>

extraction, showcasing in ReprGesture [23], QPGesture [24], and DiffuseStyleGesture [2].

Text-based co-speech gesture synthesis has seen significant contributions, such as Yoon et al.’s [25] recurrent neural network approach and Taras et al.’s [26] system, which merges acoustic and semantic speech features, employing BERT for semantic analysis [27]. Additionally, Uttaran et al. [28] utilize GloVe embeddings [29] to surpass models of similar dimensions, like Word2Vec [30] and Fast-Text [31]. Merging acoustic with semantic data offers a valuable path to enhance the relevance and context of generated gestures. Nonetheless, these modalities’ manual alignment and integration pose a challenge in effectively superior gesture synthesis.

For creating style-specific gestures, ReprGesture [23] and QPGesture [24] integrate textual data with audio features, whereas DiffuseStyleGesture [2] employs discrete labels to influence the stylistic aspects of the gestures produced. LDA [4] enables the system to generate style gestures with classifier-free guidance. Additionally, recent research has explored using textual prompts to generate stylized gestures [16]. Given that human emotions are more accurately represented on a continuous spectrum [32], [33] and emerge from a complex interplay of fuzzy factors, depending on discrete emotion labels can overly simplify the gesture generation process. This could limit the expressiveness and subtlety of the produced gestures. To address these limitations, Ghrohani et al. [7] introduced ZeroEGGS, a model that utilizes example motion clips to guide the style of gestures. Although achieving zero-shot is feasible, it still necessitates sample animation clips.

### B. Generative Approaches

DiffMotion [22], is the pioneering application of diffusion models integrating an LSTM for synthesizing diverse gestures. UnifiedGesture [5] presents a retargeting network to learn latent homeomorphic graphs to homeomorphic graphs for various gesture representations. Maximizing the transformer architecture’s potential, Alexanderson et al. [4] enhanced DiffWave by replacing dilated convolutions. Conformers [34] implementing classifier-free guidance to improve style expression. GestureDiffuCLIP [16] propose a network based on the transformer and AdaIN layers to incorporate style guidance into the diffusion model. LivelySpeaker [35] depends on contrastive learning to create a joint embedding space between gestures and transcripts. DiffuseStyleGesture (DSG) [2] and DSG+[3], integrating cross-local attention and layer normalization within transformers. Conversely, these methodologies face difficulties in achieving an optimal balance between gesture and speech synchronization, resulting in gestures that may appear either underrepresented or overly matched.

In this study, we employ a fuzzy feature inference strategy to implicitly capture fuzzy features in speech audio, synthesizing natural, personalized co-speech gestures solely relying on raw speech audio without additional modalities. Furthermore, we employ an AdaLN transformer architecture to enhance the

model’s capacity to capture the intricate relationship between speech and gestures.

## III. SYSTEM OVERVIEW

Persona-Gestor, as an end-to-end architecture, processes raw speech audio as its sole input, synthesizing personalized gestures that adeptly balance naturalness with synchronized alignment to speech.

### A. Problem Formulation

We introduce the challenge of co-speech gesture generation by framing it as a sequence-to-sequence problem, where the objective is to translate a sequence of speech audio features into a corresponding sequence of gestures. We denote the sequence of full-body gesture features and the sequence of the audio signal as  $g^0 = g_{1:T}^0 \in [g_1^0, \dots, g_t^0, \dots, g_T^0] \in \mathbb{R}^{T \times (D+3+3)}$  and  $a = a_{1:T} \in [a_1, \dots, a_t, \dots, a_T] \in \mathbb{R}^T$ .  $g_t^0 = \mathbb{R}^{(D+3+3)}$  symbolizes the representation of 3D joint angles, along with the root positional and rotational velocity at frame  $t$ , where  $D$  denoting the number of channels for these joints. The superscript indicates the diffusion time step  $n$ . Here,  $a_t$  refers to the current subsequence audio waveform signal at frame  $t$ , while  $T$  denotes the sequence length. Let us define  $p_\theta(\cdot)$  as the Probability Density Function (PDF), which aims to approximate the actual distribution of gesture data  $p(\cdot)$  and enables easy sampling. The objective is to generate a non-autoregressive whole pose sequence ( $g^0$ ) from its conditional probability distribution given audio signal ( $a$ ) as covariate:

$$g^0 \sim p_\theta(g^0|a) \approx p(\cdot) := p(g^0|a) \quad (1)$$

where the  $p_\theta(\cdot)$  aims to approximate  $p(\cdot)$  trained by the denoising diffusion model.

### B. Model Architecture

The architecture of Persona-Gestor is depicted in Fig. 2. It comprises four primary components: (1) a Fuzzy Feature Extractor, (2) an AdaLN Transformer, (3) a Gesture Encoder and Decoder, and (4) a diffusion network.

1) *Fuzzy Feature Extractor*: This module utilizes a fuzzy inference strategy, meaning it does not generate explicit classification outputs. Instead, it offers implicit, continuous, fuzzy feature information, automatically learning and inferring the global style and details directly from raw speech audio. The module, showcased in Figs. 2(b) and 3, is a dual-component extractor that integrates both global and local extractors. The local extractor leverages the WavLM large-scale pre-trained model [36] to convert the audio sequence into tokens. We chose WavLM for its adeptness at extracting the complex features of speech audio to capture universal audio latent representations, denoted as  $z_a$ .

We observe that the local extractor alone falls short of fully capturing the array of stylistic features and ensuring style consistency across sequences. To overcome this, we integrate a global style extractor, employing a depthwise separable convolution 1D layer [37] across the  $z_a$ . This global extractor is designed to

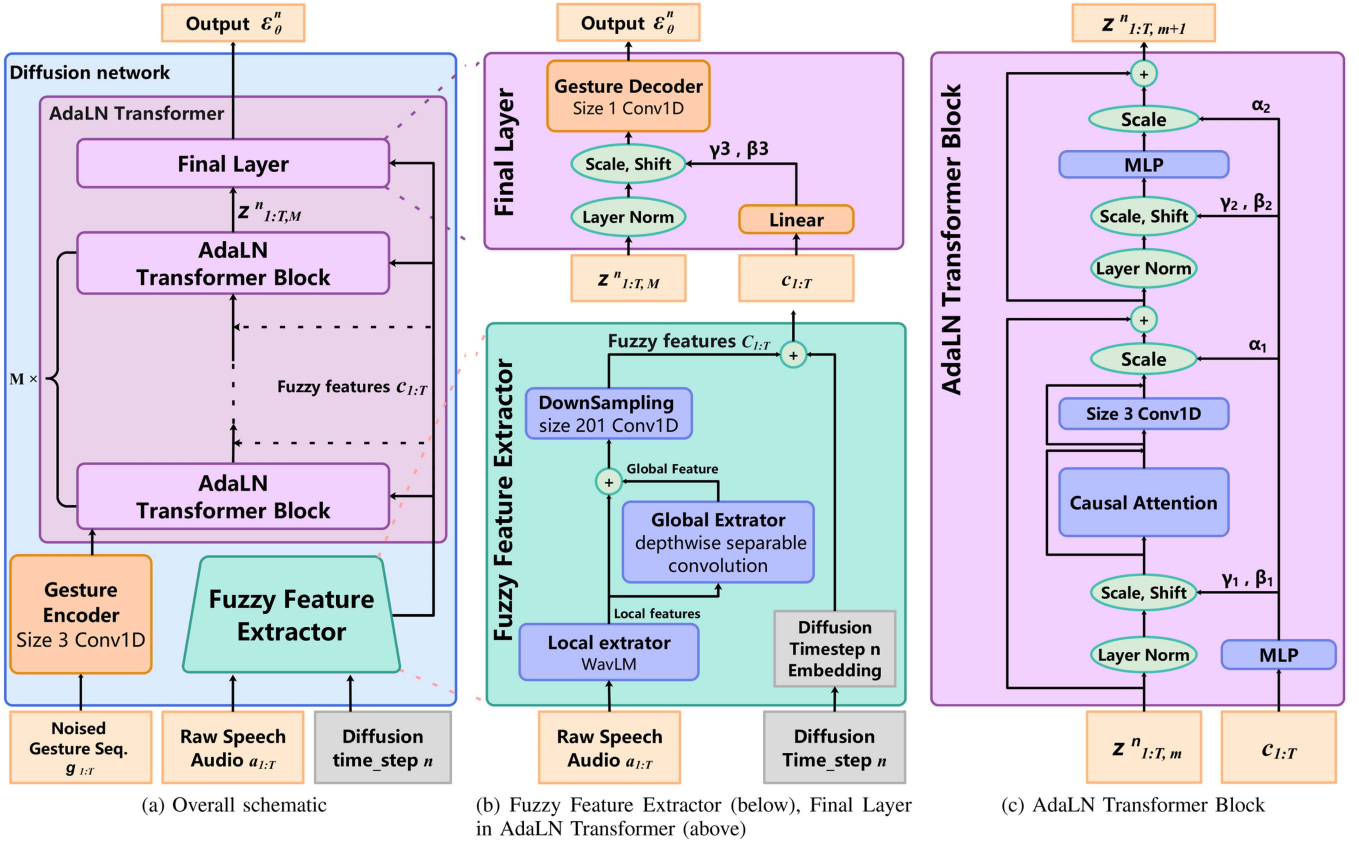


Fig. 2. The Architecture of Persona-Gestor mainly integrates a fuzzy feature extractor and an adaptive layer normalization (AdaLN) transformer diffusion architecture (DiTs-based). The fuzzy feature extractor comprises a dual-component framework to capture the fuzzy style and detail-oriented audio features comprehensively. These features, as unified latent features, are subsequently fed into the AdaLN transformer to model the relationship with the accompanist gesture, facilitating the estimation of diffusion noise for the diffusion model. (a) Overall Schematic. (b) Fuzzy Feature Extractor. (c) AdaLN Transformer Block.

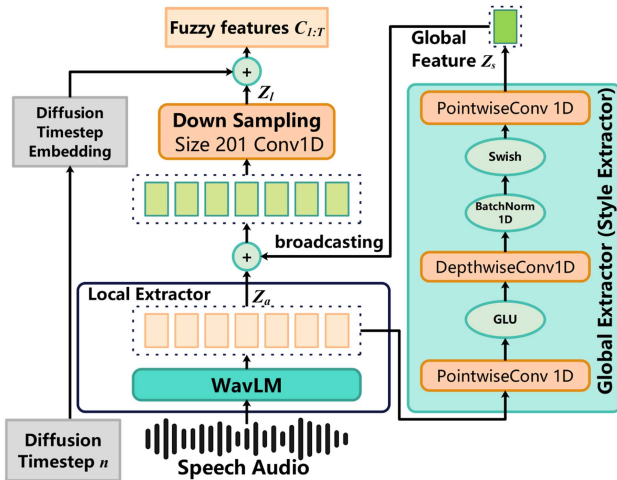


Fig. 3. An overview of the fuzzy inference condition extractor.

automatically capture and embed global fuzzy style information from  $z_a$  into a token  $z_s \in \mathbb{R}^{1 \times D'}$ . This token is then broadcasted and combined with the universal audio latent representations  $z_a \in \mathbb{R}^{T' \times D'}$  to form a unified latent representation  $z_l \in \mathbb{R}^{T \times D''}$ .

We merge local and global insights for co-speech gesture generation to enhance the sequence's overall representational fidelity. Subsequently, the unified latent representation is directed to the downsampling module for further processing.

The downsampling module is integrated into the condition extractor to ensure alignment between each latent representation and its corresponding sequence of encoded gestures. In our exploration, we experimented with linear alignment like DSG [2] and DSG+[3], but noted an issue of foot-skating arising from these methods. On the contrary, We adopt a Conv1D layer with a kernel size of 201 for this module that maps every 201-length target token output from WavLM to one gesture frame. Finally, the fuzzy feature extractor outputs  $c_{1:T}$ , representing a unified latent representation that combines encoded audio features and diffusion time step  $n$ . The condition extractor can be formalized by:

$$\begin{aligned}
 z_a &= LE(a) & z_a &\in \mathbb{R}^{T' \times D'} \\
 z_s &= GE(z_a) & z_s &\in \mathbb{R}^{1 \times D'} \\
 z_l &= DS(z_a + z_s) & z_l &\in \mathbb{R}^{T \times D''} \\
 n' &= DTE(n) & n &\in \mathbb{R}, n' \in \mathbb{R}^{1 \times D''}
 \end{aligned}$$

$$c_{1:T} = z_l + n' \quad c_{1:T} \in \mathbb{R}^{T \times D'} \quad (2)$$

where  $LE(\cdot)$  and  $GE(\cdot)$  denote the local extractor (WavLM) and the global extractor.  $DS(\cdot)$  represents the down sampling process.  $DTE(\cdot)$  signifies the diffusion time step embedding. The final output of the fuzzy feature extractor is denoted as  $c_{1:T}$ . Here,  $T'$ ,  $D'$ , and  $D''$  refer to the WavLM output token length, feature dimensionality of WavLM's output token, and the feature ( $h$ ) dimensionality of the proposed model's hidden state, respectively.  $a$  is the input raw speech audio waveform.  $Z_a$  and  $Z_s$  are extracted by the local extractor(WavLM model) and the global extractor.  $Z_l$  is the unified latent representation.  $n$  is the diffusion time step,  $n'$  is the embedded diffusion time step feature.

2) *AdaLN Transformer*: The AdaLN's fundamental purpose is to incorporate a conditional mechanism that uniformly applies a specific function across all tokens, thereby significantly improving the model's capacity for representing both conditional and output features with enhanced efficiency. It offers a more sophisticated and nuanced approach to modeling, enabling the system to capture and articulate the complex dynamics between various input conditions and their corresponding outputs. Consequently, this leads to an improvement in the model's predictive accuracy and ability to generate outputs more aligned with the given conditions.

Diffusion Transformers (DiTs) [17] represent an advanced transformer-based backbone for diffusion models, surpassing previous U-Net models in performance. By incorporating AdaLN within transformer blocks for text-to-image synthesis, DiTs achieve lower Frchet Inception Distance (FID) [38] scores, indicating superior image quality. Recently, this framework has been used for text-conditional video generation. Despite Diffusion Transformers (DiTs) success in handling discrete text prompts conditional inputs, their effectiveness in speech-driven gesture generation, a sequence-to-sequence task, necessitates a thorough investigation.

Distinctively with the DiTs, our approach utilizes continuous fuzzy features as conditional input tokens. Further, it is without any patchy for spatial input, resulting in the output being the latent feature of a sequence of gestures.

The module involves regressing the dimensionwise scale and shift parameters ( $\gamma$  and  $\beta$ ), which are derived from the fuzzy feature extractor output  $c_{1:T}$ , instead of directly learning  $\gamma$  and  $\beta$ , as depicted in Fig. 2(c). In each AdaLN transformer, a latent feature denoted as  $z_{1:T,m}^n$  is generated by fusing condition information and gesture using AdaLN and causal self-attention. Here,  $1 \leq m \leq M$ , where  $M$  represents the total number of AdaLN transformer stacks. In addition, the final layer, as illustrated in Fig. 2(b), fed the same fuzzy features but with additional scale and shift operation.

This method facilitates the creation of detailed gesture sequences solely from speech audio, eliminating the requirement for discrete style labels or supplementary inputs. Consequently, it significantly improves the model's capacity to generate personalized and closely aligned gestures with the context of the speech, offering a more refined and context-sensitive gesture synthesis capability.

3) *Gesture Encoder and Decoder*: The architecture of the gesture encoder and decoder is designed to encode and decode the gesture sequence, as illustrated in Fig. 2(a) and (b). The gesture encoder comprises a Convolution1D with a kernel size of 3. It encodes the initial sequence of gestures  $g$  into a hidden state  $h \in \mathbb{R}^{T \times D''}$ . Our experimental results revealed that employing a kernel size of 1 resulted in animation jitter. Conversely, a kernel size of 3 is instrumental in mitigating this issue by effectively capturing the spatial-temporal relationships inherent in gesture sequences.

The gesture decoder reduces the feature dimension of the output from the transformer  $D''$  to the original dimension  $D$ , corresponding to the number of channels representing skeleton joints. Result in outputting the predicted noise ( $\epsilon_\theta$ ). We utilize a size of 1 conv1D By convolving a 1D kernel in the input sequence, our model can effectively extract meaningful features and relationships between adjacent joint channels.

### C. Training and Inferencing With Denoising Diffusion Probabilistic Model

The diffusion process in this architecture aims to reconstruct the conditional probability distribution between gestures and fuzzy features. This entails employing a systematic approach to sample from this restored distribution, thereby enabling the generation of diverse gestures.

Following our previous work, Diffmotion [22], incorporating the Denoising Diffusion Probabilistic Model (DDPM) into our approach. However, we employ a non-autoregressive transformer to generate the entire sequence of gestures instead of frame-by-frame. The form is represented by  $p_\theta := \int p_\theta(g^{0:N}) dg^{1:N}$ , where  $g^1, \dots, g^N$  are latent of the same dimensionality as the data  $g^n$  at the  $n$ -th diffusion time stage.

The model contains two processes: the diffusion process and the generation process. At training time, the diffusion process gradually converts the original gesture data( $g^0$ ) to white noise( $g^N$ ) by optimizing a variational bound on the data likelihood. At inference time, the generation process recovers the data by reversing this noising process through the Markov chain using Langevin sampling [39]. The Markov chains in the diffusion process and the generation process are:

$$p(g^n | g^0) = \mathcal{N}(g^n; \sqrt{\bar{\alpha}^n} g^0, (1 - \bar{\alpha}^n) I) \quad \text{and} \\ p_\theta(g^{n-1} | g^n, g^0) = \mathcal{N}(g^{n-1}; \tilde{\mu}^n(g^n, g^0), \tilde{\beta}^n I), \quad (3)$$

where  $\alpha^n := 1 - \beta^n$  and  $\bar{\alpha}^n := \prod_{i=1}^n \alpha^i$ . As shown by [40],  $\beta^n$  is an increasing variance schedule  $\beta^1, \dots, \beta^N$  with  $\beta^n \in (0, 1)$ , and  $\tilde{\beta}^n := \frac{1 - \bar{\alpha}^{n-1}}{1 - \bar{\alpha}^n} \beta^n$ .

The training objective is to optimize the parameters  $\theta$  that minimizes the Negative Log-Likelihood (NLL) via Mean Squared Error (MSE) loss between the true noise  $\epsilon \sim \mathcal{N}(0, I)$  and the predicted noise  $\epsilon_\theta$ :

$$\mathbb{E}_{g_{1:T}, \epsilon, n} [\|\epsilon - \epsilon_\theta(\sqrt{\bar{\alpha}^n} g^0 + \sqrt{1 - \bar{\alpha}^n} \epsilon, a_{1:T}, n)\|^2], \quad (4)$$

Here  $\epsilon_\theta$  is a neural network (see Fig. 2(a)), which uses input  $g_t^0$ ,  $a_{t-1}$  and  $n$  that to predict the  $\epsilon$ , and contains the similar

**Algorithm 1:** Training for the Whole Sequence Gesture.

---

**Input:** data  $g_{1:T}^0 \sim p(g^0|a_{1:T})$  and  $a_{1:T}$   
**repeat**  
 Initialize  $n \sim \text{Uniform}(1, \dots, N)$  and  $\epsilon \sim \mathcal{N}(0, I)$   
 Take the gradient step on  

$$\nabla_{\theta} \|\epsilon - \epsilon_{\theta}(\sqrt{\bar{\alpha}_n} g_{1:T}^0 + \sqrt{1 - \bar{\alpha}_n} \epsilon, a_{1:T}, n)\|^2$$
  
**until converged**;

---

**Algorithm 2:** Sampling  $g_{1:T}^0$  via Annealed Langevin Dynamics.

---

**Input:** noise  $g_{1:T}^N \sim \mathcal{N}(0, I)$  and raw audio waveform  $a_{1:T}$   
**for**  $n = N$  **to**  $1$  **do**  
**if**  $n > 1$  **then**  
 $z \sim \mathcal{N}(0, I)$   
**else**  
 $z = 0$   
**end if**  
 $g_{1:T}^{n-1} = \frac{1}{\sqrt{\alpha^n}}(g_{1:T}^n - \frac{\beta^n}{\sqrt{1 - \bar{\alpha}^n}} \epsilon_{\theta}(g_{1:T}^n, a_{1:T}, n)) + \sqrt{\sigma_{\theta}} z$   
**end for**  
**Return**  $g_{1:T}^0$

---

architecture employed in [41]. The complete training procedure is outlined in Algorithm 1.

After training, we utilize variational inference to generate the whole sequence of new gestures matching the original data distribution ( $g^0 \sim p_{\theta}(g^0, a)$ ). We followed the sampling procedure in Algorithm 2 to obtain the whole sequence of the sample  $g^0$ . The  $\sigma_{\theta}$  is the standard deviation of the  $p_{\theta}(g^{n-1}|g^n)$ . We choose  $\sigma_{\theta} := \tilde{\beta}^n$ .

During inferencing, we send the whole sequence of the raw audio to the condition extractor component. Then, the component output is fed to the Diffusion Model to generate the whole sequence of the accompanying gesture ( $g^0$ ).

## IV. EXPERIMENTS

To validate our approach, we utilized three co-speech gesture datasets (Trinity [42], ZEGGS [7], and BEAT [43]). Our experiments concentrated on producing full 3D body gestures (including finger motions and locomotion). This choice presented a greater challenge than merely focusing on upper body motions due to the expanded output dimensionality and the need to overcome visual complexities, such as foot-skating, the naturalness of finger movements, and locomotion.

### A. Dataset and Data Processing

1) *Datasets:* The Trinity dataset focuses on individual spontaneous speech, the ZEGGS dataset encompasses a wide range of emotional expressions, and the BEAT dataset consists of personalized movements exhibited by various individuals. Further details are elaborated in Table II found in Appendix A, available online.

2) *Speech Audio Data Process:* In the Trinity dataset, the audio was recorded at a sampling rate of 44 kHz, while 48 kHz

in ZEGGS and BEAT. However, due to the pre-training of the WavLM large model on speech audio sampled at 16 kHz, we uniformly resample all audio to match this frequency.

3) *Gesture Data Process:* We focus solely on full-body gestures, adopting the data processing techniques outlined by Alexanderson et al. [20]. Given the variability in data quality and structure across motion datasets, we tailor our approach by selecting specific joints for analysis in each dataset. We omit hand skeleton data for the Trinity Gesture Dataset due to its inferior quality. For the ZEGGS and BEAT datasets, our analysis includes finger joints, which are considered to be the same set of joints in the Trinity dataset. All data capture translational and rotational velocities to detail the root’s trajectory and orientation. The datasets are uniformly downsampled to a frame rate of 20 fps. To ensure accurate and continuous representation of joint angles, we apply the exponential map technique [44]. All data are segmented into 20-second clips for training and validation purposes. As for the user evaluation, we segment the generated gesture sequence into 10 seconds to improve the efficiency of the evaluation.

### B. Model Settings

Our experiments employ 12 causal attention blocks, each comprising 16 attention heads (as depicted in Fig. 2(a)). The encoding process transforms each frame of the gesture sequence into hidden states  $h \in \mathbb{R}^{1280}$ . For the WavLM model, we utilize the pre-trained WavLM Large model.<sup>3</sup> To ensure temporal translation invariance, we employ a translation-invariant self-attention (TISA) mechanism [45].

The quaternary variance schedule of diffusion model starts from  $\beta_1 = 1 \times 10^{-4}$  till  $\beta_N = 5 \times 10^{-5}$  with linear beat schedule. The number of diffusion steps is  $N = 1000$ . The training batch size is 32 per GPU.

The model was developed using the Torch Lightning framework and tested on an Intel i9 processor with an A100 GPU. Training durations were approximately 4 hours for Trinity and ZeroEGGS and 21 hours for BEAT.

### C. Visualization Results

Our system excels in creating personalized gestures that align with the speech context, leveraging the fuzzy inference strategy to autonomously derive fuzzy features directly from speech audio. Furthermore, it showcases remarkable generalization and robustness by utilizing in-the-wild speech.

Fig. 4 depicts the visual outcomes of gestures aligned with the emotional valence conveyed by the audio. For instance, the system produces gestures of joy in response to happy audio cues (refer to Fig. 4(a)) and gestures of sadness for sorrowful audio (as depicted in Fig. 4(b)). The system can also infer age-related characteristics or other nuanced states from the speech audio (as illustrated in Fig. 4(c) and (d)).

Fig. 5 shows the system’s ability to generate personalized gestures Predicated upon individuals’ unique speech traits. For

<sup>3</sup><https://github.com/microsoft/unilm/tree/master/wavlm>



Fig. 4. Samples of gestures corresponding to different emotions. The left side of the subfigure displays ground truth gestures, while the right side showcases gestures generated by our architecture.




Fig. 5. Samples of gestures corresponding to different personalities. The left side of the subfigure displays ground truth gestures, while the right side showcases gestures generated by our architecture.

example, Ayana’s gestures, with hands together and palms facing, denote reserved expressiveness. In contrast, Jaime’s ”palm up” gestures imply openness, and Luqi’s alternating hand movements add dynamic variability. These results highlight the system’s adeptness at depicting a wide range of personality-specific gestures.

Interestingly, as shown in Fig. 6, the system can produce gestures with certain semantic relevance even without explicit semantic constraints. For instance, Carla’s remark about her mother is matched with a self-referential gesture. Likewise, Lawrence’s reference to distance is visually enhanced by a gesture that emphasizes the semantic essence of his speech.




Fig. 6. Samples of gestures corresponding to semantic.




Fig. 7. Sample of gesture including finger movements and locomotion.<sup>4</sup>



Fig. 8. Samples of gestures corresponding to in-the-wild speech audio collected from TED Talks.

Further, finger movements and locomotion are included, as shown in Fig. 7, highlighting the system’s proficiency in creating realistic, character-specific animations, thereby increasing the virtual interactions’ believability and immersion.

Fig. 8 showcases gesture outcomes generated from in-the-wild speech audio, like TED talks, to demonstrate the system’s ability to create lifelike and style movements directly from unstructured real-world audio without additional prompts or labels. This highlights the system’s robust generalization capabilities. Testing in noisy environments with background music, applause, and urban sounds further revealed the system’s strong anti-interference performance, emphasizing its resilience. This efficiency simplifies the input process, enabling effortless generation of dynamic character animations from raw audio, thus enhancing user experience and system accessibility.

<sup>4</sup>Due to the inherent challenges in retargeting finger motion to the avatar, please refer to the support video for more details.

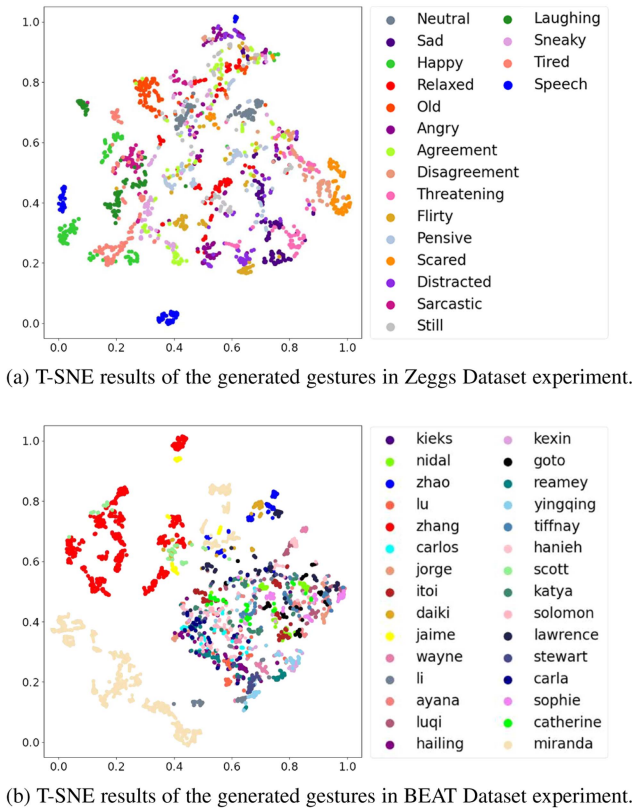


Fig. 9. The T-SNE clustering visualization displays a variety of gestures distinguished by color-coded styles, revealing distinct regions for gestures tied to specific emotions or speakers, albeit with some boundary overlaps. This highlights our approach’s capability to produce distinct style gestures through a fuzzy feature inference strategy, relying solely on speech audio.

Finally, we represent the visualizes (Fig. 9) of the distribution of generated gestures corresponding to different emotional states (Fig. 9(a)) and personalities (Fig. 9(b)) using the t-SNE method. The figure illustrates distinct separations between certain states, while others exhibit a degree of similarity yet remain distinguishable. These findings demonstrate the capability of our proposed method to generate nuanced and discernible gestures solely from raw speech audio without relying on labels or manual annotations.

#### D. Subjective and Objective Evaluation

Consistent with the prevailing practices in gesture generation research, we conducted a series of subjective and objective evaluations to evaluate the co-speech gestures generated by our proposed Persona-Gestor (PG) model.

We adopted slightly varied baselines for different datasets. For the Trinity dataset, we employed LDA [4] and Taming [15]. In addition to LDA and Taming, for the ZEGGS dataset, we also incorporated DiffuseStyleGesture (DSG) [2] and ZeroEGGS [7]. Furthermore, for the BEAT dataset, we utilized the same baseline models as in ZEGGS but replaced DSG with DSG+[3] and introduced GestureDiffuCLIP (GDC) [16] as an additional baseline model.

In our experiments with the ZEGGS and BEAT datasets, we extended the original LDA, DSG, and DSG+<sup>5</sup> models to cover all styles within these datasets. Originally, the Taming model, trained exclusively on the TED dataset, focused on upper-body gestures. We have since augmented it to support full-body gestures across the three datasets. Efforts to adapt LDA to include finger motions were met with challenges, leading to unsatisfactory outcomes in gesture generation. Consequently, we utilized LDA-generated gestures, excluding finger movements, for our analysis. For more implementation details of these baselines, please refer to Appendix C, available online.

1) *Subjective Evaluation*: The goal of speech-driven gesture generation is to produce gestures that are both natural and convincing. However, exclusive reliance on objective metrics may not adequately reflect human subjective quality assessments [20], [46], [47]. This study prioritizes subjective evaluations to gauge human perception, complemented by objective evaluations detailed in Section IV-D2.

For thorough subjective evaluations, we utilize three metrics: human likeness, appropriateness, and style appropriateness. Human likeness gauges the naturalness and resemblance of gestures to real human movements independent of speech. Appropriateness examines the temporal alignment of gestures with speech rhythm, intonation, and semantics, ensuring natural fluidity. Style-appropriateness evaluates the similarity between generated and original gestures.

We conducted a user study with pairwise comparisons, as recommended by [48]. In each trial, participants were shown two 10-second video clips generated by different models (including the Ground Truth (GT)) side by side for direct comparison. The videos were accompanied by instructions for participants to select their preferred clip based on their evaluations. Preferences were quantified on a 0 to 2 scale, with the unselected clip in each pair receiving an inverse score (e.g., a -2 score for the non-chosen clip if the chosen one received 2). A score of zero indicated no preference. Attention checks were included in the study to ensure engagement. Further details are available in Appendix B, available online.

Considering the extensive range of styles in ZEGGS (19) and BEAT (30), individual evaluations for each style were deemed impractical. Consequently, we utilized a random selection method to assign a subset of 5 styles from ZEGGS and 6 characters from BEAT to each participant. For the Trinity dataset, we chose Record\_008 and Record\_015. The training or validation sets include none of the selected audio clips.

A total of thirty volunteer participants, 17 males and 13 females aged between 19 and 31, were recruited for this study. Twenty-two participants were Chinese nationals, while the remaining eight were international students from the USA and U.K. All participants in this study exhibited a high level of English proficiency.

One-way ANOVA and post-hot Tukey multiple comparison tests were conducted to determine if the models’ scores differed

<sup>5</sup>The authors have expanded their coverage to include all types in the BEAT dataset, as originally presented in the project of that study.



TABLE I  
THE SUBJECT MEAN PERCEPTUAL RATING SCORE

Dataset	Methods		Subject Evaluation Metric			Objective Evaluation Metric		
	Model	With Fingers	Human $\uparrow$ likeness	appropriateness $\uparrow$	Style $\uparrow$ appropriateness	FGD $\downarrow$ on feature space	FGD $\downarrow$ on raw data space	BeatAlign $\uparrow$
Trinity	GT	Y	0.51 $\pm$ 1.73	0.66 $\pm$ 1.24	/	/	/	/
	LDA [4]	N	-0.22 $\pm$ 0.98	-0.39 $\pm$ 1.08	-0.18 $\pm$ 1.10	349.53	6008.37	0.68
	Taming [15]	N	-0.48 $\pm$ 0.96	-0.47 $\pm$ 1.07	-0.23 $\pm$ 1.05	3970.14	52196.87	0.68
	(Proposed)PG	N	<b>0.12<math>\pm</math>1.09</b>	<b>0.19<math>\pm</math>1.12</b>	<b>0.20<math>\pm</math>1.06</b>	<b>289.42</b>	<b>5080.57</b>	<b>0.69</b>
	(Ours)PGNSE	N	0.08 $\pm$ 1.06	0.13 $\pm$ 1.18	0.15 $\pm$ 1.03	8002.65	70617.28	0.68
	(Ours)PGCA	N	-0.17 $\pm$ 1.02	-0.01 $\pm$ 1.11	-0.03 $\pm$ 1.02	2566.65	26124.39	0.68
	(Ours)PGCF	N	-0.08 $\pm$ 1.06	-0.13 $\pm$ 1.05	0 $\pm$ 1.02	12540.17	156895.12	0.68
	(Ours)PGICC	N	0.07 $\pm$ 1.01	0.06 $\pm$ 1.12	0.08 $\pm$ 1.01	125921.53	1940124.44	0.67
ZEGGS	GT	Y	0.95 $\pm$ 1.13	1.19 $\pm$ 1.03	/	/	/	/
	LDA [4]	N	-0.73 $\pm$ 1.12	0.02 $\pm$ 1.31	0.53 $\pm$ 1.41	124.55	50996.33	0.66
	DSG [2]	Y	-0.47 $\pm$ 1.14	-0.71 $\pm$ 1.11	-0.61 $\pm$ 1.08	66.77	33297.50	0.63
	Taming [15]	Y	-1.08 $\pm$ 1.01	-0.91 $\pm$ 1.09	-0.98 $\pm$ 0.97	1419.76	293245.12	0.67
	ZeroEGGS [7]	Y	0.38 $\pm$ 1.11	0.29 $\pm$ 1.35	0.49 $\pm$ 1.31	37.19	26666.85	0.66
	(Proposed)PG	Y	<b>0.42<math>\pm</math>1.17</b>	<b>0.48<math>\pm</math>1.29</b>	<b>0.76<math>\pm</math>1.34</b>	<b>28.13</b>	<b>26193.92</b>	<b>0.68</b>
	(Ours)PGNSE	Y	0.33 $\pm$ 1.15	0.35 $\pm$ 1.29	0.59 $\pm$ 1.38	125.40	49081.55	0.66
	(Ours)PGOnehot	Y	0.25 $\pm$ 1.19	0.33 $\pm$ 1.31	0.51 $\pm$ 1.28	122.56	50259.95	0.63
	(Ours)PGCA	Y	-0.36 $\pm$ 1.22	-0.62 $\pm$ 1.14	-0.66 $\pm$ 1.09	807.12	156686.01	0.67
	(Ours)PGCF	Y	0.27 $\pm$ 1.20	-0.03 $\pm$ 1.28	-0.23 $\pm$ 1.26	97.57	39256.55	0.67
(Ours)PGICC	Y	0.04 $\pm$ 1.20	-0.39 $\pm$ 1.23	-0.40 $\pm$ 1.21	407.89	89893.96	0.67	
BEAT	GT	Y	0.65 $\pm$ 1.16	0.96 $\pm$ 1.04	/	/	/	/
	LDA [4]	N	-1.65 $\pm$ 0.73	-1.59 $\pm$ 0.74	-1.35 $\pm$ 1.05	276.25	3584.95	0.66
	DSG+ [3]	Y	-0.28 $\pm$ 1.17	-0.49 $\pm$ 1.15	-0.40 $\pm$ 1.24	23811.46	2384465.64	0.43
	GDC [16]	N	0.54 $\pm$ 1.12	0.47 $\pm$ 1.25	0.30 $\pm$ 1.27	432.15	93215.56	<b>0.69</b>
	Taming [15]	Y	-0.42 $\pm$ 1.14	-0.52 $\pm$ 1.14	-0.32 $\pm$ 1.24	1251.56	46828.23	0.66
	(Proposed)PG	Y	<b>0.56<math>\pm</math>1.14</b>	<b>0.63<math>\pm</math>1.10</b>	<b>0.66<math>\pm</math>1.16</b>	<b>264.06</b>	<b>3471.26</b>	0.68
	(Ours)PGNSE	Y	0.09 $\pm$ 1.16	0.27 $\pm$ 1.23	0.46 $\pm$ 1.31	1514.94	51077.98	0.66
	(Ours)PGOnehot	Y	-0.01 $\pm$ 1.16	0.18 $\pm$ 1.31	0.32 $\pm$ 1.36	1863.69	63872.78	0.63
	(Ours)PGCA	Y	0.35 $\pm$ 1.09	0.17 $\pm$ 1.26	0.28 $\pm$ 1.33	703.83	18990.56	0.66
	(Ours)PGCF	Y	0.14 $\pm$ 1.01	0.15 $\pm$ 1.22	0.30 $\pm$ 1.31	1160.63	48899.63	0.66
(Ours)PGICC	Y	0.02 $\pm$ 1.10	-0.24 $\pm$ 1.17	-0.25 $\pm$ 1.25	2057.31	78754.92	0.66	

Bold fonts were utilized to emphasize the best results for each metric among the different methods, except for the GT.

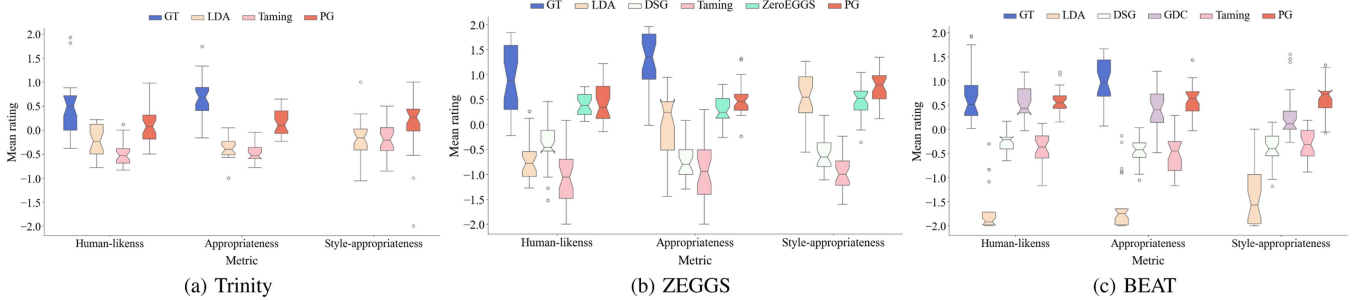


Fig. 10. The mean rating of each metric for each approach across the three datasets in comparative experiments.

on the three evaluation aspects. The results are shown in Table I and Fig. 10. The post-hoc analysis information is provided in the Appendix B, available online.

The results indicate that the GT achieves the highest scores (0.51  $\pm$  1.73 and 0.95  $\pm$  1.13) in the Trinity and ZEGGS datasets, exhibiting statistically significant differences ( $p < 0.001$ ) in human-likeness evaluations when compared to model-generated gestures. The GT is characterized by a diverse yet limited array of gestures, each with distinct traits that enhance movement realism. However, these gestures belong to the dataset's long-tail distribution, challenging the models' learning capabilities. Additionally, these unique gestures impact the appropriateness and style-appropriateness scores. Conversely, while the GT achieves higher scores (0.65  $\pm$  1.16), no significant differences were observed compared with the PG (0.56  $\pm$  1.14) and GDC (0.54  $\pm$  1.12) in the BEAT dataset analysis. This suggests that

these models are more closely aligned with GT benchmarks in this dataset.

The experiments on the Trinity dataset show our proposed model (0.12  $\pm$  1.09, 0.138  $\pm$  1.12, and 0.203  $\pm$  1.06) outperforming both LDA (-0.22  $\pm$  0.98, -0.39  $\pm$  1.08, and -0.18  $\pm$  1.10) and Taming (-0.48  $\pm$  0.96, -0.47  $\pm$  1.07, and -0.23  $\pm$  1.05) architectures significantly ( $p < 0.001$ ) across all metrics. This superior performance is due to the more natural and relaxed gestures produced by our model, PG, enhancing its effectiveness compared to the LDA and Taming models, which fall short in accurately capturing the acoustic rhythm.

Evaluation of the ZEGGS dataset showed statistically significant differences ( $p < 0.001$ ) between our method (0.42  $\pm$  1.17, 0.48  $\pm$  1.29, and 0.76  $\pm$  1.34) and others across all three metrics. However, there was no statistically significant difference ( $p > 0.05$ ) between our method and ZeroEGGS (0.38  $\pm$  1.11) in

terms of human likeness, though our method achieved a slightly higher score. These findings suggest that both our proposed model and ZeroEGGS can generate vivid gestures. Our advantage lies in the ability to synthesize emotional gestures solely through audio input, without relying on any reference example animations or labels.

In the BEAT dataset experiments, our PG model exhibited significant improvements ( $0.56 \pm 1.14$ ,  $0.63 \pm 1.10$ , and  $0.66 \pm 1.16$ ) in three metrics compared to DSG+ ( $-0.28 \pm 1.17$ ,  $-0.49 \pm 1.15$ , and  $-0.40 \pm 1.24$ ), LDA ( $-1.65 \pm 0.73$ ,  $-1.59 \pm 0.74$ , and  $-1.35 \pm 1.05$ ), and Taming ( $-0.41 \pm 1.14$ ,  $-0.52 \pm 1.14$ , and  $-0.32 \pm 1.24$ ), reflecting the degradation in synthesis quality observed in DSG+ and LDA when incorporating all styles. While human-likeness metrics were comparable to GDC, the appropriateness metric of PG achieved higher scores than GDC ( $0.47 \pm 1.25$ ) despite the GDC’s ability to better align with speech rhythm. Users reported that GDC-generated gestures overly emphasized prosodic cues, resulting in unnatural and frequent gestures. Furthermore, these gestures often displayed repetitive patterns with limited stylistic diversity, clearly indicating their origin from the GDC model and leading to a lower score ( $0.30 \pm 1.27$ ) in the style-appropriateness metric. This finding implies that the quality of gestures is not solely determined by accurately matching the audio rhythm.

2) *Objective Evaluation*: We introduce three objective evaluation metrics: Fr chet Gesture Distance (FGD) in both feature and raw data spaces [49], and BeatAlign [50]. FGD, inspired by the Fr chet Inception Distance (FID) [38], assesses the quality of generated gestures and demonstrates a moderate correlation with human-likeness ratings, outperforming other objective metrics [47]. Additionally, BeatAlign measures gesture-audio synchrony by calculating the Chamfer Distance between audio and gesture beats, providing insights into the temporal alignment of generated gestures with speech rhythms.

Table I displays our results, highlighting the state-of-the-art performance of our method in objective evaluations using FGD and BeatAlign metrics. Our model outperforms (289.42 for Trinity, 28.13 for ZEGGS, and 264.06 for BEAT) other architectures in FGD, effectively generating gestures that align closely with the Ground Truth (GT). It also achieves superior BeatAlign scores (0.69 for Trinity, 0.68 for ZEGGS, and 0.68 for BEAT) compared to other models, except for GDC (0.69 for BEAT), demonstrating its efficacy in producing co-speech gestures that synchronize accurately with speech rhythms. Although GDC scores highest in BeatAlign, corroborating user feedback, its overemphasis on prosodic cues leads to frequent high-frequency gestures. While technically accurate, this diminishes gesture naturalness.

### E. Ablation Studies

Ablation studies were performed to evaluate the impact of key components on our model’s efficacy, specifically targeting the global fuzzy feature extractor and Adaptive Layer Normalization (AdaLN).

1) *Ablation of Global Fuzzy Feature Extractor*: For the global fuzzy feature extractor, we explored the outcomes of removing this component (we call it: No Style Encoding, PGNSE) and replacing it with One-hot embedding (PGOnehot) for discrete feature extraction. PGOnehot was not applied to the Trinity dataset due to its limited style variability.

Our analysis of the global fuzzy feature extractor shows no significant differences ( $p > 0.05$ ) between PG and PGNSE on the Trinity dataset in three subjective metrics, likely due to its limited range of styles. However, the ZEGGS dataset reveals significant variances in three metrics between PG ( $0.42 \pm 1.17$ ,  $0.48 \pm 1.29$ , and  $0.76 \pm 1.34$ ) and PGOnehot ( $0.25 \pm 1.19$ ,  $0.33 \pm 1.31$ , and  $0.51 \pm 1.28$ ). In contrast, no notable differences in human-likeness and appropriateness metrics are observed between PG and PGNSE ( $0.35 \pm 1.29$ ). PG ( $0.76 \pm 1.34$ ) outperforms PGNSE ( $0.59 \pm 1.38$ ) in style-appropriateness, likely because PGNSE cannot ensure a consistent style throughout the sequence. Conversely, the BEAT dataset exhibits significant differences ( $p < 0.001$ ) between PG ( $0.56 \pm 1.14$ ,  $0.63 \pm 1.10$ , and  $0.66 \pm 1.16$ ) and the other methods, indicating the superior capability of the global fuzzy feature inference mechanism in capturing stylistic nuances. Moreover, while PGOnehot can capture various logo gesture styles, it may compromise the naturalness of the movements.

2) *Ablation of AdaLN*: We integrated Cross-Attention (PGCA), In-Context Conditions (PGICC), and Concatenation of Features (PGCF) into our analysis to evaluate AdaLN’s effectiveness. This structured approach enabled a comprehensive assessment of each component’s contribution to the model’s overall performance and its role in audio-based gesture generation. The implementations of Cross-Attention and In-Context Conditions follow the designs in [17], while Feature Concatenation combines gesture and encoded audio features along the feature axis, a technique proven effective in related studies [51]. Separate user studies were conducted for each component, with findings presented in Table I and Fig. 11.

In the ablation studies concerning AdaLN, replacing the AdaLN module with alternative architectural frameworks precipitated a significant degradation in performance across all metrics. This reduction in efficacy can be ascribed to the deficiency of alternative architectures in synchronizing speech rhythm and capturing stylistic nuances with precision. This outcome underscores the pivotal role of a uniform mechanism that applies an identical function across all attention layers throughout the sequence.

### F. Generalization and Robustness

In addition, we test our method’s generalization capabilities. We utilized in-the-wild speech audio collected from TED talks. Our system adeptly generates consistent gestures from dataset types and seamlessly produces gestures from untagged, in-the-wild audio. It also showcases remarkable robustness against various auditory disturbances, such as background music, applause, urban noise, and decorative sounds. This adaptability highlights the system’s ability to handle a broad spectrum of audio inputs, ensuring the creation of naturalistic gestures despite significant

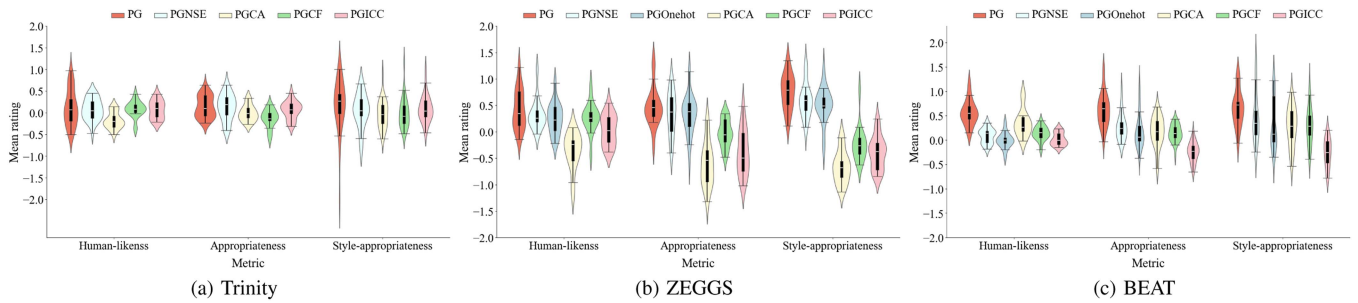


Fig. 11. The mean metric rating for each approach across the three datasets in ablation experiments.

noise interference. Such resilience emphasizes the system’s suitability for real-world applications. Yet, we encountered certain inherent challenges in assessing the generalizability and robustness of alternative models during our experimentation. More details can be found in Appendix D, available online, and supporting videos.

## V. DISSCUSSION AND CONCLUSION

In this work, we introduce *Persona-Gestor*, a novel network architecture designed to generate personality gestures, leveraging solely raw speech audio. At its core, *Persona-Gestor* combines a fuzzy feature extractor and an AdaLN transformer diffusion architecture.

The fuzzy feature extractor utilizes a fuzzy feature inference strategy in the dual-component module to implicitly infer both fuzzy stylistic features and specific details embedded within the audio data autonomously. These elements are combined into a unified latent representation, facilitating the generation of speaker-aware personalized 3D full-body gestures. This approach incorporates a highly influential feature into the capability to synthesize personality gestures through automatically inferred fuzzy features, removing the necessity for explicit style labels or additional features. This advancement facilitates the end-to-end generation of gestures that resonate with the speaker’s unique characteristics, directly from raw speech audio. Integrating fuzzy feature inference ensures a seamless and intuitive creation process that enhances generalization and user accessibility.

The AdaLN mechanism is a conditional mechanism that uniformly applies a specific function across all sequence tokens. This strategic incorporation significantly augments the model’s proficiency in accurately capturing and representing both conditional dependencies and output characteristics more efficiently. We demonstrate that AdaLN also facilitates a refined understanding and processing of the complex interplay between the continuous fuzzy features conditional input and the resultant gesture synthesis, leading to enhanced model performance and output fidelity. Ultimately, *Persona-Gestor* utilizes diffusion mechanism for producing a diverse spectrum of gesture outputs.

Our approach presents multiple benefits: 1) It exclusively uses raw speech audio to synthesize speaker-aware personalized gestures, bypassing the requirement for extra inputs, which

enhances user-friendliness. 2) It achieves the full-body (including finger motions and locomotion) gestures’ superior synchronization with speech, capturing rhythm, intonation, and certain semantics without compromising naturalness. 3) It showcases improved generalization and robustness, adapting effectively across varied conditions.

Our study highlights key areas for enhancement: First, the model’s sole dependence on speech audio may limit its effectiveness in capturing style features within segments of minimal speech. Second, the lack of control over the movement path and orientation of the digital human could lead to unintended gestures. Third, our model may not effectively replicate certain gestures crucial for expressing specific states. These observations underscore the necessity for improvements to broaden the model’s ability to accurately convey a wide range of human gestures.

## REFERENCES

- [1] U. Bhattacharya, E. Childs, N. Rewkowski, and D. Manocha, “Speech2AffectiveGestures: Synthesizing co-speech gestures with generative adversarial affective expression learning,” in *Proc. 29th ACM Int. Conf. Multimedia*, 2021, pp. 2027–2036.
- [2] S. Yang et al., “DiffuseStyleGesture: Stylized audio-driven co-speech gesture generation with diffusion models,” 2023, *arXiv:2305.04919*.
- [3] S. Yang et al., “The DiffuseStyleGesture+ entry to the GENE challenge 2023,” 2023, *arXiv:2308.13879*.
- [4] S. Alexanderson, R. Nagy, J. Beskow, and G. E. Henter, “Listen, denoise, action! audio-driven motion synthesis with diffusion models,” *ACM Trans. Graph.*, vol. 42, no. 4, pp. 1–20, 2023.
- [5] S. Yang et al., “UnifiedGesture: A unified gesture synthesis model for multiple skeletons,” in *Proc. 31st ACM Int. Conf. Multimedia*, 2023, pp. 1033–1044.
- [6] J. Li et al., “Audio2Gestures: Generating diverse gestures from audio,” *IEEE Trans. Vis. Comput. Graph.*, to be published, doi: [10.1109/TVCG.2023.3276973](https://doi.org/10.1109/TVCG.2023.3276973).
- [7] S. Ghorbani, Y. Ferstl, D. Holden, N. F. Troje, and M.-A. Carbonneau, “ZeroEGGS: Zero-shot example-based gesture generation from speech,” *Comput. Graph. Forum*, vol. 42, no. 1, pp. 206–216, 2023.
- [8] L. A. Zadeh, “Fuzzy sets,” *Inf. Control*, vol. 8, no. 3, pp. 338–353, 1965.
- [9] S. Vashishtha and S. Susan, “Unsupervised fuzzy inference system for speech emotion recognition using audio and text cues (workshop paper),” in *Proc. IEEE 6th Int. Conf. Multimedia Big Data*, 2020, pp. 394–403.
- [10] N. M. Patil and M. U. Nemade, “Content-based audio classification and retrieval using segmentation, feature extraction and neural network approach,” in *Proc. Int. Conf. Comput. Commun. Comput. Sci.*, Springer, 2019, pp. 263–281.
- [11] R. A. Calvo, S. D’Mello, J. M. Gratch, and A. Kappas, *The Oxford Handbook of Affective Computing*. London, U.K.: Oxford Univ. Press, 2015.

- [12] M. Goudbeek and K. Scherer, "Beyond arousal: Valence and potency/control cues in the vocal expression of emotion," *J. Acoustical Soc. Amer.*, vol. 128, no. 3, pp. 1322–1336, 2010.
- [13] J. Hirschberg and C. D. Manning, "Advances in natural language processing," *Science*, vol. 349, no. 6245, pp. 261–266, 2015.
- [14] K. Campbell-Kibler, "Intersecting variables and perceived sexual orientation in men," *Amer. Speech*, vol. 86, no. 1, pp. 52–68, 2011.
- [15] L. Zhu, X. Liu, X. Liu, R. Qian, Z. Liu, and L. Yu, "Taming diffusion models for audio-driven co-speech gesture generation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2023, pp. 10 544–10 553.
- [16] T. Ao, Z. Zhang, and L. Liu, "GestureDiffuCLIP: Gesture diffusion model with CLIP latents," 2023, *arXiv:2303.14613*.
- [17] W. Peebles and S. Xie, "Scalable diffusion models with transformers," 2022, *arXiv:2212.09748*.
- [18] J. Windle, D. Greenwood, and S. Taylor, "UEA digital humans entry to the GENEA challenge 2022," in *Proc. Int. Conf. Multimodal Interact.*, 2022, pp. 771–777.
- [19] L. Yu, H. Xie, and Y. Zhang, "Multimodal learning for temporally coherent talking face generation with articulator synergy," *IEEE Trans. Multimedia*, vol. 24, pp. 2950–2962, Jun. 2022.
- [20] S. Alexanderson, G. E. Henter, T. Kucherenko, and J. Beskow, "Style-controllable speech-driven gesture synthesis using normalising flows," *Comput. Graph. Forum*, vol. 39, pp. 487–496, 2020.
- [21] S. Taylor, J. Windle, D. Greenwood, and I. Matthews, "Speech-driven conversational agents using conditional flow-VAEs," in *Proc. Eur. Conf. Vis. Media Prod.*, 2021, pp. 1–9.
- [22] F. Zhang, N. Ji, F. Gao, and Y. Li, "DiffMotion: Speech-driven gesture synthesis using denoising diffusion model," in *Proc. 29th Int. Conf. MultiMedia Model.*, Springer, 2023, pp. 231–242.
- [23] S. Yang et al., "The ReprGesture entry to the GENEA challenge 2022," in *Proc. Int. Conf. Multimodal Interact.*, 2022, pp. 758–763.
- [24] S. Yang et al., "QPGesture: Quantization-based and phase-guided motion matching for natural speech-driven gesture generation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2023, pp. 2321–2330.
- [25] Y. Yoon, W.-R. Ko, M. Jang, J. Lee, J. Kim, and G. Lee, "Robots learn social skills: End-to-end learning of co-speech gesture generation for humanoid robots," in *Proc. Int. Conf. Robot. Automat.*, 2019, pp. 4303–4309.
- [26] T. Kucherenko et al., "Gesticulator: A framework for semantically-aware speech-driven gesture generation," in *Proc. Int. Conf. Multimodal Interact.*, 2020, pp. 242–250.
- [27] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in *Proc. Conf. North Amer. Chapter Assoc. Comput. Linguistics: Hum. Lang. Technol.*, 2019, pp. 4171–4186.
- [28] U. Bhattacharya, N. Rewkowski, A. Banerjee, P. Guhan, A. Bera, and D. Manocha, "Text2Gestures: A transformer-based network for generating emotive body gestures for virtual agents," in *Proc. IEEE Virtual Reality 3D User Interfaces*, 2021, pp. 1–10.
- [29] J. Pennington, R. Socher, and C. D. Manning, "GloVe: Global vectors for word representation," in *Proc. Conf. Empirical Methods Natural Lang. Process.*, 2014, pp. 1532–1543.
- [30] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean, "Distributed representations of words and phrases and their compositionality," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2013, pp. 3111–3119.
- [31] P. Bojanowski, E. Grave, A. Joulin, and T. Mikolov, "Enriching word vectors with subword information," *Trans. Assoc. Comput. Linguistics*, vol. 5, pp. 135–146, 2017.
- [32] J. Russell, "A circumplex model of affect," *J. Pers. Social Psychol.*, vol. 39, pp. 1161–1178, Dec. 1980.
- [33] E. Cambria, A. Livingstone, and A. Hussain, "The hourglass of emotions," in *Cognitive Behavioural Systems: COST 2102 International Training School, Dresden, Germany, February 21–26, 2011, Revised Selected Papers*. Berlin, Germany: Springer, 2012, pp. 144–157.
- [34] A. Gulati et al., "Conformer: Convolution-augmented transformer for speech recognition," in *Proc. 21st Annu. Conf. Int. Speech Commun. Assoc.*, 2020, pp. 5036–5040.
- [35] Y. Zhi et al., "LivelySpeaker: Towards semantic-aware co-speech gesture generation," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2023, pp. 20 807–20 817.
- [36] S. Chen et al., "WavLM: Large-scale self-supervised pre-training for full stack speech processing," *IEEE J. Sel. Topics Signal Process.*, vol. 16, no. 6, pp. 1505–1518, Oct. 2022.
- [37] L. Kaiser, A. N. Gomez, and F. Chollet, "Depthwise separable convolutions for neural machine translation," 2017, *arXiv: 1706.03059*.
- [38] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochreiter, "GANs trained by a two time-scale update rule converge to a local Nash equilibrium," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2017, pp. 6629–6640.
- [39] L. Paul, "sur la théorie du mouvement brownien," *Comptes Rendus Acad. Sci.*, vol. 65, no. 11, pp. 530–533, 1908.
- [40] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2020, pp. 6840–6851.
- [41] K. Rasul, C. Seward, I. Schuster, and R. Vollgraf, "Autoregressive denoising diffusion models for multivariate probabilistic time series forecasting," in *Proc. Int. Conf. Mach. Learn.*, 2021, pp. 8857–8868.
- [42] Y. Ferstl and R. McDonnell, "Investigating the use of recurrent motion modelling for speech gesture generation," in *Proc. 18th Int. Conf. Intell. Virtual Agents*, 2018, pp. 93–98.
- [43] H. Liu et al., "BEAT: A large-scale semantic and emotional multi-modal dataset for conversational gestures synthesis," in *Proc. Eur. Conf. Comput. Vis.*, 2022, pp. 612–630.
- [44] F. S. Grassia, "Practical parameterization of rotations using the exponential map," *J. Graph. Tools*, vol. 3, no. 3, pp. 29–48, 1998.
- [45] U. Wennberg and G. E. Henter, "The case for translation-invariant self-attention in transformer-based language models," in *Proc. 59th Annu. Meeting Assoc. Comput. Linguistics 11th Int. Joint Conf. Natural Lang. Process.*, 2021, pp. 130–140.
- [46] P. Wolfert, N. Robinson, and T. Belpaeme, "A review of evaluation practices of gesture generation in embodied conversational agents," *IEEE Trans. Human-Mach. Syst.*, vol. 52, no. 3, pp. 379–389, Jun. 2022.
- [47] T. Kucherenko et al., "Evaluating gesture-generation in a large-scale open challenge: The GENEA challenge 2022," 2023, *arXiv:2303.08737*.
- [48] P. Wolfert, J. M. Girard, T. Kucherenko, and T. Belpaeme, "To rate or not to rate: Investigating evaluation methods for generated co-speech gestures," in *Proc. Int. Conf. Multimodal Interact.*, 2021, pp. 494–502.
- [49] Y. Yoon et al., "Speech gesture generation from the trimodal context of text, audio, and speaker identity," *ACM Trans. Graph.*, vol. 39, no. 6, pp. 1–16, 2020.
- [50] R. Li, S. Yang, D. A. Ross, and A. Kanazawa, "AI choreographer: Music conditioned 3D dance generation with AIST++," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2021, pp. 13 401–13 412.
- [51] M. Zhao, J. Ning, J. Hu, and T. Li, "Attention-driven dual feature guidance for hyperspectral super resolution," *IEEE Trans. Geosci. Remote Sens.*, vol. 61, 2023, Art. no. 5525116.



**Fan Zhang** currently working toward the PhD degree in digital media with the Faculty of Humanities and Arts, Macau University of Science and Technology. He holds a position with the College of Media Engineering, the Communication University of Zhejiang, China. His research interests include artificial intelligence, human-computer interaction, and digital entertainment.



**Zhaohan Wang** is currently working toward the master's degree in digital art with the Communication University of China. His works have centered on 3D engines, digital media art theory, and virtual reality imaging. His research interests lie in the areas of virtual/mixed reality interaction, digital virtual humans, artificial intelligence, and other related fields.



**Xin Lyu** is a professor and PhD supervisor with the School of Animation and Digital Arts, Communication University of China, and serves as the director of the Digital Human Research Institute. His main research areas include virtual digital humans and intelligent product design.



**Hui Du** received the MSc degree from the Zhejiang University of Technology, and the PhD degree from the State Key Lab of CAD & CG, Zhejiang University. He is currently a faculty member with the Communication University of Zhejiang, China. His research interests include image processing and computer vision.



**Siyuan Zhao** received the PhD degree in digital media from the Macau University of Science and Technology. Currently employed with the Macau University of Science and Technology. Research interests include esports, digital media, artificial intelligence, and digital art innovation.



**Fuxing Gao** received the master's degree in art from the China Academy of Art. He currently a lecturer with the Communication University of Zhejiang. He is mainly engaged in teaching and applied research in the fields of digital interactive art, virtual human, and digital game art.



**Mengjian Li** received the PhD degree in computer science and technology from Zhejiang University, China. He is an engineering specialist with the Research Center for Data Hub and Security, Zhejiang Lab, China. His research is multimodal intelligent computing and large language models for digital humans.



**Hao Wu** is currently working toward the undergraduate degree with the Communication University of Zhejiang, China. His academic interests are focused on large language models and the AIGC multimodal.



**Weidong Geng** received the PhD degree from the Computer Science and Engineering Department, Zhejiang University, China. He is a professor with the College of Computer Science & Technology, Zhejiang University, China. His current research focuses on digital human, perceptual user interface, interactive media, and AIGC.



**Shunman Li** received the master's degree in management from South China Normal University and holds a national Level II certificate in psychological counseling. Her area of expertise is cognitive psychology.



**Naye Ji** received the PhD degree from the Institute of Computing Technology, Chinese Academy of Sciences, China. She is an associate professor with the Communication University of Zhejiang, China. Her research interests include image and video stylization, visual media editing and synthesis, and computer vision.