

Generating Holistic 3D Human Motion from Speech

Hongwei Yi^{1*} Hualin Liang^{2*} Yifei Liu^{2*} Qiong Cao^{3†}
Yandong Wen¹ Timo Bolkart¹ Dacheng Tao³ Michael J. Black^{1†}
¹Max Planck Institute for Intelligent Systems, Tübingen, Germany
²South China University of Technology ³JD Explore Academy
{hongwei.yi, yandong.wen, timo.bolkart, black}@tuebingen.mpg.de
{hualinliang3, yifei9697, mathqiong2012, dacheng.tao}@gmail.com



Figure 1. **Speech-to-motion translation example.** Given a speech signal as input, our approach generates realistic, coherent, and diverse holistic body motions; that is, the body motion together with facial expressions and hand gestures. From top to bottom: the input audio, the corresponding transcript, video frames, and the generated motions. Note that the audio is the only input to our approach, while the transcript and video frames are just shown for reference.

Abstract

This work addresses the problem of generating 3D holistic body motions from human speech. Given a speech recording, we synthesize sequences of 3D body poses, hand gestures, and facial expressions that are realistic and diverse. To achieve this, we first build a high-quality dataset of 3D holistic body meshes with synchronous speech. We then define a novel speech-to-motion generation framework in which the face, body, and hands are modeled separately. The separated modeling stems from the fact that face articulation strongly correlates with human speech, while body poses and hand gestures are less correlated. Specifically, we employ an autoencoder for face motions, and a compositional vector-quantized variational autoencoder (VQ-VAE) for the body and hand motions. The compositional VQ-VAE is key to generating diverse results. Additionally, we propose a cross-conditional autoregressive model that generates body poses and hand gestures, leading to coherent and

realistic motions. Extensive experiments and user studies demonstrate that our proposed approach achieves state-of-the-art performance both qualitatively and quantitatively. Our dataset and code are released for research purposes at <https://talkshow.is.tue.mpg.de/>.

1. Introduction

From linguistics and psychology we know that humans use body language to convey emotion and use gestures in communication [25, 32]. Motion cues such as facial expression, body posture and hand movement all play a role. For instance, people may change their gestures when shifting to a new topic [59], or wave their hands when greeting an audience. Recent methods have shown rapid progress on modeling the translation from human speech to body motion, and can be roughly divided into rule-based [43] and learning-based [23, 24, 26, 36, 37, 62] methods. Typically, the body motion in these methods is represented as the motion of a 3D mesh of the face/upper-body [5, 18, 30, 50, 51], or 2D/3D landmarks of the face with 2D/3D joints of the

*Equal Contribution.

†Joint Corresponding Authors.

hands and body [24, 26, 62]. However, this is not sufficient to understand human behavior. Humans communicate with their bodies, hands and facial expressions together. Capturing such coordinated activities as well as the full 3D surface in tune with speech is critical for virtual agents to behave realistically and interact with listeners meaningfully.

In this work, we focus on generating the expressive 3D motion of person, including their body, hand gestures, and facial expressions, from speech alone; see Fig. 1. To do this, we must learn a cross-modal mapping between audio and 3D holistic body motion, which is very challenging in practice for several reasons. First, datasets of 3D holistic body meshes and synchronous speech recordings are scarce. Acquiring them in the lab is expensive and doing so in the wild has not been possible. Second, real humans often vary in shape, and their faces and hands are highly deformable. It is not trivial to generate both realistic and stable results of 3D holistic body meshes efficiently. Lastly, as different body parts correlate differently with speech signals, it is difficult to model the cross-modal mapping and generate realistic and diverse holistic body motions.

We address the above challenges and learn to model the conversational dynamics in a data-driven way. Firstly, to overcome the issue of data scarcity, we present a new set of 3D holistic body mesh annotations with synchronous audio from in-the-wild videos. This dataset was previously used for learning 2D/3D gesture modeling with 2D body keypoint annotations [24] and 3D keypoint annotations of the holistic body [26] by applying existing models separately. Apart from facilitating speech and motion modeling, our dataset can also support broad research topics like realistic digital human rendering. Then, to support our data-driven approach to modeling speech-to-motion translation, an accurate holistic body mesh is needed. Existing methods have focused on capturing either the body shape and pose isolated from the hands and face [9, 20, 29, 38, 54, 64, 65, 68], or the different parts together, which often produces unrealistic or unstable results, especially when applied to video sequences [21, 46, 69]. To solve this, we present SHOW, which stands for “Synchronous Holistic Optimization in the Wild”. Specifically, SHOW adapts SMPLify-X [46] to the videos of talking persons, and further improves it in terms of stability, accuracy, and efficiency through careful design choices. Figure 9 shows example reconstruction results.

Lastly, we investigate the translation from audio to 3D holistic body motion represented as a 3D mesh (Fig. 1). We propose TalkSHOW, the first approach to autoregressively synthesize realistic and diverse 3D body motions, hand gestures and facial expression of a talking person from speech. Motivated by the fact that the face (i.e. mouth region) is strongly correlated with the audio signal, while the body and hands are less correlated, or even uncorrelated, TalkSHOW designs separate motion generators for different parts and

gives each part full play. For the face part, to model the highly correlated nature of phoneme-to-lip motion, we design a simple encoder-decoder based face generator that encodes rich phoneme information by incorporating the pre-trained wav2vec 2.0 [7]. On the other hand, to predict the non-deterministic body and hand motions, we devise a novel VQ-VAE [57] based framework to learn a compositional quantized space of motion, which efficiently captures a diverse range of motions. With the learned discrete representation, we further propose a novel autoregressive model to predict a multinomial distribution of future motion, cross-conditioned between existing motions. From this, a wide range of motion modes representing coherent poses can be sampled, leading to realistic looking motion generation.

We quantitatively evaluate the realism and diversity of our synthesized motion compared to ground truth and baseline methods and ablations. To further corroborate our qualitative results, we evaluate our approach through an extensive user study. Both quantitative and qualitative studies demonstrate the state-of-the-art quality of our speech-synthesized full expressive 3D character animations.

2. Related work

2.1. Holistic Body Reconstruction

Recent work addresses the problem of 3D holistic body mesh recovery [14, 28, 46, 61, 69]. SMPLify-X [46] fits the parametric and expressive SMPL-X model [46] to 2D keypoints obtained by off-the-shelf detectors (e.g. OpenPose [10]). PIXIE [21] directly regresses SMPL-X parameters using moderators that estimate the confidence of part-specific features. These features are fused and fed to independent regressors. PyMAF-X [69] improves the body and hand estimation with spatial alignment attention. In this work, we adapt the optimization-based SMPLify-X to videos of talking persons, and improve the stability and accuracy with several good engineering practices in terms of initialization, data term design, and regularization.

2.2. Speech-to-Motion Datasets

The existing speech-to-motion datasets can be roughly categorized as in-house and in-the-wild. The annotations of in-house datasets [15, 19, 23, 55, 60] are accurate but are limited in scale since the multi-camera systems used for data capture are expensive and labor intensive. Moreover, these datasets only provide annotations of the head [15, 19, 60] or body [23, 55], and thus do not support whole-body generation. To learn richer and more diverse speaking styles and emotions, [66, 67] propose to use in-the-wild videos. The annotations are pseudo ground truth (p-GT) given by advanced reconstruction approaches, e.g. [10]. However, these released datasets use either 2D keypoints or 3D keypoints with 3D head mesh to represent the body. This disconnected

Dataset	Head	Hand	Body	Holistic Body Connection	In-the-wild	Length	Annotations
Multiface [60]	3D mesh	✗	✗	✗	✗	-	multi-camera
BIWI [19]	3D mesh	✗	✗	✗	✗	-	3D-scanner
VOCASET [15]	3D mesh	✗	✗	✗	✗	-	4D-scan
Takeuchi et.al [55]	✗	✗	3D keypoint	✗	✗	5h	MoCap
Trinity [23]	✗	✗	3D keypoint	✗	✗	4h	MoCap
Yoon et.al [66, 67]	✗	✗	3D keypoint	✗	✓	52h	p-GT
Speech2Gesture [24]	✗	2D keypoint	2D keypoint	✗	✓	144h	p-GT
Habibie et.al [26]	3D mesh	3D keypoint	3D keypoint	✗	✓	33h	p-GT
Ours	3D mesh	3D mesh	3D mesh	✓	✓	27h	p-GT

Table 1. Comparison of different speech-to-motion datasets.

representation limits the possible applications of the generated talking motions. In contrast to the aforementioned work, our dataset, reconstructed by SHOW, consists of holistic body meshes and synchronized speech, covering a wide range of body poses, hand gestures, and facial expressions. More details can be found in Table 1.

2.3. Holistic Body Motion Generation from Speech

Holistic body motion generation from speech consists of three body parts motion generation, i.e., faces, hands, and bodies. Existing 3D talking face generation methods [15, 18, 50, 70] rely heavily on 4D face scan datasets for training [15, 19, 50]. There are many attempts to perform body motion generation, and these can be divided into rule-based and learning-based methods. Rule-based methods [11, 33, 37, 47] map the input speech to pre-collected body motion “units” with manually designed rules. They are explainable and controllable but it is expensive to create complex, realistic, motion patterns. Learning-based body motion generation approaches [1, 8, 24, 34, 35, 40, 66] have advanced significantly in part due to publicly released synchronous speech and body motion datasets [24, 26, 39, 55, 66, 67]. However, they only consider parts of the human body rather than the holistic body. Most related to our work, Habibie et al. [26] propose to generate 3D facial meshes and 3D keypoints of the body and hands from speech, but the generated faces, bodies and hands are disjoint. Also, these methods are deterministic, can not generate diverse motions when given the same speech recording. There are a few attempts to incorporate the diversity into motion generation using GANs [2, 41, 67], VAEs [39, 49, 62], VQ-VAEs [5, 63], or normalising-flows [3]. Nevertheless, the diversity of motions produced by these methods is inadequate.

In contrast, TalkSHOW generates holistic body motions and models different body parts separately according to their natures: the face part is more correlated to the speech signal than body parts. TalkSHOW develops a simple deterministic encoder-decoder structure for mapping acoustic signals to facial expressions. TalkSHOW adopts two VQ-VAEs to generate more diverse body and hand motions. This novel de-

sign allows the learned quantized space to be compositional and more expressive for conversational gestures. Compared with previous VQ-VAE-based methods [53, 63], we design a cross-conditional autoregressive model to generate different body-part motions, which are more fluid and natural.

3. Dataset

In this section, we introduce a high-quality audiovisual dataset, which consists of expressive 3D body meshes at 30fps, and their synchronized audio at a 22K sample rate. The 3D body meshes are reconstructed from in-the-wild monocular videos and are used as our pseudo ground truth (p-GT) in speech-to-motion generation. We provide detailed descriptions of this dataset in Sec. 3.1 and highlight several good practices for obtaining more accurate p-GT from videos in Sec. 3.2. Our experiments show that this dataset is effective for training speech-to-motion models.

3.1. Dataset Description

The dataset is built from the in-the-wild talking videos of different people with various speaking styles. We use the same video sources from [24] for straightforward comparisons with the previous work. To facilitate the subsequent 3D body reconstruction, we manually filter out videos if they are in any following cases: (i) low resolution (<720p), (ii) occluded hand(s), or (iii) invalid download link. The filtering leads to a high-quality dataset of 26.9 hours from 4 speakers. For the mini-batch processing, the raw videos are cropped into short clips (<10 seconds). Direct comparisons to the existing datasets can be found in Table 1.

Expressive 3D whole-body meshes are reconstructed from these videos and used as the p-GT. Specifically, the 3D holistic body meshes consist of face, hands, and bodies in a connected way, which is achieved by adopting a well-designed 3D topology from SMPL-X [46]. As a result, we represent the p-GT of the dataset as SMPL-X parameters. Given a video clip of T frames, the p-GT comprises parameters of a shared body shape $\beta \in \mathbb{R}^{300}$, poses $\{\theta_t | \theta_t \in \mathbb{R}^{156}\}_{t=1}^T$, a shared camera pose $\theta^c \in \mathbb{R}^3$ and translation $\epsilon \in \mathbb{R}^3$, and facial expressions $\{\psi_t | \psi_t \in \mathbb{R}^{100}\}_{t=1}^T$.



Figure 2. The 3D holistic body reconstruction results from SMPLify-X, PIXIE, PyMAF-X, and ours. Compared to other methods, ours produces more accurate and stable results with details.

Here the pose θ_t includes the jaw pose $\theta_t^{jaw} \in \mathbb{R}^3$, the body pose $\theta_t^b \in \mathbb{R}^{63}$, and the hand pose $\theta_t^h \in \mathbb{R}^{90}$.

We note that this dataset can not only be used in speech and motion modeling, but also supports broad research topics like realistic digital human rendering and learning-based holistic body recovery from videos, etc.

3.2. Good Practices for Improving p-GT

In this section, we present SHOW, which adapts SMPLify-X [46] to the videos of talking persons with several good practices, to improve the stability, accuracy, and efficiency in 3D whole-body reconstruction. In the following, we briefly summarize our efforts for improving the p-GT. See more details in the supplemental material.

Initialization. A good initialization can significantly accelerate and stabilize the SMPLify-X optimization. We apply several advanced regression-based approaches to the videos, and use the resulting predictions as the initial parameters of SMPLify-X. Specifically, PIXIE [21], PyMAF-X [69], and DECA [22] are used to initialize θ^b , θ^h , and θ^f , respectively. The camera is assumed to be static, and its parameters θ^c and ϵ are estimated by PIXIE [21] as well.

Data Term. The joint re-projection loss is the most important data objective function in SMPLify-X, as it optimizes the difference between joints extracted from the SMPL-X model, projected into the image, with joints predicted with OpenPose [10]. Here we extend the data term by incorporating body silhouettes from DeepLab V3, facial landmarks from MediaPipe [31], and facial shapes from MICA [72]. Further, we use a photometric loss between the rendered faces and the input image to better capture facial details.

Regularization. Different regularization terms in SMPLify-X prevent the reconstruction of unrealistic bodies. To derive more reasonable regularizations, we explicitly take information about the video into account and make the following assumptions. First, the speaker in each video clip remains the same. This is further verified by a face recognition pipeline using the ArcFace model [16]. So we can use consistent shape parameters β to represent the holistic body shape. Second, the holistic body pose, facial expression, and environmental lighting in video clips change smoothly over time. This temporal smoothness assumption has proven useful in many previous approaches [65, 72], and we observe similar improvements in our experiments. Third, the person’s surface does self-penetrate, which should be self-evident in the real world.

Overall, as shown in Figure 9, the p-GT can be significantly improved by incorporating the aforementioned practices. See more results in the supplemental video.

4. Method

Given a speech recording, our goal is to generate conversational body poses, hand gestures as well as facial expressions that match the speech in a plausible way. Motivated by the fact that the face motion is highly correlated to the speech signal, while the body and hand parts are less correlated, we propose TalkSHOW, a novel framework that can model speech and different human parts separately. In the following, we present an encoder-decoder based face generator in Sec. 4.2, and a body and hand generator in Sec. 4.3.

4.1. Preliminary

Let $M_{1:T} = \{m_t\}_{t=1}^T$ be a p-GT holistic motion (i.e., a temporal sequence of the holistic body poses $m_t = \{\theta_t, \psi_t\}$) provided in Sec. 3. We denote the motion of the face, body and hands as $M_{1:T}^f$, $M_{1:T}^b$ and $M_{1:T}^h$ respectively. In particular, a facial motion $M_{1:T}^f$ is represented as a sequence of jaw poses and facial expression parameters $\{\theta_t^{jaw}, \psi_t\}_{t=1}^T$. And the body motion $M_{1:T}^b$ and the hands motion $M_{1:T}^h$ are denoted as a sequence of body poses $\{\theta_t^b\}_{t=1}^T$ and hand poses $\{\theta_t^h\}_{t=1}^T$, respectively.

4.2. Face Generator

Given a raw audio signal $A_{1:T}$ and speaker identity I , our face generator G_F aims to generate an expressive facial motion $\widehat{M}_{1:T}^f = (\widehat{m}_1^f, \dots, \widehat{m}_T^f) \in \mathbb{R}^{103 \times T}$ close to $M_{1:T}^f \in \mathbb{R}^{103 \times T}$.

Figure 3 (A) illustrates our idea. In order to produce synchronized mouth motions [18], we leverage a pretrained speech model, wav2vec 2.0 [7]. Specifically, the encoder consists of an audio feature extractor and a transformer encoder [58], leading to a 768-dimensional speech representation. A linear projection layer is added on top of the encoder

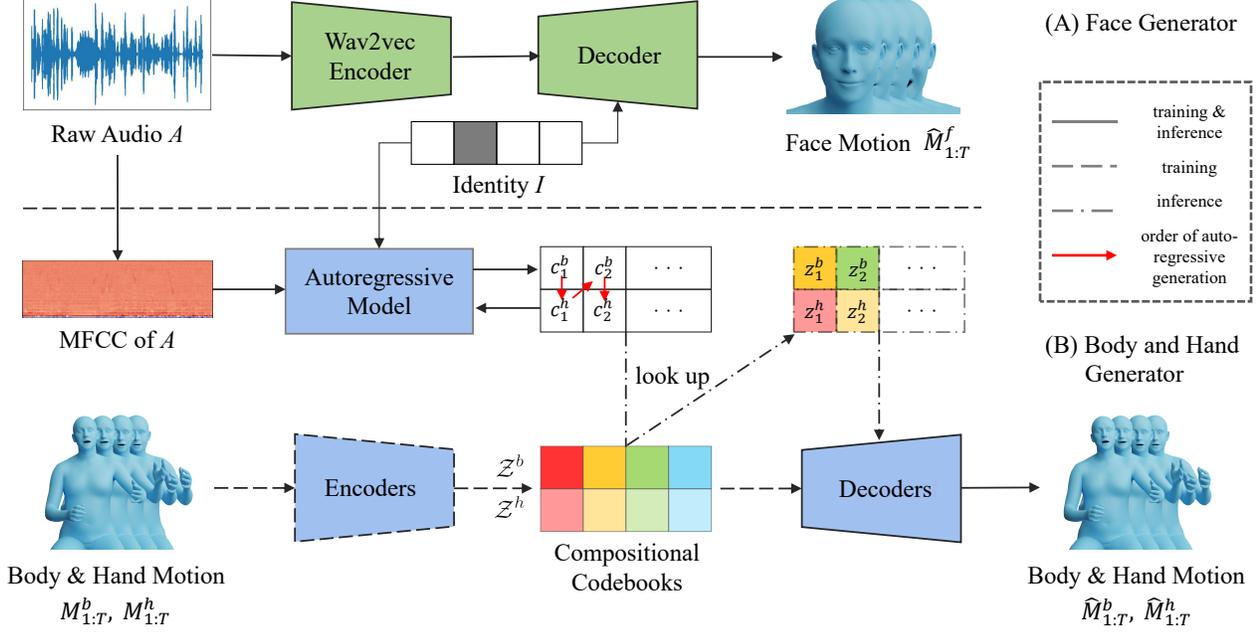


Figure 3. Overview of the proposed TalkSHOW. We employ a simple encoder-decoder model for face motions, and a novel framework for body and hand motions. Specifically, this framework first learns VQ-VAEs on each piece separately to obtain a compositional quantized space. Then, we autoregressively predict the code indices of the body or hand motion following the red arrows orders. Our predictor is designed to be cross-conditioned between the body and hand motions to keep the synchronization of the holistic body. Lastly, we look up codes in the codebook according to indices and decode them to obtain the body and hand motion. Colors are utilized to differentiate codes within the same codebook, while various shades of these colors are used to separate body and hand codes. Best viewed in color.

to reduce the feature dimension to 256. We then concatenate the audio feature with the speaker identity and feed them to the decoder. The speaker identity is represented as a one-hot vector $I \in \{0, 1\}^{N_I}$, where N_I is the number of speakers. Our decoder comprises six layers of temporal convolutional networks (TCNs) followed by a fully-connected layer. We train the encoder and decoder with an Mean Square Error (MSE) loss.

4.3. Body and Hand Generator

Given an audio input, we aim to generate a realistic and diverse motion for the body and hands, *i.e.* $\widehat{M}_{1:T}^b = (\widehat{m}_1^b, \dots, \widehat{m}_T^b) \in \mathbb{R}^{63 \times T}$ and $\widehat{M}_{1:T}^h = (\widehat{m}_1^h, \dots, \widehat{m}_T^h) \in \mathbb{R}^{90 \times T}$, respectively. Figure 3 (B) illustrates our idea. Instead of learning a direct mapping from audio to motion, we leverage the recent advances of VQ-VAE [57] to learn a multi-mode distribution space for the body and hand motions. Specifically, we first encode and quantize the body and hand motions into two finite codebooks, from which we can sample a wide range of plausible body and hand combinations. Then, we introduce a novel cross-conditional autoregressive model over the learned codebooks, which allows us to predict diverse body and hand motions. Our predictor is designed to be cross-conditioned between the body and hands to keep the synchronization of the holistic

body. Lastly, we obtain the future body/hand motion by decoding codebook indices sampled from the distribution.

Representation. We use 64-dimensional MFCC features [52] as the audio representation for body and hand motion generation, *i.e.*, $A_{1:T} = (a_1, \dots, a_T) \in \mathbb{R}^{64 \times T}$. Since body and hand gestures are more correlated to the rhythm and beat instead of phonemes, low-dimensional MFCC features are sufficient to produce plausible gestures from audio. Besides, considering that speakers often present different motion styles, we also leverage the modality of speaker identity I to differentiate those styles.

Compositional Quantized Motion Codebooks. The vanilla VQ-VAE learns a discrete codebook $\mathcal{Z} = \{z_i\}_{i=1}^{|\mathcal{Z}|}$ consisting of multiple vectors $z_i \in \mathbb{R}^{d_z}$ to quantize the latent space of the input. To further expand the range that the learned codebook can represent, we divide the body and hand motion into compositional pieces, *i.e.*, body and hands, and learn VQ-VAEs on each piece separately. By doing this, the body and hand motions are encoded and quantized into two separate finite codebooks $\mathcal{Z}^b = \{z_i^b\}_{i=1}^{|\mathcal{Z}^b|}$ and $\mathcal{Z}^h = \{z_j^h\}_{j=1}^{|\mathcal{Z}^h|}$, where $z_i^b, z_j^h \in \mathbb{R}^{d_z}$ with lengths $|\mathcal{Z}^b|$ and $|\mathcal{Z}^h|$ respectively. This approach enables us to obtain $|\mathcal{Z}^b| \times |\mathcal{Z}^h|$ different body-hand pose code pairs (z_i^b, z_j^h) and expand the range of motion diversity. In this

scheme, given the body and hand motions $M_{1:T}^b \in \mathbb{R}^{63 \times T}$ and $M_{1:T}^h \in \mathbb{R}^{90 \times T}$ as input, we first encode them into the feature sequence $E_{1:\tau}^b = (e_1^b, \dots, e_\tau^b) \in \mathbb{R}^{64 \times \tau}$ and $E_{1:\tau}^h = (e_1^h, \dots, e_\tau^h) \in \mathbb{R}^{64 \times \tau}$. Here, $\tau = \frac{T}{w}$, where w is the temporal window size, and w consecutive poses correspond to one feature embedding. In this paper, we set $w = 4$ to achieve a balance between inference speed and quality. Then, we quantize the embedding by mapping it into the nearest code in the corresponding codebook:

$$\begin{aligned} z_t^b &= \arg \min_{z_k^b \in \mathcal{Z}^b} \|e_t^b - z_k^b\| \in \mathbb{R}^{64}, \\ z_t^h &= \arg \min_{z_k^h \in \mathcal{Z}^h} \|e_t^h - z_k^h\| \in \mathbb{R}^{64}. \end{aligned} \quad (1)$$

Finally, the quantized features $Z_{1:\tau}^b = (z_1^b, \dots, z_\tau^b) \in \mathbb{R}^{64 \times \tau}$ and $Z_{1:\tau}^h = (z_1^h, \dots, z_\tau^h) \in \mathbb{R}^{64 \times \tau}$ are fed into the decoder for the synthesis.

We train the encoder, decoder, and codebook simultaneously with the following loss function:

$$\begin{aligned} \mathcal{L}_{VQ} &= \mathcal{L}_{rec}(M_{1:T}, \widehat{M}_{1:T}) + \|\text{sg}[E_{1:T}] - Z_{1:T}\| \\ &+ \beta \|E_{1:T} - \text{sg}[Z_{1:T}]\|, \end{aligned} \quad (2)$$

where \mathcal{L}_{rec} is an MSE reconstruction loss, sg is a stop gradient operation [13] that is used to calculate codebooks loss, and the third part is a ‘‘commitment’’ loss with a trade-off β .

Cross-Conditional Autoregressive Modeling. After we learn the compositional quantized codebooks, any body and hand motions can be represented as a sequence of codebook vectors via the encoder and quantization, which is denoted as $C_{1:\tau}^b = (c_1^b, \dots, c_\tau^b) \in \mathbb{R}^{|\mathcal{Z}^b| \times \tau}$ and $C_{1:\tau}^h = (c_1^h, \dots, c_\tau^h) \in \mathbb{R}^{|\mathcal{Z}^h| \times \tau}$.

Now, with the quantized motion representation, we design a temporal autoregressive model over it to predict the distribution of the possible next code, given the input audio embedding A and existing motions. Besides, we enable the modality input of identity I to distinguish different gesture styles. Because we model the body and hands independently, to keep the consistency of the holistic body and thus predict realistic gestures, we exploit the mutual information and design our model to be cross-conditioned between the body and hand motions. Specifically, following Bayes’ Rule, we model the joint probability of $C_{1:\tau}^b$ and $C_{1:\tau}^h$ as follows:

$$\begin{aligned} p(C_{1:\tau}^b, C_{1:\tau}^h | A_{1:\tau}, I) &= \prod_{t=1}^{\tau} p\left(c_t^b | c_{<t}^b, c_{<t}^h, a_{\leq t}, I\right) \\ &p\left(c_t^h | c_{\leq t}^b, c_{<t}^h, a_{\leq t}, I\right). \end{aligned} \quad (3)$$

Note that our cross-condition modeling between the body and hand motions makes the most of mutual information in two ways: (1) the current body/hand motion (i.e. c_t^b/c_t^h) depend on past body/hand motion information (i.e. $c_{<t}^b/c_{<t}^h$); (2) we argue that the current body motion c_t^b is also responsible for predicting the distribution of current hand motion.

Method	Face	
	L2 ↓	LVD ↓
Habibie et al. [26]	0.139	0.257
TalkSHOW (Ours)	0.130	0.248
Method	Body&Hands	
	RS ↑	Variation ↑
Habibie et al. [26]	0.146	0
Audio Encoder-Decoder	0.214	0
Audio VAE	0.182	0.044
Audio+Motion VAE	0.240	0.176
TalkSHOW (Ours)	0.414	0.821

Table 2. Comparison to Habibie et al. [26] and several baselines. ↑ indicates higher is better and ↓ indicates lower is better.

Such modeling guarantees the coherence of the body and hand motions as a whole and thus achieves realistic gestures. Gated PixelCNN [56] is adopted to model these quantities, in which the convolutional kernel is masked to make sure the model cannot read future information. During the training phase, the quantized body/hand motion representation concatenated with the audio and identity features is used for training. A teacher-forcing scheme and cross-entropy loss are adopted for the optimization. At inference, the model predicts multinomial distributions of the future body and hand motions, from which we can sample to acquire codebook indices for each motion. A codebook lookup is then conducted to retrieve the corresponding quantized element of motion, which we feed into the decoder for the final synthesis. Figure 3 (B) illustrates the pipeline. More training details are given in the supplemental material.

5. Experiments

We evaluate the ability of our method in generating body motions (i.e. a sequence of poses) from the speech on the created dataset both quantitatively and qualitatively. Specifically, we choose video sequences longer than 3s and split them into 80%/10%/10% for the train/val/test set. Several metrics are used to measure the realism and diversity of the generated motions including facial expression and poses. Furthermore, we conduct perceptual studies to assess the performance of our method.

5.1. Experimental Setup

Evaluation Metrics. Because we model face motion as a deterministic task and the body and hand motions as a non-deterministic task, we assess the generated motion in terms of the realism and the synchronization of face motion, and the realism and the diversity of body and hand motions. Specifically, the following metrics are adopted:

- *L2*: L2 distance between p-GT and generated facial landmarks, including jaw joints and lip shape [44, 71].

- *LVD*: Landmark Velocity Difference calculates the velocity difference between p-GT and generated facial landmarks, which measures the synchronization between the input speech and the facial expression [71].
- *RS*: Score on the realism of the generated body and hand motions. Following [4, 62], we trained a binary classifier to discriminate real samples from fake ones and the prediction represents the realistic score.
- *Variation*: As used in [44], diversity is measured by the variance across 16 samples of body and hand motions.

Compared Methods. We compare TalkSHOW to Habibie et al. [26], a SOTA speech-to-motion method. Also, we compare several baselines for modeling body and hand motions when using the same face generator as ours:

- **Audio Encoder-Decoder.** It encodes input audio and outputs motions; this is used by [24, 26].
- **Audio VAE.** Given the input audio, the VAE-like structure encodes audio into a Gaussian distribution, and then the sampled audio is fed into the decoder, which transforms the sample into motions.
- **Audio+Motion VAE.** Given the input motion and audio, it adopts a VAE-like structure with two encoders to encode motion and audio into Gaussian distributions, respectively, and then the sampled motion and audio are concatenated and fed into the decoder for the synthesis.

5.2. Quantitative Analysis

Table 2 shows the comparison results. We see that our method outperforms Habibie et al. [26] across all metrics. Particularly, our method surpasses it in terms of L2 and LVD, which demonstrates the effectiveness of our face generator for generating realistic facial expressions. Also, our method significantly outperforms it in terms of variation, which demonstrates the powerful capacity to generate diverse body and hand motions resulting from our proposed compositional quantized motion representation. Moreover, regarding the realism (RS) for body and hand motions, we surpass Habibie et al. [26] considerably, which confirms the effectiveness of our cross-conditioned autoregressive model in generating realistic motion.

On the other hand, compared to VAE-based models, our method achieves large gains in both realism and diversity. In particular, we obtain much higher diversity. This indicates the advantage of the learned compositional quantized motion codebooks, which effectively memorize multiple motion modes of the body and hands and thus boost the diversity of the generated body and hand gestures.

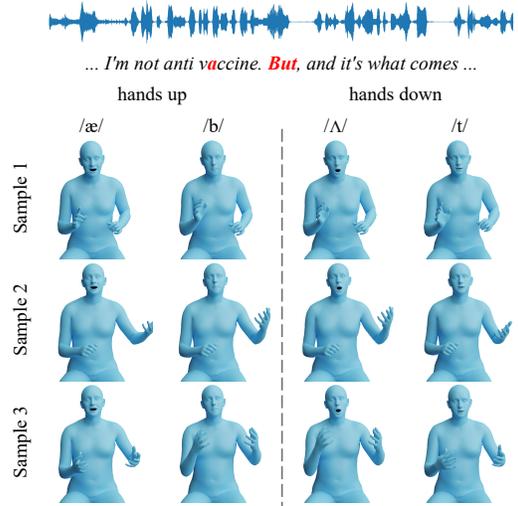


Figure 4. Our method generates diverse motions consistent with the rhythm of the input audio. For instance, we can generate different movements of hands corresponding to the strengthening tone of “*But*” in the speech, e.g. using left hand only (top), right hand only (middle), or both hands (bottom).

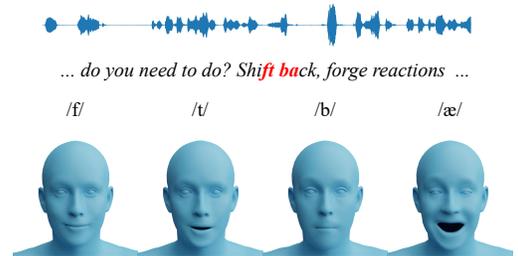


Figure 5. Given speech audio as input, our method generates facial expressions with accurate lip shapes.

5.3. Qualitative Analysis

Figure 4 shows examples of our generated 3D holistic body motion from speech. We see that given the word “*But*” from the speech which represents a strengthening tone of voice, our method generates plausible holistic body motions with hands up before saying “*But*” and hands down after saying “*But*”. Notably, the generated motions are diverse in many aspects, e.g. the range of motion and which hands to use, as shown in three different generated samples.

Figure 5 illustrates the qualitative performance of our face generator. Our approach generates realistic face motions including consistent lip motions with the corresponding phonemes such as /f/, /t/, /b/, and /æ/. Furthermore, our method exhibits a robust generalization ability to unseen languages and various audio types, e.g. French and songs. Additional interesting examples can be found in our supplemental video.

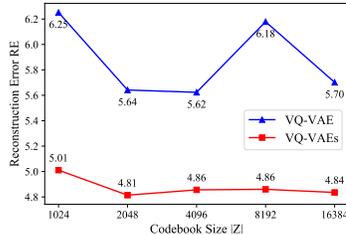


Figure 6. The comparison of VQ-VAE and VQ-VAEs with compositional codebooks.

5.4. Model Ablation

Effect of Wav2vec Feature in Face Generation. We evaluate the effect of the wav2vec feature used in face generation compared to the MFCC feature. We add an extra encoder to increase the dimension of the MFCC feature from 64 to 256 for a fair comparison. The wav2vec-based model outperforms the MFCC-based model in both metrics (0.130 vs. 0.165 in L2 and 0.251 vs. 0.277 in LVD) due to its larger capacity for modeling the relationship between audio and phonemes. Moreover, we experimentally find that the wav2vec-based model can generalize well to unseen identities; see supplementary for more details.

Effect of Compositional Quantized Motion Codebooks. We analyze the capability of capturing the diverse motion modes represented in motion data by our proposed compositional quantized motion codebooks of VQ-VAEs. To this end, we compare VQ-VAE with a single codebook. Reconstruction Error RE is adopted as the metric, in which a lower reconstruction loss indicates a higher capacity. Figure 6 illustrates the results. We see that compared to VQ-VAE with a single codebook, VQ-VAEs with compositional codebooks yield consistently lower RE across different codebook sizes. This demonstrates the effectiveness of the proposed compositional codebooks in modeling the diverse motion modes.

Effect of Cross-Conditional Modeling. In contrast to cross-conditional modeling (w/ c-c), the model without cross-condition (w/o c-c) generates body and hand motions independently. Our method w/ c-c yields a higher realistic score than that w/o c-c (0.414 vs. 0.409), benefiting from the cross-conditional modeling between the body and hand motions which leads to generating more coherent and realistic motions. Our method w/ c-c attains a slight reduction in diversity (0.821 vs. 0.922 in variance), however the method w/o c-c leads to implausible body and hands combination.

5.5. Perceptual Study

We conduct perceptual studies with Google Forms to evaluate our reconstruction and generation results, respectively. We randomly sample 40 videos in total with 10 videos from each speaker. Ten participants took part in the study.

Method	face	body	hands	holistic body
PyMAF-X [69]	0.323	0.500	0.438	0.193
SHOW (ours)	0.898	0.738	0.800	0.768

Table 3. Perceptual study results on reconstruction. For each method, we report the average percentage of answers that the reconstructed results match the input video.

Method	face	body and hands	holistic body
[26] vs. p-GT	0.153	0.141	0.169
TalkSHOW (Ours) vs. p-GT	0.478	0.464	0.458
TalkSHOW (Ours) vs. [26]	0.888	0.910	0.913

Table 4. Perceptual study of motion generation. We use A/B testing and report the percentage of answers where A is preferred over B.

Reconstruction. We assess the quality of our holistic body reconstruct results against PyMAF-X [69], compared with the ground truth. Participants are asked to answer the following questions with Yes or No: Does the reconstructed face/hands/body/full-body match the input video? Table 3 reports the average percentage of answers that the reconstructed results match the input video. We see that our method outperforms PyMAF-X by a large margin.

Holistic Body Motion Generation. We use A/B testing to evaluate our generation results, compared to the p-GT and Habibie et al. [26]. Specifically, participants are asked to answer the following questions with A or B: For the face/body&hands/overall region, which one is a better match with the given speech? Table 4 reports the average preference percentage of answers. We see that participants favor our method over Habibie et al. in terms of all the regions. Not surprisingly, participants perceive the p-GT better over both methods, with our method preferred by many more users.

6. Conclusion

In this work, we propose TalkSHOW, the first approach to generate 3D holistic body meshes from speech. We devise a simple and effective encoder-decoder for realistic face generation with accurate lip shape. For body and hands, we enable diverse generation and coherent prediction with compositional VQ-VAE and cross-conditional modeling, respectively. Moreover, we contribute a new set of accurate 3D holistic body meshes with synchronous audios from in-the-wild videos. The annotations are obtained by an empirical approach designed for videos. Experimental results demonstrate that our proposed approach achieves state-of-the-art performance both qualitatively and quantitatively.

Acknowledgments. We thank W. Zielonka and J. Thies for helping us incorporate MICA into SHOW, C. Ding, H. Jiang, Y. Feng, Z. Liu, and W. Liu for insightful discussions, and B. Pellkofer for IT support. This work was supported by the German Federal Ministry of Education and Research (BMBF): Tübingen AI Center, FKZ: 01IS18039B.

Disclosure. https://files.is.tue.mpg.de/black/CoL_CVPR_2023.txt

References

- [1] Chaitanya Ahuja, Dong Won Lee, Ryo Ishii, and Louis-Philippe Morency. No gestures left behind: Learning relationships between spoken language and freeform gestures. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1884–1895, 2020. 3
- [2] Chaitanya Ahuja, Dong Won Lee, Yukiko I Nakano, and Louis-Philippe Morency. Style transfer for co-speech gesture animation: A multi-speaker conditional-mixture approach. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XVIII 16*, pages 248–265. Springer, 2020. 3
- [3] Simon Alexanderson, Gustav Eje Henter, Taras Kucherenko, and Jonas Beskow. Style-controllable speech-driven gesture synthesis using normalising flows. In *Computer Graphics Forum*, volume 39, pages 487–496. Wiley Online Library, 2020. 3
- [4] Sadeqh Aliakbarian, Fatemeh Sadat Saleh, Mathieu Salzmann, Lars Petersson, and Stephen Gould. A stochastic conditioning scheme for diverse human motion prediction. In *Computer Vision and Pattern Recognition (CVPR)*, pages 5223–5232, 2020. 7
- [5] Tenglong Ao, Qingzhe Gao, Yuke Lou, Baoquan Chen, and Libin Liu. Rhythmic gesticulator: Rhythm-aware co-speech gesture synthesis with hierarchical neural embeddings. *arXiv preprint arXiv:2210.01448*, 2022. 1, 3
- [6] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv*, 2016. 14
- [7] Alexei Baevski, Yuhao Zhou, Abdelrahman Mohamed, and Michael Auli. wav2vec 2.0: A framework for self-supervised learning of speech representations. *Conference on Neural Information Processing Systems (NeurIPS)*, 33:12449–12460, 2020. 2, 4
- [8] Uttaran Bhattacharya, Elizabeth Childs, Nicholas Rewkowski, and Dinesh Manocha. Speech2affectivegestures: Synthesizing co-speech gestures with generative adversarial affective expression learning. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 2027–2036, 2021. 3
- [9] Federica Bogo, Angjoo Kanazawa, Christoph Lassner, Peter Gehler, Javier Romero, and Michael J Black. Keep it smpl: Automatic estimation of 3d human pose and shape from a single image. In *European conference on computer vision*, pages 561–578. Springer, 2016. 2
- [10] Zhe Cao, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Realtime multi-person 2d pose estimation using part affinity fields. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7291–7299, 2017. 2, 4, 13
- [11] Justine Cassell, Hannes Högni Vilhjálmsón, and Timothy Bickmore. Beat: the behavior expression animation toolkit. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 477–486, 2001. 3
- [12] Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking atrous convolution for semantic image segmentation. *arXiv*, 2017. 13, 14
- [13] Xinlei Chen and Kaiming He. Exploring simple siamese representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15750–15758, 2021. 6
- [14] Vasileios Choutas, Georgios Pavlakos, Timo Bolkart, Dimitrios Tzionas, and Michael J. Black. Monocular expressive body regression through body-driven attention. In *European Conference on Computer Vision (ECCV)*, pages 20–40, 2020. URL <https://expose.is.tue.mpg.de>. 2
- [15] Daniel Cudeiro, Timo Bolkart, Cassidy Laidlaw, Anurag Ranjan, and Michael Black. Capture, learning, and synthesis of 3D speaking styles. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 10101–10111, 2019. 2, 3
- [16] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *Computer Vision and Pattern Recognition (CVPR)*, pages 4690–4699, 2019. 4
- [17] Yu Deng, Jiaolong Yang, Sicheng Xu, Dong Chen, Yunde Jia, and Xin Tong. Accurate 3d face reconstruction with weakly-supervised learning: From single image to image set. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2019. 14
- [18] Yingruo Fan, Zhaojiang Lin, Jun Saito, Wenping Wang, and Taku Komura. Faceformer: Speech-driven 3d facial animation with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18770–18780, 2022. 1, 3, 4
- [19] Gabriele Fanelli, Juergen Gall, Harald Romsdorfer, Thibaut Weise, and Luc Van Gool. A 3-D Audio-Visual

- Corpus of Affective Communication. *IEEE Transactions on Multimedia*, 12(6), October 2010. ISSN 1941-0077. doi: 10.1109/TMM.2010.2052239. 2, 3
- [20] Qi Fang, Kang Chen, Yinghui Fan, Qing Shuai, Jiefeng Li, and Weidong Zhang. Learning analytical posterior probability for human mesh recovery. In *Computer Vision and Pattern Recognition (CVPR)*, June 2023. 2
- [21] Yao Feng, Vasileios Choutas, Timo Bolkart, Dimitrios Tzionas, and Michael J Black. Collaborative regression of expressive bodies using moderation. In *2021 International Conference on 3D Vision (3DV)*, pages 792–804. IEEE, 2021. 2, 4, 13, 14
- [22] Yao Feng, Haiwen Feng, Michael J Black, and Timo Bolkart. Learning an animatable detailed 3d face model from in-the-wild images. *ACM Transactions on Graphics (TOG)*, 40(4):1–13, 2021. 4, 13, 14
- [23] Ylva Ferstl and Rachel McDonnell. Investigating the use of recurrent motion modelling for speech gesture generation. In *Proceedings of the 18th International Conference on Intelligent Virtual Agents*, pages 93–98, 2018. 1, 2, 3
- [24] S. Ginosar, A. Bar, G. Kohavi, C. Chan, A. Owens, and J. Malik. Learning individual styles of conversational gesture. In *Computer Vision and Pattern Recognition (CVPR)*. IEEE, June 2019. 1, 2, 3, 7
- [25] Susan Goldin-Meadow. The role of gesture in communication and thinking. *Trends in cognitive sciences*, 3(11):419–429, 1999. 1
- [26] Ikhsanul Habibie, Weipeng Xu, Dushyant Mehta, Lingjie Liu, Hans-Peter Seidel, Gerard Pons-Moll, Mohamed Elgharib, and Christian Theobalt. Learning speech-driven 3d conversational gestures from video. In *Proceedings of the 21st ACM International Conference on Intelligent Virtual Agents*, pages 101–108, 2021. 1, 2, 3, 6, 7, 8, 15
- [27] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International Conference on Machine Learning (ICML)*, pages 448–456, 2015. 15
- [28] Hanbyul Joo, Tomas Simon, and Yaser Sheikh. Total capture: A 3d deformation model for tracking faces, hands, and bodies. In *Computer Vision and Pattern Recognition (CVPR)*, pages 8320–8329, 2018. 2
- [29] Angjoo Kanazawa, Michael J Black, David W Jacobs, and Jitendra Malik. End-to-end recovery of human shape and pose. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7122–7131, 2018. 2
- [30] Tero Karras, Timo Aila, Samuli Laine, Antti Herva, and Jaakko Lehtinen. Audio-driven facial animation by joint end-to-end learning of pose and emotion. *Transactions on Graphics (TOG)*, 36(4):1–12, 2017. 1
- [31] Yury Kartynnik, Artsiom Ablavatski, Ivan Grishchenko, and Matthias Grundmann. Real-time facial surface geometry from monocular video on mobile gpus. *arXiv*, 2019. 4, 13, 14
- [32] Adam Kendon. *Gesture: Visible action as utterance*. Cambridge University Press, 2004. 1
- [33] Stefan Kopp and Ipke Wachsmuth. Synthesizing multimodal utterances for conversational agents. *Computer animation and virtual worlds*, 15(1):39–52, 2004. 3
- [34] Taras Kucherenko, Dai Hasegawa, Gustav Eje Henter, Naoshi Kaneko, and Hedvig Kjellström. Analyzing input and output representations for speech-driven gesture generation. In *Proceedings of the 19th ACM International Conference on Intelligent Virtual Agents*, pages 97–104, 2019. 3
- [35] Taras Kucherenko, Patrik Jonell, Sanne Van Waveren, Gustav Eje Henter, Simon Alexandersson, Iolanda Leite, and Hedvig Kjellström. Gesticulator: A framework for semantically-aware speech-driven gesture generation. In *Proceedings of the 2020 International Conference on Multimodal Interaction*, pages 242–250, 2020. 3
- [36] Sergey Levine, Christian Theobalt, and Vladlen Koltun. Real-time prosody-driven synthesis of body language. pages 1–10. 2009. 1
- [37] Sergey Levine, Philipp Krähenbühl, Sebastian Thrun, and Vladlen Koltun. Gesture controllers. pages 1–11. 2010. 1, 3
- [38] Jiefeng Li, Siyuan Bian, Qi Liu, Jiasheng Tang, Fan Wang, and Cewu Lu. NIKI: Neural inverse kinematics with invertible neural networks for 3d human pose and shape estimation. In *Computer Vision and Pattern Recognition (CVPR)*, June 2023. 2
- [39] Jing Li, Di Kang, Wenjie Pei, Xuefei Zhe, Ying Zhang, Zhenyu He, and Linchao Bao. Audio2gestures: Generating diverse gestures from speech audio with conditional variational autoencoders. In *Computer Vision and Pattern Recognition (CVPR)*, pages 11293–11302, 2021. 3, 16
- [40] Miao Liao, Sibozhang, Peng Wang, Hao Zhu, Xinxin Zuo, and Ruigang Yang. Speech2video synthesis with 3d skeleton regularization and expressive body poses. In *Proceedings of the Asian Conference on Computer Vision*, 2020. 3

- [41] Xian Liu, Qianyi Wu, Hang Zhou, Yinghao Xu, Rui Qian, Xinyi Lin, Xiaowei Zhou, Wayne Wu, Bo Dai, and Bolei Zhou. Learning hierarchical cross-modal association for co-speech gesture generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10462–10472, 2022. 3, 15
- [42] Andrew L Maas, Awni Y Hannun, Andrew Y Ng, et al. Rectifier nonlinearities improve neural network acoustic models. In *International Conference on Machine Learning (ICML)*, volume 30, page 3, 2013. 14, 15
- [43] Stacy Marsella, Yuyu Xu, Margaux Lhommet, Andrew Feng, Stefan Scherer, and Ari Shapiro. Virtual character performance from speech. In *Proceedings of the 12th ACM SIGGRAPH/Eurographics symposium on computer animation*, pages 25–35, 2013. 1
- [44] Evonne Ng, Hanbyul Joo, Liwen Hu, Hao Li, Trevor Darrell, Angjoo Kanazawa, and Shiry Ginosar. Learning to listen: Modeling non-deterministic dyadic facial motion. In *Computer Vision and Pattern Recognition (CVPR)*, pages 20395–20405, 2022. 6, 7
- [45] Jorge Nocedal and Stephen J Wright. Nonlinear equations. *Numerical Optimization*, pages 270–302, 2006. 14
- [46] Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed AA Osman, Dimitrios Tzionas, and Michael J Black. Expressive body capture: 3d hands, face, and body from a single image. In *Computer Vision and Pattern Recognition (CVPR)*, pages 10975–10985, 2019. 2, 3, 4, 13, 14
- [47] Isabella Poggi, Catherine Pelachaud, F de Rosis, Valeria Carofiglio, and B De Carolis. Greta. a believable embodied conversational agent. In *Multimodal intelligent information presentation*, pages 3–25. Springer, 2005. 3
- [48] Sergey Prokudin, Michael J Black, and Javier Romero. Smp1pix: Neural avatars from 3d human models. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1810–1819, 2021. 16
- [49] Shenhan Qian, Zhi Tu, Yihao Zhi, Wen Liu, and Shenghua Gao. Speech drives templates: Co-speech gesture synthesis with learned templates. In *Computer Vision and Pattern Recognition (CVPR)*, pages 11077–11086, 2021. 3
- [50] Alexander Richard, Michael Zollhöfer, Yandong Wen, Fernando De la Torre, and Yaser Sheikh. Meshtalk: 3d face animation from speech using cross-modality disentanglement. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1173–1182, 2021. 1, 3
- [51] Najmeh Sadoughi, Yang Liu, and Carlos Busso. Meaningful head movements driven by emotional synthetic speech. *Speech Communication*, 95:87–99, 2017. 1
- [52] Md Sahidullah and Goutam Saha. Design, analysis and experimental evaluation of block based transformation in mfcc computation for speaker recognition. *Speech communication*, 54(4):543–565, 2012. 5
- [53] Li Siyao, Weijiang Yu, Tianpei Gu, Chunze Lin, Quan Wang, Chen Qian, Chen Change Loy, and Ziwei Liu. Bailando: 3d dance generation by actor-critic gpt with choreographic memory. In *Computer Vision and Pattern Recognition (CVPR)*, pages 11050–11059, 2022. 3
- [54] Yu Sun, Qian Bao, Wu Liu, Tao Mei, and Michael J Black. Trace: Temporal regression of 5d avatars with dynamic cameras in 3d environments. In *Computer Vision and Pattern Recognition (CVPR)*, June 2023. 2
- [55] Kenta Takeuchi, Souichirou Kubota, Keisuke Suzuki, Dai Hasegawa, and Hiroshi Sakuta. Creating a gesture-speech dataset for speech-based automatic gesture generation. In *International Conference on Human-Computer Interaction*, pages 198–202. Springer, 2017. 2, 3
- [56] Aaron Van den Oord, Nal Kalchbrenner, Lasse Espeholt, Oriol Vinyals, Alex Graves, et al. Conditional image generation with pixelcnn decoders. *Conference on Neural Information Processing Systems (NeurIPS)*, 29, 2016. 6, 15
- [57] Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. In *Conference on Neural Information Processing Systems (NeurIPS)*, volume 30, 2017. 2, 5
- [58] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 4
- [59] Petra Wagner, Zofia Malisz, and Stefan Kopp. Gesture and speech in interaction: An overview, 2014. 1
- [60] Cheng-hsin Wu, Ningyuan Zheng, Scott Ardisson, Rohan Bali, Danielle Belko, Eric Brockmeyer, Lucas Evans, Timothy Godisart, Hyowon Ha, Alexander Hypes, Taylor Koska, Steven Krenn, Stephen Lombardi,

- Xiaomin Luo, Kevyn McPhail, Laura Millerschoen, Michal Perdoch, Mark Pitts, Alexander Richard, Jason Saragih, Junko Saragih, Takaaki Shiratori, Tomas Simon, Matt Stewart, Autumn Trimble, Xinshuo Weng, David Whitewolf, Chenglei Wu, Shoou-I Yu, and Yaser Sheikh. Multiface: A dataset for neural face rendering. *arXiv*, 2022. doi: 10.48550/ARXIV.2207.11243. URL <https://arxiv.org/abs/2207.11243>. 2, 3
- [61] Hongyi Xu, Eduard Gabriel Bazavan, Andrei Zanfir, William T Freeman, Rahul Sukthankar, and Cristian Sminchisescu. Ghum & ghuml: Generative 3d human shape and articulated pose models. In *Computer Vision and Pattern Recognition (CVPR)*, pages 6184–6193, 2020. 2
- [62] Jing Xu, Wei Zhang, Yalong Bai, Qibin Sun, and Tao Mei. Freeform body motion generation from speech. *arXiv*, 2022. 1, 2, 3, 7
- [63] Payam Jome Yazdian, Mo Chen, and Angelica Lim. Gesture2vec: Clustering gestures using representation learning methods for co-speech gesture generation. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3100–3107. IEEE, 2022. 3
- [64] Hongwei Yi, Chen Li, Qiong Cao, Xiaoyong Shen, Sheng Li, Guoping Wang, and Yu-Wing Tai. Mmface: A multi-metric regression network for unconstrained face reconstruction. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 7663–7672, 2019. 2
- [65] Hongwei Yi, Chun-Hao P. Huang, Dimitrios Tzionas, Muhammed Kocabas, Mohamed Hassan, Siyu Tang, Justus Thies, and Michael J. Black. Human-aware object placement for visual environment reconstruction. In *Computer Vision and Pattern Recognition (CVPR)*, 2022. 2, 4
- [66] Youngwoo Yoon, Woo-Ri Ko, Minsu Jang, Jaeyeon Lee, Jaehong Kim, and Geehyuk Lee. Robots learn social skills: End-to-end learning of co-speech gesture generation for humanoid robots. In *International Conference on Robotics and Automation. (ICRA)*, pages 4303–4309. IEEE, 2019. 2, 3
- [67] Youngwoo Yoon, Bok Cha, Joo-Haeng Lee, Minsu Jang, Jaeyeon Lee, Jaehong Kim, and Geehyuk Lee. Speech gesture generation from the trimodal context of text, audio, and speaker identity. *Transactions on Graphics (TOG)*, 39(6):1–16, 2020. 2, 3, 15
- [68] Hongwen Zhang, Yating Tian, Xinchu Zhou, Wanli Ouyang, Yebin Liu, Limin Wang, and Zhenan Sun. Pymaf: 3d human pose and shape regression with pyramidal mesh alignment feedback loop. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11446–11456, 2021. 2, 13, 14
- [69] Hongwen Zhang, Yating Tian, Yuxiang Zhang, Mengcheng Li, Liang An, Zhenan Sun, and Yebin Liu. Pymaf-x: Towards well-aligned full-body model regression from monocular images. *arXiv*, 2022. 2, 4, 8
- [70] Yang Zhou, Zhan Xu, Chris Landreth, Evangelos Kalogerakis, Subhransu Maji, and Karan Singh. Visemenet: Audio-driven animator-centric speech animation. *ACM Transactions on Graphics (TOG)*, 37(4): 1–10, 2018. 3
- [71] Yang Zhou, Xintong Han, Eli Shechtman, Jose Echevarria, Evangelos Kalogerakis, and Dingzeyu Li. Makelttalk: speaker-aware talking-head animation. *ACM Transactions On Graphics (TOG)*, 39(6): 1–15, 2020. 6, 7
- [72] Wojciech Zielonka, Timo Bolkart, and Justus Thies. Towards metrical reconstruction of human faces. In *European Conference on Computer Vision (ECCV)*, 2022. 4, 13, 14

Appendices

A. Dataset

A.1. Dataset Description

Our dataset is built from the in-the-wild talking videos of four persons with various poses. The dataset contains high-quality 3D holistic body mesh annotations that are reconstructed from video clips of 26.9 hours in total. Each clip is less than 10 seconds. Fig. 7 illustrates the distributions of video durations from different characters.

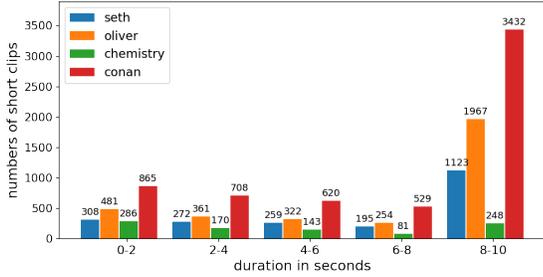


Figure 7. The distribution of the number of short clips for each character (0-10 seconds) of different speakers.

A.2. Good Practices for Improving p-GT

Preliminary. The 3D holistic body meshes consist of face, hands, and body, which is achieved by adopting SMPL-X [46]. It uses standard vertex-based linear blend skinning with learned corrective blend shapes and has $N = 10475$ vertices and $K = 67$ joints. Let W be the linear blend skinning function, the predicted mesh vertices can be represented as $v = W(\theta, \psi, \beta) \in \mathbb{R}^{N \times 3}$. Let $V = \{v_t | v_t \in \mathbb{R}^{N \times 3}\}_{t=1}^T$ and $J = \{j_t | j_t \in \mathbb{R}^{67}\}_{t=1}^T$ be the temporal sequence of mesh vertices and its 3D joint locations regressed from a linear regressor. We also denote $P^b = \{p_t^b | p_t^b \in \mathbb{R}^{32}\}_{t=1}^T$ and $P^h = \{p_t^h | p_t^h \in \mathbb{R}^{24}\}_{t=1}^T$ as the temporal sequence of the coefficients of the latent space of VPoser and low dimensional pose space after principal component analysis (PCA) for the body and hands respectively. For time interval $[1 : t]$, $V_{1:t} = (v_1, \dots, v_t)$, $J_{1:t} = (j_1, \dots, j_t)$, $P_{1:t}^b = (p_1^b, \dots, p_t^b)$ and $P_{1:t}^h = (p_1^h, \dots, p_t^h)$ represent segments of mesh vertices, 3D joints, body pose, and hand pose, respectively. Note that we use a fixed pose (sitting or standing) for the invisible lower body. And in a temporal sequence of the p-GT holistic motions m_i , at each time step t , the facial representation $m_t^f = [\theta_t^f, \psi_t] \in \mathbb{R}^{103}$ is a concatenation of jaw orientation and expression, and the body and hand motions are respectively represented by their poses $m_t^b = \theta_t^b \in \mathbb{R}^{63}$ and $m_t^h = \theta_t^h \in \mathbb{R}^{90}$.

Initialization. Since optimization-based methods are often slow and sensitive to the initialization. In contrast, regression-based methods tend to give a reasonable, but not well pixel-aligned results. Therefore, we use the results from PIXIE [21] and PyMAF-X [68] to initialize the parameters of body and hand pose, respectively. Results from DECA [22] are used to initialize the parameters of jaw pose and facial expression.

Data terms. We extend the data term by incorporating body silhouettes, facial landmarks, facial shapes, and facial details.

Firstly, to deal with the imperfect 2D landmarks by Openpose [10], we introduce the silhouette constraint to encourage the rendered SMPL-X body to be inside the human body mask. Ground-truth person segmentations are expensive to obtain for in-the-wild datasets. Hence, we employ an off-the-shelf segmentation model, Deeplab V3 [12] to generate p-GT person semantics mask $M_{sil} \in \mathbb{R}^{T \times h \times w}$, where H and W are the height and width of the input image. Pytorch3D is used as the differential renderer to process the rendered pixels of all mesh triangles, leading to the predicted semantics mask $\widehat{M}_{sil} \in \mathbb{R}^{T \times h \times w}$. The silhouette loss term is given by:

$$\mathcal{L}_{sil} = \sum ||d(\widehat{M}_{sil}) \odot d_{edt}(g(M_{sil}))||_2, \quad (4)$$

where $g(x) = \text{MaxPool}(x) - x$ is a function for detecting the edge of the binary mask. d_{edt} is a distance function to calculate the smallest Euclidean distance from the background point to the silhouette boundary.

Secondly, to get a better facial geometry in SMPL-X, we minimize the difference between the facial shape in SMPL-X and the reconstructed facial shape from MICA [72]. We term this as a facial shape objective \mathcal{L}_{FS} given by:

$$\mathcal{L}_{FS} = ||M_{g1}(V_{SMPL-X}) - M_{g2}(V_{MICA} + t_{FS})||_2, \quad (5)$$

where $V_{SMPL-X} \in \mathbb{R}^{N \times 3}$ is the SMPL-X vertices at neutral pose (i.e. $\theta = 0, \psi = 0$). $V_{MICA} \in \mathbb{R}^{5023 \times 3}$ is the MICA shape, and $t_{FS} \in \mathbb{R}^3$ is the offset of V_{MICA} from V_{SMPL-X} . M_{g1} and M_{g2} are functions that maps the original mesh vertices of V_{SMPL-X} and V_{MICA} to the corresponding 1787 vertices of frontal face part, respectively.

Thirdly, to get better facial expression, we use MediaPipe [31] to extract 105 of 468 dense 2D facial landmarks for each image. The loss term \mathcal{L}_{FE} is calculated as:

$$\mathcal{L}_{FE} = \sum_t ||U_{1:t} - \widehat{U}_{1:t}||_2, \quad (6)$$

where $U_{1:t}$ and \widehat{U}_i are temporal segments of landmarks from MediaPipe [31] and the 2D projection of the corresponding 3D joints $J_{1:t}$, respectively.

Lastly, to obtain high-frequency resolution facial details, we employ face expression tracking to monocular RGB

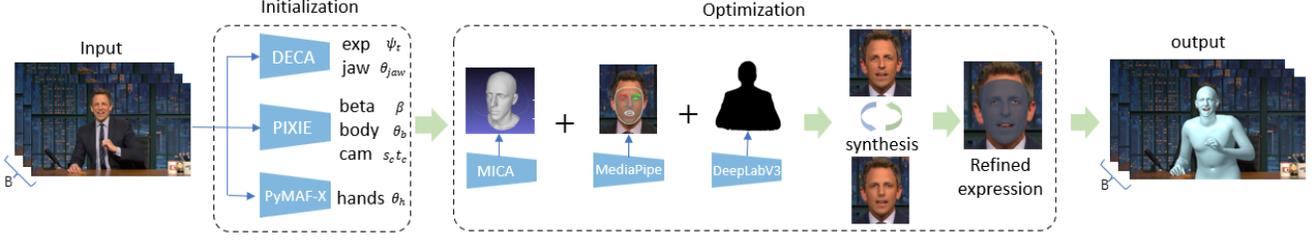


Figure 8. The architecture of SHOW. It consists of initialization and optimization modules. Specifically, given an input the image sequence, firstly, PIXIE [21], DECA [17] and PyMAF-X [68] are used to initialize the parameters of SMPL-X. Secondly, the optimization routine incorporates body silhouettes from DeepLab V3 [12], facial landmarks from MediaPipe [31], and facial shapes from MICA [72]. Then, it uses a photometric loss between the rendered faces and the input image to better capture facial details. Lastly, SHOW outputs the final results.

images in a self-supervised fashion. Specially, we follow [22, 72] to reconstruct the face jointly with an illumination model based on spherical harmonics and a Lambertian material assumption:

$$\mathcal{L}_{FR} = \sum_t \|I_r(\mathcal{M}_{S2F}(V_{1:t})) - I_{1:t}^{head}\|_2, \quad (7)$$

where \mathcal{M}_{S2F} is a function that selects the head part of $V_{1:t}$. I_r is the forward pass of differential rendering. $I_{1:t}^{head}$ is the cropped head image from input image. Note that we choose different scales (e.g. 256, 512, 1024) for different stages in the optimization procedure.

Regularization. Different regularization terms in SMPLify-X prevent the reconstruction of unrealistic bodies. To derive more reasonable regularization terms, we explicitly take the video prior into account.

To reduce the jittery results caused by the noisy 2D detected keypoints, we introduce a smooth term for body and motion poses (P^b and P^h). They are defined as:

$$\mathcal{M}_b = \sum_t \|P_{2:t}^b - P_{1:t-1}^b\|_2, \quad (8)$$

$$\mathcal{M}_h = \sum_t \|P_{2:t}^h - P_{1:t-1}^h\|_2. \quad (9)$$

We also add constant-velocity smooth term \mathcal{M}_j on J :

$$\mathcal{M}_j = \sum_t \|J_{3:t} + J_{1:t-3} - 2 \times J_{2:t-2}\|_2, \quad (10)$$

Furthermore, to prevent the inter-penetration of two hands, we use Collision Penalizer [46] and denote this loss term as L_{pen} .

Training Losses. The final objective function is given by:

$$E(\beta, \{\theta\}_{t=1}^T, \{\psi\}_{t=1}^T, \psi_{light}, \psi_{lbs}, t_{FS}) = \sum_{t=1}^T (E_{SMPLify-X}(t)) + \lambda_{FE} \mathcal{L}_{FE} + \lambda_{FS} \mathcal{L}_{FS} + \lambda_{FR} \mathcal{L}_{FR} + \lambda_{mb} \mathcal{M}_b + \lambda_{mh} \mathcal{M}_h + \lambda_{mj} \mathcal{M}_j + \lambda_{sil} \mathcal{L}_{sil} + \lambda_{pen} \mathcal{L}_{pen}, \quad (11)$$

where $\psi_{light} \in R^3$ is the spherical harmonic coefficients representing the environmental illumination. $\psi_{lbs} \in R^{128}$ is the linear blend skinning parameters of albedo model. $E_{SMPLify-X}(t)$ is the basic prior on single image as describe in [46]. Weights λ steer the influence of each term.

Optimization. Following [46], we adopt the Limited-memory BFGS [45] with strong wolfe line search for optimization. An iterative fitting routine is used for better fitting. With proper initialization, we minimize the objective function using a five-stage fitting procedure to avoid the local minima trap and reduce the optimization time. The learning rate is set to 1. As the required GPU memory increases dramatically with the image batch size for neural rendering, we use a mini-batch of 50 on NVIDIA Tesla V100.

B. Network Architecture Details

B.1. Face Generator

The raw audio input is normalized to zero mean and unit variance, and then is fed to encoder, which consists of an audio feature extractor, a transformer encoder, and a full-connected layer. The audio feature extractor is followed by an interpolation operation, in which the audio feature is re-sampled into target frames. For the decoder, it comprises six temporal convolution layers (with a kernel size, stride and padding of 3, 1 and 1 respectively) and a full-connected layer. Each temporal convolution layer is followed by layer normalization [6] and a Leaky RELU activation function [42]. We adopt SGD with momentum and a learning rate of

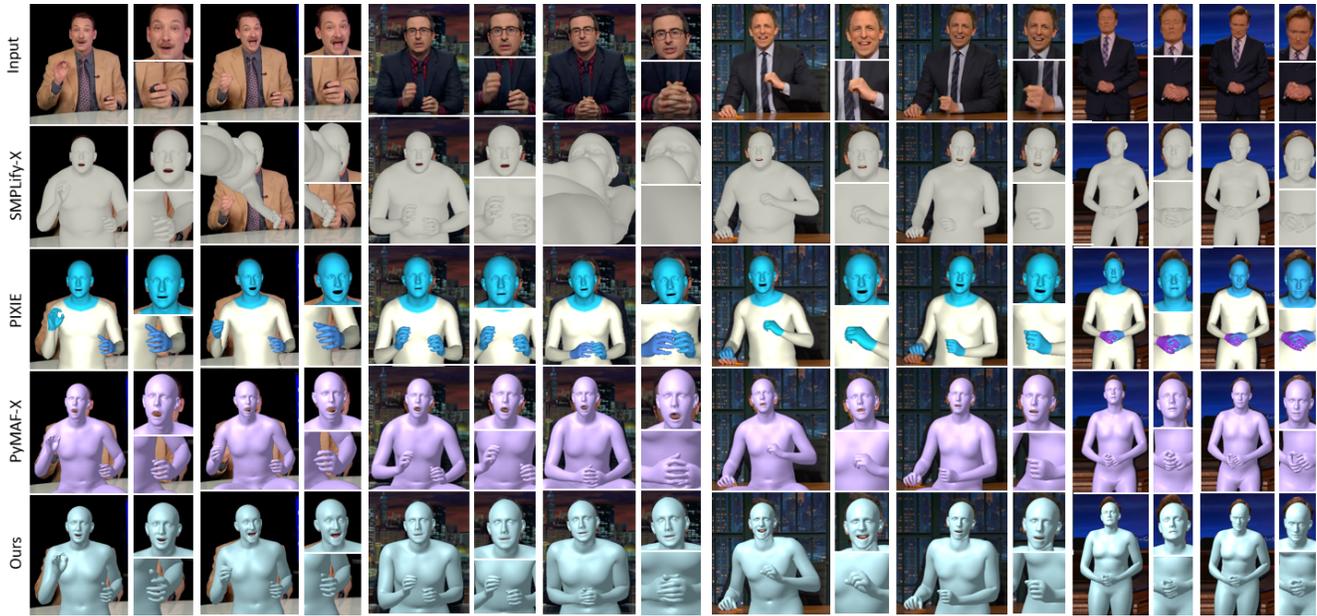


Figure 9. The 3D holistic body reconstructions of four subjects from SMPLify-X, PIXIE, PyMAF-X, and ours. Compared to other methods, ours produces more accurate and stable results with details.

0.001 as the optimizer. The face generator is trained with batchsize of 1 for 100 epochs, in which each batch contains a full-length audio and corresponding facial motions.

B.2. Body and Hand Generator

VQ-VAEs Details. The VQ-VAE takes body or hand motions as input. The encoder of each VQ-VAE is composed of three residual layers, which includes three temporal convolution layers (with a kernel size, stride and padding of 3, 1 and 1 respectively) followed by batch normalization [27] and a Leaky RELU activation function [42]. The encoder is interleaved with a temporal convolution layer with a kernel size, stride and padding of 4, 2 and 1 respectively after every residual layer except the last so that the temporal window size w is equal to 4. On the top of the encoder, a full-connected layer is added to reduce the dimension before quantization. The decoder is symmetric with the encoder. We adopt Adam with $\beta_1 = 0.9$, $\beta_2 = 0.999$ and a learning rate of 0.0001 as the optimizer. The commitment loss weight β is set to 0.25. The VQ-VAEs are trained with a batchsize of 128 and a sequence length of 88 frames for 100 epochs.

Autoregressive Model Details. The autoregressive model consists of an audio encoder and a Gated PixelCNN [56]. The audio encoder, which has the same structure as the VQ-VAE encoder, takes MFCC feature as input. Then we concatenate the output of the audio encoder and VQ-VAEs encoders and feed it to the Gated PixelCNN. The Gated PixelCNN has 15 gated convolution layers conditioned on identity, in which the convolution kernel is masked to make

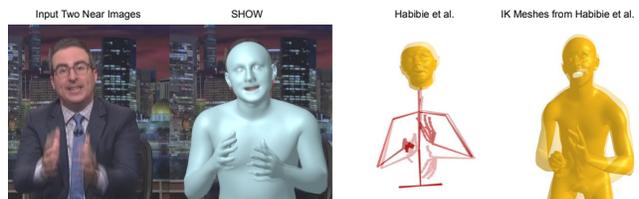


Figure 10. Holistic body reconstruction compared to Habibie et al.

sure the model cannot read future information. We adopt Adam with $\beta_1 = 0.9$, $\beta_2 = 0.999$ and a learning rate of 0.0001 as the optimizer. The autoregressive model is trained with a batchsize of 128 and a sequence length of 88 frames for 100 epochs.

C. More Comparison

Habibie et al. [26] v.s. SHOW. Habibie et al. [26] represent body, hands, and face separately. The lack of connection between body and face/hands results in unnatural poses of the face/hands w.r.t. the body. Fig. 10 a) shows that the hand and head poses of the body mesh, reconstructed from their estimated 3D skeleton, are less accurate than ours. Generated video results are further jittery. In contrast, SHOW generates more stable and accurate holistic body meshes.

Experimental Results. We compare our method with more other approaches and more metrics in Tab. 5. Specifically, We add Frechet Gesture Distance (FGD) [67] to measure the motion realism and beat consistency (BC) [41] to measure

Method	Habibie	Audio VAE	Audio+Motion VAE	Audio2Gesture[39]	Ours w/o c-c	Ours w/ c-c
FGD ↓	239.32	121.01	166.65	203.99	147.81	74.88
Variance ↑	0	0.044	0.176	0.240	0.922	0.821
BC (GT 0.868)	0.948	0.746	0.822	0.943	0.851	0.872

Table 5. More experimental results.

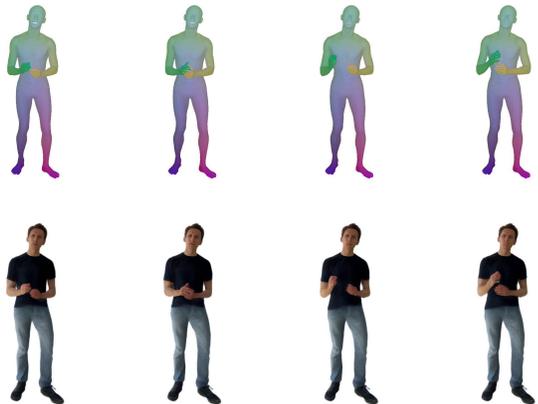


Figure 11. The application with SMPLpix to create photo-realistic neural avatars. Top row (input): the mesh vertices provided by TalkSHOW and their colors projected onto the image plane, bottom row: rendered output.

the alignment between the generated body motion and input audio, and compare with another audio-to-body motion baseline [39]. Our method outperforms the baselines in all these metrics and generates more diverse body motions, which are better aligned with the input audio.

D. Application

One application of our speech-to-motion generation is to create the photo-realistic neural avatars through neural renderers such as SMPLpix [48]. Given the mesh vertices provided by TalkSHOW and their colors, we first project them onto the image plane. Then, with the projected mesh vertices, SMPLpix allows us to efficiently synthesise photo-realistic images of humans. As TalkSHOW can produce continuous yet diverse motions, integrating SMPLpix with our motion generation framework enables us generate human avatars under different poses (see Fig. 11), leading to end-to-end photo-realistic video generation.

E. Discussions

Reconstruction. SHOW is based on SMPLify-X whose supervision signal is obtained from 2D keypoint reprojection. Thus, it is sensitive to severe hand shape deformation and heavy occlusion. A future direction would be to leverage advanced hand model with rich shape and pose space. Besides,

SHOW can only handle static camera cases currently. In the future, we plan to extend it to moving cameras.

Audio2motion. While we have demonstrated that TalkSHOW can generate realistic, coherent, and diverse holistic body motion with facial expression, body, and hand motions, it is subject to a limitation that can be addressed in the future. For the face generator, we mainly focus on facial motion (e.g. lip motion) and might not handle the very complex facial movements caused by emotions. In the future, we plan to extend to model this sort of part.

F. Risks and Potential Misuse

This work is intended for studying the translation from human speech to holistic body motion, helping building virtual agents to behave realistically and interact with listeners meaningfully. Since our techniques can generate a realistic and diverse 3D talking humans from audio, there is a risk that such technique could be potentially misused for fake video generation. For instance, a fake speech could be used to construct highly realistic 3D holistic body motion while it never happened. Thus, we should use such technology responsibly and carefully. We hope to raise the public’s awareness about a safe use of such technology.