

Diverse 3D Hand Gesture Prediction from Body Dynamics by Bilateral Hand Disentanglement

Xingqun Qi^{1,4}, Chen Liu², Muyi Sun³, Lincheng Li⁴, Changjie Fan⁴, Xin Yu^{1,2*}

¹ AAIL, University of Technology Sydney ² The University of Queensland

³ CRIPAC, NLPR, Institute of Automation, Chinese Academy of Sciences

⁴ Netease Fuxi AI Lab

{xingqunqi, yenanliu36}@gmail.com, muyi.sun@cripac.ia.ac.cn

{lilincheng, fanchangjie}@corp.netease.com, xin.yu@uq.edu.au




Figure 1. Diverse exemplary clips sampled by our method from our newly collected TED Hands Dataset. The vital frames are visualized to demonstrate that the hand gestures change with the upper body dynamics, synchronously. Best view on screen.

Abstract

Predicting natural and diverse 3D hand gestures from the upper body dynamics is a practical yet challenging task in virtual avatar creation. Previous works usually overlook the asymmetric motions between two hands and generate two hands in a holistic manner, leading to unnatural results. In this work, we introduce a novel bilateral hand disentanglement based two-stage 3D hand generation method to achieve natural and diverse 3D hand prediction from body dynamics. In the first stage, we intend to generate natural hand gestures by two hand-disentanglement branches. Considering the asymmetric gestures and motions of two hands, we introduce a Spatial-Residual Memory (SRM) module to model spatial interaction between the body and each hand by residual learning. To enhance the coordination of two hand motions wrt. body dynamics holistically, we then present a Temporal-Motion Memory (TMM) module. TMM can effectively model the temporal association between body dynamics and two hand motions. The second stage is built upon the insight that 3D hand predictions should be non-deterministic given the sequential body postures. Thus, we further diversify our 3D hand predictions based on the initial output from the stage one. Concretely, we propose a Prototypical-Memory Sampling Strategy (PSS) to generate the non-deterministic hand gestures by gradient-

based Markov Chain Monte Carlo (MCMC) sampling. Extensive experiments demonstrate that our method outperforms the state-of-the-art models on the B2H dataset and our newly collected TED Hands dataset. The dataset and code are available at: <https://github.com/XingqunQi-lab/Diverse-3D-Hand-Gesture-Prediction>.

1. Introduction

Given a sequence of upper body skeletons, our task aims at predicting natural and diverse 3D hand gestures. Such non-verbal body-hand coordination plays an important role in various virtual avatar scenarios, including human-agent interface [13, 17, 35, 40, 41], co-speech gesture synthesis [5, 20, 24, 46, 47], holoportation [29].

However, it is quite difficult to predict the natural and diverse 3D hand gestures due to three major challenges: **(1) Spatially asymmetric motions:** Asymmetric motions of two hands have been overlooked by previous works [28], e.g., when one hand moves, the other could be static or moving slowly. **(2) Temporal consistency wrt. body dynamics:** Predicted sequential hand gestures should be temporally consistent with respect to the body dynamics. **(3) Non-deterministic hand prediction:** Given the upper body dynamics, various 3D hand gestures can match the body postures rather than a deterministic result. Since the existing dataset has only a few avatar identities, the generated hands

*Corresponding author.

often lack diversity.

Due to complex body-hand interaction and asymmetric movements of two hands, directly predicting the natural and diverse 3D hand is rather difficult. Therefore, we propose to address the aforementioned challenges in a prediction followed by a diversification paradigm. To be specific, we propose a novel bilateral hand disentanglement based two-stage method, to generate natural and diverse 3D hand gestures from upper body dynamics.

In the first stage, we aim to predict the natural 3D hand gestures. Our key insight is to establish the bilateral hand disentanglement for asymmetric motions of two hands. We construct two hand-disentangled branches that interact with a body-specific branch to initially predict natural 3D hands. Here, we leverage a single-hand autoencoder to extract each hand feature respectively, thus achieving the effects of two hands disentanglement. Furthermore, we propose a Spatial-Residual Memory (SRM) module to model the spatial relation between the body and each hand via residual learning. Specifically, we employ a spatial memory bank to store the spatial residual deformation representations of each hand. Then, we leverage the current upper body embedding as a *query* to retrieve the most relevant spatial residual deformation. The queried spatial residual deformation and the synchronized hand embedding vector are used to generate the next step hand representation.

To ensure the motions of both hands are temporally consistent with the upper body sequence, we present a Temporal-Motion Memory (TMM) module. TMM models the correlation between body dynamics and hand gestures. Similar to SRM, we utilize a temporal memory bank to store the motion features of each hand at the sequence level. Moreover, we employ a motion encoder to acquire the input body dynamics representation, and then use the body dynamics representation as a *query* to produce temporally consistent hand motion features. In this fashion, we can effectively synchronize hand motions with body dynamics.

In the second stage, we develop a Prototypical-Memory Sampling Strategy (PSS) to diversify our 3D hand predictions based on the initial prediction at stage one. Here, we leverage an external memory that stores realistic hand prototypes, and then search a prototype that is closest to our initial predicted hands. Once we obtain the prototype, we project it into the feature space. To increase the diversity while preserving the authenticity of generated hand gestures, we perturb the prototype feature via gradient-based Markov Chain Monte Carlo (MCMC) sampling [23]. Specifically, a random noise is first sampled from a Gaussian distribution as a prior perturbation and then we update its posterior via Langevin dynamics based MCMC [4, 27]. Then, the perturbation is concatenated with the prototype feature to produce realistic new hand gestures.

Moreover, the existing 3D hand prediction dataset [28]

contains less than 10 avatar identities, resulting in insufficient diversity of gestures. Therefore, we newly collect a large-scale 3D hand gestures dataset (dubbed TED Hands) with more than 1.7K avatar identities from in-the-wild scenarios. Our dataset contains around 100 hours of TED talk speeches, enabling research on diverse 3D hand gesture predictions. Extensive experiments conducted on the B2H dataset [28] and our TED Hands dataset demonstrate our method outperforms the state-of-the-art.

Overall, our contributions are summarized as follows:

- We propose a novel bilateral hand disentanglement based two-stage 3D hand generation method to predict diverse 3D hand gestures from body dynamics in a prediction followed by diversification paradigm.
- We present a Spatial-Residual Memory (SRM) module to predict authentic hand poses from disentangled hand representations and a Temporal-Motion Memory (TMM) module to ensure temporal consistency of generated 3D hands.
- We design a Prototypical-Memory Sampling Strategy (PSS) to diversify our initial 3D hand prediction, thus obtaining natural and diverse hand gestures.
- We collect a new large-scale sequential 3D hand dataset from 1.7K persons' hand gestures, significantly facilitating research on diverse 3D hands generation.

2. Related Work

Since our method predicts the 3D hand motions from upper body skeletons, we briefly review the approaches of 3D hand pose estimation and 3D human motion prediction.

3D hand pose estimation: 3D hand estimation has received impressive attention for its wide applications in virtual reality, augmented reality, and robotics. Recently, parametric 3D hand model MANO [33] and various mocap-captured 3D hand datasets [2, 9, 26, 52] significantly improve the performance of 3D hand pose estimation from the single image [6, 18, 32, 37, 49, 51]. Among various image-based 3D hand estimation approaches, transformer-based models [10, 18, 21] have demonstrated outstanding performance. For instance, Li *et al.* [18] propose an interacting attention graph transformer to estimate two-hand poses from a single image. Shreyas *et al.* [10] present a hand keypoint based transformer combined with convolutional networks to establish the association among different joints.

Since our task aims to predict the 3D hand motions from upper body skeletons, these single-image-based methods [10, 18, 21] cannot be directly employed. However, their excellent performance in 3D hand estimation inspires us to model the body-hand interaction via transformer-based architectures. Besides, Ng *et al.* [28] generate the hand gestures from body dynamics in a holistic manner. As the

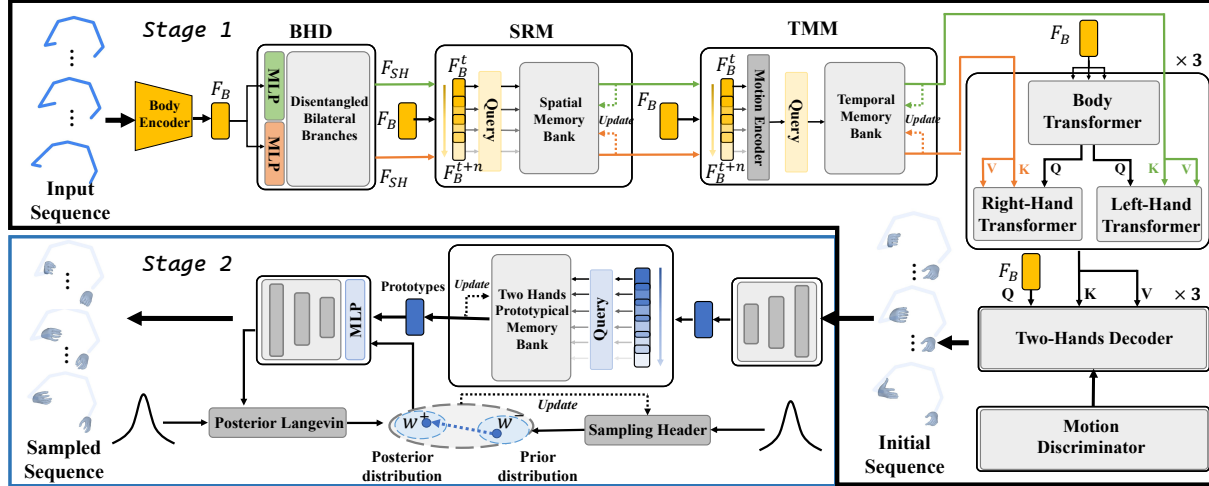


Figure 2. The pipeline of our proposed method. In stage one, we generate the natural initial 3D hands from sequential upper body skeletons. In stage two, we diversify our 3D hand predictions based on the initial gestures at stage one.

asymmetric motions of two hands, their method might produce some unnatural hand gestures. This also motivates us to disentangle two hands in hand pose estimation and then synchronize two-hands motions.

3D human motion prediction: Predicting human future motions from existing 3D poses has been widely studied [1, 8, 19, 25, 39]. Wang *et al.* [39] leverage a multi-range transformer that attends learn local-global features for multi-person motion prediction. Bouazizi *et al.* [1] introduce a multi-layer perception based human motion prediction network to learn spatio-temporal pose relation. Due to the similar nature between 3D human motion prediction and hand pose prediction, these models [1, 8, 19, 25, 39] can be easily adopted as our baselines. Moreover, we notice that memory based networks [3, 42, 43, 48] achieve promising performance in long-sequence modeling tasks. Motivated by this, we introduce external memories into our network to facilitate our long-range hand pose generation.

3. Proposed Method

3.1. Problem Formulation

Given a sequence of 3D upper body skeletons $\mathcal{B} = \{b^t\}_{t=0}^T$, our goal is to predict the natural and diverse 3D hand poses $\mathcal{H} = \{h^t\}_{t=0}^T$, where T is the sequence length. Specifically, we define each step body skeleton (*i.e.*, collars, shoulders, elbows, and wrists) $b^t \in \mathbb{R}^{24}$ as 8 joints. Each joint is represented by a 3D axis-angle representation. Meanwhile, the 3D hand pose at time t , $h^t \in \mathbb{R}^{90}$, is described as 30 joints with 3D axis-angle representation.

3.2. Stage One: Natural Hand Prediction

Due to complex body-hand interaction and asymmetric motions of two hands, we leverage the bilateral hand disen-

tanglement branches and a body-specific branch to generate the natural 3D hands. Specifically, as depicted in Fig. 2 (stage one), our method mainly consists of a Bilateral Hand Disentanglement (BHD) module, a Spatial-Residual Memory (SRM) module, a Temporal-Motion Memory (TMM) module, and a transformer-based backbone. Firstly, we leverage a body encoder to obtain the body features $F_B = \{F_B^t\}_{t=0}^T$, where $F_B^t \in \mathbb{R}^{1 \times C}$ denotes a body feature at time t , and C is the feature dimension. Then, the body features are disentangled by BHD to acquire the single hand features $F_{SH} = \{F_{SH}^t\}_{t=0}^T$, where $F_{SH}^t \in \mathbb{R}^{1 \times C}$. Considering the asymmetry motions of two hands, we exploit SRM to model the spatial interaction between the body and each hand, respectively. Besides, we further develop TMM to keep the two hand gestures temporal consistent with body dynamics.

Bilateral Hand Disentangle Module: Due to the asymmetric motions of two hands, we employ a single hand autoencoder as the feature extractor to disentangle two hands. We firstly train an autoencoder on 3D single hands. As shown in Fig. 3 (a), we then utilize a Multilayer Perceptron (MLP) to project the body features into the single hand features F_{SH} . To ensure these single hand features represent each disentangled hand separately, we conduct the feature level and 3D representation level constraints by the single hand autoencoder, as follows:

$$\mathcal{L}_{dis} = \left\| \widetilde{F}_{SH}^t - F_{SH}^t \right\|_1 + \left\| \mathbb{D} \left(\widetilde{F}_{SH}^t \right) - \mathbb{D} \left(F_{SH}^t \right) \right\|_1, \quad (1)$$

where \widetilde{F}_{SH}^t denotes reconstructed hand feature at time t , and \mathbb{D} is the pre-trained decoder.

Spatial-Residual Memory Module: After obtaining the disentangled single hand features, we introduce a Spatial-




Figure 3. Details of three sub-modules in the first stage. For brevity, we only show them in terms of single-branch view.

Residual Memory (SRM) module to model the spatial interaction between the body and each hand by residual learning. As illustrated in Fig. 3 (b), we leverage a spatial memory bank to store the spatial residual deformation representation between the current hand feature and the next step hand feature. Initially, each slot of the spatial memory bank is randomized as $m_j \in \mathbb{R}^{1 \times C}$, where j denotes slot index. To model the body-hand correlation, the current body feature is exploited as a *query* to retrieve the most relevant memory slot by calculating the cosine similarity.

In particular, we represent the relevance between the query and memory slots via cosine similarity. Then we leverage the argmax operation to acquire the most relevant memory slot m_r . Intuitively, the retrieved most relevant memory slot is leveraged as the spatial residual deformation between the current feature and the next step hand feature. However, it is intractable for argmax operation to back-propagate gradients. Thus, we conduct the softmax operation on the above calculated cosine similarity matrix to obtain the slot-wise affinity matrix. Then we acquire the learnable spatial residual deformation $F_{SH}^{\Delta t}$ by aggregation between the slot-wise affinity matrix and each memory slot. We apply the above reading process to allow gradient back-propagation. In the training phase, we utilize the current hand feature to *update* the retrieved memory slot as $m_r = \gamma m_r + (1 - \gamma) F_{SH}^t$, where $\gamma = 0.8$ in our experiments.

Different from obtaining the next hand feature by simple addition between the spatial residual deformation and current hand feature, we claim the spatial residual deformation refers to local spatial dependency from the current feature to the next time feature. This spatial dependency S is formulated as:

$$S_t^{t+1} = \text{softmax}(F_{SH}^{\Delta t}{}' \otimes F_{SH}^t), \quad (2)$$

where \otimes indicates matrix multiplication, and $'$ indicates the transpose operation. With the help of the spatial dependency, the next step hand feature is represented as $F_{SH}^{t+1} = F_{SH}^t + F_{SH}^t \otimes S_t^{t+1}$. Through SRM, we obtain the spatial-enhanced single-hand feature.

Temporal-Motion Memory Module: We further propose a Temporal-Motion Memory (TMM) module to ensure the single hand motions consistent with the body dynamics,

globally. As depicted in Fig. 3 (c), we develop a motion encoder to obtain the sequence-aware body motion embedding as $F_{B,M} \in \mathcal{R}^{T \times 1}$. Here, we use body motion embedding to represent temporal changes among body skeletons in a sequence. Then, we leverage the body motion embedding as a *query* to retrieve the most consistent single hand motion embedding $F_{SH,M} \in \mathcal{R}^{T \times 1}$ stored in a temporal memory bank. The initialization and reading processes of the temporal memory bank are similar to SRM. We utilize the body motion embedding to *update* the retrieved slot of the temporal memory bank in the same manner as SRM.

Afterward, we utilize the retrieved hand motion embedding to boost the body-hand temporal consistency. Specifically, we leverage the retrieved temporal embedding to represent the temporal correlation of hand gesture changes in a sequence. Then, the temporal-consistent enhanced sequential hand gestures are attained by: $F_{SH} = F_{SH} + F_{SH} \odot \text{softmax}(F_{SH,M})$, where \odot means dot product.

Transformer-based Backbone: As illustrated in Fig. 2 (stage one), to further enhance the body-hand interaction, we construct the bilateral hand disentanglement branches that interact with a body-specific branch to predict 3D hands. In particular, we introduce a body-specific transformer to generate body features as query Q . Then, through 3 times Multi-Head Attention (MHA) [38], we use Q to match the key features K and value features V in each hand transformer branch, respectively.

Furthermore, we feed the concatenation of bilateral hand features into an MLP, thus obtaining the merged two-hand features. Similarly, we conduct 3 times MHA in the transformer-based decoder. This encourages the merged two-hand features to be more temporal consistent with body dynamics, holistically. At the end of the decoder, we adopt two fully-connected layers to obtain the 3D two hands sequence. Additionally, similar to [28], we employ a motion discriminator to ensure the natural and continuous initial prediction at stage one.

3.3. Stage Two: Diverse Sampling

We leverage a Prototypical-Memory Sampling Strategy (PSS) to achieve hands diversification based on our initial predictions from stage one. Specifically, we first con-

struct a prototypical memory to store the realistic two hands prototypes. Then, we obtain the retrieved realistic prototype representation based on the initial prediction at stage one. Meanwhile, as shown in Fig. 2 (stage two), we propose a sampling header to obtain the hand-informative perturbation via gradient-based Markov Chain Monte Carlo (MCMC) sampling. Finally, we concatenate the retrieved realistic hand prototype representation and perturbation to generate the diverse 3D two hands.

Hands Prototypical Memory Construction: Inspired by [50], we reformulate the two hands diversification as a generation task. Instead of directly sampling the prior of the perturbation from the isotropic Gaussian distribution [7, 11, 16, 36], we assume the sampling space follows a learnable realistic prior distribution [30, 31, 50]. Since the existing 3D hand prediction dataset is noisy due to the automated annotation process, we propose a prototypical memory bank to store the realistic hand prototype representations encoded from real 3D hands. These 3D hands are captured from a studio-based mocap-captured dataset, named BEAT [22]. The reading and updating strategies of hands prototypical memory are the same as TMM.

Gradient-based MCMC Sampling: We exploit the retrieved hand prototype representation concatenated with sampled perturbation vector w to predict the diverse hand gestures. Given the observed 3D two hands $h \in \mathbb{R}^{90}$, the \tilde{h} is expressed as:

$$\tilde{h} = R_\theta(h, w) + \epsilon, w \sim p_\alpha(w), \epsilon \sim \mathcal{N}(0, \sigma_\epsilon^2 I), \quad (3)$$

where $p_\alpha(w)$ is the prior distribution of perturbation with parameters α , and $\epsilon \sim \mathcal{N}(0, \sigma_\epsilon^2 I)$ is the observation noise of hand gestures with σ_ϵ being given, R_θ is the generation model with parameters θ .

Inspired by [11], we leverage an MLP-based sampling header $S_\alpha(w)$ to model the diversification sampling process, as illustrated in Fig. 2 (stage two). The prior distribution of perturbation is initialized from an isotropic Gaussian reference distribution, expressed as:

$$p_\alpha(w) \propto \exp \left[-S_\alpha(w) - \frac{1}{2\sigma_w^2} \|w\|^2 \right], \quad (4)$$

where the $M_\alpha(w) = S_\alpha(w) + \frac{1}{2\sigma_w^2} \|w\|^2$ is defined as the whole sampling function, and α is the learnable parameters of sampling header. The hyperparameter σ_w denotes the standard deviation. For notation simplicity, let $\beta = \{\theta, \alpha\}$. For the i th sample in a training mini-batch with size n , the log-likelihood function of β is defined as:

$$L(\beta) = \sum_{i=1}^n \log \left[\int p_\alpha(w_i) p_\theta(\tilde{h}_i | h_i, w_i) dw_i \right]. \quad (5)$$

Thus the gradient of $L(\beta)$ is computed as:

$$\nabla L(\beta) = E_{p_\beta(w|\tilde{h}, h)} \left[\nabla_\alpha \log p_\alpha(w) + \nabla_\theta \log p_\theta(\tilde{h} | h, w) \right]. \quad (6)$$

The intractable expectation term $E_p(\cdot)$ in Eq. (6) is approximately solved by a gradient-based MCMC (Langevin dynamics) [4, 27]. In particular, the initial state of perturbation $p_0(w)$ is sampled from the Gaussian distribution. Through MCMC, the perturbation from the prior distribution and posterior distribution are represented as $\{w_i^-\}$ and $\{w_i^+\}$, respectively. The detailed calculation process is in [27]. Therefore, the gradients of the generation model and sampling header are updated by :

$$\nabla_\theta R_\theta = \frac{1}{n} \sum_{i=1}^n \left[\frac{1}{\sigma_\epsilon^2} (\tilde{h}_i - R_\theta(h_i, w_i^+)) \nabla_\theta R_\theta(h_i, w_i^+) \right], \quad (7)$$

$$\nabla_\alpha S_\alpha = \frac{1}{n} \sum_{i=1}^n [\nabla_\alpha S_\alpha(w_i^-)] - \frac{1}{n} \sum_{i=1}^n [\nabla_\alpha S_\alpha(w_i^+)]. \quad (8)$$

3.4. Objective Functions

In stage one, we apply the following loss functions to constrain the initial prediction process of hand gestures.

Reconstruction Loss: We use the ground truth to supervise the predicted two hands as $\mathcal{L}_{rec} = \|\mathcal{H} - \tilde{\mathcal{H}}\|_1$, where $\tilde{\mathcal{H}}$ denotes generated hands.

Perceptual Loss: Inspired by [15], we leverage the studio-based 3D hands dataset BEAT to pre-train an MLP-based autoencoder as a feature extractor $\phi(\cdot)$. We utilize the $L1$ loss to constrain the difference between the extracted features of \mathcal{H} and $\tilde{\mathcal{H}}$, as formulated: $\mathcal{L}_{perc} = \|\phi(\mathcal{H}) - \phi(\tilde{\mathcal{H}})\|_1$.

Adversarial Learning: Following the configuration of [28], we employ the adversarial training loss as:

$$\mathcal{L}_{adv} = \mathbb{E}_{\mathcal{H}} [\log D(\mathcal{H})] + \mathbb{E}_{\mathcal{B}} [\log(1 - (G(\mathcal{B})))] , \quad (9)$$

where D denotes the discriminator and G means hands generator. Finally, the overall objective in stage one is:

$$\min_G \max_D \mathcal{L}_{total}^{stage1} = \mathcal{L}_{rec} + \mathcal{L}_{adv} + \mathcal{L}_{perc} + 0.5\mathcal{L}_{dis}, \quad (10)$$

where \mathcal{L}_{dis} denotes disentanglement constraint in Eq. ((1)).

In the second stage, the overall objective to train our model is defined as follows:

$$\mathcal{L}_{total}^{stage2} = \|h - R_\theta(h, w^+)\|_1, \quad (11)$$

where w^+ is the posterior distribution of the perturbation. We conduct the frame-level diversification based on the initial prediction at stage one. To ensure temporal smoothness, we further apply temporal smoothness [45] to the diversified 3D hand gestures as a post processing.

Table 1. Comparison with the state-of-the-art methods on **B2H dataset** and **our proposed TED Hands dataset** of the first stage initial prediction. \downarrow indicates the lower the better.

Methods	B2H Dataset [28]			TED Hands Dataset		
	L2 \downarrow	FHD \downarrow	MPJRE \downarrow	L2 \downarrow	FHD \downarrow	MPJRE \downarrow
Body2Hands [28] _{CVPR'21}	3.760	1.103	13.148	2.551	1.174	11.371
MRT [39] _{NeurIPS'21}	3.634	0.973	12.576	2.325	0.877	10.314
BTM [8] _{WACV'23}	3.657	1.263	12.696	2.350	1.111	10.440
LTD [25] _{JCCV'19}	3.755	1.073	13.315	2.482	1.367	11.078
MotionMixer [1] _{JCAI'22}	3.616	0.986	12.702	2.324	0.910	10.427
SPGSN [19] _{ECCV'22}	3.656	0.962	12.757	2.435	0.990	10.887
Ours (stage one)	3.420	0.471	11.406	2.037	0.258	8.888

Table 2. Sampling experiment results on **B2H dataset**. \downarrow indicates the lower the better, and \uparrow indicates the higher the better. \pm means 95% confidence interval.

Methods	FHD \downarrow	Diversity \uparrow
Body2hands [28] w/ PSS	2.504	0.497 \pm 0.028
MRT [39] w/ PSS	1.982	0.698 \pm 0.031
BTM [8] w/ PSS	2.379	0.703 \pm 0.038
LTD [25] w/ PSS	1.893	0.577 \pm 0.028
MotionMixer [1] w/ PSS	2.169	0.869 \pm 0.055
SPGSN [19] w/ PSS	2.004	0.874 \pm 0.046
Ours w/ VAE [36]	1.617	1.001 \pm 0.061
Ours (stage two)	1.239	1.324\pm0.066

Table 3. Sampling experiment results on **our proposed TED Hands dataset**. \downarrow indicates the lower the better, and \uparrow indicates the higher the better. \pm means 95% confidence interval.

Methods	FHD \downarrow	Diversity \uparrow
Body2hands [28] w/ PSS	2.005	2.705 \pm 0.153
MRT [39] w/ PSS	1.188	4.954 \pm 0.190
BTM [8] w/ PSS	1.224	3.632 \pm 0.150
LTD [25] w/ PSS	1.385	2.851 \pm 0.101
MotionMixer [1] w/ PSS	1.613	3.487 \pm 0.208
SPGSN [19] w/ PSS	1.565	3.812 \pm 0.219
Ours w/ VAE [36]	1.033	4.018 \pm 0.190
Ours (stage two)	0.938	6.426\pm0.300

4. Experiments

4.1. Datasets and Experimental Setting

B2H Dataset: The B2H dataset [28] is automatically annotated by a monocular-based 3D pose estimation model [44], including the gestures of 8 speakers from in-the-wild talk show videos. In the original B2H dataset, there are more than 139K sequences. Each sequence includes 64 frames, with 32 frames overlapping between every two sequences. However, due to the noisiness and unavailable videos, we obtain 120,188 sequences finally. In our work, we adopt the same division criteria as [28] to split the training and testing dataset.

TED Hands Dataset: Due to the limited avatar identities in the existing 3D hand prediction dataset, we newly collect a large-scale 3D hands dataset (dubbed TED Hands) based on the TED Gesture [46, 47]. The original TED Gesture dataset only contains 10 upper body joints without elaborate fingers of two hands. Hence, we leverage a state-of-the-art 3D pose estimator FrankMocap [34] to extract the whole body-hand joints as pseudo ground truth from the video links provided by TED Gesture. Concretely, each joint of our dataset is normalized as the 3D axis-angle representation in a fixed kinematic structure [33]. In this fashion, our dataset is invariant to root joint motion and body shape. Finally, we acquired the 134,456 sequences from 1,755 different speakers, with 99.6 hours of TED talking speeches. Each sequence includes 64 frames without over-

lap. We randomly split the dataset into the 70% training set, 10% validation set, and 20% testing set. Please refer to the supplementary material for more details.

Implementation Details: Our model is implemented on the PyTorch platform with 2 NVIDIA RTX 2080Ti GPUs. We train the model utilizing Adam optimizer with an initial learning rate of 0.003. The whole training takes 30 epochs with a batch size of 64. We set $\sigma_\epsilon = 1$ in Eq. (3), and $\sigma_w = 1$ in Eq. (4).

Evaluation Metrics: To fully evaluate the naturalness of initial predictions at stage one and the diversity of sampled hand gestures at stage two, we employ various metrics:

- **L2 Distance:** Distance between the generated whole two hand poses and pseudo ground truth.
- **FHD:** Inspired by FGD [46, 47], we propose a hand-specific metric (named FHD) to measure the Fréchet distance between the features of generated hands and pseudo ground truth. The feature extractors of FHD are pre-trained by B2H and TED Hands datasets, respectively.
- **MPJRE:** We leverage the Mean Per Joint Rotation Error [o] (MPJRE) [14] to measure the absolute distance between predicted 3D representation joints and pseudo ground truth.
- **Diversity:** To verify the diversity of sampled hand gestures, we calculate the average feature distance between 500 random combined sequential pairs [20, 24]. The feature extractor is the same as FHD.

Table 4. Ablation study on different loss functions and different components of the proposed method. \downarrow indicates the lower the better. \checkmark indicates the employment of a certain module or loss function. # means the temporal-motion memory module is adopted before the spatial-residual memory module. \dagger denotes the weights of the two hands transformers are shared.

Model Variations						B2H Dataset [28]			TED Hands Dataset		
Baseline	Full Model	\mathcal{L}_{dis}	\mathcal{L}_{perc}	SRM	TMM	L2 \downarrow	FHD \downarrow	MPJRE \downarrow	L2 \downarrow	FHD \downarrow	MPJRE \downarrow
\checkmark						3.650	1.091	12.690	2.366	0.977	10.531
	\checkmark					3.604	0.773	12.453	2.303	0.731	10.250
	\checkmark	\checkmark				3.552	0.682	12.037	2.216	0.595	9.786
	\checkmark		\checkmark			3.563	0.657	12.085	2.261	0.699	10.026
	\checkmark	\checkmark	\checkmark			3.504	0.640	11.828	2.161	0.495	9.508
	\checkmark	\checkmark	\checkmark	\checkmark		3.457	0.580	11.687	2.100	0.393	9.209
#	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	3.484	0.614	11.758	2.124	0.463	9.311
\dagger	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	3.476	0.598	11.748	2.156	0.380	9.467
	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	3.420	0.471	11.406	2.037	0.258	8.888

4.2. Quantitative Results

Comparisons with the state-of-the-art: We compare our method with the pioneering 3D hands prediction work, named Body2hands [28]. Moreover, to fully verify the superiority of our method, we implement various state-of-the-art (SOTA) human motion prediction models: MTR [39], BTM [8], LTD [25], MotionMixer [1], and SPGSN [19]. For fair comparisons, all methods are implemented by source codes or pre-trained models released by authors. Since the above counterparts are designed without the diversification setting, the comparison is divided into two parts.

In the first part, we leverage our stage one initial prediction to compare with other competitors. We adopt the $L2$ distance, FHD, and MPJRE to provide a well-rounded view of comparisons. As reported in Tab. 1, our method outperforms all the counterparts by a large margin. Remarkably, our method achieves the lowest FHD and MPJRE scores on both datasets. Concretely, on FHD scores of the TED Hands dataset, our method performs even 70.6% lower than the optimal results of other competitors (e.g., $(0.877 - 0.258)/0.877 \approx 70.6\%$).

To demonstrate the effectiveness of the diversification process in our method, we compare the diversified results based on the various predictions from above mentioned competitive methods. Meanwhile, we replace our MCMC-based sampling strategy with CVAE [36] to diversify our initial prediction at stage one. We exploit the FHD and Diversity metrics in the second part. As shown in Tab. 2 and Tab. 3, our method achieves superior performance from the perspective of diversity. Although the FHD is slightly worse than stage one due to the sampled perturbation, our method still achieves the optimal result. Compared with CVAE, our method effectively avoids the potential posterior collapse problem [12]. Moreover, we observe that the diversity scores on our TED Hands dataset are much higher than the B2H dataset with all the counterparts. It applies that our newly collected TED Hands dataset includes more diverse hand gestures.

Ablation Study: We further conduct the ablation study based on stage one to verify the effectiveness of each module and different loss functions, as indicated in Tab. 4. The **baseline model** is implemented by a simple transformer-based pipeline without bilateral hand branches. Our **full model** adopts the bilateral hand branches interacting with a body-specific transformer to produce the 3D hand gestures.

As shown in Tab. 4, all combinations of the different loss functions and components have positive impacts on 3D hand predictions. The most complete combination attains the best results. Even though we directly construct the bilateral branches without disentanglement constraints (i.e., \mathcal{L}_{dis}) as **full model**, our method achieves comparable performance against the simple **baseline**. This indicates the global body-hand interaction are effectively modeled by bilateral hand transformers and a body-specific transformer. Then, we observe that after adding the **loss functions** \mathcal{L}_{dis} and \mathcal{L}_{perc} , the performance of our method is significantly improved, especially in MPJRE scores on both datasets (e.g., $10.250 \rightarrow 9.508$ on TED Hands).

To verify the effectiveness of **SRM** and **TMM**, we conduct plenty of experiments on both datasets. As reported in Tab. 4, when we first employ the SRM, our method achieves the lower $L2$ distance and FHD scores on both datasets. Then, after adding TMM, our method achieves the best performance. This result proves that TMM effectively models the body-hand temporal association. Besides, we find that when TMM is adopted before SRM, the model performance decreases slightly (**the third row from bottom to top** in Tab. 4). This corresponds to both TMM and transformer-based backbone (i.e. body-transformer, bilateral hand branches) focus on modeling the body-hand temporal correlation. Both of them encourage the two hands to be more temporal consistent with body dynamics. Employing the SRM between them leads to worse spatial-temporal modeling for body-hand interaction.

Moreover, we conduct experiments that the parameter weights of bilateral hand transformers are shared. As shown in Tab. 4 (**the second row from bottom to top**), the per-

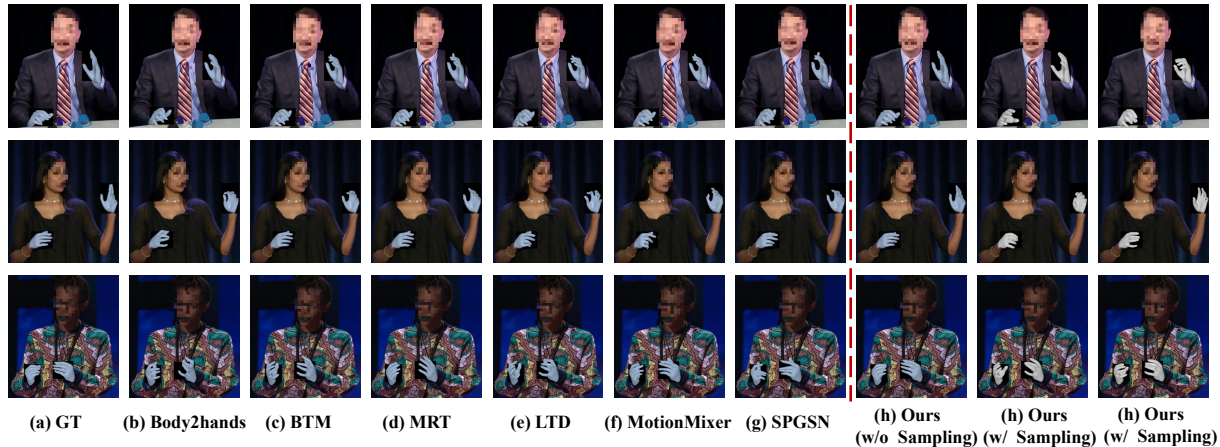


Figure 4. Visualization of our predicted 3D hand gestures against various state-of-the-art methods [1, 8, 19, 25, 28, 39]. From top to bottom, samples of the first row are from the B2H dataset, and the others are from our TED Hands dataset. Best view on screen.




Figure 5. User study on the motion naturalness, smoothness, and body-hand temporal consistency. ↑ indicates the higher the better.

formance is decreased due to the spatial-temporal interaction between the body and single-hand-specific branch being merged here. This supports our key technical insight on bilateral hand disentanglement. Additionally, to verify the practical ability of our method, we conduct experiments on model complexity and inference time, more details are presented in the supplementary material.

4.3. Qualitative Evaluation

User Study: To further analyze the visual quality of the predicted 3D hand gestures by different methods, we recruit 15 volunteers to participate in the user study. We take the naturalness, smoothness, and body-hand temporal consistency as three evaluation perceptions and the range of grades is from 0 to 5 (the higher, the better). The statistical results are reported in Fig. 5. Our method shows the best performance in all three metrics. Especially in the naturalness and body-hand temporal consistency, our method achieves remarkable advantages among all the methods due to the body-hand interaction modeling by SRM and TMM.

Visualization: To fully verify the performance of our method, we show the visualization of generated 3D hands from the first and second stages, respectively. As depicted in Fig. 4, we compare our initial predictions (from stage one) with various SOTA approaches. The avatars are ren-

dered from ground truth speakers in B2H and TED Hands datasets. We observe that Body2hands, BTM, and LTD generate unreasonable hand gestures (*e.g.*, in the third row of Body2hands, left-hand gestures of avatar). MRT, MotionMixer, and SPGSN predict the natural results, but the finger gestures are less consistent with the upper body dynamics. In contrast, our method successfully predicts the perceptually convincing 3D hand gestures from the body dynamics. This supports our key insight to build the bilateral hand disentanglement branches since the asymmetric motions of two hands. We further visualize the diversification results generated by our method at stage two. As shown in Fig. 1 and Fig. 4, with the same upper body dynamics as input, our method obtains the natural and diverse 3D hand gestures on B2H and TED Hands datasets, respectively. Please refer to the supplementary material for more visualization results.

5. Conclusion

In this paper, we present a novel bilateral hand disentanglement based two-stage method towards natural and diverse 3D hand prediction from body dynamics. In the first stage, upon the bilateral hand branches, we fully take advantage of the spatial-temporal body-hand interaction for generating natural initial 3D hand gestures. Then, in the second stage, we further diversify our natural initial 3D hand gestures by gradient-based markov chain monte carlo sampling. Extensive experiments conducted on B2H and our newly collected TED Hands datasets demonstrate the superiority of our method. As our method aims at generating diverse 3D avatar hand postures, in future work, we will investigate sampling diverse 3D hand gestures with temporal smoothness, instead of frame-wise sampling.

Acknowledgements. This research is funded in part by ARC-Discovery grant (DP220100800 to XY) and ARC-DECRA grant (DE230100477 to XY).

References

- [1] Arij Bouazizi, Adrian Holzbock, Ulrich Kressel, Klaus Dietmayer, and Vasileios Belagiannis. Motionmixer: Mlp-based 3d human body pose forecasting. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22*, pages 791–798. International Joint Conferences on Artificial Intelligence Organization, 7 2022. [3](#), [6](#), [7](#), [8](#)
- [2] Yu-Wei Chao, Wei Yang, Yu Xiang, Pavlo Molchanov, Ankur Handa, Jonathan Tremblay, Yashraj S Narang, Karl Van Wyk, Umar Iqbal, Stan Birchfield, et al. Dexycb: A benchmark for capturing hand grasping of objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9044–9053, 2021. [2](#)
- [3] Ho Kei Cheng and Alexander G Schwing. Xmem: Long-term video object segmentation with an atkinson-shiffrin memory model. In *European conference on computer vision*. Springer, 2022. [3](#)
- [4] Kumar Avinava Dubey, Sashank J Reddi, Sinead A Williamson, Barnabas Poczos, Alexander J Smola, and Eric P Xing. Variance reduction in stochastic gradient langevin dynamics. *Advances in neural information processing systems*, 29, 2016. [2](#), [5](#)
- [5] Hongyu Fu, Chen Liu, Xingqun Qi, Beibei Lin, Lincheng Li, Li Zhang, and Xin Yu. Sign spotting via multi-modal fusion and testing time transferring. In *Computer Vision – ECCV 2022 Workshops*, pages 271–287, Cham, 2023. Springer Nature Switzerland. [1](#)
- [6] Liuhaog Ge, Zhou Ren, Yuncheng Li, Zehao Xue, Yingying Wang, Jianfei Cai, and Junsong Yuan. 3d hand shape and pose estimation from a single rgb image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10833–10842, 2019. [2](#)
- [7] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020. [5](#)
- [8] Wen Guo, Yuming Du, Xi Shen, Vincent Lepetit, Xavier Alameda-Pineda, and Francesc Moreno-Noguer. Back to mlp: A simple baseline for human motion prediction. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, January 2023. [3](#), [6](#), [7](#), [8](#)
- [9] Shreyas Hampali, Mahdi Rad, Markus Oberweger, and Vincent Lepetit. Honnotate: A method for 3d annotation of hand and object poses. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3196–3206, 2020. [2](#)
- [10] Shreyas Hampali, Sayan Deb Sarkar, Mahdi Rad, and Vincent Lepetit. Keypoint transformer: Solving joint identification in challenging hands and object interactions for accurate 3d pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11090–11100, June 2022. [2](#)
- [11] Tian Han, Yang Lu, Song-Chun Zhu, and Ying Nian Wu. Alternating back-propagation for generator network. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 31, 2017. [5](#)
- [12] Junxian He, Daniel Spokoyny, Graham Neubig, and Taylor Berg-Kirkpatrick. Lagging inference networks and posterior collapse in variational autoencoders. In *International Conference on Learning Representations*, 2018. [7](#)
- [13] Ann Huang, Pascal Knierim, Francesco Chioffi, Lewis L Chuang, and Robin Welsch. Proxemics for human-agent interaction in augmented reality. In *CHI Conference on Human Factors in Computing Systems*, pages 1–13, 2022. [1](#)
- [14] Jiayi Jiang, Paul Strelci, Huajian Qiu, Andreas Fender, Larissa Laich, Patrick Snape, and Christian Holz. Avatarposer: Articulated full-body pose tracking from sparse motion sensing. In *European conference on computer vision*. Springer, 2022. [6](#)
- [15] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European conference on computer vision*, pages 694–711. Springer, 2016. [5](#)
- [16] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. [5](#)
- [17] Hema Swetha Koppula and Ashutosh Saxena. Anticipating human activities for reactive robotic response. In *IROS*, page 2071. Tokyo, 2013. [1](#)
- [18] Mengcheng Li, Liang An, Hongwen Zhang, Lianpeng Wu, Feng Chen, Tao Yu, and Yebin Liu. Interacting attention graph for single image two-hand reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2761–2770, 2022. [2](#)
- [19] Maosen Li, Siheng Chen, Zijing Zhang, Lingxi Xie, Qi Tian, and Ya Zhang. Skeleton-parted graph scattering networks for 3d human motion prediction. In *European conference on computer vision*. Springer, 2022. [3](#), [6](#), [7](#), [8](#)
- [20] Yuanzhi Liang, Qianyu Feng, Linchao Zhu, Li Hu, Pan Pan, and Yi Yang. Seeg: Semantic energized co-speech gesture generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10473–10482, 2022. [1](#), [6](#)
- [21] Kevin Lin, Lijuan Wang, and Zicheng Liu. End-to-end human pose and mesh reconstruction with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1954–1963, 2021. [2](#)
- [22] Haiyang Liu, Zihao Zhu, Naoya Iwamoto, Yichen Peng, Zhengqing Li, You Zhou, Elif Bozkurt, and Bo Zheng. Beat: A large-scale semantic and emotional multi-modal dataset for conversational gestures synthesis. In *European conference on computer vision*. Springer, 2022. [5](#)
- [23] Jun S Liu and Jun S Liu. *Monte Carlo strategies in scientific computing*, volume 10. Springer, 2001. [2](#)
- [24] 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. [1](#), [6](#)
- [25] Wei Mao, Miaomiao Liu, Mathieu Salzmann, and Hongdong Li. Learning trajectory dependencies for human motion pre-

- diction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9489–9497, 2019. 3, 6, 7, 8
- [26] Gyeongsik Moon, Shoou-I Yu, He Wen, Takaaki Shiratori, and Kyoung Mu Lee. Interhand2.6m: A dataset and baseline for 3d interacting hand pose estimation from a single rgb image. In *European Conference on Computer Vision*, pages 548–564. Springer, 2020. 2
- [27] Radford M Neal et al. Mcmc using hamiltonian dynamics. *Handbook of markov chain monte carlo*, 2(11):2, 2011. 2, 5
- [28] Evonne Ng, Shiry Ginosar, Trevor Darrell, and Hanbyul Joo. Body2hands: Learning to infer 3d hands from conversational gesture body dynamics. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11865–11874, June 2021. 1, 2, 4, 5, 6, 7, 8
- [29] Sergio Orts-Escolano, Christoph Rhemann, Sean Fanello, Wayne Chang, Adarsh Kowdle, Yury Degtyarev, David Kim, Philip L Davidson, Sameh Khamis, Mingsong Dou, et al. Holoportation: Virtual 3d teleportation in real-time. In *Proceedings of the 29th annual symposium on user interface software and technology*, pages 741–754, 2016. 1
- [30] Bo Pang, Tian Han, Erik Nijkamp, Song-Chun Zhu, and Ying Nian Wu. Learning latent space energy-based prior model. *Advances in Neural Information Processing Systems*, 33:21994–22008, 2020. 5
- [31] Bo Pang and Ying Nian Wu. Latent space energy-based model of symbol-vector coupling for text generation and classification. In *International Conference on Machine Learning*, pages 8359–8370. PMLR, 2021. 5
- [32] JoonKyu Park, Yeonguk Oh, Gyeongsik Moon, Hongseok Choi, and Kyoung Mu Lee. Handocnet: Occlusion-robust 3d hand mesh estimation network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1496–1505, June 2022. 2
- [33] Javier Romero, Dimitrios Tzionas, and Michael J Black. Embodied hands: modeling and capturing hands and bodies together. *ACM Transactions on Graphics (TOG)*, 36(6):1–17, 2017. 2, 6
- [34] Yu Rong, Takaaki Shiratori, and Hanbyul Joo. Frankmocap: A monocular 3d whole-body pose estimation system via regression and integration. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops*, pages 1749–1759, October 2021. 6
- [35] Maha Salem, Stefan Kopp, Ipke Wachsmuth, Katharina Rohlfing, and Frank Joubin. Generation and evaluation of communicative robot gesture. *International Journal of Social Robotics*, 4(2):201–217, 2012. 1
- [36] Kihyuk Sohn, Honglak Lee, and Xinchen Yan. Learning structured output representation using deep conditional generative models. *Advances in neural information processing systems*, 28, 2015. 5, 6, 7
- [37] Adrian Spurr, Aneesh Dahiya, Xi Wang, Xucong Zhang, and Otmar Hilliges. Self-supervised 3d hand pose estimation from monocular rgb via contrastive learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11230–11239, 2021. 2
- [38] 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
- [39] Jiashun Wang, Huazhe Xu, Medhini Narasimhan, and Xiaolong Wang. Multi-person 3d motion prediction with multi-range transformers. *Advances in Neural Information Processing Systems*, 34:6036–6049, 2021. 3, 6, 7, 8
- [40] Zijian Wang, Xingqun Qi, Kun Yuan, and Muye Sun. Self-supervised correlation mining network for person image generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7703–7712, 2022. 1
- [41] Pieter Wolfert, Nicole Robinson, and Tony Belpaeme. A review of evaluation practices of gesture generation in embodied conversational agents. *IEEE Transactions on Human-Machine Systems*, 2022. 1
- [42] Chao-Yuan Wu, Yanghao Li, Karttikeya Mangalam, Haoqi Fan, Bo Xiong, Jitendra Malik, and Christoph Feichtenhofer. Memvit: Memory-augmented multiscale vision transformer for efficient long-term video recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13587–13597, 2022. 3
- [43] Zhuojie Wu, Xingqun Qi, Zijian Wang, Wanting Zhou, Kun Yuan, Muye Sun, and Zhenan Sun. Showface: Coordinated face inpainting with memory-disentangled refinement networks. In *33rd British Machine Vision Conference 2022, BMVC 2022, London, UK, November 21-24, 2022*. BMVA Press, 2022. 3
- [44] Donglai Xiang, Hanbyul Joo, and Yaser Sheikh. Monocular total capture: Posing face, body, and hands in the wild. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10965–10974, 2019. 6
- [45] Xiangyu Xu, Li Siyao, Wenxiu Sun, Qian Yin, and Ming-Hsuan Yang. Quadratic video interpolation. In *NeurIPS*, 2019. 5
- [46] 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. *ACM Transactions on Graphics (TOG)*, 39(6):1–16, 2020. 1, 6
- [47] 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 *2019 International Conference on Robotics and Automation (ICRA)*, pages 4303–4309. IEEE, 2019. 1, 6
- [48] Jiyang Yu, Jingen Liu, Liefeng Bo, and Tao Mei. Memory-augmented non-local attention for video super-resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17834–17843, 2022. 3
- [49] Baowen Zhang, Yangang Wang, Xiaoming Deng, Yinda Zhang, Ping Tan, Cuixia Ma, and Hongan Wang. Interacting two-hand 3d pose and shape reconstruction from single color image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11354–11363, 2021. 2

- [50] Jing Zhang, Jianwen Xie, Nick Barnes, and Ping Li. Learning generative vision transformer with energy-based latent space for saliency prediction. *Advances in Neural Information Processing Systems*, 34:15448–15463, 2021. [5](#)
- [51] Yuxiao Zhou, Marc Habermann, Weipeng Xu, Ikhsanul Habibie, Christian Theobalt, and Feng Xu. Monocular real-time hand shape and motion capture using multi-modal data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5346–5355, 2020. [2](#)
- [52] Christian Zimmermann, Duygu Ceylan, Jimei Yang, Bryan Russell, Max Argus, and Thomas Brox. Freihand: A dataset for markerless capture of hand pose and shape from single rgb images. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 813–822, 2019. [2](#)