

RAM-Avatar: Real-time Photo-Realistic Avatar from Monocular Videos with Full-body Control

Xiang Deng¹, Zerong Zheng², Yuxiang Zhang¹, Jingxiang Sun¹, Chao Xu²,
 Xiaodong Yang³, Lizhen Wang¹, Yebin Liu^{1*}

¹ Department of Automation, Tsinghua University ²NNKosmos Technology ³Li Auto



Figure 1. We present RAM-Avatar, a real-time photo-realistic human avatar learning method based on monocular videos, which not only achieves high-fidelity rendering with full-body control including the face and hands but also supports real-time animation.

Abstract

*This paper focuses on advancing the applicability of human avatar learning methods by proposing RAM-Avatar, which learns a **Real-time, photo-realistic Avatar** that supports full-body control from **Monocular videos**. To achieve this goal, RAM-Avatar leverages two statistical templates responsible for modeling the facial expression and hand gesture variations, while a sparsely computed dual attention module is introduced upon another body template to facilitate high-fidelity texture rendering for the torsos and limbs. Building on this foundation, we deploy a lightweight yet powerful StyleUnet along with a temporal-aware discriminator to achieve real-time realistic rendering. To enable robust animation for out-of-distribution poses, we propose a Motion Distribution Align module to compensate for the discrepancies between the training and testing motion distribution. Results and extensive experiments conducted*

in various experimental settings demonstrate the superiority of our proposed method, and a real-time live system is proposed to further push research into applications. The training and testing code will be released for research purposes.

1. Introduction

The automatic learning of animatable, high-quality human avatars from videos holds significant value in diverse applications, including movie production, human-computer interaction, and immersive telepresence. Intuitively, an ideal human avatar modeling method should exhibit: 1) the ability to support real-time, high-resolution, photo-realistic animation with minimal data capture effort; 2) the capability to model not only the torsos and limbs but also achieve fine-grained control of the face and hands.

The ambition to achieve this goal has led to an in-

* Corresponding author

tensive exploration of neural radiance fields (NeRF) [45] for acquiring a 3D neural representation of human avatars [5, 21, 23, 38, 39, 49, 57, 66, 68, 69]. These techniques, however, typically rely on multi-view data and struggle to achieve photo-realistic real-time rendering, which limits their widespread application. By leveraging the power of generative adversarial networks (GANs) [17, 32, 33], researchers have proposed several effective image-to-image translation approaches to reconstruct human avatars from monocular videos and achieve photo-realistic rendering [31, 51, 72]. Despite their advancements, these methods tend to generate blurry results when driven by poses that are different from the training ones. Moreover, most methods only model the torsos and the limbs and leave the fine-grained body parts like the face and hands ignored. The lack of attention to these body parts significantly impairs their expressive ability and realistic appearance.

To overcome the aforementioned limitations, we introduce RAM-Avatar, a neural texture [62] based image-to-image translation framework which learns Real-time, photo-realistic Avatar supports full-body control from Monocular videos. RAM-Avatar comprises three essential components that collaborate to model the subtle yet important details of the face and hands, achieve high-fidelity texture rendering for the torsos and limbs, and improve animate robustness and generalization, respectively.

Firstly, to faithfully reproduce the facial expressions and hand gestures, we integrate two sophisticated statistical templates: FaceVerse [64] for facial dynamics and MANO [53] for hand articulations. Leveraging these, we deploy convolutional neural networks to distill characteristic features from the facial expression manifold of FaceVerse. It is an optimal strategy since the face has limited shape variation and contains many subtle details. Additionally, we find that a hand skeleton map indicating the hand key point locations is crucial for deriving more photo-realistic and distinct finger rendering results. Diverging from prior work [31, 51] that adopt a basic U-net architecture for rendering, our method incorporates a more powerful and computationally efficient variant, Style-Unet [65], which allows us to achieve superior results by leveraging the rich, high-level information present in our input priors. However, initial iterations revealed two primary shortcomings: a deficiency for temporal inconsistencies and unrealistic animations, stemming from ignoring temporal integration and misalignment of the body templates; and a limitation in the robustness and stability of animation results, bound by the scope of training data and thus struggling with poses beyond the training distribution.

To compensate for the first deficiency, our system embeds a dual attention module upon the body feature map rasterized from the body template. The dual attention module can be divided into spatial and temporal parts, where the

key insight is to enhance the differentiation between areas proximal to the body template and those further away while integrating motion trends to capture temporal variations. To facilitate real-time animation, we decouple the dense calculations inherent in the attention mechanism [63] into a series of sparse operations [26]. In this way, we augment the high-level texture information and improve the consistency of body feature maps without compromising computational efficiency.

In response to the second shortcoming, we innovate with a *Motion Distribution Align* (MDA) module, which is designed to transfer out-of-distribution poses to the target avatar pose distribution at test time. To implement this module, we employ a two-stage training process. First, a variational auto-encoder (VAE) [34] is trained on a large-scale pose dataset to establish a robust latent representation of human motions. Following this, a conditional decoder [56] is trained exclusively on the pose data of the target avatar. The MDA module, by conditioning on previous poses, consistently adjusts the out-of-distribution pose to align with the avatar pose distribution, where the robust encoder ensures the maintenance of semantic integrity, and the distribution-specific decoder ensures the semantic content will be properly translated into the target pose distribution. The integration of the MDA module thus empowers our system to reliably animate out-of-distribution poses while preserving the inherent human motion semantics.

Compared to existing human avatar approaches, our method excels in multiple key aspects, including animation speed, visual fidelity, fine-grained controllability and robustness. We believe our pipeline and results will serve as a catalyst for further research in similar domains. In summary, our contributions can be summarized as:

- A real-time animatable monocular framework that supports photo-realistic full-body animation, including body poses, hand gestures, and facial expressions.
- A dual attention module that facilitates stable realistic clothed body details by considering the geometric misalignment of body template and motion trend.
- A motion distribution alignment module, which bridges the domain gap between out-of-distribution poses and avatar pose domain to improve robustness.

2. Related Works

Human Avatars Reconstructing photo-realistic animatable human avatars is one of the most popular research topics in computer vision. Benefiting from the high-quality data captured by the dense array of cameras, researchers have successfully reconstructed fine-grained human avatars by fusing observations from dense views [11, 13, 20, 27, 40, 69, 70]. However, the expensive investment needed for construction and operation significantly restrict their application, and researchers have devoted great efforts in reduc-

ing the number of views needed for reconstructing realistic human avatars. For example, Peng *et al.* [49] introduce a neural blend weight field, which recovers animatable human models by combining NeRF and 3D human skeletons. Kwon *et al.* [36] also leverages a 3D body motion prior to learn generalizable neural radiance representations. Li *et al.* [39] use a joint-structured pose embeddings to encode high-frequency details and improve the representation ability of pose embeddings through feature lines, while Zheng *et al.* [78] learn high-quality full-body avatars that can be animated and rendered in real-time based on a compositional representation modeling the hands, face, and body with independent implicit fields.

With the rapid development of flexible and powerful representations, methods based on monocular videos have achieved great progress [12, 14, 19, 29, 30, 41, 68, 72, 73]. Pioneer works can only generate some reasonable but blurred results by decoding images from input pose skeleton [42, 43]. Some methods adopt a parametric body template to improve generalization and robustness [2, 3, 6]. The rasterization pipeline and texture map are further replaced by neural textures to store more high-level information [50, 51, 55]. Temporal constraints have also been proven to be essential for generating stable results [9, 31, 72, 80]. To overcome the inevitable tracking error of template fitting, researchers explore to build animatable full-body avatars directly from videos [38, 57–59]. Nevertheless, most of the existing monocular human avatar methods fail to achieve real-time photo-realistic animation and typically ignore the face and hands, which restricts their applications.

Neural Texture The core idea of Neural Texture [62] is replacing traditional RGB texture maps with learnable feature planes, where the rich signal contained in high-dimensional neural features serves as a more accurate encoding of appearance than RGB values. Since its proposal, neural textures have been adopted in many areas and proved its superiority [18, 42, 43, 50, 55]. As for modeling high-quality avatars, ANR [51] uses neural texture to account for geometric misalignment and pose-dependent surface deformation, while Next3D [60] adopts neural textures as the facial deformation representation to achieve photo-realistic portrait generation with fine-grained control. Thanks to the large-scale dataset of human images and scans, Dinar [61] builds realistic animatable avatars from a single reference image. In this paper, we propose a dual attention module upon the neural texture based feature map, which not only compensates for the misalignment of the body template but also improves the temporal consistency.

Domain Adaption Enhancing the robustness of the learning model against the out-of-distribution data is a challenging task [16, 37, 47, 54, 71]. Although significant progress

has been achieved in human avatar modeling, state-of-the-art methods may still generate artifacts or even collapse when driven by poses quite different from the training data [30, 77, 78]. To overcome this deficiency, a common practice is constructing motion graphs [4, 35] based on the training sequence. Some methods [7, 8, 25] synthesis human videos depending on reconstructed human meshes based on multi-view datasets, while [79] split and re-assemble video clips to generate a video with gestures matching a target speech audio. Zhang *et al.* [74] combine motion graph with local neural radiance fields to ensure the rendering quality. However, adopting motion graphs means that the generated video is strictly restricted by training pose sequences even though existing human avatar modeling methods can generalize to novel poses that fall in the training pose domain. In this work, we propose to transfer driving poses into avatar pose space by a motion distribution align module. This module can faithfully reconstruct the poses that falls into the avatar pose domain, and properly adjust out-of-distribution poses without losing the motion semantics.

3. Method

This section presents RAM-Avatar, a novel method that enables the learning of high-fidelity full-body avatars from monocular videos with real-time animation capabilities. The architecture of our system is depicted in Figure 2. We begin by presenting a methodology for capturing realistic facial and hand details. Building on this foundation, our system leverages a lightweight StyleUnet and a time-aware discriminator to facilitate photo-realistic real-time rendering. In addition, we introduce a dual attention module designed to guarantee temporal consistency and enhance the image realism of modeled clothed body elements. Lastly, we discuss our *Motion Distribution Align* module, which improves the robustness and generalization at test time, and outline the training objectives in our work.

3.1. Achieving Full-body Control

Neural Textures [51, 62] aim to improve the expressive capability of traditional texture maps by learning a set of high-dimensional feature maps on top of template meshes, which can be further interpreted by a neural renderer. Specifically, the neural textured feature map is generated by rasterizing the mesh to image space according to the standard graphics pipeline and texturing it with the high-dimensional neural texture. ANR [51] presents a neural texture based method to generate high-resolution images by learning a neural texture plane on top of coarse body templates and forming a deformation-aware body feature map for further rendering. However, it ignores the fine-grained body parts including facial expressions and hand gestures, which serve as significant forms of nonverbal communication and transmit critical social cues among humans. To tackle this challenge, we




Figure 2. The pipeline of our method. Our monocular high-fidelity avatar learning framework can be divided into three main parts: 1) the dual attention module aims to enhance the clothed body details; 2) the motion distribution align module aims to bridge the distribution gap between training and testing motions to improve robustness; 3) the face and hands modeling part aims to achieve full-body control. The green lines only activate at testing stage while the yellow lines only activate at training stage.

leverage FaceVerse [64] to obtain the relevant facial expression parameters. Considering the limited variation of facial shapes and the requirement of computational efficiency, we utilize a convolutional network to extract distinguishing features from the synthesized face mesh, thereby generating the t -th face feature map $F_t^{face} \in \mathbb{R}^{H \times W \times C_1}$. Acknowledging the intricate and versatile characteristics of finger movements, we additionally construct a hand pose feature map $F_t^{hand} \in \mathbb{R}^{H \times W \times C_2}$ by establishing connections between key points on the hands, which serves as an auxiliary input condition. Subsequently, we concatenate F_t^{hand} , F_t^{face} , along with the neural texture based body feature map $F_t^{body} \in \mathbb{R}^{H \times W \times C_3}$, to form a comprehensive full-body feature map $F_t^{full} \in \mathbb{R}^{H \times W \times C}$, where C equals to $C_1 + C_2 + C_3$.

Given the high-level full-body feature maps, we adopt a lightweight Style-Unet [65] architecture to render photo-realistic results. More precisely, a 2D generator G is employed to interpret feature maps F_t^{full} into RGB images $G(F_t^{full})$. Similar to [9, 10], we concatenate the t -th generated image $G(F_t^{full})$ with adjacent ground truth images and feed them into the discriminator. Such a method forces the discriminator to determine both the difference in realism and temporal coherence between the “real” sequence (I_{t-1}^{gt} , I_t^{gt} , I_{t+1}^{gt}) and “fake” sequence (I_{t-1}^{gt} , $G(F_t^{full})$, I_{t+1}^{gt}). In this way, our method enables the generation of high-quality, photo-realistic animation results while improving temporal consistency across neighboring frames.

3.2. Dual Attention Module

Although high-level texture information has been embedded in neural textures, naively synthesized body feature maps based on per-frame neural textures neglects temporal dependencies and suffers from the misalignment of the body template, leading to unsatisfactory results. To address this limitation, we propose a dual attention module. It is

inspired by [15], which adopts two types of attention mechanism on top of scene feature maps to capture contextual dependencies. Our dual attention module consists of spatial and temporal parts, which aim to compensate for the misalignments associated with the body template and enhance the temporal consistency of body feature maps.

Specifically, rather than simply generating a body feature map $F_t \in \mathbb{R}^{H \times W \times C_3}$ by rasterizing the body template mesh corresponding to the t -th frame, we pull outwards the vertices in the current mesh along their normal directions, and render it into an expanded body feature map $F_t^e \in \mathbb{R}^{H \times W \times C_3}$. The original body feature map F_t is then refined based on the distance and similarity between F_t and F_t^e . This refinement process is designed to facilitate network to distinguish the regions near the template and those far away. To ensure computational efficiency, we adopt a variant of criss-cross attention [26] as the central attention mechanism in our framework, where pixel-level features are updated based on features sharing the same row or column.

Concretely, two convolutional layers with 1×1 filters on F_t^e and F_t are deployed to get two feature maps, denoted as Q_t and K_t , each with dimensions $\mathbb{R}^{H_1 \times W_1 \times C_4}$. H_1 , W_1 , and C_4 are less than H , W , and C_3 for computing efficiency. For the sake of brevity, we omit the symbol t in the subsequent equations. For each position u in feature map Q , we obtain a vector $Q_u \in \mathbb{R}^{C_4}$ and its corresponding vector set $\Omega_u \in \mathbb{R}^{(H_1+W_1-1) \times C_4}$, where vectors in Ω_u are derived from K and are in the same row or column with position u . The correlation degree $a_{i,u}$ is computed as follows:

$$a_{i,u} = d(i, u) Q_u \Omega_{i,u} \quad (1)$$

where $\Omega_{i,u} \in \mathbb{R}^{C_4}$ represents the i -th element in Ω_u , $a_{i,u} \in \mathbf{A}$ and $\mathbf{A} \in \mathbb{R}^{(H_1+W_1-1) \times H_1 \times W_1}$, and $d(i, u)$ represents the distance weight function between the position of $\Omega_{i,u}$ and position u .

We then apply another convolutional layer with 1×1 filters on F to get $V \in \mathbb{R}^{H_1 \times W_1 \times C_3}$ for aggregation. Simi-




Figure 3. Visualization of a single iteration of criss-cross attention in spatial part of dual attention module.

larly, we obtain a vector $\mathbf{V}_u \in \mathbb{R}^{C_3}$ and corresponding vector set $\Phi_u \in \mathbb{R}^{(H_1+W_1-1) \times C_3}$, where vectors in Φ_u come from V and share the same row or column with position u . The output of the spatial part is calculated as follows:

$$\mathbf{F}_u^{spa_1} = \alpha \sum_{i=1}^{H_1+W_1-1} \mathbf{A}'_{i,u} \Phi_{i,u} \quad (2)$$

where α is learned during the training process, \mathbf{A}' represent the softmax version of \mathbf{A} , $\Phi_{i,u} \in \mathbb{R}^{C_3}$ represents the i -th element in Φ_u , and $\mathbf{F}_u^{spa_1} \in \mathbb{R}^{H_1 \times W_1 \times C_3}$ represents the result feature map. Note the above procedure is performed twice in order to capture more contextual information and we adopt a skip connection [22] to improve stability, thus generating the feature map $\mathbf{F}^{spa} \in \mathbb{R}^{H \times W \times C_3}$ of spatial part.

For temporal coherence, we generate an additional feature map $\mathbf{F}^{tem} \in \mathbb{R}^{H \times W \times C_3}$ in an analogy way with Q generated from the body feature map of the previous frame. Next, we pass \mathbf{F}^{spa} and \mathbf{F}^{tem} through two separate convolution layers, after which we perform an element-wise summation for feature fusion and obtain the aggregated body feature map $\mathbf{F}^{body} \in \mathbb{R}^{H \times W \times C_3}$. By doing so, we augment the pixel-level features with rich contextual information, thereby benefiting the subsequent rendering process.

3.3. Motion Distribution Alignment

Although the above pipeline can render promising photo-realistic animation results, it suffers from the domain-shift issue: it tends to generate blurry results when the input pose deviates significantly from the training pose [30, 74]. Therefore, the key challenge of this task is how to adjust the driven poses to match the training pose domain while maintaining semantic meaning, *i.e.*, manipulating the range of motion while sustaining the underlying trend.

To overcome this challenge, we propose a motion distribution align module to seamlessly transfer the driven pose to the avatar pose domain without semantic degradation. Specifically, to learn a robust latent pose representation to indicate inherent human motion meaning, we follow

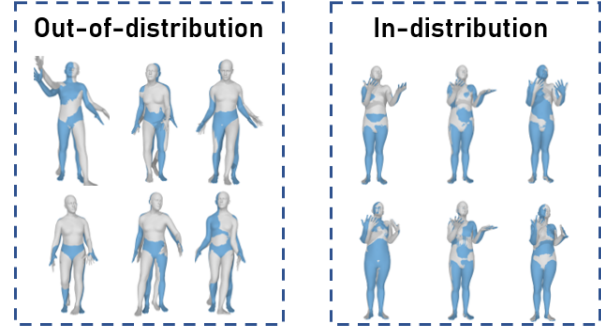


Figure 4. Visualization results of motion distribution align module. Given a training video mainly contains stand motions, *e.g.* speech video, our motion distribution align module faithfully constructs upper body motions while transferring the lower body motions to stand motions, thus avoiding blurry results. The white meshes represent the original motion, while the blue meshes represent the transferred motion.

[48, 52] to train a variational autoencoder [34] on several large-scale human motion capture datasets [1, 28, 44, 46]. To coherently translate the latent pose representation to specific distribution, we implement an extra conditional decoder trained on the avatar pose data, which acquires the latent representation of the current pose and the former pose as input. The robust pose encoder is frozen when training along with the conditional decoder to maintain robustness. Once deployed, we leverage the combination of the robust pose encoder and the distribution-specified pose decoder to bridge the gap between the distributions of the training and testing pose. In this way, our motion distribution align module inherits the best of both worlds: the robust pose encoder is responsible for extracting the intrinsic semantics of human motions, while the distribution-specified pose decoder transforms these semantics into continuous human motion sequences.

3.4. Training objectives

We adopt the common GAN loss, L1 loss, and perceptual loss with a VGG19 to train the rendering framework. We also provide a hand-loss item to regularize the hand regions based on perceptual loss.

$$\mathcal{L} = \mathcal{L}_{GAN} + \mathcal{L}_1 + \mathcal{L}_{percep} + \mathcal{L}_{hand} \quad (3)$$

$$\begin{aligned} \mathcal{L}_{GAN} = & \mathbb{E}_t[\log D(\mathbf{I}_{t-1}^{gt}, \mathbf{I}_t^{gt}, \mathbf{I}_{t+1}^{gt})] \\ & + \mathbb{E}_t[\log(1 - D(\mathbf{I}_{t-1}^{gt}, G(\mathbf{F}_t^{full}), \mathbf{I}_{t+1}^{gt}))] \end{aligned} \quad (4)$$

As for training the robust variational autoencoder, we adopt the Kullback-Leibler loss, reconstruction loss, and regularized loss as in [48]. After that, we train the distribution-specified decoder with reconstruction loss.




Figure 5. Exemplary results of our framework. We display the novel pose synthesis results of four avatars, where two pairs of avatars perform the same body pose, hand gesture, and facial expressions.




Figure 6. Comparisons for novel in-distribution pose animation against ANR [51], HFMT [31], and InstantAvatar [30]. Our method generates a more realistic appearance and outperforms other methods for novel in-distribution poses.

4. Experiments

In this section, we evaluate the performance of our methods for learning high-quality human avatars from monocular training videos that achieve high-resolution photo-realistic rendering in real time according to monocular driving videos. We begin by describing the experimental settings of our evaluation and then make comparisons with state-of-the-art methods. Lastly, we assess the influence of key components in our framework. Some exemplary results are shown in Figure 1 and Figure 5.

Table 1. Novel in-distribution pose animation quantitative comparison against state-of-the-art methods. The red numbers indicate the best results in different metrics, while the blue numbers indicate the sub-optimal results. Best viewed in color.

	SSIM \uparrow	LPIPS \downarrow	tLPIPS \downarrow	FID \downarrow
ANR [51]	0.871	0.098	0.405	54.15
HFMT [31]	0.895	0.053	0.469	14.48
InstantAvatar [30]	0.918	0.105	0.886	44.48
Ours	0.943	0.036	0.353	12.79

Table 2. Novel large rotation pose animation quantitative comparison against state-of-the-art methods. Notations in this table have the same meanings as those in Table 1.

	SSIM \uparrow	LPIPS \downarrow	tLPIPS \downarrow	FID \downarrow
ANR [51]	0.940	0.039	0.174	74.35
HFMT [31]	0.938	0.042	0.455	63.73
InstantAvatar [30]	0.930	0.039	0.314	26.97
Ours	0.961	0.020	0.155	32.09

4.1. Experimental Settings

Dataset. We evaluate the performance of our method on the “HFshirt” sequence from prior work [31], which contains 12,099 frames for training and 640 frames for testing. We also evaluate the model capability for handling large-scale rotational movements on the “female-4-sport” sequence from *People-Snapshot* dataset [3], where we leverage 300 frames for training and 100 frames for testing. To assess the robustness of the model against out-of-distribution poses, we further collect a sequence “Male1”, which contains 2,000 frames for training and 600 frames for testing. Note motions in the test set of “Male1” are beyond the scope of the training set, and this sequence is shot with natural backgrounds to accommodate real scenarios *e.g.* training with limited data.

Comparison baselines. To validate the superiority of our RAM-Avatar, we compare with 1) ANR [51], which also adopts neural texture to learn a 2D avatar, 2) HFMT [31], which adopts recurrent deep neural networks to predict plausible motion-dependent shape and appearance from 2D keypoints, 3) InstantAvatar [30], which proposes an efficient neural radiance field variant to reconstruct avatars. We reimplement ANR under the guidance of the authors and run HFMT and InstantAvatar based on the official codes.

Implementation Details. We adopt three statistic templates to capture the status of corresponding parts, *i.e.*, SMPL-X [48] for the body, FaceVerse [64] for the face, and MANO [53] for the hands. We adopt the method in [64, 76] to estimate the template parameters. We adopt a variant of Style-UNet [65] to render 1024×1024 resolution images, where the original discriminator is replaced by a temporal-aware discriminator to improve consistency. We encode the body




Figure 7. Comparisons for novel large rotation pose animation against ANR [51], HFMT [31], and InstantAvatar [30]. Our method preserves the body details in comparison with other methods for novel large rotation poses.




Figure 8. Comparisons for out-of-distribution pose animation against ANR [51], HFMT [31], and InstantAvatar [30]. Our method generates high-quality animation results by transferring out-of-distribution motions to the avatar motion domain while preserving inherent motion trends, resulting in improved robustness compared to competing methods for novel out-of-distribution poses.

template with a neural texture map $1024 \times 1024 \times 48$. For the face and hands, we adopt two separate shallow Unets to extract characteristic features, resulting in two $1024 \times 1024 \times 16$ feature maps respectively. The structure of cVAE is a variant of Vposer [48].

4.2. Results and Comparisons

Evaluation metrics. To evaluate the quality of the generated images, we employ a range of widely used metrics, including LPIPS [75], SSIM [67], and FID [24], which assess the perceptual distance of neural network features, pixel-space structural similarity, and the Fréchet inception distance between two datasets, respectively. Additionally, to evaluate the animation quality in videos, we compute and report tLPIPS [72] which quantifies temporal plausibility by comparing perceptual changes across consecutive frames.

By utilizing these metrics, we can provide a comprehensive assessment of animation quality in our evaluation.

Comparison settings. In our comparative analysis, we consider three distinct settings. Firstly, we evaluate methods in the context of novel in-distribution pose animation, where testing poses belong to the same distribution as the training poses. Secondly, we assess methods in the scenario of novel large rotation pose animation where models are driven by novel unseen rotation poses. Lastly, we investigate the performance of our method in the case of novel out-of-distribution pose animation. In this setting, the models are driven by novel unseen poses that do not conform to the distribution of training poses. We remove the facial expression and hand gesture modeling part in our framework and model body only for a fair comparison. The motion distribution align module also only activate in novel out-of-




Figure 9. Qualitative results of ablation study for face and hand modeling.

distribution pose animation setting.

Evaluation results. The in-distribution comparison results are shown in Table 1 and Figure 6. It can be seen that our animation results own photo-realistic appearance while other methods generate a few artifacts. Our method also obtains the best performance in all metrics, which demonstrates the superiority of our framework in image quality. From Figure 7 and Table 2, we observe that our method achieves better performance than 2D methods ANR and HFMT. Our method also outperforms the 3D method InstantAvatar in most metrics and achieves comparable performance in comparison with FID. Note although InstantAvatar performs good results when training with 300 frames only, it tends to produce average or even blurry results when training with a larger dataset (see Figure 6). As for driving by out-of-distribution poses in Figure 8, both ANR and InstantAvatar fail to render reasonable results, while our method successfully adjusts the range of body motion without changing the inherent meaning and avoids collapse. We attribute the improved robustness of our framework to the special design of the motion distribution align module, which effectively mitigates the gap between the training and testing pose distributions, resulting in improved synthesis quality and robustness.

4.3. Ablation Study

After demonstrating the superiority of our method, we conduct an ablation study to evaluate the influence of our key components. Specifically, we consider four framework variations, *i.e.*, our framework without spatial part in dual attention module (w/o Spa), our framework without temporal part in dual attention module (w/o Tem), our framework without hand modeling component (w/o Hand), and our framework without face modeling component (w/o Face). As shown in Table 3, we can conclude that each component contributes to the final quality, where removing the temporal part hampers the temporal consistency, and removing the spatial part leads to degenerate image quality. In Figure 9, ablation studies confirm that each face and hands modeling component positively impacts the overall output quality.

Table 3. Ablation study for novel in-distribution pose animation.

	SSIM \uparrow	LPIPS \downarrow	tLPIPS \downarrow	FID \downarrow
Ours (w/o Spa)	0.938	0.038	0.354	13.23
Ours (w/o Tem)	0.942	0.038	0.380	13.72
Ours	0.943	0.036	0.353	12.79

5. Discussion and Conclusion

Limitations Despite achieving better performance compared to other state-of-the-art human avatar methods and supporting real-time animation, the proposed RAM-Avatar method has several limitations. Firstly, it relies on an image-to-image translation system, which is constrained by the available training data. Although the motion distribution align module enhances robustness by transferring the animation pose to the avatar pose domain, the generated images may still appear blurry when encountering extreme poses, *e.g.* driven by rotation poses while training video only captures the front side of the subject. Secondly, the system generates rendering results based on parametric template tracking, which does not accurately describe body poses or facial expressions. As a result, the full-body control is imperfect, impacting the overall performance of our framework.

Social Impact Our method enables learning human avatars that can be animated by another person, resulting in a threat to be used to generate fake videos, which needs to be addressed before deploying.

Conclusion This paper introduces RAM-Avatar, a method for learning real-time, photo-realistic avatars that supports full-body control from monocular videos. Our approach is building upon two statistic templates to model the facial expressions and hand gestures while leveraging dual attention augmented Neural Texture feature maps to model realistic and consistent body details, followed by a lightweight StyleUnet to generate high-quality animation results in real-time. To enhance the robustness of our method, we propose a motion distribution align module, which transfers the animation pose to the avatar pose domain while preserving semantic meaning. Through comprehensive experiments, we demonstrate that our method achieves state-of-the-art performance in human avatar learning based on monocular videos. We believe our real-time framework will boost a wide range of downstream applications, including augmented reality and virtual reality.

Acknowledgements. This paper is supported by National Key R&D Program of China (2022YFF0902200), the NSFC project No.62125107, the Beijing Municipal Science & Technology Z231100005923030.

References

- [1] Ijaz Akhter and Michael J Black. Pose-conditioned joint angle limits for 3d human pose reconstruction. In *CVPR*, pages 1446–1455, 2015. 5
- [2] Thiemo Alldieck, Marcus Magnor, Weipeng Xu, Christian Theobalt, and Gerard Pons-Moll. Detailed human avatars from monocular video. In *3DV*, pages 98–109, 2018. 3
- [3] Thiemo Alldieck, Marcus Magnor, Weipeng Xu, Christian Theobalt, and Gerard Pons-Moll. Video based reconstruction of 3d people models. In *CVPR*, pages 8387–8397, 2018. 3, 6
- [4] Okan Arıkan and David A Forsyth. Interactive motion generation from examples. *ACM Transactions on Graphics*, pages 483–490, 2002. 3
- [5] Timur Bagautdinov, Chenglei Wu, Tomas Simon, Fabian Prada, Takaaki Shiratori, Shih-En Wei, Weipeng Xu, Yaser Sheikh, and Jason Saragih. Driving-signal aware full-body avatars. *ACM Transactions on Graphics*, pages 1–17, 2021. 2
- [6] Andrei Burov, Matthias Nießner, and Justus Thies. Dynamic surface function networks for clothed human bodies. In *ICCV*, pages 10754–10764, 2021. 3
- [7] Dan Casas, Marco Volino, John Collomosse, and Adrian Hilton. 4d video textures for interactive character appearance. In *Computer Graphics Forum*, pages 371–380, 2014. 3
- [8] Dan Casas, Christian Richardt, John Collomosse, Christian Theobalt, and Adrian Hilton. 4d model flow: Precomputed appearance alignment for real-time 4d video interpolation. In *Computer Graphics Forum*, pages 173–182, 2015. 3
- [9] Caroline Chan, Shiry Ginosar, Tinghui Zhou, and Alexei A Efros. Everybody dance now. In *ICCV*, pages 5933–5942, 2019. 3, 4
- [10] Mengyu Chu, You Xie, Jonas Mayer, Laura Leal-Taixé, and Nils Thuerey. Learning temporal coherence via self-supervision for gan-based video generation. *ACM Transactions on Graphics*, pages 75–85, 2020. 4
- [11] Alvaro Collet, Ming Chuang, Pat Sweeney, Don Gillett, Dennis Evseev, David Calabrese, Hugues Hoppe, Adam Kirk, and Steve Sullivan. High-quality streamable free-viewpoint video. *ACM Transactions on Graphics*, pages 1–13, 2015. 2
- [12] Zijian Dong, Chen Guo, Jie Song, Xu Chen, Andreas Geiger, and Otmar Hilliges. Pina: Learning a personalized implicit neural avatar from a single rgb-d video sequence. In *CVPR*, pages 20470–20480, 2022. 3
- [13] Mingsong Dou, Sameh Khamis, Yury Degtyarev, Philip Davidson, Sean Ryan Fanello, Adarsh Kowdle, Sergio Orts Escolano, Christoph Rhemann, David Kim, Jonathan Taylor, et al. Fusion4d: Real-time performance capture of challenging scenes. *ACM Transactions on Graphics*, pages 1–13, 2016. 2
- [14] Yao Feng, Jinlong Yang, Marc Pollefeys, Michael J. Black, and Timo Bolkart. Capturing and animation of body and clothing from monocular video. In *SIGGRAPH Asia*, pages 1–7, 2022. 3
- [15] Jun Fu, Jing Liu, Haijie Tian, Yong Li, Yongjun Bao, Zhiwei Fang, and Hanqing Lu. Dual attention network for scene segmentation. In *CVPR*, pages 3146–3154, 2019. 4
- [16] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *ICML*, pages 1180–1189, 2015. 3
- [17] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NIPS*, 2014. 2
- [18] Artur Grigorev, Karim Iskakov, Anastasia Ianina, Renat Bashirov, Ilya Zakharkin, Alexander Vakhitov, and Victor Lempitsky. Stylepeople: A generative model of fullbody human avatars. In *CVPR*, pages 5151–5160, 2021. 3
- [19] Chen Guo, Tianjian Jiang, Xu Chen, Jie Song, and Otmar Hilliges. Vid2avatar: 3d avatar reconstruction from videos in the wild via self-supervised scene decomposition. In *CVPR*, pages 12858–12868, 2023. 3
- [20] Kaiwen Guo, Peter Lincoln, Philip Davidson, Jay Busch, Xueming Yu, Matt Whalen, Geoff Harvey, Sergio Orts-Escolano, Rohit Pandey, Jason Dourgarian, et al. The relightables: Volumetric performance capture of humans with realistic relighting. *ACM Transactions on Graphics*, pages 1–19, 2019. 2
- [21] Marc Habermann, Lingjie Liu, Weipeng Xu, Michael Zollhoefer, Gerard Pons-Moll, and Christian Theobalt. Real-time deep dynamic characters. *ACM Transactions on Graphics*, pages 1–16, 2021. 2
- [22] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 5
- [23] Tong He, Yuanlu Xu, Shunsuke Saito, Stefano Soatto, and Tony Tung. Arch++: Animation-ready clothed human reconstruction revisited. In *ICCV*, pages 11046–11056, 2021. 2
- [24] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *NIPS*, 2017. 7
- [25] Peng Huang, Margara Tejera, John Collomosse, and Adrian Hilton. Hybrid skeletal-surface motion graphs for character animation from 4d performance capture. *ACM Transactions on Graphics*, pages 1–14, 2015. 3

- [26] Zilong Huang, Xinggang Wang, Lichao Huang, Chang Huang, Yunchao Wei, and Wenyu Liu. Cc-net: Criss-cross attention for semantic segmentation. In *ICCV*, pages 603–612, 2019. 2, 4
- [27] Mustafa Işık, Martin Rünz, Markos Georgopoulos, Taras Khakhulin, Jonathan Starck, Lourdes Agapito, and Matthias Nießner. Humanrf: High-fidelity neural radiance fields for humans in motion. *ACM Transactions on Graphics*, pages 1–12, 2023. 2
- [28] Catalin Ionescu, Dragos Papava, Vlad Olaru, and Cristian Sminchisescu. Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *PAMI*, pages 1325–1339, 2013. 5
- [29] Boyi Jiang, Yang Hong, Hujun Bao, and Juyong Zhang. Selfrecon: Self reconstruction your digital avatar from monocular video. In *CVPR*, pages 5605–5615, 2022. 3
- [30] Tianjian Jiang, Xu Chen, Jie Song, and Otmar Hilliges. Instantavatar: Learning avatars from monocular video in 60 seconds. In *CVPR*, pages 16922–16932, 2023. 3, 5, 6, 7
- [31] Moritz Kappel, Vladislav Golyanik, Mohamed Elgharib, Jann-Ole Henningson, Hans-Peter Seidel, Susana Castillo, Christian Theobalt, and Marcus Magnor. High-fidelity neural human motion transfer from monocular video. In *CVPR*, pages 1541–1550, 2021. 2, 3, 6, 7
- [32] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *CVPR*, pages 4396–4405, 2019. 2
- [33] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of StyleGAN. In *CVPR*, 2020. 2
- [34] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. 2014. 2, 5
- [35] Lucas Kovar, Michael Gleicher, and Frédéric Pighin. Motion graphs. In *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*, pages 723–732. 2023. 3
- [36] Youngjoong Kwon, Dahun Kim, Duygu Ceylan, and Henry Fuchs. Neural human performer: Learning generalizable radiance fields for human performance rendering. In *NIPS*, pages 24741–24752, 2021. 3
- [37] Mengxue Li, Yi-Ming Zhai, You-Wei Luo, Peng-Fei Ge, and Chuan-Xian Ren. Enhanced transport distance for unsupervised domain adaptation. In *CVPR*, pages 13936–13944, 2020. 3
- [38] Ruilong Li, Julian Tanke, Minh Vo, Michael Zollhöfer, Jürgen Gall, Angjoo Kanazawa, and Christoph Lassner. Tava: Template-free animatable volumetric actors. In *ECCV*, pages 419–436, 2022. 2, 3
- [39] Zhe Li, Zerong Zheng, Yuxiao Liu, Boyao Zhou, and Yebin Liu. Posevocab: Learning joint-structured pose embeddings for human avatar modeling. In *SIGGRAPH*, 2023. 2, 3
- [40] Zhe Li, Zerong Zheng, Lizhen Wang, and Yebin Liu. Animatable gaussians: Learning pose-dependent gaussian maps for high-fidelity human avatar modeling. In *CVPR*, 2024. 2
- [41] Hongyu Liu, Xintong Han, ChengBin Jin, Huawei Wei, Zhe Lin, Faqiang Wang, Haoye Dong, Yibing Song, Jia Xu, and Qifeng Chen. Human motionformer: Transferring human motions with vision transformers. In *ICLR*, 2023. 3
- [42] Lingjie Liu, Weipeng Xu, Michael Zollhoefer, Hyeonwoo Kim, Florian Bernard, Marc Habermann, Wenping Wang, and Christian Theobalt. Neural rendering and reenactment of human actor videos. *ACM Transactions on Graphics*, pages 1–14, 2019. 3
- [43] Lingjie Liu, Weipeng Xu, Marc Habermann, Michael Zollhöfer, Florian Bernard, Hyeonwoo Kim, Wenping Wang, and Christian Theobalt. Neural human video rendering by learning dynamic textures and rendering-to-video translation. *IEEE Transactions on Visualization and Computer Graphics*, pages 1–1, 2020. 3
- [44] Naureen Mahmood, Nima Ghorbani, Nikolaus F Troje, Gerard Pons-Moll, and Michael J Black. Amass: Archive of motion capture as surface shapes. In *ICCV*, pages 5442–5451, 2019. 5
- [45] B Mildenhall, PP Srinivasan, M Tancik, JT Barron, R Ramamoorthi, and R Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020. 2
- [46] Aron Monszpart, Paul Guerrero, Duygu Ceylan, Ersin Yumer, and Niloy J Mitra. imapper: interaction-guided scene mapping from monocular videos. *ACM Transactions on Graphics*, pages 1–15, 2019. 5
- [47] Pau Panareda Busto and Juergen Gall. Open set domain adaptation. In *ICCV*, pages 754–763, 2017. 3
- [48] 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 *CVPR*, pages 10975–10985, 2019. 5, 6, 7
- [49] Sida Peng, Junting Dong, Qianqian Wang, Shangzhan Zhang, Qing Shuai, Xiaowei Zhou, and Hujun Bao. Animatable neural radiance fields for modeling dynamic human bodies. In *ICCV*, pages 14314–14323, 2021. 2, 3
- [50] Sergey Prokudin, Michael J. Black, and Javier Romero. Smplpix: Neural avatars from 3d human models. In *WACV*, pages 1809–1818, 2021. 3

- [51] Amit Raj, Julian Tanke, James Hays, Minh Vo, Carsten Stoll, and Christoph Lassner. Anr: Articulated neural rendering for virtual avatars. In *CVPR*, pages 3722–3731, 2021. 2, 3, 6, 7
- [52] Davis Rempe, Tolga Birdal, Aaron Hertzmann, Jimei Yang, Srinath Sridhar, and Leonidas J Guibas. Humor: 3d human motion model for robust pose estimation. In *ICCV*, pages 11488–11499, 2021. 5
- [53] Javier Romero, Dimitrios Tzionas, and Michael J. Black. Embodied hands: Modeling and capturing hands and bodies together. *ACM Transactions on Graphics*, 2017. 2, 6
- [54] Kuniaki Saito, Shohei Yamamoto, Yoshitaka Ushiku, and Tatsuya Harada. Open set domain adaptation by backpropagation. In *ECCV*, pages 153–168, 2018. 3
- [55] Aliaksandra Shysheya, Egor Zakharov, Kara-Ali Aliev, Renat Bashirov, Egor Burkov, Karim Isakov, Aleksei Ivakhnenko, Yury Malkov, Igor Pasechnik, Dmitry Ulyanov, et al. Textured neural avatars. In *CVPR*, pages 2387–2397, 2019. 3
- [56] Kihyuk Sohn, Honglak Lee, and Xinchen Yan. Learning structured output representation using deep conditional generative models. In *NIPS*, 2015. 2
- [57] Shih-Yang Su, Frank Yu, Michael Zollhöfer, and Helge Rhodin. A-nerf: Articulated neural radiance fields for learning human shape, appearance, and pose. In *NIPS*, 2021. 2, 3
- [58] Shih-Yang Su, Timur Bagautdinov, and Helge Rhodin. Danbo: Disentangled articulated neural body representations via graph neural networks. In *ECCV*, pages 107–124, 2022.
- [59] Shih-Yang Su, Timur Bagautdinov, and Helge Rhodin. Npc: Neural point characters from video. In *ICCV*, 2023. 3
- [60] Jingxiang Sun, Xuan Wang, Lizhen Wang, Xiaoyu Li, Yong Zhang, Hongwen Zhang, and Yebin Liu. Next3d: Generative neural texture rasterization for 3d-aware head avatars. In *CVPR*, pages 20991–21002, 2023. 3
- [61] David Svitov, Dmitrii Gudkov, Renat Bashirov, and Victor Lempitsky. Dinar: Diffusion inpainting of neural textures for one-shot human avatars. In *ICCV*, pages 7062–7072, 2023. 3
- [62] Justus Thies, Michael Zollhöfer, and Matthias Nießner. Deferred neural rendering: Image synthesis using neural textures. *Acm Transactions on Graphics*, pages 1–12, 2019. 2, 3
- [63] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NIPS*, 2017. 2
- [64] Lizhen Wang, Zhiyuan Chen, Tao Yu, Chenguang Ma, Liang Li, and Yebin Liu. Faceverse: a fine-grained and detail-controllable 3d face morphable model from a hybrid dataset. In *CVPR*, pages 20333–20342, 2022. 2, 4, 6
- [65] Lizhen Wang, Xiaochen Zhao, Jingxiang Sun, Yuxiang Zhang, Hongwen Zhang, Tao Yu, and Yebin Liu. Styleavatar: Real-time photo-realistic portrait avatar from a single video. In *SIGGRAPH*, 2023. 2, 4, 6
- [66] Shaofei Wang, Katja Schwarz, Andreas Geiger, and Siyu Tang. Arah: Animatable volume rendering of articulated human sdfs. In *ECCV*, pages 1–19. Springer, 2022. 2
- [67] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, pages 600–612, 2004. 7
- [68] Chung-Yi Weng, Brian Curless, Pratul P Srinivasan, Jonathan T Barron, and Ira Kemelmacher-Shlizerman. Humannerf: Free-viewpoint rendering of moving people from monocular video. In *CVPR*, pages 16210–16220, 2022. 2, 3
- [69] Donglai Xiang, Fabian Prada, Timur Bagautdinov, Weipeng Xu, Yuan Dong, He Wen, Jessica Hodgins, and Chenglei Wu. Modeling clothing as a separate layer for an animatable human avatar. *ACM Transactions on Graphics*, pages 1–15, 2021. 2
- [70] Yuelang Xu, Benwang Chen, Zhe Li, Hongwen Zhang, Lizhen Wang, Zerong Zheng, and Yebin Liu. Gaussian head avatar: Ultra high-fidelity head avatar via dynamic gaussians. In *CVPR*, 2024. 2
- [71] Zhenhui Ye, Ziyue Jiang, Yi Ren, Jinglin Liu, Jinzheng He, and Zhou Zhao. Geneface: Generalized and high-fidelity audio-driven 3d talking face synthesis. In *ICLR*, 2023. 3
- [72] Jae Shin Yoon, Duygu Ceylan, Tuanfeng Y Wang, Jingwan Lu, Jimei Yang, Zhixin Shu, and Hyun Soo Park. Learning motion-dependent appearance for high-fidelity rendering of dynamic humans from a single camera. In *CVPR*, pages 3407–3417, 2022. 2, 3, 7
- [73] Zhengming Yu, Wei Cheng, xian Liu, Wayne Wu, and Kwan-Yee Lin. MonoHuman: Animatable human neural field from monocular video. In *CVPR*, 2023. 3
- [74] He Zhang, Fan Li, Jianhui Zhao, Chao Tan, Dongming Shen, Yebin Liu, and Tao Yu. Controllable free viewpoint video reconstruction based on neural radiance fields and motion graphs. *IEEE Transactions on Visualization and Computer Graphics*, 2022. 3, 5
- [75] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, pages 586–595, 2018. 7

- [76] Yuxiang Zhang, Hongwen Zhang, Liangxiao Hu, Jijun Zhang, Hongwei Yi, Shengping Zhang, and Yebin Liu. Proxycap: Real-time monocular full-body capture in world space via human-centric proxy-to-motion learning. In *CVPR*, 2024. 6
- [77] Zerong Zheng, Han Huang, Tao Yu, Hongwen Zhang, Yandong Guo, and Yebin Liu. Structured local radiance fields for human avatar modeling. In *CVPR*, 2022. 3
- [78] Zerong Zheng, Xiaochen Zhao, Hongwen Zhang, Boning Liu, and Yebin Liu. Avatarrex: Real-time expressive full-body avatars. *ACM Transactions on Graphics*, 2023. 3
- [79] Yang Zhou, Jimei Yang, Dingzeyu Li, Jun Saito, Deepali Aneja, and Evangelos Kalogerakis. Audio-driven neural gesture reenactment with video motion graphs. In *CVPR*, pages 3418–3428, 2022. 3
- [80] Long Zhuo, Guangcong Wang, Shikai Li, Wayne Wu, and Ziwei Liu. Fast-vid2vid: Spatial-temporal compression for video-to-video synthesis. In *ECCV*, pages 289–305, 2022. 3