

# CLIP-Actor: Text-Driven Recommendation and Stylization for Animating Human Meshes

Kim Youwang<sup>1\*</sup>  Kim Ji-Yeon<sup>2\*</sup>  Tae-Hyun Oh<sup>1,2†</sup> 

Department of <sup>1</sup>EE & <sup>2</sup>CiTE, POSTECH, Korea  
{youwang.kim, jiyeon.kim, taehyun}@postech.ac.kr  
<https://clip-actor.github.io>

**Abstract.** We propose CLIP-Actor, a text-driven motion recommendation and neural mesh stylization system for human mesh animation. CLIP-Actor animates a 3D human mesh to conform to a text prompt by recommending a motion sequence and optimizing mesh style attributes. We build a text-driven human motion recommendation system by leveraging a large-scale human motion dataset with language labels. Given a natural language prompt, CLIP-Actor suggests a text-conforming human motion in a coarse-to-fine manner. Then, our novel zero-shot neural style optimization detailizes and texturizes the recommended mesh sequence to conform to the prompt in a temporally-consistent and pose-agnostic manner. This is distinctive in that prior work fails to generate plausible results when the pose of an artist-designed mesh does not conform to the text from the beginning. We further propose the spatio-temporal view augmentation and mask-weighted embedding attention, which stabilize the optimization process by leveraging multi-frame human motion and rejecting poorly rendered views. We demonstrate that CLIP-Actor produces plausible and human-recognizable style 3D human mesh in motion with detailed geometry and texture solely from a natural language prompt.

**Keywords:** mesh animation, mesh stylization, text-driven editing

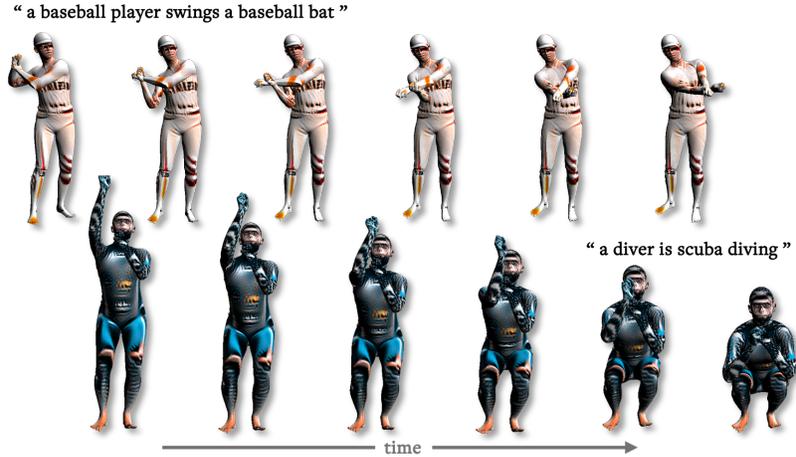
## 1 Introduction

Manual generation of animatable and detailed 3D avatars is cumbersome and requires time-consuming efforts with intensive labor and pain of creation. To reduce such burdens, many attempts have been introduced to automate such processes [5, 9, 10, 12, 16, 21, 29, 38, 50]. Furthermore, highly deformable human bodies make it more challenging to design temporally-consistent detailed geometries and textures. This process may be fully automated by text-guided 3D avatar generation, *i.e.*, making a machine understand the human text prompt to create a 3D avatar amenable to the prompt. Text-guided 3D avatar generation can be widely applied to machine-created media, such as virtual human animation [50], language-driven robot task planning [52, 59], and movie script visualization [22].

---

\* Authors contributed equally to this work.

† Joint affiliated with Yonsei University, Korea.



**Fig. 1. CLIP-Actor.** Given an input text prompt, CLIP-Actor recommends the best matching human meshes in motion and iteratively stylizes them by synthesis-through-optimization. CLIP-Actor can detail and texturize not just a single mesh frame but a short action clip by optimizing temporal-consistent and pose-agnostic style attributes.

Our key intuition is from the text-visual coupled understanding of humans. For example, when an actor reads a script for a play, the actor brings up an image of gestures, tone of speech, and clothes to her/his mind following the context described in the script. We believe such text-visual coupled imagination can be a breakthrough for accelerating machine-created media, *e.g.*, stylized 3D humans in motion. We can embody it to the machine by leveraging the text-image joint embedding space of CLIP [45]. With the representational power of the CLIP embedding space, the similarity measure between text and image provides concrete signals in building text-to-3D human meshes in motion.

In this work, we propose **CLIP-Actor**, an automated framework of text-driven recommendation and stylization of animating 3D human meshes. Given a text prompt describing human action and style, CLIP-Actor crafts a short clip of animated human meshes conforming to the prompt (see Fig. 1). Our method is free from extra artist-designed 3D mesh inputs since it searches meshes in motion from a database that strongly correlates with the given query text. The CLIP-Actor then detailizes and texturizes the mesh sequence by optimizing our proposed Decoupled Neural Style Fields (DNSF) in a pose-agnostic manner. The objective of the optimization is to maximize the correlation between the input text prompt and 2D rendered images of the stylized 3D mesh. We optimize DNSF with spatio-temporally augmented rendered images and provide an initial content mesh with a multi-modal sampling strategy. Moreover, we propose mask-weighted embedding attention for stable neural optimization. We demonstrate that CLIP-Actor can stylize visually and physically plausible 3D human meshes in motion with various text descriptions in zero-shot.

We summarize our main contributions as follows:

- We propose CLIP-Actor, a text-driven animated human mesh synthesis system.
- We propose a hierarchical text-driven human motion recommendation module that utilizes fine-grained textual semantic matching to capture visual and textual cues within the text prompt.
- Our novel decoupled neural style field (DNSF) learns style attributes of the human meshes in motion in a temporally consistent and pose-agnostic manner.
- We further develop novel methods to improve the convergence of the text-driven neural DNSF optimization: multi-modal content mesh sampling, spatio-temporal view augmentation, and mask-weighted embedding attention.

## 2 Related Work

Our work is closely related to text-driven 3D object contents and style manipulation. Multi-modal object stylization has been mainly studied using learned multi-modal embedding space, such as CLIP [45] and 3D content/style manipulation methods. We briefly review these lines of research.

**Text-driven visual data manipulation.** Recent advances in learned text and image joint embedding space [13, 45] have lit a fire in research about the style manipulation of images and 3D objects. CLIP embedding space is learned with abundant natural images and texts and was originally developed for zero-shot image and language analysis tasks. Interestingly, its representation turns out powerful enough to manipulate visual data with intuitive text guides. For images, text conditional image generation [15, 17, 27, 28, 36] has been notably advanced by CLIP. A representative work, StyleCLIP [39], manipulates an input image by optimizing over its latent code of a pre-trained generative model given a natural language text-prompt. CLIPDraw [15] synthesizes images with text guidance by optimizing the parameters of a set of curves via gradient descent.

Analogous to the image domain, several works [23, 34] extend the manipulation target domain to 3D objects by leveraging the advances in differentiable rendering [3, 26, 35, 37, 46]. The differentiable rendering technique enables seamless gradient flow from 2D rendered images to their 3D objects, allowing CLIP to bridge between language and 3D modalities through 2D images. Dream Fields [23] generates a 3D structure using implicit representation in free space, given a text prompt. It exploits no structural prior knowledge to learn or manipulate 3D contents. This allows flexible content exploration with novel styles but often results in abstract visual contents. As another concurrent work, Michel *et al.* [34] propose Text2Mesh, a CLIP-guided optimization method to manipulate the given fixed source mesh styles to conform to the target text condition prompt. In contrast to Dream Fields, since Text2Mesh stylizes a 3D object over the displacement and its texture map defined on a fixed, T-posed template human mesh, it imposes strong structural prior. It demonstrates plausible and interesting styles and textures of meshes given a text prompt. However, we observed that when the given template mesh is hard to conform to the given text prompt, it produces undesirable stylization; *e.g.*, the text containing detailed human action

produces a failure of stylization when the pose of the given human template mesh and the action are not conformed to each other.

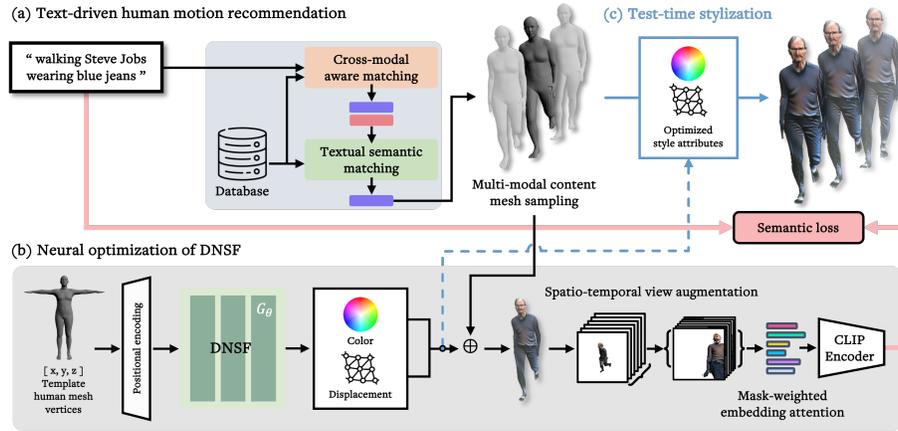
We focus on animating human meshes with details and styles according to the input text prompt. We exploit the parametric human mesh model to disentangle style from geometric contents, *i.e.*, pose. Such disentanglement enables the pose, detail, and style of human mesh to conform to the input text sequentially. This enables to stably manipulate 3D human objects, better conforming to the input text prompt from action to style.

**Text-driven human motion manipulation.** Many recent approaches have been studied to generate human body motion with given natural language descriptions. One line of work [1, 2, 18, 30, 43] guides the machine to translate natural language descriptions in a sequential manner and generate human skeletal motions using recurrent neural models. Another line of work [19, 20, 41] generates human motion conditioned on the limited number of closed-set action categories. CLIP-Actor focuses on textual and visual semantics in a whole sentence and can tackle various natural language descriptions.

Recently, MotionCLIP [55] and TEMOS [42] propose to learn the natural language conditioned mesh motion generation. MotionCLIP learns the human motion autoencoder and makes its latent space compatible with CLIP text and image space using semantic similarity. Similarly, TEMOS learns generative human mesh motion latent space with transformer-VAE [14, 25, 41] and aligns it with natural language latent space via DistilBERT [51], thus composing the cross-modal motion latent space. While both methods focus on the latent space to capture textual and visual semantics of natural language descriptions, CLIP-Actor directly maps the descriptions to realistic motion using a recommendation system. Moreover, our detailed volumetric meshes are stylized with appearance attributes much more expressive than those of the aforementioned methods.

**Texture and geometric stylization of human mesh in motion.** Aside from 3D mesh pose, recent work has added different levels of details to bare human meshes, *e.g.*, cloth modeling or texture color. The separate modeling of human and cloth meshes [5, 21, 29], the neural extension of the parametric human mesh model [4, 11, 32], the neural parametric approach [9, 10, 38], and the neural implicit approach [48, 49, 58] show promising clothed human mesh results from the given human scans, but without surface colors. Those works deal with texture and geometric styles separately. Recently, Saito *et al.* [50] propose a weakly-supervised way to recover both texture and geometric styles.

None of these methods can generate diverse color and cloth details of human motion in a zero-shot manner, *e.g.*, with only a text guide. We present a novel text-driven recommendation, detailization, and texturization of animating human meshes in zero-shot, where human meshes in motion with texture and geometric details are generated from the machine’s imagination without the task dataset.

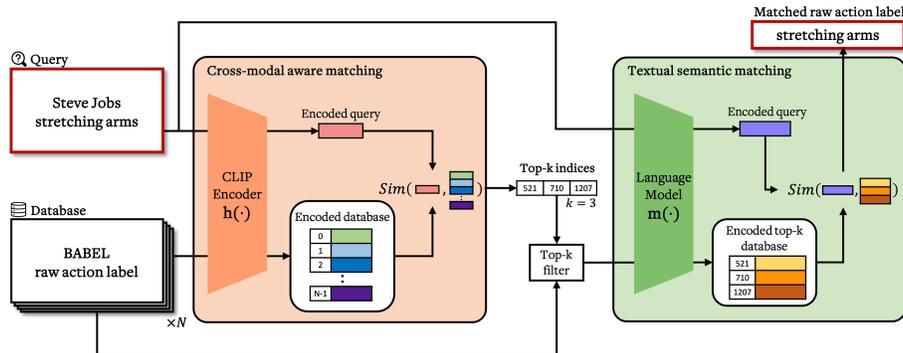


**Fig. 2. Overall architecture.** Given a text description of the human action, the text-driven human motion recommendation module finds the best semantically matched motion sequence from the motion database [44]. Content meshes are then sampled in a multi-modal context. Decoupled Neural Style Field takes a T-posed human mesh and learns text-driven style attributes, which are then applied to the content meshes. We apply spatio-temporal view augmentation and weight rendered images to guide the neural optimization with similarity among rendered images and the text.

### 3 CLIP-Actor: An Overview

Our goal is to visualize 3D motion that conforms to the input description by stylizing mesh with the color and displacement of its vertices. For example, consider a natural language prompt, “walking Steve Jobs wearing blue jeans.” Instead of preparing extra fixed 3D mesh inputs, our model obtains a sequence of 3D meshes that conforms to the input prompt, *i.e.*, walking, by retrieving a motion sequence from a dataset, *e.g.*, BABEL [44]. The retrieved mesh sequence becomes the “*content*” of our mesh stylization. We then grant the characteristics, *e.g.*, cloth, hair, to the meshes by optimizing the neural model to learn the color and displacement of the mesh vertices. Finally, our model generates a short clip of walking Steve jobs wearing blue jeans (see Fig. 2).

Formally, given a text prompt  $y$ , we retrieve a sequence of pose parameters  $\mathbf{R}_{1:T} = [\mathbf{R}_1, \dots, \mathbf{R}_T]$  of SMPL [31, 40, 47] for duration  $T$ . In a single frame  $t$ , mesh vertices  $\mathbf{M}_t$  can be acquired with a linear mapping as:  $\mathbf{M}_t = \mathcal{M}(\mathbf{R}_t, \beta_t)$ ,  $\forall t \in \{1, \dots, T\}$ , where  $\mathbf{R}_t$  denotes the pose parameters, and  $\beta_t$  the shape parameters for a human mesh. Then, a single mesh at frame  $t$  is represented by the faces  $F$ , and the 3D mesh vertices  $\mathbf{M}_t \in \mathbb{R}^{V \times 3}$ , where  $V$  is the number of vertices. Since SMPL mesh faces  $F$  for every frame are identical with given triangulation, we represent a single mesh using the mesh vertices,  $\mathbf{M}_t$ . Hence,  $\mathbf{M}_{1:T} = [\mathbf{M}_1, \dots, \mathbf{M}_T]$  denotes a full sequence of human meshes and is taken to our decoupled neural style field (DNSF) as “*content*.” The DNSF then learns “*style*,” *i.e.*, color and displacement, of mesh vertices and produces a sequence of textured mesh  $\mathbf{M}_{1:T}^*$ .



**Fig. 3. Hierarchical multi-modal motion retrieval.** Given text prompt as a query, our retrieval module finds the most relevant raw action label from the database [44]. First, the query and all of the raw action labels are encoded by CLIP text encoder  $h(\cdot)$ , and we measure similarity between them. Top- $k$  indices of raw action labels are selected, and corresponding raw action labels are retrieved by the top- $k$  filter. The language model encoder  $m(\cdot)$  vectorizes the query and top- $k$  action labels. The highest-scored raw action label is retrieved as the final matched result for the input text prompt.

## 4 Text-driven Human Motion Recommendation

In this section, we propose a motion recommendation module to obtain a motion sequence that conforms to the text prompt. We recommend the motion by retrieving visually and textually relevant action labels from the dataset.

**Hierarchical multi-modal motion retrieval.** We propose a hierarchical multi-modal motion retrieval module to obtain a motion sequence corresponding to the given prompt (see Fig. 3). We utilize a large scale human motion dataset with language labels, BABEL [44], containing frame-level aligned SMPL pose parameters and raw action labels. Given the text prompt as a query, the raw action label that is visually and linguistically associated with the query is matched through our motion recommendation system. We design a two-stage retrieval; cross-modal aware matching and textual semantic matching. The hierarchical matching enables CLIP-Actor to catch visual-language (cross-modal) aware contexts and linguistic semantics. These comprehensive matching modules hand over good initial content to the subsequent neural mesh stylization.

**Cross-modal aware matching.** Cross-modal aware matching finds the action labels similar to the input text prompt on the joint image-text space. We prepare the database, *i.e.*, a set of raw action labels  $\mathcal{A}$ , gathered from BABEL [44]. We retrieve a set  $\mathcal{A}_k \subset \mathcal{A}$  of the raw action label  $a_i \in \mathcal{A}$ , given a text prompt  $y$  as:

$$\mathcal{A}_k = \text{top-k}[S(\mathbf{h}(a_i), \mathbf{h}(y))], \quad \text{where } S(\mathbf{x}, \mathbf{y}) = \frac{\mathbf{x}^\top \mathbf{y}}{\|\mathbf{x}\|_2 \|\mathbf{y}\|_2}, \quad (1)$$

$\mathbf{h}(\cdot)$  is the pre-trained CLIP text encoder, and  $\text{top-k}[\cdot]$  denotes a function that returns  $k$  best matches. The similarity is measured by the cosine similarity.

Specifically, consider an input prompt “a man walking backwards.” The cross-modal aware matching vectorizes the input prompt and the action labels using the CLIP text encoder and computes the similarity between them. The set of matched action labels  $\mathcal{A}_k$  is determined as {“walking in place,” “walking backward,” “walking laterally”} and the top-1 matched label is “walking in place.” Since the CLIP text encoder is learned to focus on words that appear visually, it catches visual semantics (*i.e.*, walking), and all the elements in  $\mathcal{A}_k$  are closely related to the input prompt in visual space. However, the text encoder of CLIP is trained with still images instead of videos so that it cannot distinguish fine-grained action (*i.e.*, “walking in place” vs. “walking backwards”; because both appear to be the same in a still image). Thus, we propose textual semantic matching as the following step to compensate for the single-stage retrieval.

**Textual semantic matching.** The textual semantic matching finds the most relevant action label with the input prompt by capturing textual semantics in the sentence. We utilize the language expert, MPNet [53], so that our two-stage module can distinguish textual semantics and grammatical structures. The best matching label  $a_*$  is retrieved as:

$$a_* = \arg \max_{a_j \in \mathcal{A}_k} \mathcal{S}(\mathbf{m}(a_j), \mathbf{m}(y)), \quad (2)$$

where  $\mathbf{m}(\cdot)$  denotes the pre-trained MPNet encoder. Again, consider our above example of “a man walking backwards.” The top- $k$  action labels are re-ranked, and the most similar action label “walking backward” is retrieved as a final result. The sequence of meshes  $\mathbf{M}_{1:T}$  associated with the retrieved action label is passed to the following neural mesh stylization pipeline as the content mesh sequence.

## 5 Decoupled Stylization of Human Meshes in Motion

We represent a stylized human mesh with the content mesh and the style attributes. In practice, we denote the content mesh as  $\mathbf{M}_i \in \mathbb{R}^{V \times 3}$  sampled from the retrieved human motion sequence of  $T$  frames,  $\mathbf{M}_{1:T}$ . The mesh’s surface style attributes  $(\mathbf{c}, \mathbf{d}) \in \mathbb{R}^{V \times 3} \times \mathbb{R}^V$  are interpreted as the per-vertex RGB color and per-vertex displacement, which are applied over surfaces via given triangulation. The Neural Style Field [34] takes a fixed static mesh as input and learns the style attributes with a multi-layer perceptron (MLP). Michel *et al.* [34] claim that this implicit formulation tightly couples the style field to the source mesh. However, since Neural Style Field takes a single posed mesh at a time, a significant number of MLPs are required to stylize a sequence of human meshes.

**Decoupled Neural Style Field.** Instead, we introduce *Decoupled Neural Style Field* (DNSF). We propose to rather decouple the style field from the content mesh so that we need only one neural network to learn style attributes for the meshes in motion. Specifically, we first map the style attributes from the template human mesh  $\mathbf{M}_c$  and merge it with the content meshes  $\mathbf{M}_i, \forall i \in \{1, \dots, T\}$ , right before rendering (refer to Fig. 2). DNSF can achieve the same mesh stylization as the basic Neural Style Field while effectively decoupling the style from the content mesh. In practice, parameterized as an MLP  $G_\theta$ , DNSF maps the vertices

on the template human mesh, *i.e.*, T-posed SMPL  $\mathbf{M}_c$ , to style attributes  $\mathbf{c}$  and  $\mathbf{d}$ , in a pose-agnostic manner as:

$$\text{DNSF: } G_\theta(\mathbf{M}_c) \mapsto \{\mathbf{c}, \mathbf{d}\}. \quad (3)$$

We also employ the Fourier feature-based positional encoding to the mesh vertices, which helps the style field to cover higher frequency details [54]. In detail, the MLP  $G_\theta$  gets the positional encoded feature as input and outputs per-vertex RGB value,  $\mathbf{c} \in [0, 1]^{V \times 3}$ , and per-vertex displacement value,  $\mathbf{d} \in [-0.1, 0.1]^V$ , along the per-vertex normal direction. The predicted style attributes are then applied to the content posed mesh  $\mathbf{M}_i$  to produce the stylized human mesh  $\mathbf{M}_i^*$ .

**Text-driven DNSF optimization.** The core of the text-driven DNSF optimization is to maximize the semantic correlation between the visual mesh observation and the input text prompt. However, we cannot directly utilize CLIP to measure the semantic correlation with the created 3D mesh itself because the CLIP visual encoder is designed and trained only for 2D images.

We leverage an intuitive idea that the observations of a 3D object can be described similarly from any viewpoint [23, 24, 57]. To utilize the representation power of CLIP as a supervision signal, we first render images of the 3D meshes for input compatibility. With randomly sampled  $N$  camera poses,  $\mathbf{p} = [\mathbf{p}_1, \dots, \mathbf{p}_N]$ , we differentially render the stylized mesh  $\mathbf{M}_i^*$  to get  $N$ -view rendered images  $\mathbf{I}_{ij}^*$ ,  $\forall j \in \{1, 2, \dots, N\}$  [56]. Thereby, our main optimization objective, *semantic loss*, is defined with the pre-trained CLIP image and text encoders,  $\mathbf{g}(\cdot)$  and  $\mathbf{h}(\cdot)$ , as:

$$\mathcal{L}_s = 1 - \frac{\bar{\mathbf{g}}(\mathbf{I}_i^*)^\top \mathbf{h}(y)}{\|\bar{\mathbf{g}}(\mathbf{I}_i^*)\|_2 \|\mathbf{h}(y)\|_2}, \quad \bar{\mathbf{g}}(\mathbf{I}_i^*) = \frac{1}{N} \sum_{j=1}^N \mathbf{g}(\mathbf{I}_{ij}^*), \quad (4)$$

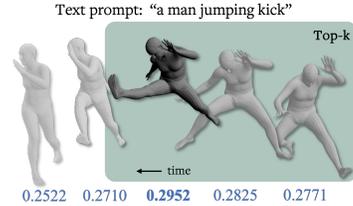
where  $y$  denotes the input text prompt, and  $\mathbf{g}(\mathbf{I}_i^*) \in \mathbb{R}^{512}$  and  $\mathbf{h}(y) \in \mathbb{R}^{512}$  the unnormalized CLIP embedding vectors for the image and the text prompt, respectively. The semantic loss is basically a cosine similarity between the normalized mean embedding vectors for  $N$  rendered images of the stylized mesh  $\mathbf{M}_i^*$  and the normalized embedding for the input text prompt  $y$ .

**Spatio-temporal view augmentation.** Prior works show that *spatial* augmentations, such as 3D viewpoint or 2D image augmentations, improve the quality of content generation [15, 23, 24, 34]. We extend it to *spatio-temporal view augmentation*, where we propose to leverage both multi-view property and human motion originating from the temporal movement. This naturally diversifies views in a combinatorial way by the spatio-temporal context of human motion.

The strength of DNSF is amplified with the spatio-temporal view augmentation. Recall that DNSF  $G_\theta$  takes a template SMPL mesh as an input, which is pose-agnostic. Therefore, the semantic loss  $\mathcal{L}_s$  can be measured with any content mesh  $\mathbf{M}_{i \in \{1, \dots, T\}}$  in the motion sequence for learning DNSF. One can sample and use the center frame or the frame that conforms best with the text prompt. This increases the chance to measure the loss with a view favorable to DNSF learning. Considering that the naïve selection of the content mesh fails to generate plausible color and geometric details when it does not conform with the text prompt, the content mesh sampling strategy is a crucial design choice.

**Multi-modal content mesh sampling.** One naïve way to choose the content mesh to stylize with is to randomly select a single mesh within a mesh sequence.

However, as careful text prompt and its semantic alignment with the mesh’s rendered image are crucial for the optimization [34, 62], we design *multi-modal content mesh sampling* that finds the best text-conforming meshes within the motion. Specifically, we render the images of the content mesh sequence,  $\mathbf{I}(\mathbf{M}_{1:T})$ , and compute each image’s CLIP similarity score with  $y$ . For example, when the text prompt is given as “a man jumping kick”, we render each mesh in retrieved motion into an image and find the semantically matching frames with the jumping kick action (refer to Fig. 4).



**Fig. 4.** We choose the top-k best matching mesh frames that conform with a given text prompt. We show the CLIP scores.

**Mask-weighted embedding attention.** Before the pre-trained CLIP encoder  $\mathbf{g}(\cdot)$  encodes the rendered images, we apply differentiable 2D image augmentations, including random crop and perspective transformations [46]. Such 2D augmentations help DNSF learn style attributes from diverse perspective images, thus can achieve better generalization in 3D contents [23, 34].

However, the problem occurs when careless random crops are applied. The prior work [34] applies extreme close-ups to crop the rendered images, which severely samples the empty renders (see Fig. 5). Such redundant images do not conform to the text prompt even for the properly stylized meshes and distract the stable DNSF optimization.



**Fig. 5.** Naïve mean embedding of the random-cropped renders may distract the optimization of DNSF.

We mitigate this issue by weighting the CLIP embedded vectors  $\mathbf{g}(\mathbf{I}_{ij}^*)$  from  $N$  different camera poses  $\{\mathbf{p}_j\}_{j=1}^N$  according to each image’s foreground pixel ratio. In other words, we reject the embedding vector  $\mathbf{g}(\mathbf{I}_{ij}^*)$ , if  $\mathbf{I}_{ij}^*$  has an extremely small portion of mesh foreground pixels in it. We call this *mask-weighted embedding attention*, and implement it by simply adding the weight  $w_{ij}$  to Eq. (4) as:

$$\bar{\mathbf{g}}(\mathbf{I}_i^*) = \sum_{j=1}^N \frac{w_{ij}}{\sum_{k=1}^N w_{ik}} \mathbf{g}(\mathbf{I}_{ij}^*), \quad w_{ij} = \frac{1}{HW} \sum_{H,W} \mathbb{1}[\mathbf{m}_{ij}(h, w) = 1], \quad (5)$$

where  $H$  and  $W$  denote the height and width of the rendered image  $\mathbf{I}_{ij}^*$  and its foreground mask  $\mathbf{m}_{ij}$ .

## 6 Experiments

In this section, we evaluate CLIP-Actor in several aspects. Since our model is in an over-fitting regime and the first approach that addresses the stylization of 3D human meshes in “motion” conditioned on the natural language, we mainly ablate our technical components and the design choices qualitatively.



**Fig. 6. Qualitative results of CLIP-Actor.** Each image shows the representative frame from the recommended motion sequence, with detailed surface geometries and textures, along with the input text prompt. CLIP-Actor shows good action and style consistency, vivid and attractive texture results.

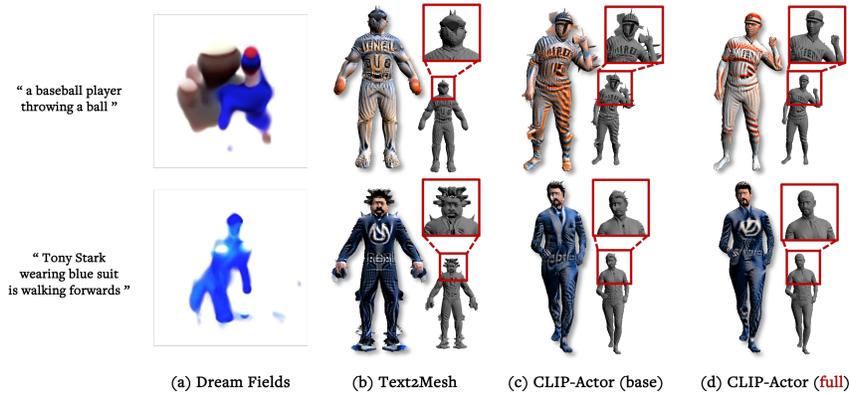
Specifically, we describe the models for our evaluation and ablation study in Sec. 6.1. We then show our qualitative and quantitative results compared with competing methods in Secs. 6.2 and 6.3. In Sec. 6.4, we empirically support our choice of the hierarchical motion retrieval. In Sec. 6.5, we show how our decoupled style representation, mesh sampling, view augmentation, and attention mechanism help us achieve better qualitative results.

## 6.1 Model Description

We define our full model, **CLIP-Actor**, as the one that uses the top-3 best matching mesh frames conforming with a given text prompt, *i.e.*, using multi-modal content mesh sampling and spatio-temporal view augmentations along with the mask-weighted embedding attention. Also, the CLIP-Actor (base) is the model that utilizes only the center frame of the retrieved motion sequence and does not utilize DNSF, *i.e.*, using posed mesh to learn the style field. Still, CLIP-Actor (base) is a strong baseline model since it at least mitigates the limitation of Text2Mesh [34] by suggesting the initial mesh to the neural optimization.

## 6.2 Qualitative Results

In Fig. 6, we visualize CLIP-Actor’s recommendation and mesh stylization results for a given text prompt. With only a single text prompt, CLIP-Actor can retrieve visually conforming motion sequences containing representative poses. Moreover, CLIP-Actor can capture the subject’s representative identities. For example, the geometric and texture details such as Spiderman’s webbed costume,



**Fig. 7. Qualitative comparison.** We compare CLIP-Actor with other competing methods [23, 34] and our strong baseline, CLIP-Actor (base). Given the same text prompts as input, (a) Dream Fields shows abstract generations, which are blurry and hard to recognize, (b) Text2Mesh shows a better generation than Dream Fields but suffers from substantial defects on the surfaces. (c) CLIP-Actor (base) shows more text-conforming meshes with human-recognizable style attributes but still suffers from surface defects. (d) CLIP-Actor shows human-recognizable and semantically conforming action, while presenting detailed color and geometry, such as hairstyle and face identities.

the iconic color of Lionel Messi’s uniform, Freddie Mercury’s hairstyle, and the robe that Gandhi wears are well-illustrated in Fig. 6.

We also evaluate CLIP-Actor with other recent competing methods, Dream Fields [23] and Text2Mesh [34], and our strong baseline model, CLIP-Actor (base). Figure 7 illustrates the visual comparison of the methods.

Given the same text prompts, Dream Fields shows blurry and non-human-recognizable renderings of the generated 3D content. We postulate that such performance degradation is due to the lack of structural prior when training the Dream Fields. Dream Fields learns the occupancy and the color of 3D points in virtual space without any structural guidance. For example, we cannot impose the human body’s physical constraints on Dream Fields when performing specific actions (refer to Fig. 7a). We found that applying only semantic supervision to such a highly unrestrained content generation process fails to handle physically constrained human motion and textures.

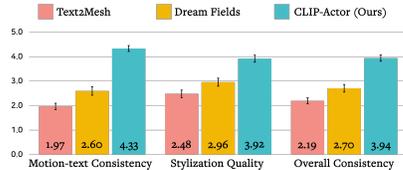
Text2Mesh shows enhanced texture generation than Dream Fields. However, it still fails since the given artist-designed human mesh is absolutely uncorrelated with the target human action. Such limitation is originated from Text2Mesh’s highly coupled style field, which learns the style field from the “posed” content mesh. Text2Mesh also clamps the per-vertex displacement to lie in a limited range, preventing style attributes from largely changing the content [34]. On the other hand, by adding our novel text-driven human motion recommendation module before Text2Mesh, and providing the text-conforming content mesh as an

initial point, *i.e.*, CLIP-Actor (base), we can significantly enhance Text2Mesh’s qualitative performance.

Finally, our full CLIP-Actor further enhances the qualitative result by capturing semantically meaningful details such as a cap on a baseball player and the hairstyle of Tony Stark (see Fig. 7d) while reducing the messy spikes. Our novel DNSF, multi-modal content mesh sampling, and spatio-temporal view augmentation enable CLIP-Actor to leverage multi-view renderings originating from multi-frame human motion; thus, results are much smoother and text-conforming. More importantly, note that all the other methods except CLIP-Actor cannot handle human motion. CLIP-Actor recommends text-conforming human motion and synthesizes temporally consistent and pose-agnostic mesh style attributes.

### 6.3 Quantitative Results

Since there are no benchmarks for our task, we conduct a user study to evaluate CLIP-Actor quantitatively. We ask 46 non-expert users to score (1-5) five random text-avatar paired results regarding motion-text consistency, stylization quality, and overall consistency with the given text. Figure 8 shows that our CLIP-Actor outperforms other competing methods in all aspects, while none of the competing methods scored higher than a neutral point (*i.e.*,  $<3$ ). The differences are significantly noticeable in the motion-text consistency, which validates our good action consistency.



**Fig. 8. User study results.** CLIP-Actor outperforms other competing methods in various aspects.

### 6.4 Evaluation on Retrieval Module

We validate the performance of our retrieval module design choice by comparing it with other variants. We use the SICK dataset [33] that contains contextually similar sentence pairs generated from image descriptions. We build module variants to simulate various retrieval scenarios. Details about the dataset and the experiment settings can be found in the supplementary material.

**Retrieval module variants.** We consider two hierarchical modules and two single-stage baselines for the design variants. Our full hierarchical retrieval module utilizes pre-trained CLIP and MPNet sequentially (see CLIP+MPNet in Sec. 4). We also design the reverse ordered hierarchical module, MPNet+CLIP. For hierarchical models, top-k candidates are matched at the first stage, and the best-matched item is selected after re-ranking. Finally, our single-stage baselines are the modules that only use either pre-trained CLIP or MPNet encoder.

**Quantitative results.** We empirically show the performance of our hierarchical motion retrieval module by comparing it with the different design variants. As shown in Table 1a, the single-stage baselines show comparable results with our hierarchical model in the SICK4.8 setting, where the sentence pairs are more related to each other (refer to Table 1b). However, CLIP shows higher precision than MPNet in SICK4.4. We postulate that CLIP catches the visual semantics,

**Table 1. Evaluation results on the retrieval module.** (a) The precision is measured as the percentage of matched sentence pairs among all pairs. (b) The samples of each range of the SICK dataset. Our hierarchical retrieval outperforms all the other variants.

(a)				(b)	
Retrieval Module	Retrieval Precision [%]			SICK ID	A: A man is slicing an onion B: An onion is being sliced by a man
	SICK4.8	SICK4.4	SICK[4.4,4.8]		
CLIP	91.94	85.21	81.62	SICK4.9	A: A skateboarder is jumping off a ramp B: A skateboarder is making a jump off a ramp
MPNet	91.94	83.56	80.55	SICK4.8	A: A man is performing with a guitar B: A man is playing a guitar
MPNet+CLIP	91.34	85.48	80.41	SICK4.4	A: A man is cutting a paper B: There is no man cutting a paper
CLIP+MPNet (Ours)	<b>92.24</b>	<b>85.75</b>	<b>81.90</b>	SICK4.3	

*e.g.*, similar context can be imagined from “performing with a guitar” and “playing a guitar”. On the other hand, the language expert, MPNet, focuses on the textual difference of description, *e.g.*, “performing with” and “playing”. Thus, it is sensitive to the textual structure. In the setting with the increased number of samples, *i.e.*, SICK[4.4,4.8], CLIP shows comparable results with ours but is still insufficient without the help of the language expert. Moreover, the MPNet+CLIP shows unstable performances over settings. In contrast, our full retrieval module consistently outperforms over all settings. We demonstrate that a coarse-to-fine matching system achieves favorable retrieval performance on natural language.

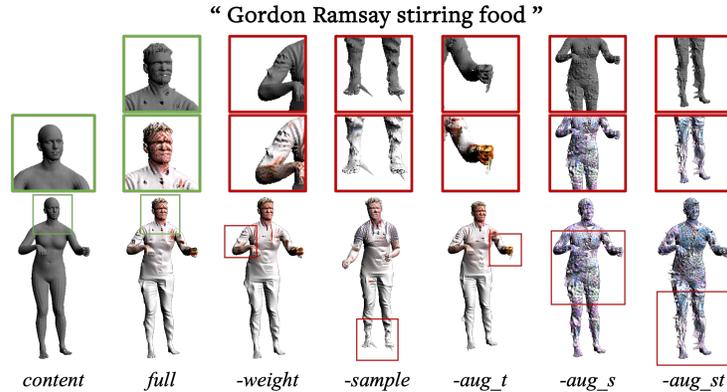
### 6.5 Ablation on Decoupled Neural Style Fields (DNSF)

We analyze CLIP-Actor by ablating each of the components of DNSF. Figure 9 shows the qualitative ablation results for our major technical components.

**Effects of temporal augmentation.** First, we remove temporal view augmentation so that DNSF utilizes only a single mesh frame (top-1). Removing the multi-frame renderings significantly degrades the visual quality, where it presents noticeable spikes on the surface and unrealistic colors ( $-aug\_t$  in Fig. 9). Since our full model utilizes top-3 relevant frames and 2D, 3D augmentations, it leverages multi-view of stylized mesh, which regularizes the model from overfitting [23, 34].

**Effects of multi-modal content mesh sampling.** We also compare the full CLIP-Actor with the model without *multi-modal content mesh sampling*. Multi-modal content mesh sampling enables DNSF to begin its optimization with better initialization that conforms with the text prompt. Naïve sampling of the content mesh yields unrecognizable face identities, degraded texture, and geometric details (see  $-sample$  in Fig. 9).

**Effects of mask-weighted embedding attention.** The *mask-weighted embedding attention* adds detailed touches to the stylized meshes. By preventing empty renderings from guiding the optimization, it enables learning fine geometric and texture details via focused gradient flow in back-propagation. When the augmented rendered images contain extreme close-ups of distal body regions, such as tiptoe or fingertips, our embedding attention method draws the DNSF’s attention to the mesh foreground pixels rather than empty space with focused



**Fig. 9. Ablation results.** We remove each of CLIP-Actor’s components to validate corresponding effects. Our *full* model shows the most smooth geometry and vivid color.

gradient flow. In Fig. 9, *-weight* shows the result when we train DNSF without our attention mechanism. Our full CLIP-Actor shows much smooth, fine-grained geometric details. We believe our novel attention mechanism can be applied to not only text-driven 3D object manipulation pipelines [23, 34, 62] but also differentiable rendering applications [24, 35, 61].

## 7 Conclusion

We present CLIP-Actor, a text-driven animated human mesh synthesis system. Leveraging multi-modal aware and semantic textual matching, CLIP-Actor recommends the best semantically matching human motion sequence with the input text prompt in a hierarchical manner. Our CLIP-Actor then stylizes the meshes of recommended motion by synthesis-through-optimization in a pose-agnostic manner via decoupled neural style fields. We further develop novel neural optimization techniques to utilize multi-modal sampling and embedding weighting, which stabilize and enhance the detailization and texturization quality.

CLIP-Actor can be extended to other parametric mesh models, such as hands and animals [6–8, 47, 60, 63, 64], enabling diverse animation of 3D objects. One promising future application of CLIP-Actor would be a dataset generation of stylized meshes in motion, paired with natural language description. We believe such multi-modal datasets can boost exciting future applications.

**Acknowledgment.** This work was supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT) (No.2022-00164860, Development of Human Digital Twin Technology Based on Dynamic Behavior Modeling and Human-Object-Space Interaction; and No.2021-0-02068, Artificial Intelligence Innovation Hub).

## References

1. Ahn, H., Ha, T., Choi, Y., Yoo, H., Oh, S.: Text2action: Generative adversarial synthesis from language to action. In: IEEE International Conference on Robotics and Automation (ICRA) (2018) 4
2. Ahuja, C., Morency, L.P.: Language2pose: Natural language grounded pose forecasting. International Conference on 3D Vision (3DV) (2019) 4
3. Barron, J.T., Mildenhall, B., Verbin, D., Srinivasan, P.P., Hedman, P.: Mip-nerf 360: Unbounded anti-aliased neural radiance fields. IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2021) 3
4. Bhatnagar, B.L., Sminchisescu, C., Theobalt, C., Pons-Moll, G.: Combining implicit function learning and parametric models for 3d human reconstruction. In: European Conference on Computer Vision (ECCV) (2020) 4
5. Bhatnagar, B.L., Tiwari, G., Theobalt, C., Pons-Moll, G.: Multi-garment net: Learning to dress 3d people from images. In: IEEE International Conference on Computer Vision (ICCV) (2019) 1, 4
6. Biggs, B., Boyne, O., Charles, J., Fitzgibbon, A., Cipolla, R.: Who left the dogs out?: 3D animal reconstruction with expectation maximization in the loop. In: European Conference on Computer Vision (ECCV) (2020) 14
7. Biggs, B., Roddick, T., Fitzgibbon, A., Cipolla, R.: Creatures great and SMAL: Recovering the shape and motion of animals from video. In: Asia Conference on Computer Vision (ACCV) (2018) 14
8. Blanz, V., Vetter, T.: A morphable model for the synthesis of 3d faces. In: SIGGRAPH (1999) 14
9. Božič, A., Palafox, P., Zollhöfer, M., Thies, J., Dai, A., Nießner, M.: Neural deformation graphs for globally-consistent non-rigid reconstruction. In: Advances in Neural Information Processing Systems (NeurIPS) (2021) 1, 4
10. Bozic, A., Palafox, P., Zollöfer, M., Dai, A., Thies, J., Nießner, M.: Neural non-rigid tracking. In: Advances in Neural Information Processing Systems (NeurIPS) (2020) 1, 4
11. Burov, A., Nießner, M., Thies, J.: Dynamic surface function networks for clothed human bodies. In: IEEE International Conference on Computer Vision (ICCV) (2021) 4
12. Canfes, Z., Atasoy, M.F., Dirik, A., Yanardag, P.: Text and image guided 3d avatar generation and manipulation. arXiv:2202.06079 (2022) 1
13. Du, Y., Collins, M.K., Tenenbaum, B.J., Sitzmann, V.: Learning signal-agnostic manifolds of neural fields. In: Advances in Neural Information Processing Systems (NeurIPS) (2021) 3
14. Fragkiadaki, K., Levine, S., Felsen, P., Malik, J.: Recurrent network models for human dynamics. IEEE International Conference on Computer Vision (ICCV) (2015) 4
15. Frans, K., Soros, L.B., Witkowski, O.: Clipdraw: Exploring text-to-drawing synthesis through language-image encoders. arXiv:2106.14843 (2021) 3, 8
16. Gafni, G., Thies, J., Zollhöfer, M., Nießner, M.: Dynamic neural radiance fields for monocular 4d facial avatar reconstruction. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2021) 1
17. Gal, R., Patashnik, O., Maron, H., Chechik, G., Cohen-Or, D.: Stylegan-nada: Clip-guided domain adaptation of image generators. In: ACM Transactions on Graphics (SIGGRAPH) (2022) 3

18. Ghosh, A., Cheema, N., Oguz, C., Theobalt, C., Slusallek, P.: Synthesis of compositional animations from textual descriptions. *IEEE International Conference on Computer Vision (ICCV)* (2021) 4
19. Guo, C., Zuo, X., Wang, S., Liu, X., Zou, S., Gong, M., Cheng, L.: Action2video: Generating videos of human 3d actions. *International Journal of Computer Vision (IJCV)* pp. 1–31 (2022) 4
20. Guo, C., Zuo, X., Wang, S., Zou, S., Sun, Q., Deng, A., Gong, M., Cheng, L.: Action2motion: Conditioned generation of 3d human motions. In: *ACM International Conference on Multimedia (MM)* (2020) 4
21. Guo, J., Li, J., Narain, R., Park, H.: Inverse simulation: Reconstructing dynamic geometry of clothed humans via optimal control. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2021) 1, 4
22. Hanser, E., Kevitt, P.M., Lunney, T.F., Condell, J.: Scenemaker: Intelligent multi-modal visualisation of natural language scripts. In: *Proceedings of the 20th Irish conference on Artificial intelligence and cognitive science* (2009) 1
23. Jain, A., Mildenhall, B., Barron, J.T., Abbeel, P., Poole, B.: Zero-shot text-guided object generation with dream fields. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2022) 3, 8, 9, 11, 13, 14
24. Jain, A., Tancik, M., Abbeel, P.: Putting nerf on a diet: Semantically consistent few-shot view synthesis. In: *IEEE International Conference on Computer Vision (ICCV)* (2021) 8, 14
25. Jiang, J., Xia, G.G., Carlton, D.B., Anderson, C.N., Miyakawa, R.H.: Transformer vae: A hierarchical model for structure-aware and interpretable music representation learning. In: *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)* (2020) 4
26. Kato, H., Beker, D., Morariu, M., Ando, T., Matsuoka, T., Kehl, W., Gaidon, A.: Differentiable rendering: A survey. *arXiv:2006.12057* (2020) 3
27. Kim, G., Ye, J.C.: Diffusionclip: Text-guided diffusion models for robust image manipulation. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2022) 3
28. Kwon, G., Ye, J.C.: Clipstyler: Image style transfer with a single text condition. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2022) 3
29. Li, Y., Tang, M., Yang, Y., Huang, Z., Tong, R., Yang, S., Li, Y., Manocha, D.: N-Cloth: Predicting 3D cloth deformation with mesh-based networks. *Computer Graphics Forum (Proceedings of Eurographics)* pp. 547–558 (2022) 1, 4
30. Lin, A.S., Wu, L., Corona, R., Tai, K.W.H., Huang, Q., Mooney, R.J.: Generating animated videos of human activities from natural language descriptions. In: *Proceedings of the Visually Grounded Interaction and Language Workshop at NeurIPS 2018* (2018) 4
31. Loper, M., Mahmood, N., Romero, J., Pons-Moll, G., Black, M.J.: SMPL: A skinned multi-person linear model. *ACM Transactions on Graphics (SIGGRAPH Asia)* 34(6), 248 (2015) 5
32. Ma, Q., Yang, J., Ranjan, A., Pujades, S., Pons-Moll, G., Tang, S., Black, M.J.: Learning to dress 3d people in generative clothing. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2020) 4
33. Marelli, M., Menini, S., Baroni, M., Bentivogli, L., Bernardi, R., Zamparelli, R.: A sick cure for the evaluation of compositional distributional semantic models. In: *International Conference on Language Resources and Evaluation (LREC)* (2014) 12

34. Michel, O., Bar-On, R., Liu, R., Benaim, S., Hanocka, R.: Text2mesh: Text-driven neural stylization for meshes. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2022) [3](#), [7](#), [8](#), [9](#), [10](#), [11](#), [13](#), [14](#)
35. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. In: European Conference on Computer Vision (ECCV) (2020) [3](#), [14](#)
36. Nichol, A., Dhariwal, P., Ramesh, A., Shyam, P., Mishkin, P., McGrew, B., Sutskever, I., Chen, M.: Glide: Towards photorealistic image generation and editing with text-guided diffusion models. In: International Conference on Machine Learning (ICML) (2022) [3](#)
37. Niemeyer, M., Mescheder, L., Oechsle, M., Geiger, A.: Differentiable volumetric rendering: Learning implicit 3d representations without 3d supervision. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2020) [3](#)
38. Palafox, P., Bozic, A., Thies, J., Nießner, M., Dai, A.: Neural parametric models for 3d deformable shapes. In: IEEE International Conference on Computer Vision (ICCV) (2021) [1](#), [4](#)
39. Patashnik, O., Wu, Z., Shechtman, E., Cohen-Or, D., Lischinski, D.: Styleclip: Text-driven manipulation of stylegan imagery. In: IEEE International Conference on Computer Vision (ICCV) (2021) [3](#)
40. Pavlakos, G., Choutas, V., Ghorbani, N., Bolkart, T., Osman, A.A.A., Tzionas, D., Black, M.J.: Expressive body capture: 3D hands, face, and body from a single image. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019) [5](#)
41. Petrovich, M., Black, M.J., Varol, G.: Action-conditioned 3D human motion synthesis with transformer VAE. In: IEEE International Conference on Computer Vision (ICCV) (2021) [4](#)
42. Petrovich, M., Black, M.J., Varol, G.: TEMOS: Generating diverse human motions from textual descriptions. In: European Conference on Computer Vision (ECCV) (2022) [4](#)
43. Plappert, M., Mandery, C., Asfour, T.: Learning a bidirectional mapping between human whole-body motion and natural language using deep recurrent neural networks. *Robotics and Autonomous Systems* **109**, 13–26 (2018) [4](#)
44. Punnakal, A.R., Chandrasekaran, A., Athanasiou, N., Quiros-Ramirez, A., Black, M.J.: BABEL: Bodies, action and behavior with english labels. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2021) [5](#), [6](#)
45. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., Sutskever, I.: Learning transferable visual models from natural language supervision. In: International Conference on Machine Learning (ICML) (2021) [2](#), [3](#)
46. Ravi, N., Reizenstein, J., Novotny, D., Gordon, T., Lo, W.Y., Johnson, J., Gkioxari, G.: Accelerating 3d deep learning with pytorch3d. arXiv:2007.08501 (2020) [3](#), [9](#)
47. Romero, J., Tzionas, D., Black, M.J.: Embodied hands: Modeling and capturing hands and bodies together. *ACM Transactions on Graphics (SIGGRAPH Asia)* **36**(6) (Nov 2017) [5](#), [14](#)
48. Saito, S., Huang, Z., Natsume, R., Morishima, S., Kanazawa, A., Li, H.: Pifu: Pixel-aligned implicit function for high-resolution clothed human digitization. In: IEEE International Conference on Computer Vision (ICCV) (2019) [4](#)
49. Saito, S., Simon, T., Saragih, J., Joo, H.: Pifuhd: Multi-level pixel-aligned implicit function for high-resolution 3d human digitization. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2020) [4](#)

50. Saito, S., Yang, J., Ma, Q., Black, M.J.: SCANimate: Weakly supervised learning of skinned clothed avatar networks. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2021) 1, 4
51. Sanh, V., Debut, L., Chaumond, J., Wolf, T.: Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. arXiv:1910.01108 (2019) 4
52. Shree, V., Asfora, B., Zheng, R., Hong, S., Banfi, J., Campbell, M.: Exploiting natural language for efficient risk-aware multi-robot sar planning. IEEE Robotics and Automation Letters 6(2), 3152–3159 (2021) 1
53. Song, K., Tan, X., Qin, T., Lu, J., Liu, T.Y.: Mpnet: Masked and permuted pre-training for language understanding. Advances in Neural Information Processing Systems (NeurIPS) (2020) 7
54. Tancik, M., Srinivasan, P.P., Mildenhall, B., Fridovich-Keil, S., Raghavan, N., Singhal, U., Ramamoorthi, R., Barron, J.T., Ng, R.: Fourier features let networks learn high frequency functions in low dimensional domains. Advances in Neural Information Processing Systems (NeurIPS) (2020) 8
55. Tevet, G., Gordon, B., Hertz, A., Bermano, A.H., Cohen-Or, D.: Motionclip: Exposing human motion generation to clip space. In: European Conference on Computer Vision (ECCV) (2022) 4
56. Tsang, C.F., Shugrina, M., Lafleche, J.F., Takikawa, T., Wang, J., Loop, C., Chen, W., Jatavallabhula, K.M., Smith, E., Rozantsev, A., Perel, O., Shen, F., Gao, J., Fidler, S., State, G., Gorski, J., Xiang, T., Li, J., Li, M., Lebedev, R.: Kaolin (2019) 8
57. Wang, C., Chai, M., He, M., Chen, D., Liao, J.: Clip-nerf: Text-and-image driven manipulation of neural radiance fields. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2022) 8
58. Wang, S., Mihajlovic, M., Ma, Q., Geiger, A., Tang, S.: Metaavatar: Learning animatable clothed human models from few depth images. In: Advances in Neural Information Processing Systems (NeurIPS) (2021) 4
59. Yoon, Y., Ko, W.R., Jang, M., Lee, J., Kim, J., Lee, G.: Robots learn social skills: End-to-end learning of co-speech gesture generation for humanoid robots. In: IEEE International Conference on Robotics and Automation (ICRA) (2019) 1
60. Youwang, K., Ji-Yeon, K., Joo, K., Oh, T.H.: Unified 3d mesh recovery of humans and animals by learning animal exercise. In: British Machine Vision Conference (BMVC) (2021) 14
61. Yu, A., Ye, V., Tancik, M., Kanazawa, A.: pixelNeRF: Neural radiance fields from one or few images. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2021) 14
62. Zhang, R., Guo, Z., Zhang, W., Li, K., Miao, X., Cui, B., Qiao, Y., Gao, P., Li, H.: Pointclip: Point cloud understanding by clip. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2021) 9, 14
63. Zuffi, S., Kanazawa, A., Berger-Wolf, T., Black, M.J.: Three-d safari: Learning to estimate zebra pose, shape, and texture from images “in the wild”. In: IEEE International Conference on Computer Vision (ICCV) (2019) 14
64. Zuffi, S., Kanazawa, A., Jacobs, D., Black, M.J.: 3D menagerie: Modeling the 3D shape and pose of animals. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2017) 14