

# MagicArticulate: Make Your 3D Models Articulation-Ready

Chaoyue Song<sup>1,2</sup>, Jianfeng Zhang<sup>†2</sup>, Xiu Li<sup>2</sup>, Fan Yang<sup>1</sup>, Yiwen Chen<sup>1</sup>, Zhongcong Xu<sup>2</sup>,  
Jun Hao Liew<sup>2</sup>, Xiaoyang Guo<sup>2</sup>, Fayao Liu<sup>3</sup>, Jiashi Feng<sup>2</sup>, Guosheng Lin<sup>†1</sup>

<sup>1</sup>Nanyang Technological University <sup>2</sup>ByteDance Seed

<sup>3</sup>Institute for Inforcomm Research, A\*STAR

## Abstract

*With the explosive growth of 3D content creation, there is an increasing demand for automatically converting static 3D models into articulation-ready versions that support realistic animation. Traditional approaches rely heavily on manual annotation, which is both time-consuming and labor-intensive. Moreover, the lack of large-scale benchmarks has hindered the development of learning-based solutions. In this work, we present MagicArticulate, an effective framework that automatically transforms static 3D models into articulation-ready assets. Our key contributions are threefold. First, we introduce Articulation-XL, a large-scale benchmark containing over 33k 3D models with high-quality articulation annotations, carefully curated from Objaverse-XL. Second, we propose a novel skeleton generation method that formulates the task as a sequence modeling problem, leveraging an auto-regressive transformer to naturally handle varying numbers of bones or joints within skeletons and their inherent dependencies across different 3D models. Third, we predict skinning weights using a functional diffusion process that incorporates volumetric geodesic distance priors between vertices and joints. Extensive experiments demonstrate that MagicArticulate significantly outperforms existing methods across diverse object categories, achieving high-quality articulation that enables realistic animation. Project page: <https://chaoyuesong.github.io/MagicArticulate>.*

## 1. Introduction

The rapid advancement of 3D content creation has led to an increasing demand for articulation-ready 3D models, especially in gaming, VR/AR, and robotics simulation. Converting static 3D models into articulation-ready versions traditionally requires professional artists to manually place skeletons, define joint hierarchies and specify

skinning weights, which is both time-consuming and demands significant expertise, making it a major bottleneck in modern content creation pipelines.

To address these issues, various automatic approaches for skeleton extraction have been proposed, which can be categorized into template-based [3, 22] and template-free methods [2, 17, 42, 43]. Template-based methods, like Pinocchio [3], fit predefined skeletal templates to input shapes. While they achieve satisfactory results for specific categories like human characters, they struggle to generalize to objects with varying structural patterns. Moreover, these methods mostly rely on distance metrics between joints and vertices for skinning weight prediction, which often fail on shapes with complex topology. Many template-free methods [2, 6, 17, 24, 36] extract curve skeletons from meshes or point clouds using shape medial axis or the centerline of shapes, but often produce densely packed joints that are unsuitable for animation. Recent deep learning methods like RigNet [43] have shown promise in predicting skeletons and skinning weights directly from input shapes. However, they rely heavily on carefully crafted features and make strong assumptions about shape orientation, limiting their ability to handle diverse object categories. These limitations stem from two fundamental challenges: the lack of a large-scale, diverse dataset for training generalizable models, and the inherent difficulty in designing an effective framework capable of handling complex mesh topologies, accommodating varying skeleton structures, and ensuring the coherent generation of both accurate skeletons and skinning weights.

To overcome these challenges, we first introduce Articulation-XL, a large-scale dataset containing over 33k 3D models with high-quality articulation annotations carefully curated from Objaverse-XL [11, 12]. Built upon this benchmark, we propose MagicArticulate, a novel framework that addresses both skeleton generation and skinning weight prediction. Specifically, we reformulate skeleton generation as an auto-regressive sequence modeling task, enabling our model to naturally handle varying numbers of bones or joints within skeletons across different 3D mod-

<sup>†</sup> Corresponding authors.

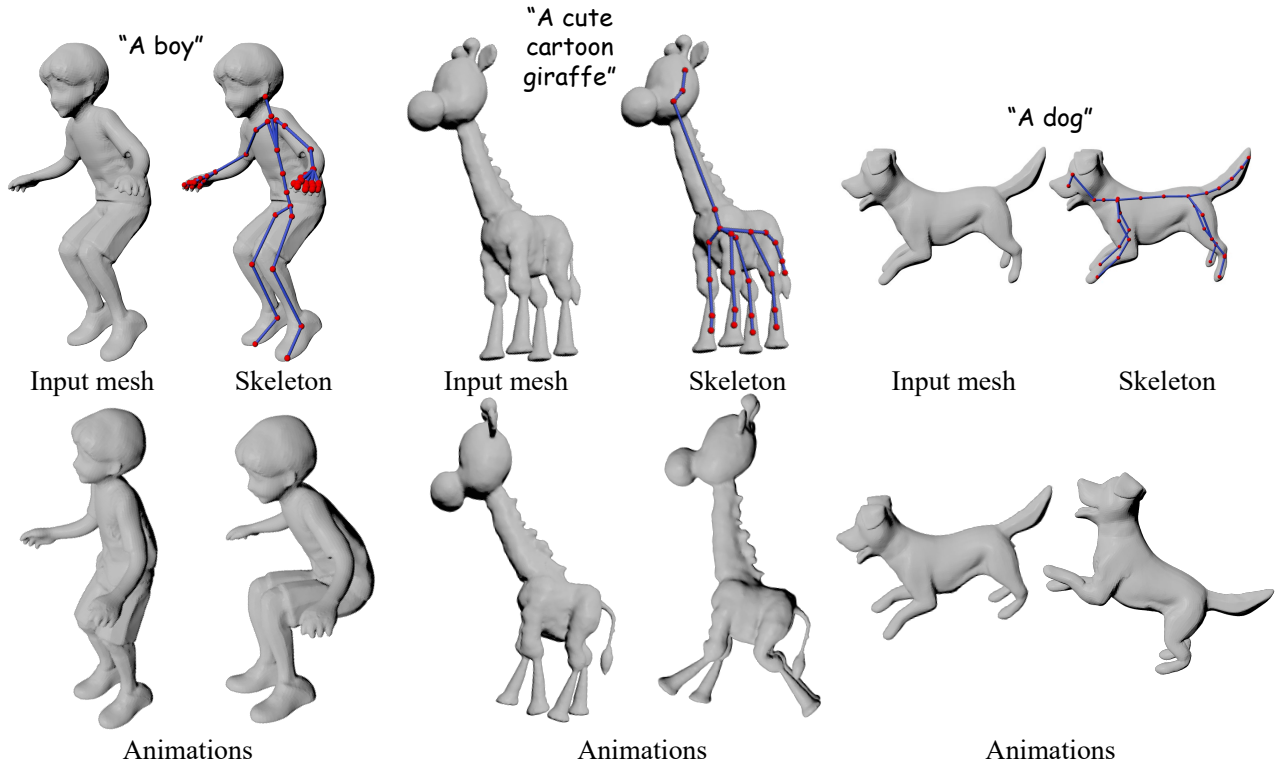


Figure 1. **Given a 3D model, MagicArticulate can automatically generate the skeleton and skinning weights, making the model articulation-ready without further manual refinement.** The input meshes are generated by Rodin Gen-1 [50] and Tripo 2.0 [1]. The meshes and skeletons are rendered using Maya Software Renderer [19].

els. For skinning weight prediction, we develop a functional diffusion framework that learns to generate smoothly transitioning skinning weights over mesh surfaces by incorporating volumetric geodesic distance priors between vertices and joints, effectively handling complex mesh topologies that challenge traditional geometric-based methods. These designs demonstrate superior scalability on large-scale datasets and generalize well across diverse object categories, without requiring assumptions about shape orientation or topology.

Extensive experiments on our Articulation-XL and ModelsResource [38] collected by Xu et al. [42, 43] demonstrate the effectiveness of MagicArticulate in both skeleton generation and skinning weight prediction. The proposed methods also generalize well to 3D models from various sources, including artist-created assets, and models generated by AI techniques. With the generated skeleton and skinning weights, our method automatically creates ready-to-animate assets that support natural pose manipulation without manual refinement (Figure 1), particularly beneficial for large-scale animation content creation.

Our key contributions include: (1) The first large-scale articulation benchmark containing over 33k models with

high-quality articulation annotations; (2) A novel two-stage framework that effectively handles both skeleton generation and skinning weight prediction; (3) State-of-the-art performance and demonstrated practicality in real-world animation pipelines.

## 2. Related works

### 2.1. Skeleton generation

There are two categories of methods for creating skeletons in 3D models. The first category relies on predefined templates [3, 22] or additional annotations [10, 18, 21, 44]. Pinocchio [3] is a pioneering method for automatically extracting an animation skeleton from an input 3D model. It fits a predefined skeleton template to the 3D model, evaluating the fitting cost for different templates and selecting the most suitable one for a given model. Li et al. [22] proposed a deep learning-based method to estimate joint positions for a given human skeletal template. However, these template-based methods are limited to rigging characters whose articulation structures are compatible with the predefined templates, making it difficult to generalize to objects with distinct structures.

There are also methods that rely on additional inputs or annotations to generate skeletons for 3D models, including point cloud sequences [44], mesh sequences [10, 21], and manual annotations [18]. Additionally, recent works [34, 35, 45, 47–49] have focused on learning the joints and bones of articulated objects directly from videos to reconstruct object motion. In contrast, our approach aims to generate skeletons using only 3D models as input.

The second category consists of template-free methods that operate without relying on predefined templates or additional annotations. Many approaches [2, 6, 17, 24, 36] are designed to extract curve skeletons from meshes or point clouds by utilizing the medial axis or the centerline of shapes. These methods often result in densely packed joints that are unsuitable for effective articulation and animation. Recent deep-learning approaches have also been developed to learn skeletons directly from input shapes without relying on predefined templates. These methods are generally trained on datasets containing thousands of rigged characters, allowing them to generate skeletons that align with articulated components. For instance, Xu et al. [42] introduced a volumetric network designed to generate skeletons for input 3D models. RigNet [43] leverages graph convolutions to learn mesh representations, thereby enhancing the accuracy of skeleton extraction. However, it relies on the strong assumption that the input training and test shapes maintain a consistent upright and front-facing orientation.

In this work, we formulate skeleton generation as an auto-regressive problem to accommodate the varying number of bones in different 3D models. By generating bones auto-regressively, our method dynamically adapts to each model’s specific requirements, ensuring flexibility and accuracy in skeleton creation.

## 2.2. Skinning weight prediction

To make 3D models ready for articulation, we also predict skinning weights conditioned on the 3D shape and corresponding skeleton, which define the influence of each joint on each vertex of the mesh.

Several geometric-based techniques have been introduced for skinning [3, 13, 14, 20]. These methods assign skinning weights based on the distance between joints and vertices. However, this distance-based assumption often fails when the 3D shape has a complex topology. Deep learning-based methods [23, 25, 27, 43], such as NeuroSkinning [25], take a skeleton template as input and predict skinning weights using a learned graph neural network. RigNet [43] utilizes intrinsic shape representations that capture geodesic distances between vertices and bones, often struggles with highly intricate mesh topologies and may require extensive feature engineering to maintain performance across varied object categories. SkinningNet [27] employs a two-stream graph neural network to compute

skinning weights directly from input meshes and the corresponding skeletons. However, the performance of these GNN-based methods can degrade when applied to datasets with highly varying orientations, such as Articulation-XL, leading to reduced accuracy and robustness in complex and varied scenarios.

In this work, we predict skinning weights in a functional diffusion process by incorporating volumetric geodesic distance priors between vertices and joints. This approach effectively handles complex mesh topologies and diverse skeletal structures without the constraints of shape orientations.

## 2.3. Auto-regressive 3D generation

Recently, auto-regressive models have been widely used in 3D mesh generation [7–9, 28, 30, 37, 41]. MeshGPT [30] models meshes as sequences of triangles and tokenizes them using a VQ-VAE [39]. It then employs an auto-regressive transformer to generate the token sequences. This approach enables the creation of meshes with varying face counts. However, most subsequent methods [7, 8, 41] are limited to generating meshes up to 800 faces, due to the computational cost of mesh tokenization. MeshAnythingV2 [9] introduces Adjacent Mesh Tokenization (AMT), doubling the maximum face count to 1,600. EdgeRunner [37] further increases this limit to 4,000 faces by enhancing mesh tokenization techniques. In this work, we explore the potential of auto-regressive models for shape-conditioned skeleton generation. To achieve this, we formulate skeletons as sequences of bones. Unlike mesh generation, which focuses on creating detailed and realistic shapes by utilizing a high number of faces, skeleton generation prioritizes accuracy over complexity. Accurate skeletons are crucial for realistic articulation and animation, and typically consist of fewer than 100 bones, as indicated by the statistics in Articulation-XL.

## 3. Articulation-XL

To facilitate large-scale learning of 3D model articulation, we present Articulation-XL, a comprehensive dataset curated from Objaverse-XL [11, 12]. Our dataset construction pipeline consists of three main stages: initial filtering, VLM-based filtering, and category annotation. **We will release our Articulation-XL to facilitate future work.**

**Initial data collection.** We begin by identifying 3D models from Objaverse-XL that contain both skeleton and skinning weight annotations. To ensure data quality and practical utility, we apply the following filtering criteria: 1) we remove duplicate data based on both skeleton and mesh similarity; 2) we exclude models with only a single joint/bone structure; 3) we filter out data with more than 100 bones, which constitute a negligible portion of the dataset. This

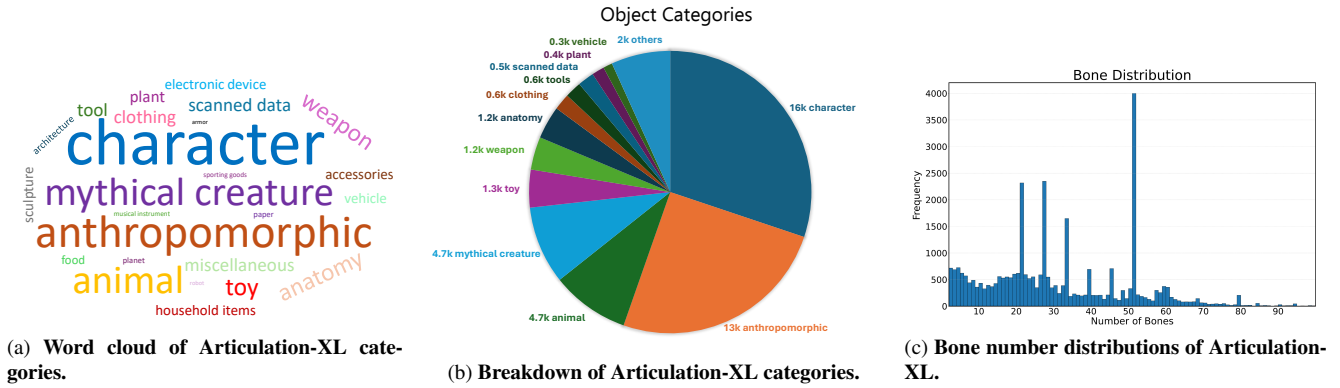


Figure 2. Articulation-XL statistics.




Figure 3. Some examples from Articulation-XL alongside examples of poorly defined skeletons that were curated out.

initial filtering yields 38.8k candidate models with articulation annotations.

**VLM-based filtering.** However, we observe that many initial candidates contain poorly defined skeletons that may impair learning (see Figure 3). To ensure dataset quality, we further implement a Vision-Language Model (VLM)-based filtering pipeline: 1) we render each object with its skeleton from four viewpoints; 2) and then utilize GPT-4o [29] to assess skeleton quality based on specific criteria (detailed in supplementary). This process results in a final collection of over 33k 3D models with high-quality articulation annotations, forming the curated dataset Articulation-XL <sup>1</sup>. The dataset exhibits diverse structural complexity: the number of bones per model ranges from 2 to 100, and the number of joints ranges from 3 to 101. The distribution of bone numbers is illustrated in Figure 2c.

**Category label annotation.** Additionally, we also leverage a Vision-Language Model (VLM) to automatically as-

<sup>1</sup>We have expanded the dataset to over 48K models in Articulation-XL2.0. For further details, please refer to <https://huggingface.co/datasets/chaoyue7/Articulation-XL2.0>.

sign category labels to each model using specific instructions. The distribution of these categories is illustrated via a word cloud and a pie chart, as shown in Figure 2a and Figure 2b, respectively. We observe a rich diversity of object categories, with human-related models forming the largest subset. Detailed statistics and distribution analyses are provided in the supplementary material.

## 4. Methods

We propose a two-stage pipeline to make 3D models articulation-ready. Given an input 3D mesh, our method first employs an auto-regressive transformer to generate a structurally coherent skeleton (Section 4.1). Subsequently, we predict skinning weights in a functional diffusion process, conditioning on both the input shape and its corresponding skeleton (Section 4.2).

### 4.1. Auto-regressive skeleton generation

In the initial stage of MagicArticulate, we generate skeletons for 3D models. Unlike previous approaches that rely on fixed templates, our method can handle the inherent structural diversity of 3D objects through an auto-regressive generation framework, as presented in Figure 5.

#### 4.1.1. Problem formulation

Given an input 3D mesh  $\mathcal{M}$ , our goal is to generate a structurally valid skeleton  $\mathcal{S}$  that captures the articulation structure of the object. A skeleton consists of two key components: a set of joints  $\mathbf{J} \in \mathbb{R}^{j \times 3}$  defining spatial locations, and bone connections  $\mathbf{B} \in \mathbb{N}^{b \times 2}$  specifying the topological structure through joint indices. Formally, we aim to learn the conditional distribution:

$$p(\mathcal{S}|\mathcal{M}) = p(\mathbf{J}, \mathbf{B}|\mathcal{M}), \quad (1)$$

where  $\mathcal{M}$  can be sourced from various inputs, including direct 3D models, text-to-3D generation, or image-based reconstruction.

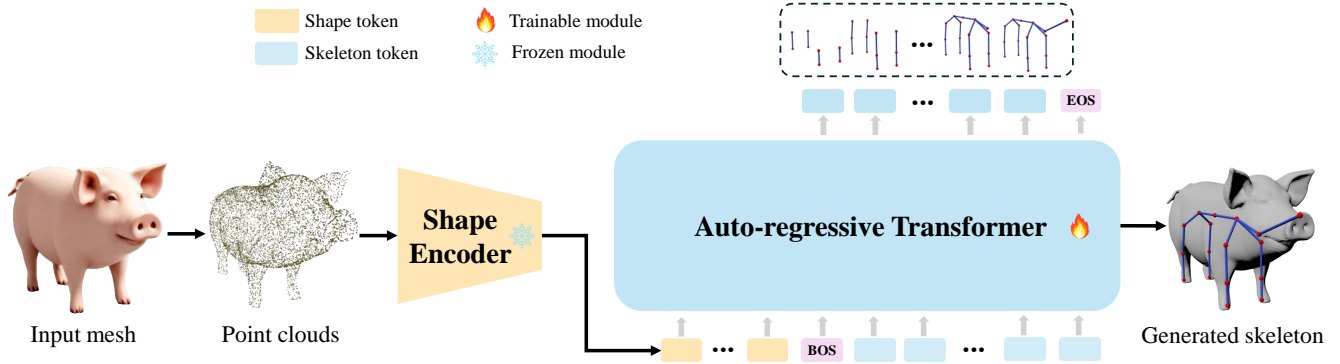


Figure 4. **Overview of our method for auto-regressive skeleton generation.** Given an input mesh, we begin by sampling point clouds from its surface. These sampled points are then encoded into fixed-length shape tokens, which are appended to the start of skeleton tokens to achieve auto-regressive skeleton generation conditioned on input shapes. The input mesh is generated by Rodin Gen-1 [50].

A key challenge in skeleton generation lies in the variable complexity of articulation structures across different objects. Traditional approaches [3, 22] often adopt predefined skeleton templates, which work well for specific categories like human bodies but fail to generalize to objects with diverse structural patterns. This limitation becomes particularly apparent when dealing with our large-scale dataset that contains a wide range of object categories.

To address this challenge, we draw inspiration from recent advances in auto-regressive mesh generation [9, 30] and reformulate skeleton generation as a sequence modeling task. This novel formulation allows us to: 1) handle varying numbers of bones or joints within skeletons across different 3D models; 2) capture the inherent dependencies between bones; 3) scale effectively to diverse object categories.

#### 4.1.2. Sequence-based generation framework

Our framework transforms the skeleton generation task into a sequence modeling problem through four key components: skeleton tokenization, sequence ordering, shape conditioning, and auto-regressive generation.

**Skeleton tokenization.** We represent each skeleton  $\mathcal{S}$  as a sequence of bones, where each bone is defined by its two connecting joints (6 coordinates in total). To ensure consistent and discrete representation, we employ a carefully designed tokenization process. We first scale and translate the input mesh and corresponding skeleton to a unit cube  $[-0.5, 0.5]^3$ , ensuring their spatial alignment. Subsequently, we map the normalized joint coordinates to a discrete  $128^3$  space, leading to a sequence length of  $6b$  for  $b$  bones. As such, the discretized coordinates are converted into tokens, which serve as input to the auto-regressive transformer. Unlike MeshGPT [30], we omit the VQ-VAE compression step based on our dataset analysis. Specifically, in Articulation-XL, most of the models have fewer




Figure 5. **Spatial sequence ordering versus hierarchical sequence ordering.** The numbers indicate the bone ordering indices.

than 100 bones (i.e., 600 tokens). Given these relatively short sequence lengths, using VQ-VAE compression would potentially introduce artifacts without significant benefits in computational efficiency.

**Sequence ordering.** In this work, we investigate two distinct ordering strategies. Our first approach follows the sequence ordering strategy from recent 3D mesh generation methods [28, 30]. In this approach, joints are initially sorted in ascending z-y-x order (with z representing the vertical axis), and the corresponding joint indices in the bones are updated accordingly. Bones are then ordered first by their lower joint index and subsequently by the higher one. Ad-

ditionally, for each bone, the joint indices are cyclically permuted so that the lower index appears first. we refer to this ordering as **spatial sequence ordering** in this paper. However, this ordering strategy disrupts the parent-child relationships among bones and does not facilitate identifying the root joint. Consequently, additional processing is required to build the skeleton’s hierarchy.

To overcome these limitations, we propose an alternative approach termed **hierarchical sequence ordering**<sup>2</sup>, which leverages the intrinsic hierarchical structure of the skeleton by processing bones layer by layer. After sorting joints in ascending z-y-x order and updating their indices in bones, we first order the bones directly connected to the root joint. When the root has several child joints, we begin with the bone linked to the child joint having the smallest index and then proceed in ascending order. For subsequent layers, bones are grouped by their immediate parent, and within each group, they are arranged in ascending order based on the child joint index. Additionally, among groups in the same layer, the group corresponding to the smallest parent joint index is processed first, followed by those with larger indices.

**Shape-conditioned generation.** Following the conventions in [8, 9], we utilize point clouds as the shape condition by sampling 8,192 points from the input mesh  $\mathcal{M}$ . We then process this point cloud through a pre-trained shape encoder [52], which transforms the raw 3D geometry into a fixed-length feature sequence suitable for transformer processing. This encoded sequence is then appended to the start of the transformer’s input skeleton sequence for auto-regressive generation. Additionally, for each sequence, we insert a  $\langle \text{bos} \rangle$  token after the shape latent tokens to signify the beginning of the skeleton tokens. Similarly, a  $\langle \text{eos} \rangle$  token is added following the skeleton tokens to denote the end of the skeleton sequence.

**Auto-regressive learning.** For skeleton generation, we employ a decoder-only transformer architecture, specifically the OPT-350M model [51], which has demonstrated strong capabilities in sequence modeling tasks. During training, we provide the ground truth sequences and utilize cross-entropy loss for next-token prediction to supervise the model:

$$\mathcal{L}_{pred} = \text{CE}(\mathbf{T}, \hat{\mathbf{T}}), \quad (2)$$

where  $\mathbf{T}$  represents the one-hot encoded ground truth token sequence, and  $\hat{\mathbf{T}}$  denotes the predicted sequence.

At inference time, the generation process begins with only the shape tokens as input, and the model sequentially generates each skeleton token, ending when the  $\langle \text{eos} \rangle$  token is produced. The resulting token sequence is then detokenized to recover the final skeleton coordinates and connectivity structure.

<sup>2</sup>Hierarchical ordering is an extension of our under review version.

## 4.2. Skinning weight prediction

The second stage focuses on predicting skinning weights, which controls how the mesh deforms with skeleton movements. In this work, we represent skinning weights as an  $n$ -dimensional function defined on mesh surfaces, which are continuous, high-dimensional, and exhibit significant variation across different skeletal structures. To address these complexities, we employ a functional diffusion framework for accurate skinning weight prediction.

### 4.2.1. Preliminary: Functional diffusion

Functional diffusion [46] extends classical diffusion models to operate directly on functions, making it particularly suitable for our task. Consider a function  $f_0$  mapping from domain  $\mathcal{X}$  to range  $\mathcal{Y}$ :

$$f_0 : \mathcal{X} \rightarrow \mathcal{Y}. \quad (3)$$

The diffusion process gradually adds functional noise  $g$  (mapping the same domain to range) to the original function:

$$f_t(x) = \alpha_t \cdot f_0(x) + \sigma_t \cdot g(x), \quad t \in [0, 1] \quad (4)$$

where  $\alpha_t$  and  $\sigma_t$  control the noise schedule. The goal is to train a denoiser  $D$  that recovers the original function:

$$D_\theta[f_t, t](x) \approx f_0(x). \quad (5)$$

This formulation naturally aligns with our task requirements. By treating skinning weights as continuous functions over the mesh surface, we can capture smoothly transitioning weights between vertices. Additionally, the framework’s flexibility allows it to adapt to diverse mesh topologies and skeletal structures.

### 4.2.2. Skinning weight prediction

Building upon the functional diffusion framework, we formulate skinning weight prediction as learning a mapping  $f : \mathbb{R}^3 \rightarrow \mathbb{R}^n$  from 3D points to their corresponding weights. Specifically, the input to our model consists of 3D points  $\mathcal{P} \in \mathbb{R}^{v \times 3}$  sampled from the surface of the mesh. The output is an  $n$ -dimensional skinning weight matrix  $\mathcal{W} \in \mathbb{R}^{v \times n}$ . Here, the ground truth skinning weights of sampled points for training are copied from their nearest vertices and will also be copied back when inference.  $n$  denotes the maximum number of joints in the dataset.

To enhance prediction accuracy, we introduce two key components. First, we condition the generation on both joint coordinates and global shape features extracted by a pre-trained encoder [52]. Second, we leverage volumetric geodesic priors calculated from [13]. Specifically, we compute the volumetric geodesic priors from each mesh vertex to each joint. We then assign these priors to sampled points based on their nearest vertices and normalize them to

match the range of skinning weights, forming a volumetric geodesic matrix  $\mathcal{G} \in \mathbb{R}^{v \times n}$ . Our model learns to predict the residual between the actual skinning weights and this geometric prior, i.e.,  $f : \mathcal{P} \rightarrow (\mathcal{W} - \mathcal{G})$ , enabling more stable predictions.

Following [46], we optimize our model using  $x_0$ -prediction with the objective:

$$\mathcal{L}_{denoise} = \|D_\theta(\{x, f_t(x)\}, t) - f_0(x)\|_2^2, \quad x \in \mathcal{P}. \quad (6)$$

We employ the Denoising Diffusion Probabilistic Model (DDPM) [16] as our scheduler. In practice, we normalize the skinning weights and volumetric geodesic priors to the range  $[-1, 1]$  before adding noise. We will conduct ablation studies on this design in Section 5.4.2.

## 5. Experiments

### 5.1. Implementation details

**Datasets.** We evaluate our method on two datasets: our proposed Articulation-XL and ModelsResource [38, 43]. Articulation-XL contains 33k samples, with 31.4k for training and 1.6k for testing. ModelsResource is a smaller dataset, containing 2,163 training and 270 testing samples. The number of joints for each object varies from 3 to 48, with an average of 25.0 joints. While the data in ModelsResource maintains a consistent upright and front-facing orientation, the 3D models in Articulation-XL exhibit varying orientations. We have verified that there are no duplications between Articulation-XL and ModelsResource.

**Training details.** Our training process consists of two stages. For skeleton generation, we train the auto-regressive transformer on 8 NVIDIA A100 GPUs for approximately two days. For skinning weight prediction, models are trained on the same hardware configuration for about one day. To enhance model robustness, we apply data augmentation including scaling, shifting, and rotation transformations. For more details, please refer to the appendix.

### 5.2. Skeleton generation results

**Metrics.** We adopt three standard metrics following [43] to evaluate skeleton quality: CD-J2J, CD-J2B, and CD-B2B. These Chamfer Distance-based metrics measure the spatial alignment between generated and ground truth skeletons by computing distances between joints-to-joints, joints-to-bones, and bones-to-bones respectively. Lower values indicate better skeleton quality.

**Baselines.** We compare our method against two representative approaches: Pinocchio [3], a traditional template-fitting method, and RigNet [43], a learning-based method using graph convolutions. All methods are evaluated on the Articulation-XL and ModelsResource datasets.

Table 1. **Quantitative comparison on skeleton generation.** We compare different methods using CD-J2J, CD-J2B, and CD-B2B as evaluation metrics on both Articulation-XL (Arti-XL) and ModelsResource (Modelres.). Lower values indicate better performance. The metrics are in units of  $10^{-2}$ . Here, \* denotes models trained on Articulation-XL and tested on ModelsResource.

	Dataset	CD-J2J	CD-J2B	CD-B2B
RigNet*		7.132	5.486	4.640
Pinocchio		6.852	4.824	4.089
Ours-hier*		4.451	3.454	2.998
RigNet	<i>ModelsRes.</i>	4.143	2.961	2.675
Ours-spatial*		4.103	3.101	2.672
Ours-hier		3.654	2.775	2.412
Ours-spatial		<b>3.343</b>	<b>2.455</b>	<b>2.140</b>
Pinocchio		8.360	6.677	5.689
RigNet	<i>Arti-XL</i>	7.478	5.892	4.932
Ours-hier		3.025	2.408	2.083
Ours-spatial		<b>2.586</b>	<b>1.959</b>	<b>1.661</b>

**Comparison results.** Qualitative comparisons are presented in Figure 6, where we compare different methods across various object categories. Pinocchio struggles with objects that differ from its predefined templates, especially obvious in non-humanoid objects (as shown in the 2nd row and the 3rd row on the right). RigNet demonstrates improved performance when tested on ModelsResource, where the data maintains a consistent upright and front-facing orientation. However, it still struggles with complex topologies (as illustrated in the 1st and 2nd rows on the left). Furthermore, RigNet performs worse on Articulation-XL, where the data exhibit varying orientations. In contrast, our method generates high-quality skeletons that closely match artist-created references across diverse object categories.

The quantitative results are shown in Table 1. Our method consistently outperforms baselines across all metrics on both datasets. Additionally, we compare our method using both spatial and hierarchical ordering strategies. The spatial ordering consistently achieves better performance, likely because the hierarchical ordering requires the model to allocate part of its capacity to learning the skeleton’s hierarchy and identifying the root joint. Results obtained using spatial ordering are well-suited for applications such as skeleton-driven pose transfer [47], whereas those derived from hierarchical ordering are more readily integrated with 3D models for animation.

**Generalization analysis.** To evaluate the generalization capability, we perform cross-dataset evaluation by training RigNet and our MagicArticulate on Articulation-XL and testing on ModelsResource. As shown in Table 1 (marked with \*), our method maintains competitive performance compared to RigNet trained directly on ModelsResource, while RigNet’s performance degrades significantly when




Figure 6. **Comparison of skeleton creation results on ModelsResource (left) and Articulation-XL (right).** Our generated skeletons more closely resemble the artist-created references, while RigNet and Pinocchio struggle to handle various object categories.

tested on unseen data distributions, performing even worse than the template-based method Pinocchio.

To further assess real-world applicability, we evaluate all methods on AI-generated 3D meshes from Tripo 2.0 [1] (Figure 7). Our method successfully generates plausible skeletons for diverse object categories, while RigNet fails to produce valid results despite being trained on our large-scale dataset. Notably, even Pinocchio’s template-based approach struggles to generate accurate skeletons for basic categories like humans and quadrupeds, highlighting the advantage of our method in handling novel object structures.

### 5.3. Skinning weight prediction results

**Metrics.** We evaluate skinning weight quality using three metrics: precision, recall, and L1-norm error. Precision and recall measure the accuracy of identifying significant joint influences (defined as weights larger than  $1e - 4$  following [43]), while the L1-norm error computes the average difference between predicted and ground truth skinning weights across all vertices. We will also report the deformation error in appendix.

**Baselines.** We compare our method against Geodesic Voxel Binding (GVB) [13], a geometric-based method available in Autodesk Maya [19] and RigNet [43]. When trained on Articulation-XL, we filter out a subset containing 28k training and 1.2k testing samples, excluding data with more than 55 joints (which constitute a small fraction of both real-world cases and Articulation-XL).

**Comparison results.** Qualitative comparisons in Figure 8 visualize the predicted skinning weights and their L1 error




Figure 7. **Skeleton creation results on 3D generated meshes.** Our method has a better generalization performance than both RigNet [43] and Pinocchio [3] across difference object categories. The 3D models are generated by Tripo 2.0 [1].

maps against artist-created references. Our method predicts more accurate skinning weights with significantly lower errors across diverse object categories. In contrast, both GVB and RigNet show larger deviations, particularly in regions around joint boundaries.

The quantitative results are shown in Table 2, which support qualitative observations, demonstrating that our method consistently outperforms baselines across most metrics on both datasets.



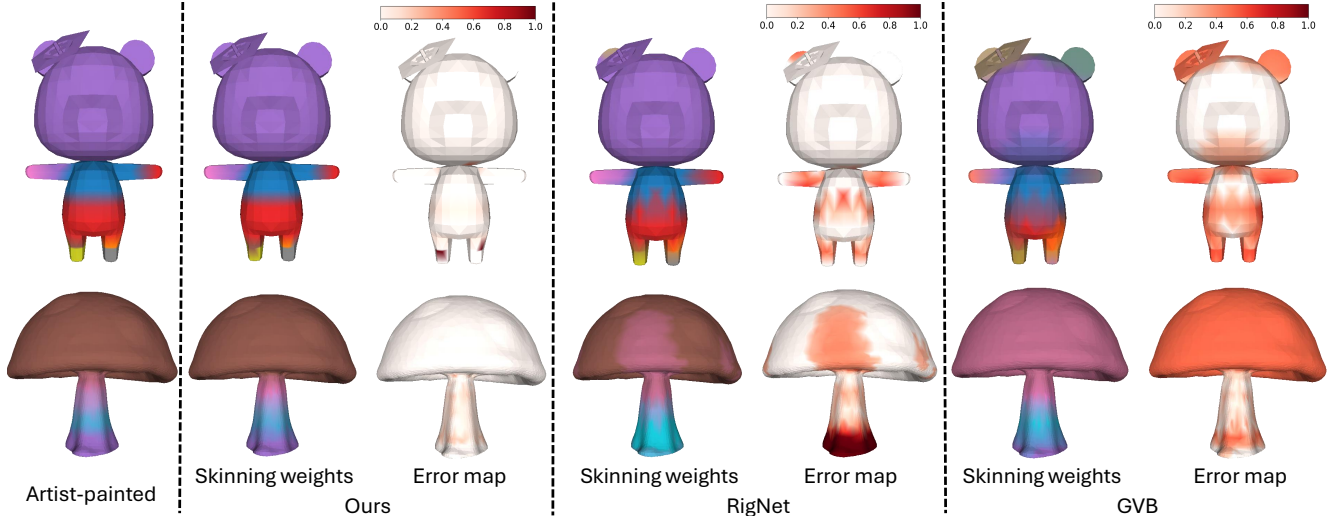


Figure 8. Comparisons with previous methods for skinning weight prediction on ModelsResource (top) and Articulation-XL (bottom). We visualize skinning weights and L1 error maps. For more results, please refer to the supplementary materials.

Table 2. **Quantitative comparison on skinning weight prediction.** We compare our method with GVB and RigNet. For Precision and Recall, larger values indicate better performance. For average L1-norm error, smaller values are preferred.

	Dataset	Precision	Recall	avg L1
GVB	<i>ModelsResource</i>	69.3%	79.2%	0.687
RigNet		77.1%	<b>83.5%</b>	0.464
Ours		<b>82.1%</b>	81.6%	<b>0.398</b>
GVB	<i>Articulation-XL</i>	75.7%	68.3%	0.724
RigNet		72.4%	71.1%	0.698
Ours		<b>80.7%</b>	<b>77.2%</b>	<b>0.337</b>

Table 3. **Ablation studies for skeleton generation.**

	CD-J2J	CD-J2B	CD-B2B
w/o data filtering	2.982	2.327	2.015
4,096 points	2.635	2.024	1.727
12,288 points	2.685	2.048	1.760
Ours (8,192)	<b>2.586</b>	<b>1.959</b>	<b>1.661</b>

## 5.4. Ablation studies

### 5.4.1. Ablation studies on skeleton generation

We conduct ablation studies to assess the impact of VLM-based data filtering and the number of sampled mesh points on skeleton generation. The results, presented in Table 3, show notable performance degradation without data filtering, highlighting the importance of high-quality training data. We also vary the number of sampled points as input to the pre-trained shape encoder [52]. As shown in Table 3, sampling 8,192 points yields superior performance.

Table 4. **Ablation studies on skinning weight prediction.**

	Precision	Recall	avg L1
w/o geodesic dist.	81.5%	77.7%	0.444
w/o weights norm	82.0%	77.9%	0.436
w/o shape features	81.4%	81.3%	0.412
Ours	<b>82.1%</b>	<b>81.6%</b>	<b>0.398</b>

### 5.4.2. Ablation studies on skinning weight prediction

We conduct ablation studies on three critical components of our skinning weight prediction framework. The quantitative results on ModelsResource are shown in Table 4. First, removing the volumetric geodesic distance initialization reduces precision by 0.6% and recall by 3.9%, demonstrating its crucial role in guiding accurate weight distribution. Second, eliminating our normalization strategy, which scales both skinning weights and geodesic distances to  $[-1, 1]$  before noise addition, leads to an 8.7% increase in L1 error. Finally, excluding global shape features from the pre-trained encoder [52] results in less accurate predictions. All these results validate our design choices and show that each component contributes notably to the final performance.

## 6. Conclusion

In this work, we present MagicArticulate to convert static 3D models into articulation-ready assets that support realistic animation. We first introduce a large-scale dataset Articulation-XL with high-quality articulation annotations, which is carefully curated from Objaverse-XL. Built upon this dataset, we develop a novel two-stage pipeline that first generates skeletons through auto-regressive sequence modeling, naturally handling varying numbers of bones or joints

within skeletons across different 3D models. Then we predict skinning weights in a functional diffusion process that incorporates volumetric geodesic distance priors between vertices and joints. Extensive experiments demonstrate our method’s superior performance and generalization ability across diverse object categories.

## **Acknowledgements**

This research is supported by the MoE AcRF Tier 2 grant (MOE-T2EP20223-0001).

# MagicArticulate: Make Your 3D Models Articulation-Ready

## Supplementary Material

### Overview

In this supplementary material, we provide additional details and experimental results for the main paper, including:

- Further details of MagicArticulate (Section 7) and Articulation-XL (Section 9);
- Additional experimental results on skeleton generation and skinning weight prediction (Section 8);
- A discussion of the limitations of our work and future works (Section 10).

## 7. More details of MagicArticulate

### 7.1. Implementation details

**Skeleton generation.** Our skeleton generation pipeline utilizes a pre-trained shape encoder [52] to process input meshes. For each mesh, we sample 8,192 points which are encoded into 257 shape tokens following MeshAnything [8]. To ensure consistent point cloud sampling across different data sources, we first extract the signed distance function from input mesh using [40], followed by generating a coarse mesh via Marching Cubes [26]. We then sample point clouds and their corresponding normals from this coarse mesh.

For training on Articulation-XL, we use 8 NVIDIA A100 GPUs for approximately two days with a batch size of 64 per GPU, resulting in an effective batch size of 512. When training on ModelsResource, we utilize 4 NVIDIA A100 GPUs for about 9 hours with a batch size of 32 per GPU, which yields an effective batch size of 128. During inference, the model generates skeleton tokens autoregressively from shape tokens until reaching the  $\langle \text{eos} \rangle$  token, followed by detokenization to recover the final skeleton coordinates in  $[-0.5, 0.5]$  range.

**Skinning weight prediction.** Our functional diffusion model employs the Denoising Diffusion Probabilistic Model (DDPM) with 1,000 timesteps and a linear beta schedule. During training, we condition the model on ground truth skeletons and supervise it with corresponding ground truth skinning weights. We add noise to the skinning weight function (the process is illustrated in Figure S10) and then feed the noised skinning weights into our denoising network (Figure S9). Following [46], our network architecture processes the noised set  $\{(x, f_t(x)) \mid x \in \mathcal{P}\}$  by splitting it into smaller subsets and handling them through multiple cross-attention stages. The time embedding at timestep  $t$  is incorporated into each self-attention layer via adaptive layer normalization. For visual clarity, Figure S9 shows only one processing stage.

We train the model on Articulation-XL using 8 NVIDIA A100 GPUs for approximately one day, with a batch size of 16 per GPU (effective batch size 128). Training on ModelsResource uses the same configuration for about 4 hours. During inference, we perform 25 denoising steps to generate predictions  $\mathcal{W} \in \mathbb{R}^{v \times n}$  in the range  $[-1, 1]$ . These results are then normalized to  $[0, 1]$ , ensuring that each row of the skinning weight matrix sums to 1. To handle varying joint counts across different models, we employ a valid joint mask during both training and testing, with a maximum joint count of 55 as discussed in the main paper (Sections 4.2 and 5.3).

### 7.2. Experimental details

For baseline comparisons, we use the implementations of RigNet [43] and Pinocchio [3] from the GitHub repositories<sup>3</sup>. The Geodesic Voxel Binding (GVB) [13] comparison is conducted using the implementation in Autodesk Maya [19]. When training RigNet on our Articulation-XL, we strictly follow the authors’ data processing pipeline and six-stage training strategy as specified in their official implementation.

### 7.3. Animation

Many recent works have explored 3D animation, including skeleton-free pose transfer [23, 31, 32], skeleton-driven pose transfer [47], and physics-driven animation [15]. In this paper, we propose a method that enables automatic articulation generation for any input 3D model, whether artist-created or AI-generated. The pipeline first generates a skeleton for the input model, then predicts skinning weights conditioned on both the model geometry and the generated skeleton. The resulting articulated model can be exported in standard formats (e.g., FBX, GLB), making it directly compatible with popular animation software such as Blender [4] and Autodesk Maya [19].

## 8. Additional experimental results

### 8.1. More results of skeleton generation

We provide additional qualitative comparisons among MagicArticulate, RigNet [43], and Pinocchio [3] for skeleton generation.

**More qualitative results on out-of-domain data.** We evaluate our method’s generalization capability on diverse out-of-domain data sources: AI-generated meshes from

<sup>3</sup><https://github.com/zhan-xu/RigNet>, <https://github.com/haozl9/Automatic-Rigging>

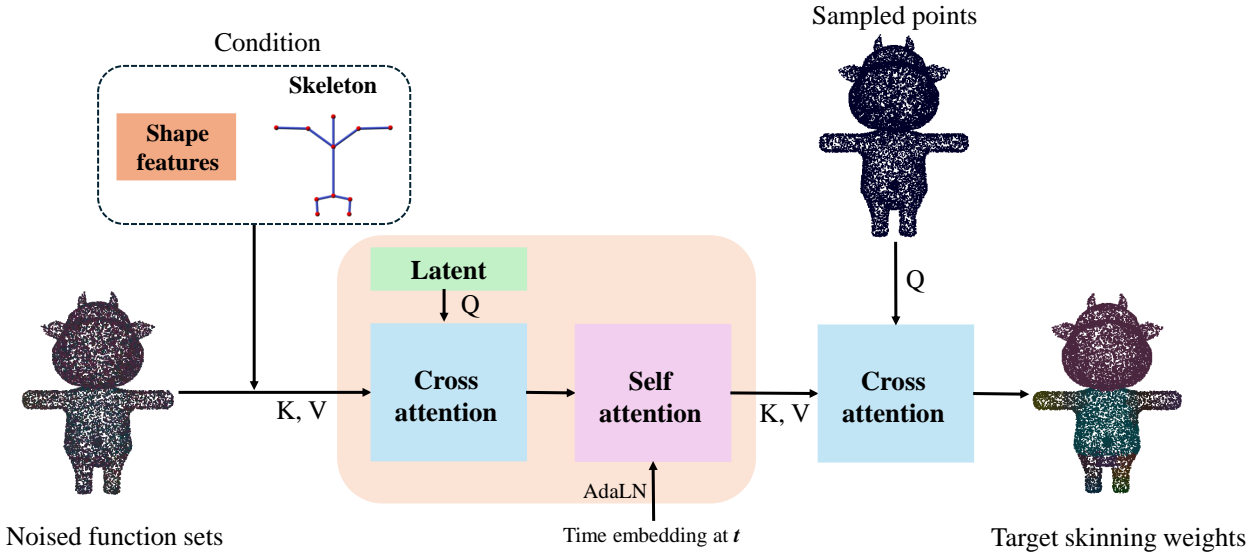


Figure S9. **Overview of the function diffusion architecture for skinning weight prediction.** Given a set of noised skinning weight functions  $\{(x, f_t(x)) \mid x \in \mathcal{P}\}$ , conditioned on skeleton and shape features from [52], we denoise the skinning weight functions to approximate the target weights.

Table S5. **Quantitative comparison on skinning weight prediction.** We compare our method with GVB and RigNet. For Precision and Recall, larger values indicate better performance. For average L1-norm error and average distance error, smaller values are preferred.

	Dataset	Precision	Recall	avg L1	avg Dist.
GVB	<i>ModelsResource</i>	69.3%	79.2%	0.687	0.0067
RigNet		77.1%	<b>83.5%</b>	0.464	0.0054
Ours		<b>82.1%</b>	81.6%	<b>0.398</b>	<b>0.0039</b>
GVB	<i>Articulation-XL</i>	75.7%	68.3%	0.724	0.0095
RigNet		72.4%	71.1%	0.698	0.0091
Ours		<b>80.7%</b>	<b>77.2%</b>	<b>0.337</b>	<b>0.0050</b>

Tripo2.0 [1], unregistered 3D scans from FAUST [5], and video-based 3D reconstructions [34]. As shown in Figure S11, while existing methods struggle with generalization (RigNet fails across all cases, and Pinocchio shows misalignments even for human bodies, see skeleton results on the 3D scan), our method maintains robust performance across different data sources and categories. Notably, for human models, our method generates more detailed skeletal structures, including accurate hand skeletons, surpassing Pinocchio’s template-based results.

**More qualitative results on Articulation-XL and ModelsResource.** We provide additional qualitative results on both Articulation-XL and ModelsResource datasets. As illustrated in Figure S12, our method consistently generates high-quality skeletons that accurately match artist-created references across diverse object categories.

**Robustness to various mesh orientations.** To further val-

idate our model’s robustness to various orientations, we include mesh rotations at multiple angles in Figure S13. These examples show that our approach remains largely rotation-stable. While minor skeleton variations may occur, all generated results maintain anatomically valid and suitable for rigging purposes.

## 8.2. More results of skinning weight prediction

**Quantitative results with deformation error.** Beyond the precision, recall, and L1-norm metrics reported in the main paper, we evaluate the practical effectiveness of predicted skinning weights through deformation error analysis. This metric computes the average Euclidean distance between vertices deformed using predicted weights and ground truth weights across 10 random poses. The comprehensive results, shown in Table S5, demonstrate our method’s superior performance across most metrics on both datasets. We also include deformation error analysis in our ablation stud-

Table S6. Ablation studies on ModelsResource for skinning weight prediction.

	Precision	Recall	avg L1	avg Dist.
w/o geodesic dist.	81.5%	77.7%	0.444	0.0046
w/o weights norm	82.0%	77.9%	0.436	0.0045
w/o shape features	81.4%	81.3%	0.412	0.0042
Ours	<b>82.1%</b>	<b>81.6%</b>	<b>0.398</b>	<b>0.0039</b>

Table S7. Object counts for each category in the Articulation-XL dataset.

Category	# Objects	Category	# Objects	Category	# Objects
character	16020	miscellaneous	584	architecture	132
anthropomorphic	13393	scanned data	546	planet	49
animal	4760	plant	382	paper	46
mythical creature	4734	accessories	293	musical instrument	25
toy	1360	vehicle	283	sporting goods	21
weapon	1257	sculpture	276	armor	13
anatomy	1227	household items	274	robot	4
clothing	595	food	206		

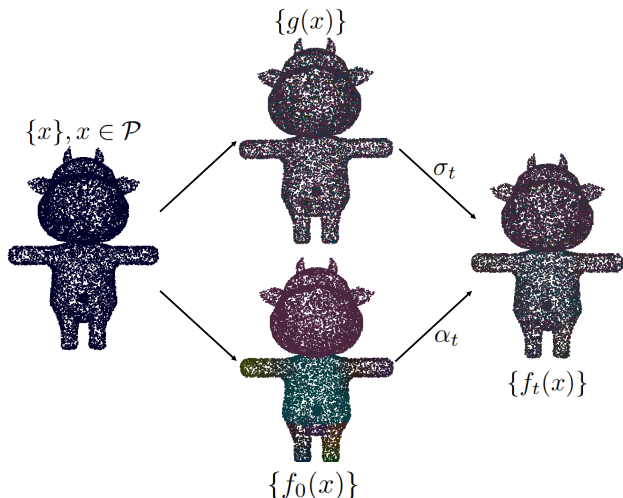


Figure S10. **Process of adding noise to the skinning weight function.** Given  $x \in \mathcal{P}$  and the original skinning weight function  $f_0(x)$ , we add the noise function  $g(x)$  to obtain the noised function  $f_t(x)$ .

ies (Table S6), further validating the effectiveness of our design choices.

**More qualitative results.** We present additional qualitative comparisons between MagicArticulate, RigNet [43], and Geodesic Voxel Binding (GVB) [13] for skinning weight prediction. Figure S14 shows both the predicted skinning weights and their L1 error maps compared to artist-created references, demonstrating our method’s superior accuracy across diverse object categories.

## 9. More details of Articulation-XL

### 9.1. Data Curation

Our dataset curation process filters out duplicates, objects with extreme joint/bone counts, and multi-component objects. A detailed category-wise object distribution is provided in Table S7.

### 9.2. Quality assessment

We employ GPT-4o [29] for quality assessment of skeleton annotations. For each model, we generate four-view renders using Pyrender<sup>4</sup> showing both the 3D model and its skeleton (Figure S17). These renders are evaluated using specific quality criteria detailed in Figure S15.

### 9.3. Category annotation

For the Visual-Language Model (VLM)-based category labeling, we render each 3D model along with its normal maps from four viewpoints using Blender [4] (see example in Figure S18). We then utilize GPT-4o [29] to classify the categories of the 3D models based on specific instructions, as outlined in Figure S16.

## 10. Limitations and future work

Despite its strong performance, our method has several notable limitations. First, our approach struggles with coarse mesh inputs, often producing inaccurate skeletons as shown in Figure S19. While we employ preprocessing techniques

<sup>4</sup><https://github.com/mmatl/pyrender>

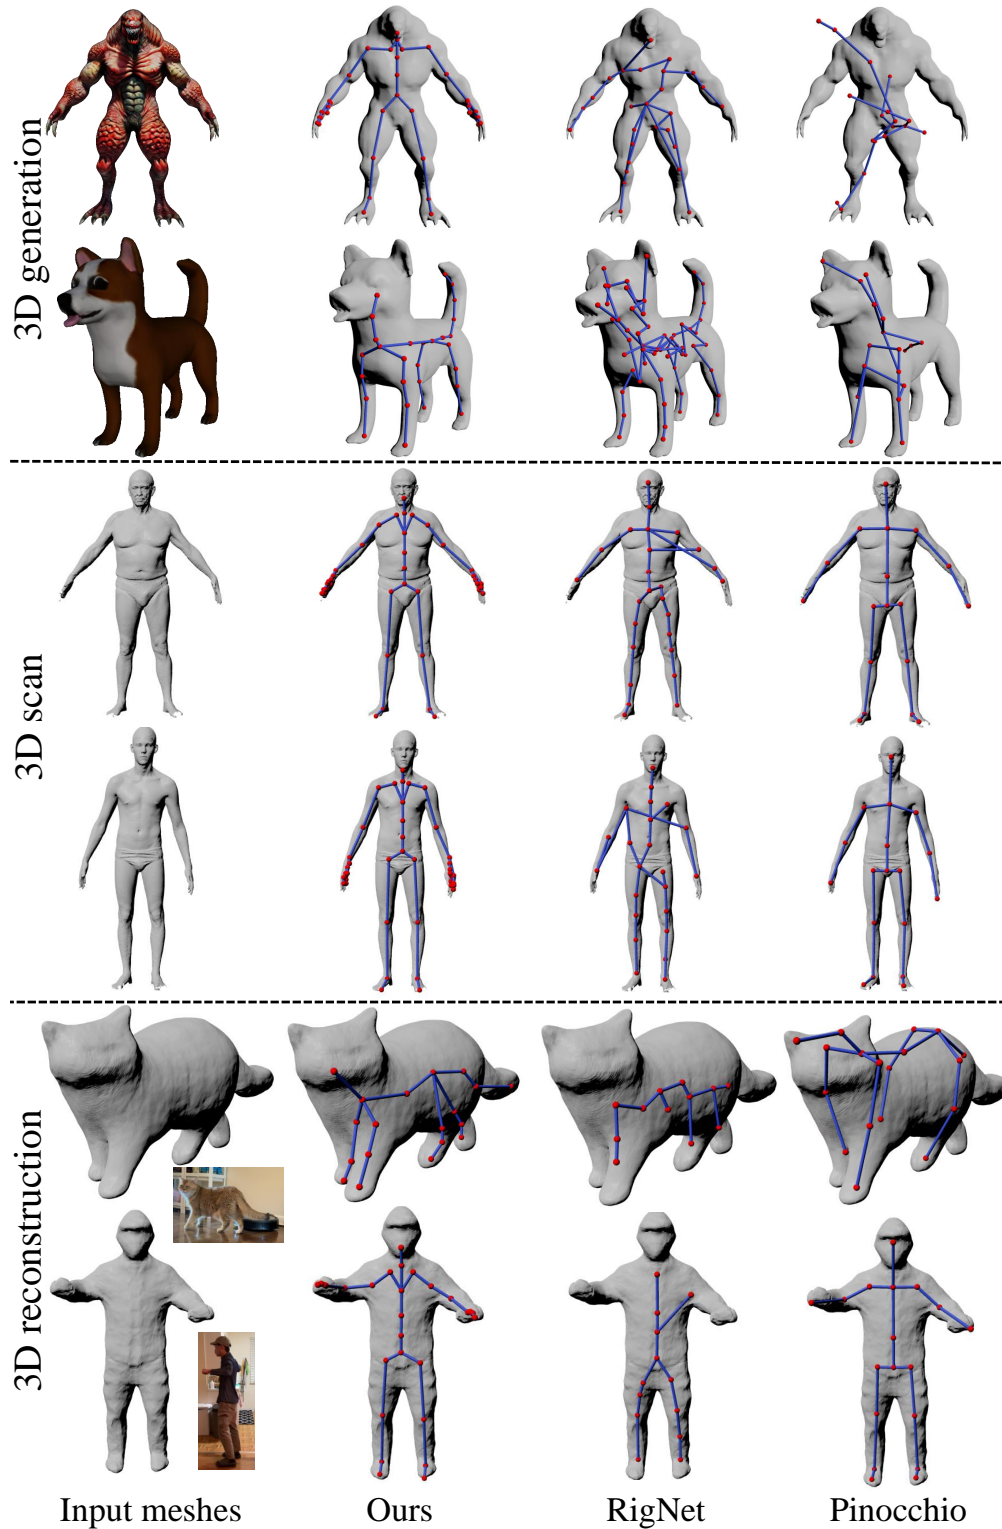


Figure S11. Comparison of skeleton generation methods on out-of-domain data. The input meshes are from 3D generation, 3D scan, and 3D reconstruction.

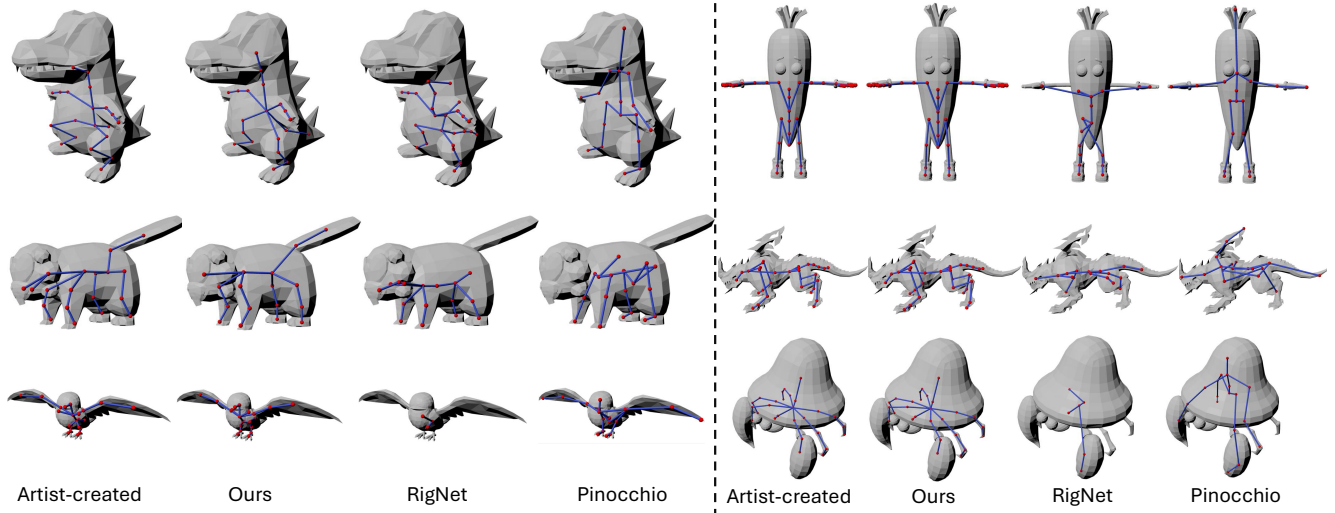


Figure S12. **Comparison of skeleton generation methods on ModelsResource (left) and Articulation-XL (right).** Our results more closely resemble the artist-created references, while RigNet and Pinocchio struggle to handle various object categories.




Figure S13. **Skeleton results on 3D models with different orientations.** Although minor differences may appear in the generated skeletons, all results maintain anatomically valid and suitable for rigging purposes.

to handle inputs from different sources, the significant domain gap between training data and coarse meshes remains challenging. Potential solutions include incorporating mesh quality augmentation during training to enhance robustness.

A second limitation lies in our dataset composition. Although Articulation-XL is large in scale, it lacks sufficient coverage of common articulated objects like laptops, staplers, and scissors, which affects our model’s generalization to these categories.

Future work will address these limitations by: 1) Developing more robust preprocessing and training strategies for handling varying mesh qualities; 2) Expanding dataset coverage to include a broader range of everyday articulated objects; 3) Exploring techniques to better bridge the domain gap between different data sources.




Figure S14. Comparison of skinning weight prediction methods on ModelsResource (first three rows) and Articulation-XL (last three rows). We visualize the predicted skinning weights alongside their corresponding L1 error maps.



## # Task Description:

Our task is to evaluate the quality of the skeleton within a 3D object using an image that contains four sub-images rendered from both the 3D mesh and its skeleton together. The four sub-images are arranged in a 2x2 grid format. These sub-images are rendered from four distinct views of the 3D object. In the images, bones are represented by blue cylinders, and joints are represented by red spheres.

## # Instructions:

1. Carefully examine the provided image, which contains four different sub-images of the 3D object and its skeleton.
2. Provide a detailed critique of the skeleton within the 3D object based on the following criteria:
  - Joint and Bone Position Relative to Mesh:**  
Identify joints or bones that are excessively protruding out of the mesh. A high-quality skeleton should have both joints and bones confined within the mesh boundary. The more joints or bones that protrude, the lower the quality. Since the joint spheres and bone cylinders have a radius, you should consider their centers without accounting for their shapes.
  - Anatomical Accuracy:**
    - Correct Joint Placement: Joints should mimic the anatomical placement. For example, shoulder joints should align with the mesh's shoulder region.
    - Natural Poses: The skeleton should maintain natural and plausible poses. Any significant deviations might indicate errors in skeletal rigging.
3. Based on your analysis, provide a rating for each criterion using the following three options. Ensure that the decision aligns with your detailed critique.
  - Poor: Poor quality, significant issues or errors that severely impact the skeleton's usability or appearance.
  - Average: Average quality, some issues or errors that moderately impact the skeleton's usability or appearance.
  - Good: Good quality, no noticeable issues/errors or minor issues/errors that slightly impact the skeleton's usability or appearance.
4. Provide an overall rating for the skeleton within the 3D object based on the ratings for each criterion, with more weight given to Joint and Bone Position Relative to Mesh. If the skeleton has good performance on this part, it can be rated as good even if it has minor issues in other criteria. Ensure that the final rating is consistent with the individual scores and the overall critique. If the provided reasons for each criterion indicate a poor quality, the rating should reflect that. Ensure that the ratings and reasons are aligned.
5. Make sure to provide your evaluation following the example output format below.

## # Example Output:

```Critique:

### **Joint and Bone Position Relative to Mesh:**

- External Joints and Bones: Good. Most joints and bones are well confined within the mesh's boundary, with minimal protrusion observed.

### **Anatomical Accuracy:**

- Correct Joint Placement: Good. Joints are well-aligned with the mesh's primary structural points.  
- Natural Poses: Good. The skeleton maintains a natural and plausible pose, consistent with the object's intended design.

**Final Rating:** Good.

```

```Critique:

### **Joint and Bone Position Relative to Mesh:**

- External Joints and Bones: Poor. Several joints and bones, particularly in the lower section, are protruding significantly out of the mesh boundary.

### **Anatomical Accuracy:**

- Correct Joint Placement: Average. While most joints are correctly placed, some joints are noticeably out of alignment.  
- Natural Poses: Poor. The skeleton maintains an unnatural pose, with significant deviations observed.

**Final Rating:** Poor.

```

```Critique:

### **Joint and Bone Position Relative to Mesh:**

- External Joints and Bones: Average. Most joints and bones are positioned correctly within the mesh, but some joints, especially in the middle section, slightly protrude, affecting the alignment.

### **Anatomical Accuracy:**

- Correct Joint Placement: Average. Joints are generally placed close to the intended structural points, but there are minor misalignments that mildly impact the skeleton's accuracy.  
- Natural Poses: Average. The skeleton maintains a generally natural pose, but some joints appear slightly off, giving a somewhat unnatural appearance.

**Final Rating:** Average.

```

Figure S15. Input instructions to VLM for data filtering.

## # Task Description:

Our task is to evaluate the categories of a 3D object based on an image that contains eight sub-images rendered from this 3D object. The eight sub-images are concatenated in a 2x4 format. The top four images are its RGB images rendered from four different angles, and the bottom four are its normal maps rendered from four different angles.

## # Instructions:

1. Carefully examine the provided image, which contains eight different sub-images of the 3D object.

2. Classify the 3D object into one or more of the following categories:

- Character:** Human or humanoid objects.
- Animal:** Any kind of animal, including pets, wild animals, and mythical creatures.
- Furniture:** Items like chairs, tables, beds, etc.
- Electronic Device:** Specifically defined as 3C devices, such as phones, Apple Watch, TVs, and other consumer electronics.
- Mythical Creature:** Strange creatures, monsters, elves, etc., including dragons and other legendary creatures.
- Anatomy:** Parts of the human body or used for medical purposes.
- Tool:** Instruments used to perform tasks, like hammers, screwdrivers, etc.
- Planet:** Celestial bodies such as planets, moons, and other astronomical objects.
- Musical Instrument:** Objects designed to produce music, such as guitars, pianos, drums, etc.
- Sculpture:** Artistic objects created by carving, modeling, or assembling materials, often for decorative or artistic purposes.
- Jewelry:** Decorative items worn for personal adornment, such as rings, necklaces, bracelets, etc.
- Accessory:** Includes fashion accessories (can overlap with jewelry) and parts or recognizable parts of a larger object.
- Paper:** Appears to be a flat texture with no thickness.
- Anthropomorphic:** Objects with human-like features, even if they are essentially other types of items. For example, Donald Duck, although essentially a duck, can be marked as anthropomorphic due to its human-like features. Anything with hands and feet counts as anthropomorphic. A bell pepper with a smiley face but no hands or feet. A soda can with human-like features.
- Toy:** Items designed for play, like dolls, action figures, etc.
- Clothing:** Wearable items like shirts, pants, shoes, etc.
- Food:** Edible items like fruits, vegetables, cooked dishes, etc.
- Scanned Data:** Objects created from 3D scans. If the object appears to be a scanned model, classify it as scanned data and rate it accordingly.
- Architecture:** Buildings and other structures.
- Vehicle:** Cars, bikes, planes, boats, etc.
- Plant:** Trees, flowers, shrubs, etc.
- Weapon:** Items designed for combat, like swords, guns, archery, etc.
- Household Item:** Common items used in daily life, like utensils, appliances, etc.
- Sporting Goods:** Items used in sports and recreational activities, like balls, bats, etc.
- Miscellaneous:** Objects that do not fit into any of the current categories but have a clear meaning.

3. Make sure to provide your evaluation following the example output format below.

## # Example Output:

\*\*\*

Categories: furniture.

\*\*\*

\*\*\*

Categories: weapon.

\*\*\*

\*\*\*

Categories: scanned data, architecture.

\*\*\*

\*\*\*

Categories: Character.

\*\*\*

Figure S16. Input instructions to VLM for category labeling.

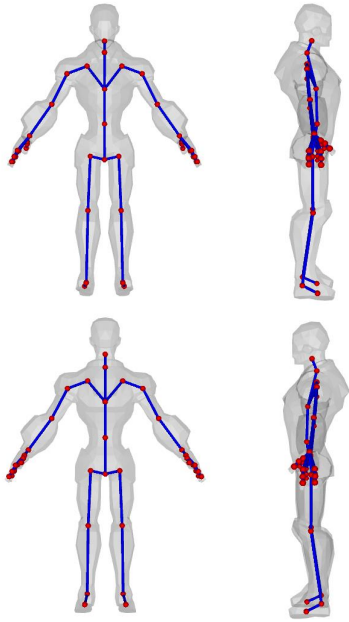


Figure S17. **Input rendered examples to VLM for data filtering.**




Figure S18. **Input rendered examples to VLM for category labeling.**




Figure S19. **Failure cases.** When input meshes possess very coarse surfaces (3D reconstruction results from [33]), our generated skeleton may exhibit inaccuracies, such as imperfect connections between the dog's trunk and legs.

## References

- [1] TriPo AI. Tripo 3d, 2023. [2](#), [8](#), [12](#)
- [2] Oscar Kin-Chung Au, Chiew-Lan Tai, Hung-Kuo Chu, Daniel Cohen-Or, and Tong-Yee Lee. Skeleton extraction by mesh contraction. *ACM transactions on graphics (TOG)*, 27(3):1–10, 2008. [1](#), [3](#)
- [3] Ilya Baran and Jovan Popović. Automatic rigging and animation of 3d characters. *ACM Transactions on graphics (TOG)*, 26(3):72–es, 2007. [1](#), [2](#), [3](#), [5](#), [7](#), [8](#), [11](#)
- [4] Blender Foundation. Blender - a 3d modelling and rendering software, 2024. Version 3.6. [11](#), [13](#)
- [5] Federica Bogo, Javier Romero, Matthew Loper, and Michael J Black. Faust: Dataset and evaluation for 3d mesh registration. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3794–3801, 2014. [12](#)
- [6] Junjie Cao, Andrea Tagliasacchi, Matt Olson, Hao Zhang, and Zhinxun Su. Point cloud skeletons via laplacian based contraction. In *2010 Shape Modeling International Conference*, pages 187–197. IEEE, 2010. [1](#), [3](#)
- [7] Sijin Chen, Xin Chen, Anqi Pang, Xianfang Zeng, Wei Cheng, Yijun Fu, Fukun Yin, Yanru Wang, Zhibin Wang, Chi Zhang, et al. Meshxl: Neural coordinate field for generative 3d foundation models. *arXiv preprint arXiv:2405.20853*, 2024. [3](#)
- [8] Yiwen Chen, Tong He, Di Huang, Weicai Ye, Sijin Chen, Jiaxiang Tang, Xin Chen, Zhongang Cai, Lei Yang, Gang Yu, et al. Meshanything: Artist-created mesh generation with autoregressive transformers. *arXiv preprint arXiv:2406.10163*, 2024. [3](#), [6](#), [11](#)
- [9] Yiwen Chen, Yikai Wang, Yihao Luo, Zhengyi Wang, Zilong Chen, Jun Zhu, Chi Zhang, and Guosheng Lin. Meshanything v2: Artist-created mesh generation with adjacent mesh tokenization. *arXiv preprint arXiv:2408.02555*, 2024. [3](#), [5](#), [6](#)
- [10] Edilson De Aguiar, Christian Theobalt, Sebastian Thrun, and Hans-Peter Seidel. Automatic conversion of mesh animations into skeleton-based animations. In *Computer Graphics Forum*, pages 389–397. Wiley Online Library, 2008. [2](#), [3](#)
- [11] Matt Deitke, Dustin Schwenk, Jordi Salvador, Luca Weihs, Oscar Michel, Eli VanderBilt, Ludwig Schmidt, Kiana Ehsani, Aniruddha Kembhavi, and Ali Farhadi. Objaverse: A universe of annotated 3d objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13142–13153, 2023. [1](#), [3](#)
- [12] Matt Deitke, Ruoshi Liu, Matthew Wallingford, Huong Ngo, Oscar Michel, Aditya Kusupati, Alan Fan, Christian Laforte, Vikram Voleti, Samir Yitzhak Gadre, et al. Objaverse-xl: A universe of 10m+ 3d objects. *Advances in Neural Information Processing Systems*, 36, 2024. [1](#), [3](#)
- [13] Olivier Dionne and Martin de Lasa. Geodesic voxel binding for production character meshes. In *Proceedings of the 12th ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, pages 173–180, 2013. [3](#), [6](#), [8](#), [11](#), [13](#)
- [14] Ana Dodik, Vincent Sitzmann, Justin Solomon, and Oded Stein. Robust biharmonic skinning using geometric fields. *arXiv preprint arXiv:2406.00238*, 2024. [3](#)
- [15] Zhoujie Fu, Jiacheng Wei, Wenhao Shen, Chaoyue Song, Xiaofeng Yang, Fayao Liu, Xulei Yang, and Guosheng Lin. Sync4d: Video guided controllable dynamics for physics-based 4d generation. *arXiv preprint arXiv:2405.16849*, 2024. [11](#)
- [16] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. [7](#)
- [17] Hui Huang, Shihao Wu, Daniel Cohen-Or, Minglun Gong, Hao Zhang, Guiqing Li, and Baoquan Chen. L1-medial skeleton of point cloud. *ACM Trans. Graph.*, 32(4):65–1, 2013. [1](#), [3](#)
- [18] Adobe Inc. Mixamo. [2](#), [3](#)
- [19] Autodesk Inc. Autodesk maya, 2024. Version 2024. [2](#), [8](#), [11](#)
- [20] Alec Jacobson, Ilya Baran, Jovan Popovic, and Olga Sorkine. Bounded biharmonic weights for real-time deformation. *ACM Trans. Graph.*, 30(4):78, 2011. [3](#)
- [21] Doug L James and Christopher D Twigg. Skinning mesh animations. *ACM Transactions on Graphics (TOG)*, 24(3):399–407, 2005. [2](#), [3](#)
- [22] Peizhuo Li, Kfir Aberman, Rana Hanocka, Libin Liu, Olga Sorkine-Hornung, and Baoquan Chen. Learning skeletal articulations with neural blend shapes. *ACM Transactions on Graphics (TOG)*, 40(4):1–15, 2021. [1](#), [2](#), [5](#)
- [23] Zhouyingcheng Liao, Jimei Yang, Jun Saito, Gerard Pons-Moll, and Yang Zhou. Skeleton-free pose transfer for stylized 3d characters. In *European Conference on Computer Vision*, pages 640–656. Springer, 2022. [3](#), [11](#)
- [24] Cheng Lin, Changjian Li, Yuan Liu, Nenglu Chen, Yi-King Choi, and Wenping Wang. Point2skeleton: Learning skeletal representations from point clouds. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4277–4286, 2021. [1](#), [3](#)
- [25] Lijuan Liu, Youyi Zheng, Di Tang, Yi Yuan, Changjie Fan, and Kun Zhou. Neuroskinning: Automatic skin binding for production characters with deep graph networks. *ACM Transactions on Graphics (ToG)*, 38(4):1–12, 2019. [3](#)
- [26] William E Lorensen and Harvey E Cline. Marching cubes: A high resolution 3d surface construction algorithm. In *Seminal graphics: pioneering efforts that shaped the field*, pages 347–353. 1998. [11](#)
- [27] Albert Mosella-Montoro and Javier Ruiz-Hidalgo. Skinningnet: Two-stream graph convolutional neural network for skinning prediction of synthetic characters. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18593–18602, 2022. [3](#)
- [28] Charlie Nash, Yaroslav Ganin, SM Ali Eslami, and Peter Battaglia. Polygen: An autoregressive generative model of 3d meshes. In *International conference on machine learning*, pages 7220–7229. PMLR, 2020. [3](#), [5](#)
- [29] OpenAI. Gpt-4o, 2023. [4](#), [13](#)
- [30] Yawar Siddiqui, Antonio Alliegro, Alexey Artemov, Tatiana Tommasi, Daniele Sirigatti, Vladislav Rosov, Angela Dai, and Matthias Nießner. Meshgpt: Generating triangle meshes with decoder-only transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19615–19625, 2024. [3](#), [5](#)

- [31] Chaoyue Song, Jiacheng Wei, Ruibo Li, Fayao Liu, and Guosheng Lin. 3d pose transfer with correspondence learning and mesh refinement. *Advances in Neural Information Processing Systems*, 34:3108–3120, 2021. 11
- [32] Chaoyue Song, Jiacheng Wei, Ruibo Li, Fayao Liu, and Guosheng Lin. Unsupervised 3d pose transfer with cross consistency and dual reconstruction. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(8):10488–10499, 2023. 11
- [33] Chonghyuk Song, Gengshan Yang, Kangle Deng, Jun-Yan Zhu, and Deva Ramanan. Total-recon: Deformable scene reconstruction for embodied view synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 17671–17682, 2023. 19
- [34] Chaoyue Song, Jiacheng Wei, Tianyi Chen, Yiwen Chen, Chuan-Sheng Foo, Fayao Liu, and Guosheng Lin. Moda: Modeling deformable 3d objects from casual videos. *International Journal of Computer Vision*, pages 1–20, 2024. 3, 12
- [35] Chaoyue Song, Jiacheng Wei, Chuan Sheng Foo, Guosheng Lin, and Fayao Liu. Reacto: Reconstructing articulated objects from a single video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5384–5395, 2024. 3
- [36] Andrea Tagliasacchi, Ibraheem Alhashim, Matt Olson, and Hao Zhang. Mean curvature skeletons. In *Computer Graphics Forum*, pages 1735–1744. Wiley Online Library, 2012. 1, 3
- [37] Jiaxiang Tang, Zhaoshuo Li, Zekun Hao, Xian Liu, Gang Zeng, Ming-Yu Liu, and Qinsheng Zhang. Edgerunner: Auto-regressive auto-encoder for artistic mesh generation. *arXiv preprint arXiv:2409.18114*, 2024. 3
- [38] The Models-Resource. The models-resource, 2019. 2, 7
- [39] Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Advances in neural information processing systems*, 30, 2017. 3
- [40] Peng-Shuai Wang, Yang Liu, and Xin Tong. Dual octree graph networks for learning adaptive volumetric shape representations. *ACM Transactions on Graphics (TOG)*, 41(4): 1–15, 2022. 11
- [41] Haohan Weng, Yikai Wang, Tong Zhang, CL Chen, and Jun Zhu. Pivotmesh: Generic 3d mesh generation via pivot vertices guidance. *arXiv preprint arXiv:2405.16890*, 2024. 3
- [42] Zhan Xu, Yang Zhou, Evangelos Kalogerakis, and Karan Singh. Predicting animation skeletons for 3d articulated models via volumetric nets. In *2019 international conference on 3D vision (3DV)*, pages 298–307. IEEE, 2019. 1, 2, 3
- [43] Zhan Xu, Yang Zhou, Evangelos Kalogerakis, Chris Landreth, and Karan Singh. Rignet: Neural rigging for articulated characters. *arXiv preprint arXiv:2005.00559*, 2020. 1, 2, 3, 7, 8, 11, 13
- [44] Zhan Xu, Yang Zhou, Li Yi, and Evangelos Kalogerakis. Morig: Motion-aware rigging of character meshes from point clouds. In *SIGGRAPH Asia 2022 conference papers*, pages 1–9, 2022. 2, 3
- [45] Gengshan Yang, Minh Vo, Natalia Neverova, Deva Ramanan, Andrea Vedaldi, and Hanbyul Joo. Banmo: Building animatable 3d neural models from many casual videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2863–2873, 2022. 3
- [46] Biao Zhang and Peter Wonka. Functional diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4723–4732, 2024. 6, 7, 11
- [47] Hao Zhang, Di Chang, Fang Li, Mohammad Soleymani, and Narendra Ahuja. Magicpose4d: Crafting articulated models with appearance and motion control. *arXiv preprint arXiv:2405.14017*, 2024. 3, 7, 11
- [48] Hao Zhang, Fang Li, Samyak Rawlekar, and Narendra Ahuja. Learning implicit representation for reconstructing articulated objects. *arXiv preprint arXiv:2401.08809*, 2024.
- [49] Hao Zhang, Fang Li, Samyak Rawlekar, and Narendra Ahuja. S3o: A dual-phase approach for reconstructing dynamic shape and skeleton of articulated objects from single monocular video. *arXiv preprint arXiv:2405.12607*, 2024. 3
- [50] Longwen Zhang, Ziyu Wang, Qixuan Zhang, Qiwei Qiu, Anqi Pang, Haoran Jiang, Wei Yang, Lan Xu, and Jingyi Yu. Clay: A controllable large-scale generative model for creating high-quality 3d assets. *ACM Transactions on Graphics (TOG)*, 43(4):1–20, 2024. 2, 5
- [51] Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, et al. Opt: Open pre-trained transformer language models. *arXiv preprint arXiv:2205.01068*, 2022. 6
- [52] Zibo Zhao, Wen Liu, Xin Chen, Xianfang Zeng, Rui Wang, Pei Cheng, Bin Fu, Tao Chen, Gang Yu, and Shenghua Gao. Michelangelo: Conditional 3d shape generation based on shape-image-text aligned latent representation. *Advances in Neural Information Processing Systems*, 36, 2024. 6, 9, 11, 12