

ReWeaver: Towards Simulation-Ready and Topology-Accurate Garment Reconstruction

Ming Li^{1,2,3}, Hui Shan^{1,2,3}, Kai Zheng³, Chentao Shen^{1,3},
Siyu Liu⁶, Yanwei Fu^{2,4}, Zhen Chen⁵, Xiangru Huang³

¹Zhejiang University ²Shanghai Innovation Institute ³Westlake University
⁴Fudan University ⁵Adobe ⁶Xidian University



Figure 1. **ReWeaver**. From as few as four input views, *ReWeaver* reconstructs high-precision sewing patterns with complex topology together with their corresponding 3D geometry. The method outputs a unified 2D–3D garment representation, where each panel and edge is explicitly linked to its associated 3D points. This enables faithful, simulation-ready garment assets to be recovered from ordinary and sparse-view photographs without a controlled capture setup. The reconstructed garment geometry and topology is precisely aligned with the input images and can be used for 3D structural perception.

Abstract

High-quality 3D garment reconstruction plays a crucial role in mitigating the sim-to-real gap in applications such as digital avatars, virtual try-on and robotic manipulation. However, existing garment reconstruction methods typically rely on unstructured representations, such as 3D Gaussian Splats, struggling to provide accurate reconstructions of garment topology and sewing structures. As a result, the reconstructed outputs are often unsuitable for high-fidelity physical simulation. We propose **ReWeaver**, a novel framework for topology-accurate 3D garment and sewing pattern reconstruction from sparse multi-view RGB images. Given as few as four input views, *ReWeaver* pre-

dicts seams and panels as well as their connectivities in both the 2D UV space and the 3D space. The predicted seams and panels align precisely with the multi-view images, yielding structured 2D–3D garment representations suitable for 3D perception, high-fidelity physical simulation, and robotic manipulation. To enable effective training, we construct a large-scale dataset **GCD-TS**, comprising multi-view RGB images, 3D garment geometries, textured human body meshes and annotated sewing patterns. The dataset contains over 100,000 synthetic samples covering a wide range of complex geometries and topologies. Extensive experiments show that *ReWeaver* consistently outperforms existing methods in terms of topology accuracy, geometry alignment and seam-panel consis-

gency. Code and data will be available at the project sii-liming.github.io/ReWeaver/.

1. Introduction

High-quality 3D garment reconstruction is central to many emerging applications in virtual try-on, digital humans, gaming, and robotic manipulation [5, 10, 22, 48]. For these tasks, it is not sufficient to recover only a visually plausible garment surface: downstream simulation, animation, and asset creation require *structured* representations that capture how garments are constructed—namely, their sewing patterns, panels, and seams.

Existing reconstruction methods typically operate on unstructured representations such as point clouds, noisy meshes, unsigned or signed distance fields, or 3D Gaussian splats [9, 16, 28, 38, 39, 44, 45]. While these formats can approximate garment geometry, they lack explicit sewing structure, making them difficult to parameterize for physics-based simulation, garment editing, or retargeting. Moreover, these representations are inherently misaligned with industry-standard garment design workflows, which rely on 2D sewing patterns as the primary medium for design and manufacturing.

Motivated by these limitations, recent works have shifted toward learning sewing pattern structures directly [6, 14, 18, 24, 26, 27, 30, 36, 41, 47]. However, these approaches have significant shortcomings. Methods relying on pre-defined topologies [27] perform well only on simple garments and fail to handle unseen layouts. Vision–language–model–based approaches [6, 36] predict tokenized JSON descriptions that are later converted into 2D patterns with GarmentCode [19] framework. While this approach generalizes to a broader range of topologies, the geometric accuracy of the reconstruction remains unsatisfactory.

To simultaneously reconstruct accurate *topology* and *geometry* of 3D garment, we propose *ReWeaver*, a unified model that reconstructs both 3D garment structure and 2D sewing patterns from as few as four input views. *ReWeaver* predicts a set of 3D patches and curves together with their 2D counterparts, establishing explicit 2D–3D correspondences. The predicted 3D patches and curves align closely with the multi-view images, enabling accurate 3D perception and robotic tasks. At the same time, the predicted sewing patterns further enable *ReWeaver* to extract simulation-ready and topology-accurate 3D garment assets from in-the-wild sparse view images.

To accommodate for sparse and unknown distribution of input views, we adopt the multi-view visual encoder from VGGT [42]. The visual encoder alternates between intra-frame attention and inter-frame attention and outputs a set of encoded tokens. We reconstruct the 3D garment geometry and topology by predicting a set of 3D curves, a set

of 3D patches, and the patch–curve connectivity. Using the predicted patch–curve connectivity, we group each 3D patch tokens with its connecting curves tokens and perform intra-group attention to flatten the 3D curves into 2D panel edges.

Compared to existing methods, *ReWeaver* achieves superior garment reconstruction quality and, for the first time, offers garment geometry and topology reconstruction in both 2D and 3D spaces, with precise alignment to the input images. This allows *ReWeaver* to be employed in various downstream applications such as 3D structural perception, physical simulation, and garment asset creation.

Our main contributions can be summarized as follows:

1. We introduce **ReWeaver**, the *first* framework to jointly reconstruct structured 3D garments and 2D sewing patterns while maintaining 2D–3D correspondence.
2. We present **GCD-TS** (GarmentCodeData with extended textures and seam annotations), an extension of the large-scale GCD dataset that adds garment and body texture assets and provides structured 3D point-cloud annotations, together with their explicit correspondences to 2D panels.
3. Our model outperforms prior methods in pattern reconstruction, achieving strong topological generalization while maintaining high geometric fidelity to the input images.

2. Related Works

2.1. Learning-based garment reconstruction and generation

With the advancement of deep learning techniques and the construction of large-scale garment datasets [5, 6, 17, 20, 30, 36], a variety of data-driven methods have been proposed for garment pattern reconstruction. NeuralTailor [18] is a pioneering work that focuses on learning-based recovery of structured garment sewing patterns from 3D point clouds, and demonstrates generalization to some unseen garment types. Going further, SewFormer [30] introduces a two-level transformer network to directly predict sewing patterns from a single image. With the rapid development of vision-language models (VLMs) [2, 15, 29], recent works [6, 14, 36] have leveraged VLMs for data annotation or fine-tuning, enabling text-conditioned and image-conditioned sewing pattern reconstruction, generation, and editing. DressCode [14] presents the first text-driven garment generation pipeline that produces high-quality sewing patterns and physically-based textures, demonstrating the feasibility of using a GPT-based autoregressive model for sewing pattern generation tasks. [6, 36] further empower the model to generate sewing patterns with different tokenization schemes by fine-tuning LLaVA [29]. Recent works explore diffusion-based denoising in a latent space and then decode to various pattern representations [24, 25,

31]. With the exception of [25], these methods focus exclusively on 2D patterns and largely overlook precise 3D geometric understanding.

2.2. Optimization-based garment reconstruction

Some optimization-based methods, as well as hybrid approaches that combine learning-based prediction with physics-guided optimization, have made notable progress in recovering simulation-ready garments by integrating differentiable physics into the reconstruction process. For example, DiffAvatar [27] leverages the XPBD [33] algorithm to drive a differentiable physics engine that simulates cloth behavior, allowing gradient-based updates to align the simulated garment with the reconstructed scan geometry. However, this method relies on a predefined garment library, which limits its ability to generalize to diverse or unseen pattern structures. Furthermore, optimizing the pattern parameters requires running the differentiable simulation multiple times, making the overall process highly time-consuming. Dress-1-to-3 [26] addresses a more challenging setting by leveraging generative models as priors and using the C-IPC [23] simulation algorithm to reconstruct simulation-ready garments from a single input image. Its initial garment patterns are estimated from a single image using SewFormer [30], but these patterns have low geometric accuracy and may even contain incorrect topology. Gaussian Garments [40] reconstructs an initial static geometry from a single template frame via multi-view stereo, surface reconstruction, and remeshing, and subsequently registers it to the video sequence. However, the acquisition of the static geometry relies on multiple traditional steps, suffering from error accumulation, and such reconstruction requires the selected template frame to possess high visibility and rich textures.

3. Method

We introduce *ReWeaver*, a framework that reconstructs accurate geometry and topology of 3D garments from multi-view images with as few as four views. In the rest of this section, we first formally define the garment reconstruction task in Section 3.1, then introduce the building blocks of *ReWeaver*: a multi-view visual encoder adopted from VGGT [42] (Section 3.2); a consecutive module that predicts 3D curve and patch using a bi-path transformer (Section 3.3); and an intra-surface attention module that implicitly learns a flattening process that converts the 3D curves and patches into their 2D counterparts (Section 3.4). Finally, we describe the loss functions used for training in Section 3.5.

Terminology. Since we operate in both 2D and 3D spaces, we adopt the following terminology for clarity. We denote 3D garment surface regions as *patches* and their 2D counterparts as *panels*. We use *curve* to refer to the 3D locus




Figure 2. Visualization of the terminologies used in this paper.

of a seam or boundary, and *edge* to denote its 2D counterpart on the panel. See Figure 2 for a visualization.

3.1. Problem Setup

ReWeaver addresses the task of garment reconstruction from multi-view RGB images. Given a set of images $\mathcal{I} = \{I_j \mid I_j \in \mathbb{R}^{H \times W \times 3}\}$ capturing a 3D garment, the goal is to reconstruct *four* groups of elements:

1. **3D curves:** $\mathcal{C} = \{C_j \mid C_j \in \mathbb{R}^{M_C \times 3}\}$, representing the garment seams in 3D space.
2. **3D patches:** $\mathcal{P} = \{P_i \mid P_i \in \mathbb{R}^{M_P \times 3}\}$, representing garment surfaces in 3D space, each 3D patch is topologically equivalent to a 2D square.
3. **2D edges:** $\mathcal{E} = \{E_{ij} \mid E_{ij} \in \mathbb{R}^{M_E \times 2}\}$, where for patch P_i and its attached 3D curves C_j , E_{ij} denotes the 2D flattened boundary edge corresponding to C_j .¹
4. **Patch–curve connectivity:** A binary matrix $\sigma_{pc} \in \{0, 1\}^{N_P \times N_C}$ where $\sigma_{pc}(i, j) = 1$ indicates patch i is attached to curve j .

Here, N_P and N_C denote the number of patches and curves, respectively, while M_P , M_C , and M_E represent the number of points used to model each patch, curve, and edge. Together, these elements provide a unified multi-level representation. The 3D patches and curves support tasks such as *3D structural perception* and shape analysis, while the 2D edges and patch–curve connectivity enable *2D asset extraction* for garment pattern design and physics-based simulation (via 2D triangulation of panel interiors bounded by $\{E_{ij}\}$).

3.2. Visual Encoder

We adopt a multi-view visual encoder design following VGGT [42]. Each input image is divided into non-overlapping patches that are embedded into tokens using a DINOv2 backbone [37]. Tokens are processed by alternating intra-frame and inter-frame self-attention layers: the intra-frame layers refine per-view features, while the inter-frame layers aggregate information across views. Following

¹A single 3D seam C_j can induce multiple 2D edges across its incident patches. Thus, we employ two indices to denote each individual edge.

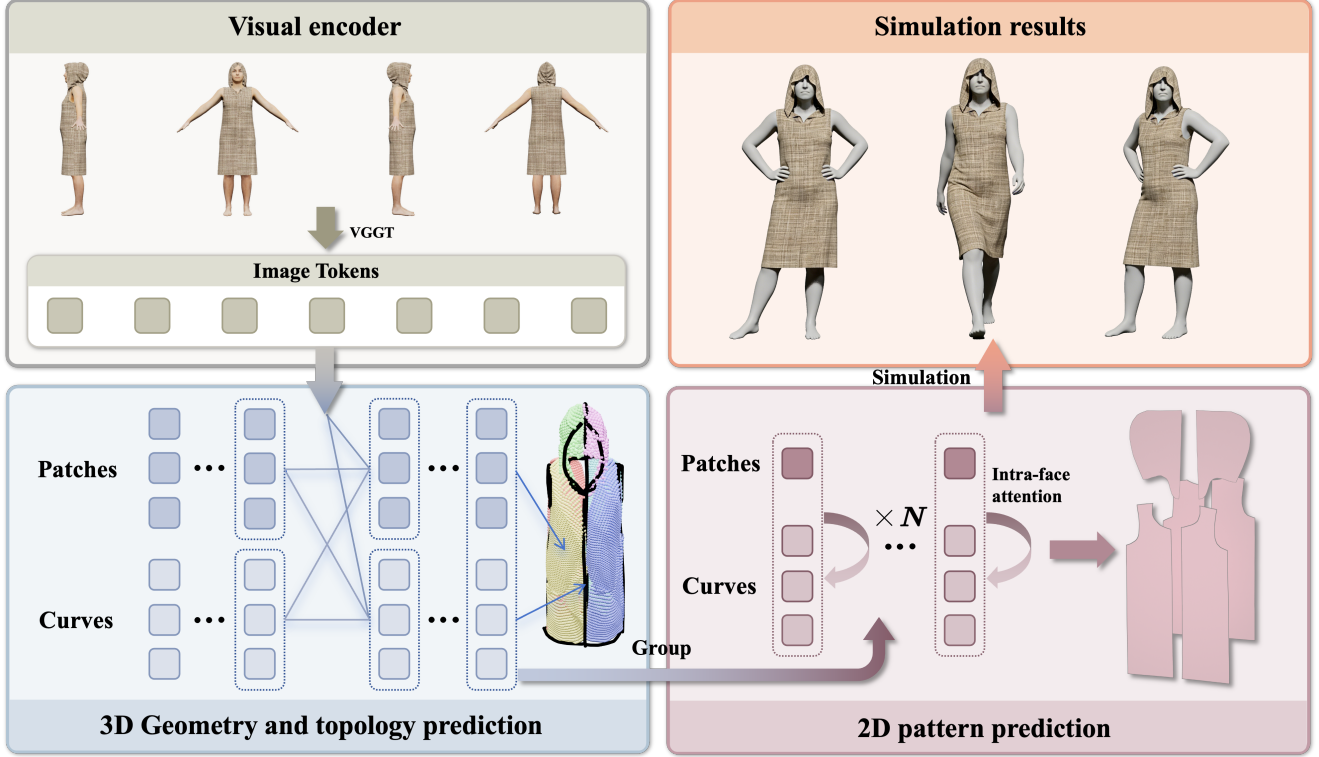


Figure 3. **Pipeline of our method.** Our VGGT-like image encoder extracts features from multi-view images (Section 3.2), which then interact with predefined patch and curve queries. In the 3D geometry and topology prediction module (Section 3.3), these queries and the image tokens pass through stacked self- and cross-attention blocks. The resulting tokens are decoded into 3D curves and patches. The same tokens are then reused in the 2D pattern prediction module (Section 3.4): guided by the refined topology, we group the patches and curves into patch-centric groups, apply intra-group attention, and finally decode the edges of the 2D panels. The decoded 2D panels can be directly used for physical simulation.

VGGT [42], these attention blocks are alternately stacked, progressively integrating local and global geometric cues into unified token representations. We concatenate the outputs of the final inter- and intra-frame layers and flatten all frame tokens into a sequence $T_i \in \mathbb{R}^{N_i \times D}$, where D is the embedding dimension.

3.3. 3D Curve and Patch Prediction

Given the token sequence T_i processed by the visual encoder, we design a topology and geometry prediction module that predicts the 3D curves, patches, and their connectivities. Following existing works in 3D object detection from multi-view images [8, 43], the prediction module uses latent patch and curve queries together with attention to make direct 3D predictions and support diverse garment topologies with an undetermined number of curves and patches.

Specifically, the prediction module takes as input the token sequence T_i and a set of learnable patch queries $Q_p \in \mathbb{R}^{N_p \times D}$ and curve queries $Q_c \in \mathbb{R}^{N_c \times D}$, and output three groups of elements: 1) the probability of each queried patch or curve being valid; 2) the predicted geometry of

each curve and patch; and 3) the predicted patch-curve connectivity. We next detail the overall architecture design and the decoding heads used for different prediction tasks.

Processing patch and curve queries. Inspired by ComplexGen [11], we adopt a bi-path Transformer architecture that models interactions both between patch and curve tokens and between these tokens and the image tokens. At each layer of the attention module, an element group (patch or curve) first performs self-attention to conduct intra-group message passing, and then applies cross-attention to retrieve contextual information from both the image tokens and the other element group. This design enables full fusion of image evidence while maintaining global structural and geometric consistency. After these attention operations, each path is passed through a layer normalization [4] and a feed-forward module, followed by a residual connection [13]. After the final layer, we obtain the refined tokens $T_p \in \mathbb{R}^{N_p \times D}$ and $T_c \in \mathbb{R}^{N_c \times D}$, which serve as the processed representations of the patch and curve queries, respectively.

Probability Prediction. The probability decoding heads separately decode tokens T_p and T_c into probability vectors for patches and curves, denoted as $\sigma_p \in [0, 1]^{N_p}$ and $\sigma_c \in [0, 1]^{N_c}$, respectively. For patches and curves, we use two separate decoding heads f_p^{prob} and f_c^{prob} , each consisting of a three-layer MLP followed by a sigmoid to normalize outputs to $[0, 1]$:

$$\sigma_p^i = \text{sigmoid}(f_p^{\text{prob}}(T_p^i)), \quad \sigma_c^i = \text{sigmoid}(f_c^{\text{prob}}(T_c^i)).$$

Low-probability elements are filtered by thresholds ϵ_p and ϵ_c , and the topology is further refined (Appendix A) to obtain binary validity masks $\sigma_p^* \in \{0, 1\}^{N_p}$ and $\sigma_c^* \in \{0, 1\}^{N_c}$.

Geometry Prediction. We represent each patch and curve as a continuous parametric function that maps canonical coordinates to 3D space. Specifically, these functions are implemented as three-layer MLPs that map $[0, 1]$ and $[0, 1]^2$ domains to \mathbb{R}^3 . Inspired by the hyper-network formulations [12, 46], we define the geometry prediction heads f_c^{geo} and f_p^{geo} as hyper-networks that generate the weights of these MLPs, conditioned on the processed query tokens T_c and T_p . Each curve or patch token thus parameterizes a unique mapping:

$$\begin{aligned} \forall u \in [0, 1], \quad g_c^i(u) &= f_c^{\text{geo}}(T_c^i)(u) \in \mathbb{R}^3 \\ \forall u, v \in [0, 1], \quad g_p^i(u, v) &= f_p^{\text{geo}}(T_p^i)(u, v) \in \mathbb{R}^3, \end{aligned} \quad (1)$$

where g_c^i and g_p^i denote the instantiated MLPs that generate 3D points along the i -th curve and patch, respectively.

During training, we uniformly sample $u \in [0, 1]$ and $(u, v) \in [0, 1]^2$, apply the corresponding MLPs to obtain the predicted point sets, and supervise them using point-based geometric losses. This implicit formulation enables differentiable supervision at arbitrary sampling densities without degrading geometric smoothness or continuity.

Connectivity Prediction. To predict the connectivity between patch i and curve j , we project the patch and curve tokens with two linear layers f_p^{adj} and f_c^{adj} , take their dot product, and apply a sigmoid to map it to $[0, 1]$ as the adjacency probability, i.e.,

$$\sigma_{pc}(i, j) = \text{sigmoid}(f_p^{\text{adj}}(T_p^i) \cdot f_c^{\text{adj}}(T_c^j)) \quad (2)$$

adjacency All these probabilities form an adjacency matrix σ_{pc} , which is first preliminarily filtered by the threshold ϵ_{adj} and then refined via a topology refinement procedure (Appendix A). The refined adjacency matrix is denoted as $\sigma_{pc}^* \in \{0, 1\}^{N_p \times N_c}$.

3.4. 2D Pattern Prediction

Given the valid patch and curve tokens (T_p, T_c) and their refined topology ($\sigma_p^*, \sigma_c^*, \sigma_{pc}^*$), we group each valid patch token with its connected curve tokens. Each group is processed by an intra-group attention module, resulting in edge tokens T_e that are then decoded into 2D panel edges.

Specifically, for each group, the curve tokens are first processed via self-attention to exchange information, followed by cross-attention with the associated patch token. We then apply the same layer normalization, feed-forward module, and residual connection as in Section 3.3. The resulting edge tokens are denoted as T_e . Let $\partial_i = \{j | \sigma_{pc}^*(i, j) = 1\}$ denote the indices of curve tokens connected to patch i , and let $T_e^{\partial_i}$ denote the corresponding edge tokens. For a connected curve $j \in \partial_i$, we use a three-layer hyper-network to generate an MLP that maps a canonical 1D parameter to normalized 2D coordinates:

$$\forall u \in [0, 1], \quad g_e^{ij}(u) = f_e^{\text{edge}}(T_e^j)(u) \in [0, 1]^2. \quad (3)$$

Here g_e^{ij} denotes the instantiated edge function that produces the 2D edge corresponding to the j -th curve connected to patch i , and f_e^{edge} denotes another hyper-network that takes in the edge token and output an MLP, similar to other previously defined hyper-networks. Because this formulation does not guarantee that the endpoints of adjacent edges meet perfectly, we apply a geometric refinement procedure (Appendix B) to enforce loop closure with a high success rate.

Since each 2D panel is predicted in a normalized coordinate system within $[0, 1]^2$, to recover the absolute panel metrics, we introduce an additional lightweight MLP f_p^{scale} that predicts a scalar scale factor s_i from the corresponding patch token T_p^i :

$$s_i = f_p^{\text{scale}}(T_p^i), \quad s_i \in \mathbb{R}. \quad (4)$$

The final scaled 2D panel is obtained by multiplying its normalized coordinates by s_i . This scale term allows the model to represent panels at realistic physical sizes and supports accurate downstream garment reconstruction.

3.5. Loss Functions

In this section, we elaborate how to compute differentiable losses for supervised learning. We first establish the correspondence between the predicted elements and the ground-truth annotations (patches, curves) using the Hungarian matching algorithm [21, 35].

Geometric Loss. For all parametric functions predicted by the hyper-networks (patches, curves, and edges), we supervise their output geometries using the Chamfer Distance (CD) loss. Let $\mathcal{G} = \{g_p^i, g_c^j, g_e^{ij}\}$ denote the set of all predicted geometric mappings, and $m(g)$ be the ground-truth

element matched to g through Hungarian matching. The geometric loss is computed over sampled point sets from each predicted mapping:

$$L_{\text{geo}} = \sum_{g \in \mathcal{G}} w_{\text{geo}}^{(g)} \text{CD}(V(g), V(m(g))), \quad (5)$$

where $V(g)$ and $V(m(g))$ are the sampled 2D or 3D point sets from the predicted and ground-truth geometries, respectively.

Classification and Connectivity Loss. For the predicted validity and connectivity probabilities ($\sigma_p, \sigma_c, \sigma_{pc}$), we apply Binary Cross-Entropy (BCE) supervision:

$$L_{\text{cls}} = \sum_{\sigma \in \{\sigma_p, \sigma_c, \sigma_{pc}\}} w_{\text{cls}}^{(\sigma)} \text{BCE}(\sigma, m(\sigma)), \quad (6)$$

where $m(\sigma)$ denotes the corresponding ground-truth label.

Scale Loss. For the 2D panel scale predictions, we use an ℓ_2 loss:

$$L_{\text{scale}} = \sum_{i=1}^{N_p} w_{\text{scale}} \|s_i - s_{m(i)}^{\text{gt}}\|_2^2. \quad (7)$$

Total Loss. The overall training objective combines all above loss functions, where the weighting coefficients $w_{\text{cls}}^{(\cdot)}$, $w_{\text{geo}}^{(\cdot)}$, and w_{scale} are empirically determined (see Appendix C for details).

4. Experiments

4.1. GCD-TS Dataset

To construct the GCD-TS dataset, we follow the garment sampling and simulation procedure of GCD [19, 20] using both male and female SMPL [32] bodies. To improve realism and generalization, we replace the original GCD textures—which contain strong seam cues and textureless bodies—with a diverse set of garment and body textures. For body textures, we adopt nearly 50 textures from BED-LAM [7]. For garment textures, we randomly sample from our collected library of tileable textures drawn from open-source sources [1] and commercial datasets [3]. Each garment-body pairing is rendered from four viewpoints (front, back, left, right), with small randomized perturbations in camera pose. In total, GCD-TS contains roughly **100,000** textured multi-view samples. Figure 4 illustrates the difference between GCD and GCD-TS textures.

4.2. Experimental Setup

Training Details. We randomly split the GCD-TS dataset into training, validation, and test sets with a ratio of 8:1:1,



Figure 4. **Texture differences between GCD and GCD-TS.** Using the same garment geometry, GCD textures reveal strong seam cues (e.g., highlighted regions), which are unrealistic and can lead to overfitting. GCD-TS replaces these with more realistic, diverse textures to improve generalization.

and trained our model on the training split. We adopted the lightweight 21 MB pre-trained weights of DINOv2 [37] to initialize the image feature extraction module. The numbers of patch and curve queries were set to 200 and 70, respectively — approximately twice the maximum counts observed in the training data. Each input sample consists of four RGB images with a resolution of 518×518 . Each curve and its corresponding edge are represented by 50 uniformly sampled points, and each patch is represented by a 20×20 point set. All tokens, including image, curve, and patch tokens, have dimension 768. See Appendix C for more implementation and training details.

Evaluation Metrics. To quantitatively evaluate our sewing pattern predictions, we define a set of evaluation metrics and also adopt several indicators from previous works [25, 36]. Before computing the 2D geometric metrics, we zero-center each panel. To assess the accuracy of topological prediction, we report: 1) **Panel Count Accuracy (Acc_p)**, the percentage of sewing patterns with a correctly predicted number of panels; 2) **Edge Count Accuracy (Acc_e)**, the percentage of correctly predicted edge counts within each correctly predicted panel; and 3) **Overall Accuracy (Acc_o)**, defined as the product of Panel Count Accuracy and Edge Count Accuracy, providing a comprehensive measure of garment topology reconstruction quality. To evaluate 2D geometric precision, we report: 4) **2D Edge CD (CD_e)**, the Chamfer Distance between the closed-loop edge point sets of the predicted and ground-truth panels.; and 5) **2D Panel IoU (IoU)**, the pixel-wise Intersection over Union between rasterized images of the predicted and ground-truth panels. To evaluate 3D geometric accuracy,




Figure 5. **Visualization on adaptive sampling density at inference time.** From left to right, the point set sampled at a fixed density of 20×20 points per patch; the point set sampled at the adaptive-density; the ground truth point cloud.

we report: 6) **Patch Point Cloud CD (CD_p)**, the Chamfer Distance between the predicted and ground-truth 3D patch point sets; and 7) **3D Curve CD (CD_c)**, the Chamfer Distance between the predicted and ground-truth 3D curve point sets. These metrics are representation-agnostic and can be applied to both our method and the baselines.

4.3. Results

Geometry Reconstruction. *ReWeaver* predicts patch geometries by mapping uniformly sampled coordinates from the unit square through hyper-networks. Although training uses a fixed 20×20 sampling density, the smoothness of the implicit mapping allows arbitrary sampling densities at inference. Because patch sizes vary significantly, a fixed sampling rate can make small patches appear overly dense and large patches overly sparse. To address this, we pre-sample a 20×20 grid and adaptively retain points based on spatial variance, producing a near-uniform point density across patches. While Li et al. [25] also predict 3D point clouds, its resolution is fixed and cannot adapt in this manner. Figure 5 compares our adaptive sampling with the ground-truth point clouds, and Table 2 reports the corresponding CD scores. For curves, we use a fixed budget of 50 samples per curve, which is sufficient to capture all curve lengths without noticeable visual artifacts.

2D Sewing Pattern Reconstruction. We choose Sewformer [30], ChatGarment [6], and AIpparel [36] as our baseline models. For fairness, we replace their monocular image input with a four-view image input. Since both GCD-TS and the GCD-MM used by AIpparel for training essentially follow the same GCD distribution, we fine-tune AIpparel on our GCD-TS dataset until the loss curve plateaus rather than training it from scratch. For the same reason, we fine-tune ChatGarment instead of training from scratch, but convert the JSON files of our GCD-TS dataset into the re-

finer version following the method described in the original paper. For the training of Sewformer, we adopt the modifications introduced in [36].

Although our task uses multi-view inputs, the garment reconstruction task is more challenging on our dataset because we employ tileable textures rather than GCD’s default textures, which contain strong seam cues (see Figure 4). A visual comparison is shown in Figure 7, and quantitative results are reported in Table 1.




Figure 6. **Topology and geometry refinement.** Without topology refinement, redundant 3D curves cause incorrect 2D edges (black segment in the leftmost image). Topology refinement removes these redundancies, producing clean panel structures, while geometry refinement enforces accurate, closed-loop boundaries suitable for triangulation.

Table 1. **Quantitative Comparison of 2D Panel Quality.** We compare the 2D panels generated by *ReWeaver* and baselines using the metrics defined in Section 4.2. *ReWeaver* outperforms on five out of six metrics.

Method	$Acc_p \uparrow$	$Acc_e \uparrow$	$Acc_o \uparrow$	$CD_e \downarrow$	IoU \uparrow
Sewformer	0.3761	0.4802	0.1806	0.1161	0.5844
ChatGarment	0.5557	0.8012	0.4452	0.0906	0.6533
AIpparel	0.4561	0.6774	0.3090	0.0648	0.7084
<i>ReWeaver</i>	0.9210	0.7175	0.6608	0.0391	0.8221

Ablation Study. We evaluate the impact of the topology and geometry refinement procedures described in Appendix A and B. Topology refinement removes duplicated or redundant edges, yielding a substantial improvement in Edge Count Accuracy (Acc_e) and enabling more accurate 2D panel reconstruction. Geometry refinement closes small gaps between edges in 2D space, producing fully closed panel boundaries. Figure 6 shows qualitative differences, and Table 2 reports the corresponding quantitative results.

5. Conclusion

We presented *ReWeaver*, an integrated framework for multi-view garment reconstruction that supports both perception and pattern recovery. *ReWeaver* produces 3D geometries that align closely with sparse-view inputs and reconstructs




Figure 7. **Comparison with AIpparel.** We compare *ReWeaver* and **AIpparel** against ground truth. Each example shows the predicted and ground-truth 2D panels along with the resulting simulated meshes. *ReWeaver* yields more accurate panels and correspondingly improved simulation results.

Table 2. **Effects of topology and geometry refinement.** Refinement removes redundant edges and enforces closed boundaries, yielding higher Edge Acc and improved geometric quality (lower Edge CD and higher Panel IoU).

	$CD_p^{\text{base}} \downarrow$	$CD_p^{\text{adapt}} \downarrow$	$CD_c \downarrow$	$Acc_p \uparrow$	$Acc_e \uparrow$	$Acc_o \uparrow$	$CD_e \downarrow$	IoU \uparrow
with refine-correction	0.0232	0.0185	0.0266	0.9210	0.7175	0.6608	0.0391	0.8221
without refine	0.0225	0.0188	0.0255	0.9101	0.5361	0.4880	0.0416	0.7775

accurate 2D sewing patterns suitable for downstream physical simulation. Our experiments demonstrate that jointly reasoning over geometric primitives in both 2D and 3D

spaces leads to notable improvements in garment topology and geometric fidelity.

References

- [1] 3D Jungle. <https://3djungle.net>. 6
- [2] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023. 2
- [3] architextures. <https://architextures.org>. 6
- [4] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016. 4
- [5] Hugo Bertiche, Meysam Madadi, and Sergio Escalera. Cloth3d: clothed 3d humans. In *European Conference on Computer Vision*, pages 344–359. Springer, 2020. 2
- [6] Siyuan Bian, Chenghao Xu, Yuliang Xiu, Artur Grigorev, Zhen Liu, Cewu Lu, Michael J Black, and Yao Feng. Chatgarment: Garment estimation, generation and editing via large language models. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 2924–2934, 2025. 2, 7
- [7] Michael J Black, Priyanka Patel, Joachim Tesch, and Jinlong Yang. Bedlam: A synthetic dataset of bodies exhibiting detailed lifelike animated motion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8726–8737, 2023. 6
- [8] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European conference on computer vision*, pages 213–229. Springer, 2020. 4
- [9] Andrés Casado-Elvira, Marc Comino Trinidad, and Dan Casas. Pergamo: Personalized 3d garments from monocular video. In *Computer Graphics Forum*, pages 293–304. Wiley Online Library, 2022. 2
- [10] H Onan Demirel, Salman Ahmed, and Vincent G Duffy. Digital human modeling: a review and reappraisal of origins, present, and expected future methods for representing humans computationally. *International Journal of Human-Computer Interaction*, 38(10):897–937, 2022. 2
- [11] Haoxiang Guo, Shilin Liu, Hao Pan, Yang Liu, Xin Tong, and Baining Guo. Complexgen: Cad reconstruction by b-rep chain complex generation. *ACM Transactions on Graphics (TOG)*, 41(4):1–18, 2022. 4
- [12] David Ha, Andrew Dai, and Quoc V Le. Hypernetworks. *arXiv preprint arXiv:1609.09106*, 2016. 5
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 4
- [14] Kai He, Kaixin Yao, Qixuan Zhang, Jingyi Yu, Lingjie Liu, and Lan Xu. Dresscode: Autoregressively sewing and generating garments from text guidance. *ACM Transactions on Graphics (TOG)*, 43(4):1–13, 2024. 2, 4
- [15] Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, et al. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*, 2024. 2
- [16] Boyi Jiang, Juyong Zhang, Yang Hong, Jinhao Luo, Ligang Liu, and Hujun Bao. Bcnet: Learning body and cloth shape from a single image. In *European Conference on Computer Vision*, pages 18–35. Springer, 2020. 2
- [17] Maria Korosteleva and Sung-Hee Lee. Generating datasets of 3d garments with sewing patterns. *arXiv preprint arXiv:2109.05633*, 2021. 2
- [18] Maria Korosteleva and Sung-Hee Lee. Neurtailor: Reconstructing sewing pattern structures from 3d point clouds of garments. *ACM Transactions on Graphics (TOG)*, 41(4):1–16, 2022. 2
- [19] Maria Korosteleva and Olga Sorkine-Hornung. Garmentcode: Programming parametric sewing patterns. *ACM Transactions on Graphics (TOG)*, 42(6):1–15, 2023. 2, 6
- [20] Maria Korosteleva, Timur Levent Kesdogan, Fabian Kemper, Stephan Wenninger, Jasmin Koller, Yuhan Zhang, Mario Botsch, and Olga Sorkine-Hornung. Garmentcodedata: A dataset of 3d made-to-measure garments with sewing patterns. In *European Conference on Computer Vision*, pages 110–127. Springer, 2024. 2, 6
- [21] Harold W Kuhn. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97, 1955. 5
- [22] Jason Lawrence, Ye Pan, Dan B Goldman, Rachel McDonnell, Julie Robillard, Carol O’Sullivan, Yaser Sheikh, Michael Zollhoefer, and Jason Saragih. State of the art in telepresence. In *ACM SIGGRAPH 2022 Courses*, pages 1–74. 2022. 2
- [23] Minchen Li, Danny M Kaufman, and Chenfanfu Jiang. Codimensional incremental potential contact. *ACM Transactions on Graphics*, 40(4), 2021. 3
- [24] Ren Li, Cong Cao, Corentin Dumery, Yingxuan You, Hao Li, and Pascal Fua. Single view garment reconstruction using diffusion mapping via pattern coordinates. In *Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers*, pages 1–11, 2025. 2
- [25] Siran Li, Chen Liu, Ruiyang Liu, Zhendong Wang, Gaofeng He, Yong-Lu Li, Xiaogang Jin, and Huamin Wang. Garmagenet: A multimodal generative framework for sewing pattern design and generic garment modeling. *arXiv preprint arXiv:2504.01483*, 2025. 2, 3, 6, 7
- [26] Xuan Li, Chang Yu, Wenxin Du, Ying Jiang, Tianyi Xie, Yunuo Chen, Yin Yang, and Chenfanfu Jiang. Dress-1-to-3: Single image to simulation-ready 3d outfit with diffusion prior and differentiable physics. *arXiv preprint arXiv:2502.03449*, 2025. 2, 3
- [27] Yifei Li, Hsiao-yu Chen, Egor Larionov, Nikolaos Sarafianos, Wojciech Matusik, and Tuur Stuyck. Diffavatar: Simulation-ready garment optimization with differentiable simulation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4368–4378, 2024. 2, 3
- [28] Zhe Li, Zerong Zheng, Lizhen Wang, and Yebin Liu. Animatable gaussians: Learning pose-dependent gaussian maps for high-fidelity human avatar modeling. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 19711–19722, 2024. 2

- [29] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *Advances in neural information processing systems*, 36:34892–34916, 2023. 2
- [30] Lijuan Liu, Xiangyu Xu, Zhijie Lin, Jiabin Liang, and Shuicheng Yan. Towards garment sewing pattern reconstruction from a single image. *ACM Transactions on Graphics (TOG)*, 42(6):1–15, 2023. 2, 3, 7
- [31] Shengqi Liu, Yuhao Cheng, Zhuo Chen, Xingyu Ren, Wenhan Zhu, Lincheng Li, Mengxiao Bi, Xiaokang Yang, and Yichao Yan. Multimodal latent diffusion model for complex sewing pattern generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 17640–17650, 2025. 3
- [32] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. Smpl: A skinned multi-person linear model. In *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*, pages 851–866. 2023. 6
- [33] Miles Macklin, Matthias Müller, and Nuttapon Chentanez. Xpbd: position-based simulation of compliant constrained dynamics. In *Proceedings of the 9th International Conference on Motion in Games*, pages 49–54, 2016. 3
- [34] Marvelous Designer. Marvelous designer official website, 2024. 4
- [35] James Munkres. Algorithms for the assignment and transportation problems. *Journal of the society for industrial and applied mathematics*, 5(1):32–38, 1957. 5
- [36] Kiyohiro Nakayama, Jan Ackermann, Timur Levent Kesdogan, Yang Zheng, Maria Korosteleva, Olga Sorkine-Hornung, Leonidas J Guibas, Guandao Yang, and Gordon Wetzstein. Aiparel: A multimodal foundation model for digital garments. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 8138–8149, 2025. 2, 6, 7, 4
- [37] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafranec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023. 3, 6, 2
- [38] Lingteng Qiu, Guanying Chen, Jiapeng Zhou, Mutian Xu, Junle Wang, and Xiaoguang Han. Rec-mv: Reconstructing 3d dynamic cloth from monocular videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4637–4646, 2023. 2
- [39] Boxiang Rong, Artur Grigorev, Wenbo Wang, Michael J Black, Bernhard Thomaszewski, Christina Tsalicoglou, and Otmar Hilliges. Gaussian garments: Reconstructing simulation-ready clothing with photorealistic appearance from multi-view video. *arXiv preprint arXiv:2409.08189*, 2024. 2
- [40] Boxiang Rong, Artur Grigorev, Wenbo Wang, Michael J Black, Bernhard Thomaszewski, Christina Tsalicoglou, and Otmar Hilliges. Gaussian garments: Reconstructing simulation-ready clothing with photorealistic appearance from multi-view video. In *2025 International Conference on 3D Vision (3DV)*, pages 1054–1063. IEEE, 2025. 3
- [41] Yuki Tatsukawa, Anran Qi, I-Chao Shen, and Takeo Igarashi. Garmentimage: Raster encoding of garment sewing patterns with diverse topologies. In *Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers*, pages 1–11, 2025. 2
- [42] Jianyuan Wang, Minghao Chen, Nikita Karaev, Andrea Vedaldi, Christian Rupprecht, and David Novotny. Vggg: Visual geometry grounded transformer. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 5294–5306, 2025. 2, 3, 4
- [43] Yue Wang, Vitor Campagnolo Guizilini, Tianyuan Zhang, Yilun Wang, Hang Zhao, and Justin Solomon. Detr3d: 3d object detection from multi-view images via 3d-to-2d queries. In *Conference on robot learning*, pages 180–191. PMLR, 2022. 4
- [44] Donglai Xiang, Timur Bagautdinov, Tuur Stuyck, Fabian Prada, Javier Romero, Weipeng Xu, Shunsuke Saito, Jingfan Guo, Breannan Smith, Takaaki Shiratori, et al. Dressing avatars: Deep photorealistic appearance for physically simulated clothing. *ACM Transactions on Graphics (TOG)*, 41(6):1–15, 2022. 2
- [45] Yuliang Xiu, Jinlong Yang, Xu Cao, Dimitrios Tzionas, and Michael J Black. Econ: Explicit clothed humans optimized via normal integration. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 512–523, 2023. 2
- [46] Dominic Zhao, Seijin Kobayashi, João Sacramento, and Johannes von Oswald. Meta-learning via hypernetworks. In *4th Workshop on Meta-Learning at NeurIPS 2020 (MetaLearn 2020)*. NeurIPS, 2020. 5
- [47] Yang Zheng, Qingqing Zhao, Guandao Yang, Wang Yifan, Donglai Xiang, Florian Dubost, Dmitry Lagun, Thabo Beeler, Federico Tombari, Leonidas Guibas, et al. Physavatar: Learning the physics of dressed 3d avatars from visual observations. In *European Conference on Computer Vision*, pages 262–284. Springer, 2024. 2
- [48] Xingxing Zou, Xintong Han, and Waikeng Wong. Cloth4d: A dataset for clothed human reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12847–12857, 2023. 2

ReWeaver: Towards Simulation-Ready and Topology-Accurate Garment Reconstruction

Supplementary Material

A. Topology Refinement

To obtain clean sewing-structure topology from the raw predictions, we apply a refinement pipeline that combines reliability-based filtering, 3D curve consolidation, and 2D loop analysis. This removes redundant or inconsistent curves and recovers stable, closed panel boundaries.

We begin by filtering patch predictions using a relatively high threshold $\epsilon_p = 0.7$. Since patch validity is already highly reliable ($\text{Acc}_p = 0.9210$), this single cutoff cleanly separates valid from invalid patch queries.

Curve validity and patch–curve adjacencies, in contrast, are filtered with more permissive thresholds, $\epsilon_c = 0.5$ and $\epsilon_{\text{adj}} = 0.5$. Empirically, these thresholds avoid false negatives but may retain redundant or overlapping curve predictions. The following refinement steps address this over-retention.

Duplicate Curve Merge We first merge curves that represent nearly the identical geometry. Two curves C_a and C_b are considered duplicates if their bidirectional Chamfer distance is small, i.e.,

$$\text{CD}(C_a \rightarrow C_b) < 0.03 \quad \text{and} \quad \text{CD}(C_b \rightarrow C_a) < 0.03.$$

For each such pair, we retain the curve with the higher predicted validity probability and transfer the discarded curve’s adjacency relations to it.

Sub-curve Removal If two curves share the same adjacency pattern, and one is geometrically contained inside the other, we treat the shorter one as a spurious “sub-curve.” Formally, if

$$\text{CD}(C_{\text{sub}} \rightarrow C_{\text{main}}) < 0.04,$$

we discard C_{sub} . This removes small fragments arising from local over-segmentation.

2D loop pruning For each panel, let $\{E_{ij}\}_{j=1}^N$ denote the incident 2D edges predicted from the associated 3D curves. Our goal is to keep only the set of edges that best forms a closed loop.

We first consider all possible orderings and orientations of the edges to form a closed boundary:

$$[E'_{ijk}] = (E'_{ij_1}, E'_{ij_2}, \dots, E'_{ij_N}).$$

For any ordering, we define the loop cost as

$$C([E'_{ijk}]) = \sum_{k=1}^N \left\| E'_{ijk}[-1] - E'_{ijk+1}[0] \right\|_2^2, \quad j_{N+1} = j_1. \quad (8)$$

The optimal cost for the unordered edge set is then

$$C(\{E_{ij}\}_{j=1}^N) = \min_{[E'_{ijk}]} C([E'_{ijk}]). \quad (9)$$

For any edge E_{ik} in the panel, we compute the loop cost after removing it:

$$C(\{E_{ij}\}_{j=1}^N \setminus \{E_{ik}\}).$$

If removal reduces the optimal cost by a sufficient margin,

$$C(\{E_{ij}\}_{j=1}^N \setminus \{E_{ik}\}) < C(\{E_{ij}\}_{j=1}^N) - \epsilon, \quad (10)$$

where $\epsilon = 10^{-9}$, then E_{ik} is deemed inconsistent with the panel loop and is pruned. When multiple edges satisfy the condition, we remove the one that yields the greatest cost reduction. The pruning is repeated until convergence.

Ablation Analysis

To isolate the effect of each refinement step, we conduct a cumulative ablation study in which the following rules are progressively enabled:

1. threshold-based filtering only;
2. thresholding + 2D loop pruning;
3. thresholding + loop pruning + sub-curve removal;
4. thresholding + loop pruning + sub-curve removal + duplicate merging (full refinement).

The results in Table 3 show that each refinement stage contributes positively to edge accuracy and panel IoU. Because topology refinement has negligible influence on the Chamfer Distance of patches and curves, we omit those metrics here.

B. 2D Geometry Refinement

Since our model does not explicitly enforce perfect edge-to-edge connectivity, the predicted panel boundaries may contain small gaps and misalignments. To obtain visually smoother and more aesthetically pleasing panels, we apply a geometry refinement stage. We first reorder and flip the edges as described in Section A, and then perform the following two geometry refinement steps:

Table 3. **More ablation on topology refinement.** Progressive rules further improve edge accuracy and panel IoU and slightly reduce edge Chamfer Distance CD_e .

Method	Acc _e ↑	CD _e ↓	IoU ↑
Threshold only	0.5361	0.0416	0.7775
+ 2D loop pruning	0.6045	0.0444	0.7584
+ sub-curve removal	0.6269	0.0418	0.7940
+ duplicate merging	0.7175	0.0391	0.8221

Replace bad edges. Consider an ordered closed loop of N edges $\{E_{ij}\}$ for panel i , where each edge $E_{ij} \in \mathbb{R}^{M_E \times 2}$ is represented by M_E sampled 2D points. For each edge E_{ij} , we denote s_j and e_j as the start and end point, and similarly e_{j-1} and s_{j+1} from its neighbors (indices modulo N). We define the local connection gaps as:

$$\text{gap}_1(j) = \|s_j - e_{j-1}\|_2, \quad \text{gap}_2(j) = \|e_j - s_{j+1}\|_2,$$

and normalize them by the edge length $\ell_j = \|e_j - s_j\|_2$. Edges with a large normalized connection error,

$$\frac{\text{gap}_1(j)}{\ell_j} + \frac{\text{gap}_2(j)}{\ell_j} > \tau_{\text{gap}}, \quad (11)$$

are marked as "bad", where $\tau_{\text{gap}} = 3.0$. For each bad edge E_{ij} , we discard its original geometry and replace it with a straight segment that linearly interpolates between e_{i-1} and s_{i+1} , resampled into M_E points. This operation preserves the loop topology while smoothing out locally inconsistent or noisy edge predictions.

Align edges to joint midpoints. After replacing obviously bad edges, we further refine the geometry by snapping all edges to a consistent set of joint targets along the loop. For a panel i with an ordered closed loop of edges $\{E_{ij}\}_{j=1}^N$, each sampled as $E_{ij} \in \mathbb{R}^{M_E \times 2}$, we let e_j and s_j denote the end point and start point of edge E_{ij} , respectively. For the joint between E_{ij} and $E_{i,j+1}$ (indices taken modulo N), we define a target vertex

$$v_j = \frac{e_j + s_{j+1}}{2}, \quad (12)$$

i.e., the midpoint of the two incident edge endpoints. This yields a set of N joint targets $\{v_j\}_{j=1}^N$ distributed along the loop.

For each edge E_{ij} , we estimate a 2D similarity transform

$$T_{ij}(p) = s_{ij}R_{ij}p + t_{ij},$$

that aligns its endpoints to the corresponding joint targets:

$$T_{ij}(s_j) = v_{j-1}, \quad T_{ij}(e_j) = v_j.$$

Here $s_{ij} \in \mathbb{R}$ denotes the scale factor, $R_{ij} \in \mathbb{R}^{2 \times 2}$ is a 2D rotation matrix, and $t_{ij} \in \mathbb{R}^2$ is the translation vector. Such a solution clearly exists and is unique in the two-dimensional space. The scale parameter is clamped to a reasonable range to avoid degenerate transforms. After solving for T_{ij} , we apply it to all M_E sampled points of E_{ij} to obtain the refined edge \tilde{E}_{ij} . This alignment step enforces consistent edge-to-edge joints, removes small gaps or overlaps, and preserves the overall shape of each edge.

C. Experiment Details

More training details. We train our model for 50 epochs on $8 \times \text{H200}$ GPUs, which takes about 120 hours in total. We use the AdamW optimizer with a unified weight decay of 1×10^{-4} , and set the base learning rates of the visual encoder, 3D prediction head, and 2D prediction head to 5×10^{-5} , 1×10^{-5} , and 5×10^{-5} , respectively. All three modules share the same iteration-wise schedule: a short linear warmup of 1 epochs that ramps the learning rate from a very small value to its base value, followed by cosine annealing down to a small floor of 5×10^{-6} . We use a total batch size of 128 (16 per GPU across 8 GPUs).

Visual Encoder. We adopt a lightweight DINO-based image encoder with 22M parameters, followed by a stack of 12 intra-frame and inter-frame attention layers. We use the DINOv2[37] backbone with a patch size of 14, and its weights are kept fully trainable during training. Inside the visual encoder, the token dimension is fixed to 384. The outputs of the last intra-frame attention layer and the last inter-frame attention layer are concatenated to form a 768-dimensional feature, which is used as the input to the subsequent modules. We resize all input images to a resolution of 518×518 and apply per-pixel normalization by subtracting the dataset mean and dividing by the dataset standard deviation.

3D Curve and Patch Prediction. We set the numbers of curve and patch queries to 200 and 70, respectively—about twice the maximum values in the dataset—to provide a safe margin for more complex or augmented cases while keeping the number of empty queries and the computational cost reasonably low. We assign a weight of 1 to all BCE losses and a weight of 300 to all geometric losses. We train curve/patch validity with a weighted binary cross-entropy loss: the positive class (valid curve/patch) has weight 1, and the negative class (empty query/patch) weight is chosen from the empirical ratio between valid curves and empty queries, scaled by a global factor so that both contribute roughly equally to the loss. We use a bipath transformer decoder with 12 layers, where each layer applies 1) per-path self-attention; 2) cross-attention between the two paths con-




Figure 8. Visualization of more results

ditioned on primitive-type embeddings; 3) cross-attention to the image features; and 4) a final feed-forward network.

2D Pattern Prediction. We assign a weight of 300 to the edge geometry loss and 1×10^{-2} to the scale loss. Each edge is represented by 50 points that are uniformly sampled along its arc length. For every panel, we normalize its 2D coordinates to the range $[-1, 1]$ by subtracting the panel-wise

mean and dividing by the maximum absolute coordinate value (scale). This normalization improves the numerical stability of the predictions and largely removes the ambiguity caused by global translations. The 2D pattern prediction module takes curve and patch features together with the patch–curve adjacency matrix, and first groups the corresponding tokens into panel-wise sets according to their adjacency relations. The panel-wise edge tokens are then pro-

cessed by 12 stacked transformer-style layers: each layer applies self-attention among edges of the same panel, followed by cross-attention where edge tokens attend to the corresponding panel token (edges as queries, panel token as keys/values), and finally a feed-forward network update, yielding the edge geometry and a scalar scale prediction for each panel.

D. More Results

The results of our experiments on various styles, such as strapless dresses, pencil skirts, and asymmetrical tops, are shown in Figure 8.

E. Garment Simulation

To dress the target human model, 2D garment panels are first positioned in 3D space based on their corresponding point cloud data. These panels are then imported into Marvelous Designer [34] along with the body model, where their placement is manually refined to ensure alignment with the respective body regions. Furthermore, the sewing relationships are established using the predictions from our model, which identifies and stitches together the edges corresponding to the same 3D curves. Subsequently, we employ the software’s built-in physics engine with default material parameters to simulate the draping process. This wraps the panels around the body under collision constraints, achieving a final dressed state that conforms to the target pose. The associated project files will be released in our code repository.

F. Discussion

Notably, compared to prior methods that predict structured pattern descriptions including seam connectivity [14, 36], our approach is inherently more flexible and potentially more generalizable to topology. Since we operate directly on panels and their 2D edge geometry rather than committing to a fixed Structured pattern representation, our method can in principle accommodate garments with additional components (e.g., sewn-on pockets) or complex topologies where seams are not strictly one-to-one at the edge level, without any architectural changes. A thorough empirical study of this generalization capability, especially on such challenging cases, is left for future work.

Unlike most previous works that focus on either 3D garment geometry or 2D panel layouts in isolation, our method reconstructs both representations jointly, which naturally enables further joint optimization to refine them simultaneously. Beyond this, one could extend our framework with additional prediction heads or differentiable cloth simulation to also estimate physical material parameters and human shape/pose from multi-view images, using the coupled 2D–3D predictions as constraints. This would allow

reconstruction of more complete and better aligned 3D assets from real-world image, and we believe this constitutes a particularly interesting direction for future work.

G. Limitations

A key limitation lies in the scarcity of high-quality 3D garments with complex topology and photorealistic textures. This limits the model’s generalization ability on out-of-distribution and real-world data. Future work could introduce greater variations in camera parameters, human shapes and poses, and backgrounds to generate more diverse synthetic data, and incorporate real-world data during training to narrow the sim-to-real gap.