SIMSEW: TOWARDS SIMULATION-READY SEWING PATTERN GENERATION

Anonymous authors
Paper under double-blind review

ABSTRACT

In this paper, we focus on sewing pattern generation for digital garments. Existing methods are primarily trained with supervised objectives on ground-truth sewing patterns. However, such objectives do not directly constrain key factors related to simulation success, such as garment simulatability and mesh quality. Tackling this issue, we present SimSew, a RL framework for simulation-ready sewing pattern generation. SimSew directly incorporates simulation-aware rewards derived from downstream garment simulation into RL, encouraging the generator to produce patterns that are more suitable for physical simulation. To provide informative group-relative learning signals, we introduce a simulability-aware reward that jointly evaluates pattern validity, stitching consistency, geometric fidelity, and mesh quality. Dense structural and geometric feedback differentiates candidate patterns even when simulation fails, while lower-tail triangle quality further distinguishes successfully constructed meshes. Extensive experiments demonstrate that SimSew improves downstream simulation readiness while preserving sewing-pattern quality. On GCD-MM, SimSew increases the simulation success rate from 71.89% to 76.92%; reward-aligned evaluation further improves One-to-one Stitch from 86.76% to 93.56% and reduces Corner MAE from $18.97°$ to $10.66°$.

1 INTRODUCTION

Digital garment design and manufacturing play an increasingly important role in apparel creation, with applications spanning garment modeling, virtual try-on, animation, and physically based clothing simulation (Pons-Moll et al., 2017; Bertiche et al., 2020; Patel et al., 2020; Ma et al., 2020; Tiwari & Bhowmick, 2021). Sewing patterns are a fundamental component of this process, specifying the panels and stitching required to construct garments. Traditionally, sewing patterns are derived from body measurements and garment design parameters by experienced professionals, or manually created in CAD systems. Such workflows demand substantial expertise and limit the accessibility and scalability of digital garment design. This makes multimodal sewing pattern generation valuable for both research and practical applications.

A generated pattern must first be triangulated into 2D panel meshes and then transformed, according to the predicted rotations and translations, into a BoxMesh that represents the garment as 2D panels arranged in 3D space. The resulting BoxMesh is finally coupled with an SMPL body model for physical simulation. Consequently, the practical usability of a sewing pattern depends not only on its similarity to the reference, but also on whether it can be successfully converted into a valid garment mesh and support subsequent physical simulation. Most existing methods are trained with reference-based supervision, encouraging the generated topology and geometry to match annotated sewing patterns (Jiang et al., 2020; Chi & Song, 2021; Xue et al., 2023; Zou et al., 2023; Nam et al., 2025). This creates a fundamental mismatch between supervised training objectives and downstream simulation requirements, as reference fidelity does not directly optimize garment constructability or simulation readiness.

In practice, this mismatch manifests in two critical failure modes along the downstream pipeline. *(i) Inconsistent stitch relations.* Missing or incompatible stitches can prevent corresponding panel boundaries from being assembled into a 3D digital garment. *(ii) Degenerate mesh geometry.* Even when stitching succeeds, distorted panel geometry can produce poorly shaped or degenerate mesh elements, compromising subsequent physical simulation.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

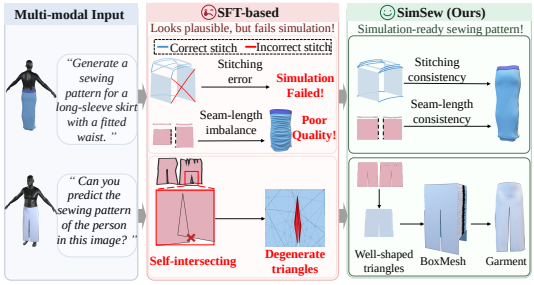

Figure 1: SFT-based methods can produce plausible sewing patterns that still fail downstream simulation due to inconsistent stitching or degenerate geometry. SimSew incorporates downstream simulation-aware reward to promote stitch consistency and simulation-ready garment meshes.

To address these challenges, we propose SimSew, a simulation-aware reinforcement-learning framework that improves downstream simulation success through dedicated reward design. Rather than optimizing only reference-level reconstruction fidelity, SimSew directly incorporates downstream simulation-aware reward into policy optimization. Specifically, the pattern-format reward first ensures that the generated sewing pattern is structurally complete and can be instantiated for downstream simulation. For stitching, each stitched boundary edge should be associated with a unique compatible counterpart, whereas conventional supervised training primarily encourages the predicted stitch relations to match their reference annotations without explicitly assessing their downstream compatibility. This can lead to missing, incorrect, or conflicting stitch assignments that prevent valid garment assembly. We therefore introduce a stitch-aware reward that evaluates stitch consistency and the geometric compatibility of paired seam boundaries, providing direct feedback on whether the predicted stitch relations can support downstream simulation. For geometry, small errors in predicted panel vertices can distort boundary edges or induce self-intersections, which may in turn produce poorly shaped or degenerate triangles during triangulation. We thus introduce a geometry-aware reward that jointly evaluates panel-level geometric fidelity and downstream mesh quality, encouraging geometrically valid panels while suppressing degeneracies that compromise triangulation and physical simulation.

Our main contributions are summarized as follows:

- We introduce SimSew, to our best knowledge the first RL framework that improves downstream simulation success through dedicated reward design, which directly incorporates downstream simulation-aware reward to bridge the gap between reference-based supervision and simulation readiness.
- We design a simulability-aware reward that explicitly targets the key requirements of downstream garment construction, including pattern-format consistency, one-to-one stitch assignment and seam compatibility, and panel geometry with mesh quality.
- Extensive experiments demonstrate that SimSew consistently improves downstream simulatability and reward-aligned pattern quality. On GCD-MM, SimSew increases the simulation success rate from 71.89% to 76.92%. Compared with the SFT baseline, it also improves All-edge-count Exact from 56.17% to 61.54% and One-to-one Stitch from 86.76% to 93.56%, while reducing Corner MAE from $18.97°$ to $10.66°$.

2 Related Work

2.1 Sewing Pattern Generation

Existing sewing pattern generation methods can be broadly categorized into regression-based, autoregressive, and diffusion-based paradigms. Regression-based methods directly predict panel geometry, transformations, and stitch relations from visual or 3D observations, typically through structured or hierarchical parameter regression (Korosteleva & Lee, 2022; Liu et al., 2023; Chen et al., 2024; Lim et al., 2024). Autoregressive methods serialize sewing patterns into structured tokens or programs and generate panels, edges, and stitch relations sequentially, enabling flexible conditioning on text, images, and editing instructions (He et al., 2024; Nakayama et al., 2025; Bian et al., 2025; Zhou et al., 2025a; Gui et al., 2025; Li et al., 2026b; Lin et al., 2026). Diffusion-based methods model the distribution of sewing patterns through iterative denoising in token, latent, or continuous representation spaces, improving generation diversity and global structural modeling (Liu et al., 2025c; Li et al., 2025; Nakayama et al., 2026).

Our work is complementary to these advances. Rather than introducing a new representation or generator, we apply RL to an SFT-trained sewing pattern model and incorporate downstream simulation-aware reward through pattern-format, stitch-aware, and geometry-aware rewards, explicitly optimizing simulatability beyond reference-based supervision.

2.2 Reinforcement Learning with 3D Generation

Reinforcement learning has been increasingly explored for 3D generation. Preference-based methods improve semantic alignment, perceptual quality, and geometric plausibility through human or learned feedback (Liu et al., 2025a;b; Leng et al., 2026; Ye et al., 2024; Zhou et al., 2025b; Li et al., 2026a; Zou et al., 2025), while structured-generation methods apply GRPO to executable CAD programs and autoregressive 3D representations (Guan et al., 2025; Yin et al., 2026; Tang et al., 2025; Zhao et al., 2025; Zhou et al., 2026). Together, these studies show that RL can improve 3D generation beyond supervised objectives (Pan & Liu, 2026; Adhikari et al., 2026; Zhao et al., 2026).

Existing RL methods for 3D generation have largely focused on perceptual, semantic, or geometric quality, rather than simulation-ready sewing patterns. SimSew instead designs simulation-aware rewards to directly optimize downstream garment simulation and simulation success through RL.

3 Preliminary

3.1 Problem Formulation

Given a multimodal condition x, which may consist of an image, a text description, an image–text pair, or an editing instruction, SimSew generates a sewing pattern $\mathcal{Y} = (\mathcal{T}, \mathcal{G})$, where $\mathcal{T}$ denotes the discrete topology and $\mathcal{G}$ the corresponding continuous geometry and transformations. Let $\mathcal{C}(\mathcal{Y})$ denote downstream garment simulation. Since reference fidelity alone does not guarantee successful simulation, SimSew trains a policy $\pi_\theta(\mathcal{Y} \mid x)$ to improve simulation-aware quality while preserving agreement with the reference pattern:

$$\theta^\star = \arg\max_\theta \mathbb{E}_{(x,\mathcal{Y}^\star)\sim\mathcal{D},\, \mathcal{Y}\sim\pi_\theta(\cdot|x)} \left[R(\mathcal{Y}, \mathcal{Y}^\star, \mathcal{C}(\mathcal{Y})) \right]. \tag{1}$$

Here, R denotes the simulation-aware reward.

3.2 SewFusion

We employ SewFusion Lin et al. (2026) as the SFT initialization. SewFusion models the heterogeneous sewing pattern representation by decomposing it into discrete topology and panel-level continuous geometry. Both are generated within the same causal multimodal Transformer: topology is generated through autoregressive next-token prediction, while continuous geometry is generated through flow matching in the latent space of a pretrained Panel Geometry VAE. Following the notation above, its conditional generation process is factorized as

$$p_{\theta_0,\psi}(\mathcal{Y} \mid x) = p_{\theta_0}(\mathcal{T} \mid x)\, p_{\theta_0,\psi}(\mathcal{G} \mid \mathcal{T}, x). \tag{2}$$

where θ_0 denotes the shared Transformer parameters of the SFT initialization and ψ denotes the parameters of the pretrained Panel Geometry VAE. SimSew initializes its policy from this SFT model and further optimizes it through reinforcement learning with simulation-aware rewards.

3.3 Group Relative Policy Optimization

Group Relative Policy Optimization (GRPO) (Shao et al., 2024) is a critic-free reinforcement learning algorithm derived from Proximal Policy Optimization (PPO) (Schulman et al., 2017). Unlike PPO, which relies on an additional value model to estimate advantages, GRPO obtains learning signals by comparing the rewards of multiple trajectories sampled for the same input.

Given an input condition x, GRPO samples a group of G trajectories $\tau_i{}_{i=1}^{G}$ from the current policy and evaluates each trajectory with a scalar reward R_i. The relative advantage of the i-th trajectory is computed by normalizing its reward within the group:

$$\widehat{A}_i = \frac{R_i - \mathrm{mean}\big(\{R_g\}_{g=1}^{G}\big)}{\mathrm{std}\big(\{R_g\}_{g=1}^{G}\big) + \epsilon}. \tag{3}$$

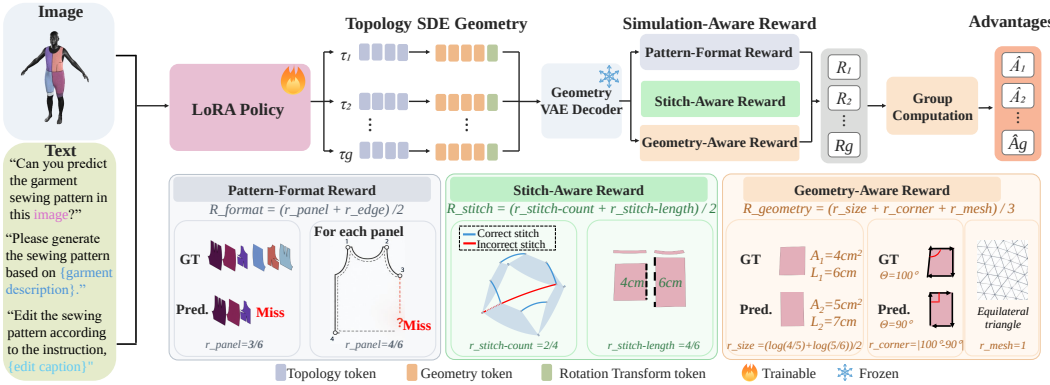

Figure 2: **Overview of SimSew.** Given multimodal garment conditions, the LoRA policy samples a group of trajectories, which are decoded into sewing patterns and evaluated by simulation-aware rewards. The reward combines pattern-format, stitch-aware, and geometry-aware terms to assess structural completeness, stitch consistency, panel geometry, and downstream mesh quality. Group-relative rewards are then normalized into advantages for GRPO-based policy optimization.

where ϵ is a small constant for numerical stability. The policy is then optimized using a PPO-style clipped surrogate objective,

$$J_{\text{GRPO}} = \frac{1}{G}\sum_{i=1}^{G}\left[\min\left(\rho_i \widehat{A}_i,\ \text{clip}\left(\rho_i, 1-\epsilon_c, 1+\epsilon_c\right)\widehat{A}_i\right) - \beta D_{\text{KL}}\left(\pi_\theta \| \pi_{\text{ref}}\right)\right], \tag{4}$$

where ρ_i denotes the policy ratio between the current and old policies, ϵ_c is the clipping threshold, and β controls the KL regularization toward the reference policy π_{ref}. By exploiting relative reward differences within each group, GRPO avoids the need for an explicit critic while providing stable learning signals for policy optimization.

4 Method

4.1 Overview

SimSew performs RL on a supervised sewing pattern generator, as illustrated in Figure 2. Starting from the supervised checkpoint, SimSew samples groups of trajectories over topology and geometry and evaluates them through downstream garment simulation. The resulting simulation-aware rewards are normalized within each group to obtain relative advantages, which jointly optimize the topology and geometry policies through a LoRA adapter. This enables direct optimization toward garment simulatability while retaining the structural and geometric priors learned during supervised training.

4.2 Simulability-Aware Reward Design

A sewing pattern can closely match the reference while still failing downstream garment simulation. This discrepancy arises because reference-level fidelity does not explicitly capture the structural and geometric conditions required by the garment-construction pipeline. In particular, incomplete panel structures can hinder pattern construction, inconsistent stitch relations can prevent valid garment assembly, and distorted panel geometry can produce poorly shaped or degenerate triangles during triangulation. Moreover, using only the final simulation outcome provides sparse supervision, as failed samples provide limited information about the source of failure.

Motivated by these observations, we design a simulability-aware reward that provides structured feedback at different stages of garment construction. Specifically, the pattern-format reward evaluates representation validity and structural completeness, the stitch-aware reward targets stitch consistency and seam compatibility, and the geometry-aware reward evaluates panel geometry together with downstream mesh quality. These complementary signals provide informative group-relative

feedback even when simulation fails, while directly encouraging patterns that are more suitable for downstream simulation.

We linearly combine the three rewards as

$$R_i = \lambda_{\text{fmt}} R_i^{\text{fmt}} + \lambda_{\text{stitch}} R_i^{\text{stitch}} + \lambda_{\text{geom}} R_i^{\text{geom}}, \tag{5}$$

where all components lie in $[0, 1]$, $\lambda_k \geq 0$, and $\sum_k \lambda_k = 1$. Unavailable components are assigned zero. For two non-negative counts, we use $a(m, n) = 1 - |m-n| / \max(m, n, 1)$, such that $a(0, 0) = 1$.

Pattern-Format Reward. The first requirement of downstream garment construction is that the generated pattern maintains a structurally complete panel representation. Missing or duplicated panels change the garment structure, while incorrect edge counts alter panel boundaries and may further affect stitching and triangulation. Therefore, we evaluate the consistency of both panel and edge counts with respect to the reference pattern.

Panel-Count Reward. Let $\widehat{\mathcal{P}}_i$ and $\mathcal{P}^\star$ denote the predicted and reference panel sets, with cardinalities $\widehat{P}_i$ and $P^\star$. We define panel-count consistency as

$$r_i^{\text{panel}} = a(\widehat{P}_i, P^\star). \tag{6}$$

Edge-Count Reward. After establishing panel correspondence $\mathcal{M}_i^{\text{panel}}$, edge-count consistency is evaluated over matched panel pairs:

$$r_i^{\text{edge}} = \frac{\sum_{(p,p^\star)\in\mathcal{M}_i^{\text{panel}}} a(\widehat{E}_{i,p}, E_{p^\star}^\star)}{D_i^{\text{panel}}}, \qquad D_i^{\text{panel}} = \max(\widehat{P}_i, P^\star, 1). \tag{7}$$

Unmatched panels contribute zero through the denominator, thereby penalizing both missing and redundant structures. The pattern-format reward is

$$R_i^{\text{fmt}} = \frac{1}{2}\left(r_i^{\text{panel}} + r_i^{\text{edge}}\right). \tag{8}$$

Detailed panel and edge correspondence procedures are provided in Appendix C.

Stitch-Aware Reward. A structurally valid pattern may still fail during garment assembly because the predicted stitch relations do not satisfy the constraints required by downstream stitching. In particular, each stitched boundary edge should be associated with a unique compatible counterpart. Missing, duplicated, or one-to-many assignments can therefore lead to a `StitchingError`. In addition, even structurally correct stitch pairs may be difficult to assemble when their seam lengths are highly inconsistent.

Stitch-Count Reward. To capture both aspects, we decompose the stitch-aware reward into stitch-count consistency and seam-length compatibility. Let $\widehat{\mathcal{S}}_i$ and $\mathcal{S}^\star$ denote the predicted and reference stitch sets, with cardinalities $\widehat{S}_i$ and $S^\star$. The stitch-count reward evaluates not only the number of predicted stitches but also whether each stitched edge participates in a unique reference-consistent pairing:

$$r_i^{\text{stitch-count}} = \begin{cases} 1, & \widehat{S}_i = S^\star = 0, \\ a(\widehat{S}_i, S^\star)\dfrac{\sum_{s\in\widehat{\mathcal{S}}_i} v_i(s)}{\max(\widehat{S}_i, 1)}, & \text{otherwise}, \end{cases} \tag{9}$$

where $v_i(s)$ indicates whether a predicted stitch satisfies one-to-one edge assignment and matches a valid reference pairing. The detailed definition of $v_i(s)$ and edge correspondence is given in Appendix C.

Seam-Length Reward. After mapping predicted stitches to reference identities, let $\mathcal{M}_i^{\text{stitch}}$ denote the matched stitch pairs. We further evaluate the compatibility of their seam-length imbalance:

$$r_i^{\text{stitch-length}} = \begin{cases} 1, & \widehat{S}_i = S^\star = 0, \\ \dfrac{\sum_{(\widehat{s},s^\star)\in\mathcal{M}_i^{\text{stitch}} \exp\left(-\frac{|\rho(\widehat{s})-\rho(s^\star)|}{\tau_{\text{seam}}}\right)}}{\max(\widehat{S}_i, S^\star, 1)}, & \text{otherwise}, \end{cases} \tag{10}$$

where $\rho(s)$ measures the length imbalance between the two stitched boundaries. Unmatched stitches contribute zero.

The final stitch-aware reward is

$$R_i^{\text{stitch}} = \frac{1}{2}\left(r_i^{\text{stitch-count}} + r_i^{\text{stitch-length}}\right). \tag{11}$$

This formulation jointly encourages structurally valid stitch assignments and geometrically compatible seam pairs.

Geometry-Aware Reward. Even when the topology and stitch relations are valid, small errors in panel geometry can be amplified during triangulation. Distorted boundaries, abnormal panel scales, or sharp local deformations may produce poorly shaped triangles or even a `DegenerateTrianglesError`, thereby compromising subsequent physical simulation. We therefore evaluate geometry at both the panel level and the resulting mesh level.

Size and Corner Reward. At the panel level, we use size and corner consistency. The size reward constrains panel area and perimeter, while the corner reward preserves local boundary structure. Their detailed distance definitions are provided in Appendix C. Each panel-level distance $d_{p,p^\star}^m$ is converted into a reward:

$$r_i^m = \frac{\sum_{(p,p^\star)\in\mathcal{M}_i^{\text{panel}}} \exp\left(-d_{p,p^\star}^m/\tau_m\right)}{D_i^{\text{panel}}}, \qquad m \in \{\text{size}, \text{corner}\}. \tag{12}$$

Non-finite distances contribute zero.

Mesh-Quality Reward. Panel-level geometric consistency alone does not directly indicate whether the resulting triangulation is well conditioned. We therefore additionally measure the quality of downstream mesh elements using the normalized mean-ratio quality Knupp (2001):

$$q_{i,f}^{(v)} = \text{clip}\left(\frac{4\sqrt{3}A_{i,f}^{(v)}}{\sum_{k=1}^{3}(l_{i,f,k}^{(v)})^2}, 0, 1\right), \tag{13}$$

where $A_{i,f}^{(v)}$ is the area of triangle f and $l_{i,f,k}^{(v)}$ denotes its three edge lengths. The metric equals one for an equilateral triangle and approaches zero as the triangle degenerates.

For the available panel-local, stitched, and rest-shape mesh views $\mathcal{V}_i$, we focus on lower-tail mesh quality:

$$R_i^{\text{mesh}} = \begin{cases} \min\limits_{v\in\mathcal{V}_i} Q_{0.05}\left(\{q_{i,f}^{(v)}\}_{f\in\mathcal{F}_i^{(v)}}\right), & \text{simulation succeeds,} \\ 0, & \text{otherwise.} \end{cases} \tag{14}$$

The fifth percentile emphasizes poorly shaped local regions, which are more likely to affect downstream simulation, while remaining more robust than a strict minimum.

Finally, the geometry-aware reward is

$$R_i^{\text{geom}} = \omega_{\text{size}} r_i^{\text{size}} + \omega_{\text{corner}} r_i^{\text{corner}} + \omega_{\text{mesh}} R_i^{\text{mesh}}, \tag{15}$$

where $\omega_k \geq 0$ and $\sum_k \omega_k = 1$. The contour terms provide dense feedback when Meshgen fails, while R_i^{mesh} directly penalizes poorly conditioned triangles after successful simulation.

Together, the three components supervise successive requirements of the garment simulation pipeline. The pattern-format reward promotes a valid and structurally complete sewing pattern representation, the stitch-aware reward promotes complete and reference-consistent seam relations, and the geometry-aware reward promotes triangulable panel geometry and well-conditioned meshes. Their linear combination provides dense feedback without routing candidates through failure-dependent reward branches.

5 Experiments

5.1 Experimental Setting

Datasets. Experiments are conducted on two sewing pattern datasets, SewFactory Liu et al. (2023) and GCD-MM Nakayama et al. (2025). SewFactory is a synthetic dataset with 8,269 valid gar-

Table 1: **Quantitative comparison with previous state-of-the-art methods on SewFactory and GCD-MM.**

Dataset	Methods	Setting	Panel Acc. (%)↑	Edge Acc. (%)↑	Stitch Acc. (%)↑	Curve Type Acc. (%)↑	Vertex L_2 ↓	Transl. L_2 ↓	Rot. L_2 ↓	Quad. Ctrl. L_2 ↓	Cubic Ctrl. L_2 ↓	Arc Ctrl. L_2 ↓	SSR (%)↑
SewFactory	SewFormer	*image*	89.60	99.06	81.06	90.53	4.5921	2.1657	0.0857	0.1935	–	–	86.96
	AIpparel		98.43	98.43	88.62	**99.88**	3.1488	1.5769	0.0275	0.0951	–	–	88.69
	SewFusion		**98.66**	99.14	90.88	**99.88**	2.6654	1.0465	**0.0231**	**0.0842**	–	–	88.72
	SimSew (Ours)		98.62	**99.17**	**91.59**	99.85	**2.4938**	**1.0024**	0.0249	0.0852	–	–	**89.86**
GCD-MM	SewFormer	*image*	68.96	47.60	36.29	45.30	12.5066	5.2676	0.0126	0.5365	0.5782	0.3209	42.34
	SewingLDM		76.75	72.26	71.58	68.48	6.0844	3.4855	0.4855	0.2497	0.2429	0.2012	56.28
	AIpparel		86.28	80.83	73.04	79.54	5.5278	2.9291	0.0084	0.2146	0.2451	0.1802	65.91
	SewFusion		92.33	85.64	81.74	82.23	4.1251	2.6933	**0.0043**	0.0896	0.1534	**0.0397**	71.89
	SimSew (Ours)		**93.25**	**87.17**	**83.52**	**84.01**	**3.6824**	**2.6563**	**0.0043**	**0.0859**	**0.1405**	0.0404	**76.92**
	SewingLDM	*text*	59.72	56.66	58.26	53.86	10.8272	8.0690	0.0316	0.4483	0.4635	0.3999	39.41
	AIpparel		61.03	56.66	54.58	55.82	7.8853	3.3294	0.0109	0.3763	0.3608	0.2767	53.65
	SewFusion		68.57	65.40	59.40	59.39	**6.8230**	**3.1491**	0.0062	0.1008	**0.0931**	0.0868	58.52
	SimSew (Ours)		**73.34**	**74.65**	**72.29**	**72.46**	7.1772	4.2009	**0.0059**	**0.0900**	0.1182	**0.0592**	**65.28**
	SewingLDM	*image-text*	77.43	75.22	76.58	58.08	6.4484	4.9755	0.0114	0.3370	0.3552	0.2959	57.05
	AIpparel		87.69	81.77	73.88	80.46	5.4643	2.8954	0.0083	0.2121	0.2423	0.1782	66.73
	SewFusion		**94.05**	**89.74**	**85.18**	**84.85**	**3.5687**	**2.5326**	**0.0043**	0.0897	**0.0506**	**0.0399**	73.24
	SimSew (Ours)		93.99	88.77	84.44	84.26	3.5850	2.7427	0.0055	**0.0859**	0.1385	0.0436	**77.49**
	AIpparel	*editing*	82.91	92.07	93.64	95.84	7.7974	3.0720	0.0088	0.2250	0.2571	0.1890	70.56
	SewFusion		**89.73**	**97.32**	**97.60**	**97.39**	**6.8309**	**2.1655**	**0.0023**	0.1135	**0.1301**	**0.0939**	71.27
	SimSew (Ours)		87.46	96.36	97.13	96.64	7.0308	2.3713	0.0025	**0.1130**	0.1321	0.0945	**73.11**

ments after preprocessing, using a 90:5:5 train/validation/test split. GCD-MM extends GarmentCodeData Korosteleva et al. (2024) with multimodal annotations and contains approximately 115K garments. Following AIpparel Nakayama et al. (2025), we adopt a 90:5:5 train/validation/test split.

Evaluation Metrics. Generation quality is evaluated at three levels: discrete structure using panel-count, edge-count, stitch, and curve-type accuracies; continuous geometry using Vertex, Quadratic Control, Cubic Control, Arc Control, Translation, and sign-invariant Rotation L2 errors; and downstream executability using Simulation Success Rate (SSR).

Implementation Details. We adopt SewFusion Lin et al. (2026) as the SFT initialization. During RL, we optimize a LoRA adapter with GRPO. For each prompt, GRPO samples 8 trajectories. Training runs for 1,000 optimizer steps on 8 H100 GPUs with a batch size of 32. Full optimization and sampling hyperparameters are provided in the appendix.

5.2 Main Results

Baseline. SimSew is compared with state-of-the-art **open-source** sewing pattern generation methods, including SewFormer Liu et al. (2023), the diffusion-based SewingLDM Liu et al. (2025c), autoregressive-based AIpparel Nakayama et al. (2025), and the SewFusion Lin et al. (2026).

Quantitative Results Table 1 reports detailed quantitative comparisons with the baselines on SewFactory and GCD-MM. Since SewFactory provides only image-conditioned inputs, we evaluate all methods under the image-conditioned setting on this dataset. Moreover, SewFactory represents all edges as single-control-point Bézier curves without explicit cubic or arc annotations; therefore, the corresponding Cubic Control and Arc Control metrics are omitted.

Overall, SimSew consistently improves downstream simulation success across datasets and input settings. On SewFactory, SimSew achieves an SSR of 89.86%, outperforming SewFusion by 1.14 percentage points while maintaining comparable or better pattern-level accuracy and geometric quality. The improvement is more pronounced on GCD-MM, where SimSew increases SSR from 71.89% to 76.92% under image conditioning, from 58.52% to 65.28% under text conditioning, from 73.24% to 77.49% under image–text conditioning, and from 71.27% to 73.11% for editing.

Notably, the gains in simulation success do not always coincide with improvements in every reference-based metric. For example, under the image–text and editing settings, SewFusion retains advantages on several topology and geometry metrics, whereas SimSew achieves substantially higher SSR. This observation supports our motivation that reference-level reconstruction fidelity alone does not fully reflect downstream simulation readiness. By incorporating simulation-aware rewards during RL post-training, SimSew improves garment simulatability while largely preserving the structural and geometric quality learned from supervised training.

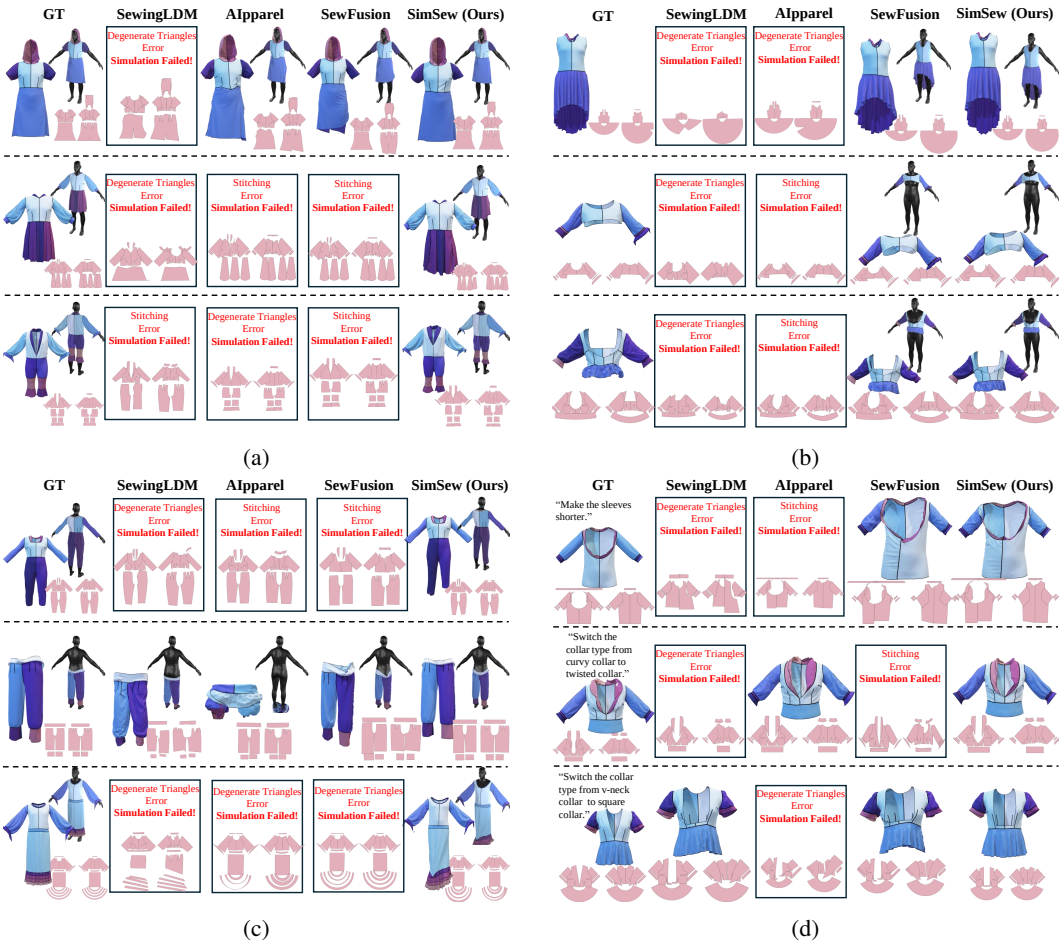

Figure 3: **Qualitative comparison on multimodal sewing pattern generation and editing.**(a) Image-conditioned generation, (b) text-conditioned generation, (c) image–text-conditioned generation, and (d) sewing pattern editing.

Qualitative Results. As shown in Figure 3, SFT-based methods frequently exhibit incorrect stitch pairings, missing panels, or distorted panel geometry, leading to `StitchingError` or `DegenerateTrianglesError` during garment simulation. SimSew effectively corrects stitch assignments and recovers missing panels, yielding more complete and structurally consistent sewing patterns. Moreover, the generated panel geometry produces better-shaped triangles and higher-quality garment meshes after triangulation. These examples demonstrate that SimSew improves both pattern-level correctness and downstream constructability, leading to more reliable physical simulation. Additional qualitative results are provided in the Appendix D.1.

Reward-Aligned Evaluation To examine whether reinforcement learning improves the properties targeted by our reward design, we compare SimSew with its SFT initialization using seven reward-aligned metrics in Fig. 4. SimSew consistently improves all four structural metrics: Panel-count Exact increases from 92.34% to 93.27% (+0.93 points), All-edge-count Exact from 56.17% to 61.54% (+5.37 points), One-to-one Stitch from 86.76% to 93.56% (+6.80 points), and Valid-pair F1 from 89.69% to 91.11% (+1.42 points). The larger gains in edge-count correctness and one-to-one stitching indicate that RL is particularly effective at enforcing discrete structural and sewing constraints that remain difficult for SFT alone. SimSew also reduces Size Distance by 9.2% (0.0994 to 0.0902) and Corner MAE by 43.8% (18.97° to 10.66°), demonstrating improved geometric fidelity. However, Absolute ρ P95 increases from 0.504 to 0.649 (+28.9%; lower is better), revealing a remaining long-tail seam-length imbalance. Overall, RL produces broad but non-uniform im-

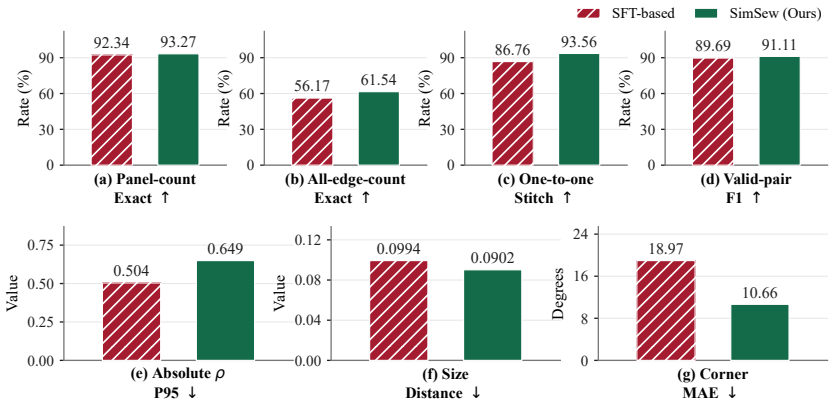

Figure 4: Reward-aligned comparison between SFT-based generation and SimSew.

Table 2: Ablation study of the proposed reward components.

Method	Panel Acc. (%)↑	Edge Acc. (%)↑	Stitch Acc. (%)↑	Curve Type Acc. (%)↑	Vertex L_2 ↓	Transl. L_2 ↓	Rot. L_2 ↓	Quad. Ctrl. L_2 ↓	Cubic Ctrl. L_2 ↓	Arc Ctrl. L_2 ↓	SSR (%)↑
Full	93.52	87.22	83.87	84.25	3.7976	2.7274	0.0053	0.0871	0.1379	0.0421	76.13
Pattern-Format Reward											
w/o Pattern-Format Reward	92.95	87.18	83.69	84.07	3.7757	2.7098	0.0053	0.0875	0.1486	0.0425	75.91
w/o panel	93.06	87.29	83.93	84.26	3.7707	2.7286	0.0053	0.0870	0.1464	0.0425	75.96
w/o edge	93.55	86.79	83.72	84.18	3.7884	2.7558	0.0052	0.0865	0.1474	0.0421	75.59
Stitch-Aware Reward											
w/o Stitch-Aware Reward	93.37	87.07	83.88	84.22	3.7596	2.8208	0.0054	0.0851	0.1363	0.0425	73.60
w/o stitch-count	93.46	87.42	82.74	84.20	3.7745	2.7331	0.0052	0.0879	0.1518	0.0426	73.63
w/o stitch-length	93.20	87.23	83.79	84.18	3.7980	2.8993	0.0056	0.0874	0.1366	0.0430	75.81
Geometry-Aware Reward											
w/o Geometry-Aware Reward	93.37	86.89	83.54	84.10	3.7989	2.7478	0.0054	0.0860	0.1471	0.0425	73.49
w/o size	93.46	87.23	83.78	84.16	3.7788	2.7603	0.0055	0.0898	0.1510	0.0423	75.20
w/o corner	93.37	87.29	83.84	84.21	3.7779	2.6590	0.0050	0.0859	0.1391	0.0426	75.57
w/o mesh	93.37	87.22	83.75	84.10	3.8197	2.6829	0.0053	0.0898	0.1530	0.0427	73.54

provements: it substantially strengthens structural validity and geometric accuracy, while extreme seam-length compatibility remains an open limitation.

5.3 Ablation Studies

Ablation of Rewards. We ablate the three reward components in Eqs 5. by removing each component wherever it contributes to the terminal reward.

Influence of Sampling Temperature. The sampling temperature T controls rollout exploration, with higher values increasing the diversity of sampled topology structures. We evaluate $T \in \{0.1, 0.3, 0.5, 0.7, 0.9, 1.0, 1.1\}$ while fixing the group size to $G = 8$.

Influence of Group Size. The group size G controls rollout diversity for group-relative advantage estimation, with larger groups providing richer comparisons at higher simulation cost. We evaluate $G \in \{2, 4, 6, 8, 10, 12\}$ at $T = 0.9$; the effects of G and sampling temperature are shown Appendix D.1, with detailed topology and geometry metrics reported in.

6 Conclusion

We presented SimSew, a simulation-aware RL framework for simulation-ready sewing pattern generation. By incorporating downstream simulation-aware rewards, SimSew bridges the gap between reference-level fidelity and garment simulatability. Experiments on SewFactory and GCD-MM show that SimSew consistently improves simulation success while largely preserving sewing pattern quality, highlighting the importance of downstream physical feedback for structured garment generation.
Limitation. Simulation success can be further improved, especially for challenging garment structures. Moreover, validation on real-world garment data remains limited.

AI USE STATEMENT

In this work, we used generative AI tools for language editing and improving the clarity and readability of the manuscript, as well as for limited assistance with code debugging and implementation. We did not use generative AI tools to generate experimental data, fabricate or modify results, or independently formulate scientific claims or conclusions. Generative AI was not used in place of human judgment in the design, evaluation, or interpretation of the proposed method. All AI-assisted text and code were carefully reviewed, revised, and verified by the authors. In particular, AI-assisted code was tested for correctness, and all technical statements, experimental results, and conclusions were independently checked by the authors. We take full responsibility for the final content of this work, including any text, code, claims, or other artifacts produced with the aid of generative AI.

REFERENCES

Saugat Adhikari, Ashok Prasad Neupane, Pramish Paudel, Ajad Chhatkuli, and Danda Pani Paudel. iarcs: Iterative agentic rl for controllable 3d scene generation. *arXiv preprint arXiv:2608.06161*, 2026.

Hugo Bertiche, Meysam Madadi, and Sergio Escalera. CLOTH3D: Clothed 3D humans. In *Computer Vision – ECCV 2020*, pp. 344–359. Springer, 2020. doi: 10.1007/978-3-030-58565-5_21.

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 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2924–2934, 2025. URL https://openaccess.thecvf.com/content/CVPR2025/html/Bian_ChatGarment_Garment_Estimation_Generation_and_Editing_via_Large_Language_Models_CVPR_2025_paper.html.

Cheng-Hsiu Chen, Jheng-Wei Su, Min-Chun Hu, Chih-Yuan Yao, and Hung-Kuo Chu. Panelformer: Sewing pattern reconstruction from 2d garment images. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pp. 454–463, 2024.

Cheng Chi and Shuran Song. GarmentNets: Category-level pose estimation for garments via canonical space shape completion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 3324–3333, 2021. doi: 10.1109/ICCV48922.2021.00331.

Yandong Guan, Xilin Wang, Ximing Xing, Jing Zhang, Dong Xu, and Qian Yu. CAD-Coder: Text-to-CAD generation with chain-of-thought and geometric reward. In *Advances in Neural Information Processing Systems*, volume 38, pp. 59765–59789, 2025. URL https://arxiv.org/abs/2505.19713.

Hongquan Gui, Zhanpeng Yang, Arjun Rachana Harish, Cheng Ren, Yishu Yang, and Ming Li. Genpattern: Dual-graph enhanced sewing pattern generation via multimodal large language model. *Journal of Manufacturing Systems*, 83:822–838, 2025. doi: 10.1016/j.jmsy.2025.11.005.

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*, 2024. URL https://arxiv.org/abs/2401.16465.

Boyi Jiang, Juyong Zhang, Yang Hong, Jinhao Luo, Ligang Liu, and Hujun Bao. BCNet: Learning body and cloth shape from a single image. In *Computer Vision – ECCV 2020*, pp. 18–35. Springer, 2020.

Patrick M. Knupp. Algebraic mesh quality metrics. *SIAM Journal on Scientific Computing*, 23(1): 193–218, 2001. doi: 10.1137/S1064827500371499.

Maria Korosteleva and Sung-Hee Lee. NeuralTailor: Reconstructing sewing pattern structures from 3d point clouds of garments. *ACM Transactions on Graphics*, 2022. URL https://arxiv.org/abs/2201.13063.

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. *arXiv preprint arXiv:2405.17609*, 2024. URL https://arxiv.org/abs/2405.17609.

Jiaqi Leng, Shuyuan Tu, Haidong Cao, Sicheng Xie, Daoguo Dong, Zuxuan Wu, and Yu-Gang Jiang. Preference score distillation: Leveraging 2d rewards to align text-to-3d generation with human preference. *arXiv preprint arXiv:2603.01594*, 2026.

Xinyu Li, Qi Yao, and Yuanda Wang. GarmentDiffusion: 3d garment sewing pattern generation with multimodal diffusion transformers. In *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence*, 2025. URL https://arxiv.org/abs/2504.21476.

Zejian Li, Jiarui Ma, Han Xu, Weiting Zheng, Yangrui Zhu, Chenye Meng, Pei Chen, Ling Yang, Zhiyuan Yang, Changyuan Yang, Guang Yang, Immanuel Koh, and Lingyun Sun. Circular-dpo: Aligning multi-stage 3d generative models via preference feedback loop. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 32592–32601, 2026a.

Zhenyang Li, Lutao Jiang, Yizhou Zhao, Ying-Cong Chen, Xin Wang, Weikai Chen, and Yifan Peng. Patterngsl: A structured specification language for template-free and simulation-ready 3d garments. *ACM Transactions on Graphics*, 2026b. SIGGRAPH 2026.

Seungchan Lim, Sumin Kim, and Sung-Hee Lee. Spnet: Estimating garment sewing patterns from a single image of a posed user. In *Eurographics 2024 – Short Papers*. The Eurographics Association, 2024. doi: 10.2312/egs.20241035.

Jiaxin Lin, Xiao Pan, Hangjie Yuan, Luyan Liang, Wan Li, and Daquan Feng. Sewfusion: Tailored generation of topology and panel-level geometry for sewing patterns, 2026. URL https://arxiv.org/abs/2609.23548.

Fangfu Liu, Junliang Ye, Yikai Wang, Hanyang Wang, Zhengyi Wang, Jun Zhu, and Yueqi Duan. Dreamreward-x: Boosting high-quality 3d generation with human preference alignment. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2025a. doi: 10.1109/TPAMI.2025.3609680.

Lijuan Liu, Xiangyu Xu, Zhijie Lin, Jiabin Liang, and Shuicheng Yan. Towards garment sewing pattern reconstruction from a single image. *ACM Transactions on Graphics*, 2023. doi: 10.1145/3618319. URL https://arxiv.org/abs/2311.04218.

Qingming Liu, Zhen Liu, Dinghuai Zhang, and Kui Jia. Nabla-r2d3: Effective and efficient 3d diffusion alignment with 2d rewards. In *Advances in Neural Information Processing Systems*, volume 38, 2025b.

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*, pp. 17640–17650, 2025c.

Qianli Ma, Jinlong Yang, Anurag Ranjan, Sergi Pujades, Gerard Pons-Moll, Siyu Tang, and Michael J. Black. Learning to dress 3D people in generative clothing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6469–6478, 2020.

Kiyohiro Nakayama, Jan Ackermann, Timur Levent Kesdogan, Yang Zheng, Maria Korosteleva, Olga Sorkine-Hornung, Leonidas J. Guibas, Guandao Yang, and Gordon Wetzstein. AIpparel: A multimodal foundation model for digital garments. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 8138–8149, 2025. URL https://openaccess.thecvf.com/content/CVPR2025/html/Nakayama_AIpparel_A_Multimodal_Foundation_Model_for_Digital_Garments_CVPR_2025_paper.html.

Kiyohiro Nakayama, I-Chao Shen, Ruofan Liu, Yiming Wang, Gordon Wetzstein, and Takeo Igarashi. Garment particles: A 2d–3d symmetric garment representation for generation and editing. In *SIGGRAPH Conference Papers '26*, 2026. doi: 10.1145/3799902.3811102.

Hyeongjin Nam, Donghwan Kim, Jeongtaek Oh, and Kyoung Mu Lee. DeClotH: Decomposable 3D cloth and human body reconstruction from a single image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5636–5645, 2025. doi: 10.1109/CVPR52734.2025.00529.

Zhenyu Pan and Han Liu. Metaspatial: Reinforcing 3d spatial reasoning in vlms for the metaverse. In *The Fourteenth International Conference on Learning Representations*, 2026.

Chaitanya Patel, Zhouyingcheng Liao, and Gerard Pons-Moll. TailorNet: Predicting clothing in 3D as a function of human pose, shape and garment style. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7365–7375, 2020.

Gerard Pons-Moll, Sergi Pujades, Sonny Hu, and Michael J. Black. ClothCap: Seamless 4D clothing capture and retargeting. *ACM Transactions on Graphics*, 36(4):73:1–73:15, 2017. doi: 10.1145/3072959.3073711.

John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms, 2017. URL https://arxiv.org/abs/1707.06347.

Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo. DeepSeekMath: Pushing the limits of mathematical reasoning in open language models, 2024. URL https://arxiv.org/abs/2402.03300.

Yiwen Tang, Zoey Guo, Kaixin Zhu, Ray Zhang, Qizhi Chen, Dongzhi Jiang, Junli Liu, Bohan Zeng, Haoming Song, Delin Qu, Tianyi Bai, Dan Xu, Wentao Zhang, and Bin Zhao. Are we ready for RL in text-to-3D generation? a progressive investigation. *arXiv preprint arXiv:2512.10949*, 2025. doi: 10.48550/arXiv.2512.10949. URL https://arxiv.org/abs/2512.10949.

Lokender Tiwari and Brojeshwar Bhowmick. DeepDraper: Fast and accurate 3D garment draping over a 3D human body. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, pp. 1416–1426, 2021.

Han Xue, Wenqiang Xu, Jieyi Zhang, Tutian Tang, Yutong Li, Wenxin Du, Ruolin Ye, and Cewu Lu. GarmentTracking: Category-level garment pose tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 21233–21242, 2023.

Junliang Ye, Fangfu Liu, Qixiu Li, Zhengyi Wang, Yikai Wang, Xinzhou Wang, Yueqi Duan, and Jun Zhu. DreamReward: Text-to-3D generation with human preference. *arXiv preprint arXiv:2403.14613*, 2024. doi: 10.48550/arXiv.2403.14613. URL https://arxiv.org/abs/2403.14613.

Xiaolong Yin, Xingyu Lu, Jiahang Shen, Jingzhe Ni, Hailong Li, Ruofeng Tong, Min Tang, and Peng Du. Rlcad: Reinforcement learning training gym for revolution involved cad command sequence generation. *Computer-Aided Design*, 192:104027, 2026. doi: 10.1016/j.cad.2025.104027.

An Zhao, Shengyuan Zhang, Zejian Li, Ling Yang, Pei Chen, Jiale Wu, Haoran Xu, Anyang Wei, Perry Pengyun Gu, and Lingyun Sun. Diffusion distillation with direct preference optimization for efficient 3d lidar scene completion. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pp. 13070–13078, 2026. doi: 10.1609/aaai.v40i15.38307.

Ruowen Zhao, Junliang Ye, Zhengyi Wang, Guangce Liu, Yiwen Chen, Yikai Wang, and Jun Zhu. Deepmesh: Auto-regressive artist-mesh creation with reinforcement learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 10612–10623, 2025.

Feng Zhou, Ruiyang Liu, Chen Liu, Gaofeng He, Yong-Lu Li, Xiaogang Jin, and Huamin Wang. Design2garmentcode: Turning design concepts to tangible garments through program synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 23712–23722, 2025a.

Zhen Zhou, Jian Liu, Biwen Lei, Jing Xu, Haohan Weng, Yiling Zhu, Zhuo Chen, Junfeng Fan, Yunkai Ma, Dazhao Du, Song Guo, Fengshui Jing, and Chunchao Guo. Mesh-pro: Asynchronous advantage-guided ranking preference optimization for artist-style quadrilateral mesh generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 34248–34258, 2026.

Zhenglin Zhou, Xiaobo Xia, Fan Ma, Hehe Fan, Yi Yang, and Tat-Seng Chua. DreamDPO: Aligning text-to-3D generation with human preferences via direct preference optimization. *arXiv preprint arXiv:2502.04370*, 2025b. doi: 10.48550/arXiv.2502.04370. URL https://arxiv.org/abs/2502.04370.

Xiandong Zou, Ruihao Xia, Hongsong Wang, and Pan Zhou. DreamCS: Geometry-aware text-to-3D generation with unpaired 3D reward supervision. *arXiv preprint arXiv:2506.09814*, 2025. doi: 10.48550/arXiv.2506.09814. URL https://arxiv.org/abs/2506.09814.

Xingxing Zou, Xintong Han, and Waikeung Wong. CLOTH4D: A dataset for clothed human reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 12847–12857, 2023. doi: 10.1109/CVPR52729.2023.01235.

APPENDIX

A ADDITIONAL IMPLEMENTATION DETAILS

During RL, we optimize a LoRA adapter with rank 16, scaling factor 32, and dropout 0.05. Geometry generation uses 12 flow steps, a 2-step stochastic sampling window, and a noise scale of 0.2. The clipping ranges for the token and flow objectives are set to 0.2 and 0.05, respectively, with a reference-policy KL coefficient of 0.01. Each GPU processes 2 prompts per step, and gradients are accumulated over 2 steps, yielding an effective batch size of 32 prompts. The learning rate is set to 2×10^{-7}.

B REWARD IMPLEMENTATION DETAILS

C REWARD IMPLEMENTATION DETAILS

This section provides additional implementation details for the three simulation-aware rewards, including panel correspondence, stitch matching, panel-level geometric distances, and reward normalization.

C.1 PATTERN-FORMAT REWARD

Let $\widehat{\mathcal{P}}_i$ and $\mathcal{P}^\star$ denote the predicted and reference panel sets, with cardinalities $\widehat{P}_i$ and $P^\star$. Panel-count consistency is defined as

$$r_i^{\text{panel}} = a(\widehat{P}_i, P^\star), \tag{16}$$

where

$$a(m, n) = 1 - \frac{|m - n|}{\max(m, n, 1)}. \tag{17}$$

Let $\mathcal{M}_i^{\text{panel}}$ denote the correspondence between predicted and reference panels. Edge-count consistency is

$$r_i^{\text{edge}} = \frac{\sum_{(p,p^\star)\in\mathcal{M}_i^{\text{panel}}} a(\widehat{E}_{i,p}, E_{p^\star}^\star)}{D_i^{\text{panel}}}, \tag{18}$$

where

$$D_i^{\text{panel}} = \max(\widehat{P}_i, P^\star, 1). \tag{19}$$

Unmatched panels therefore contribute zero through the denominator.

The resulting pattern-format reward is

$$R_i^{\text{fmt}} = \frac{1}{2}\left(r_i^{\text{panel}} + r_i^{\text{edge}}\right). \tag{20}$$

C.2 STITCH-AWARE REWARD

Let $\widehat{\mathcal{S}}_i$ and $\mathcal{S}^\star$ denote the predicted and reference stitch sets, with cardinalities $\widehat{S}_i$ and $S^\star$. For each predicted boundary edge e, its stitch multiplicity is

$$c_i(e) = \sum_{s \in \widehat{\mathcal{S}}_i} \mathbb{I}[e \in s]. \tag{21}$$

A valid stitched boundary edge should participate in exactly one stitch.

For a predicted stitch $s = (e_1, e_2)$, let $\mu_i(s)$ denote its reference stitch after edge correspondence. We define

$$v_i(s) = \mathbb{I}[c_i(e_1) = 1]\,\mathbb{I}[c_i(e_2) = 1]\,\mathbb{I}[\mu_i(s) \in \mathcal{S}^\star], \tag{22}$$

which indicates whether the predicted stitch satisfies one-to-one edge assignment and matches a valid reference pairing.

The stitch-count reward is

$$r_i^{\text{stitch-count}} = \begin{cases} 1, & \widehat{S}_i = S^\star = 0, \\ a(\widehat{S}_i, S^\star) \dfrac{\sum_{s \in \widehat{\mathcal{S}}_i} v_i(s)}{\max(\widehat{S}_i, 1)}, & \text{otherwise.} \end{cases} \tag{23}$$

The count-similarity term penalizes missing or excessive stitches, while $v_i(s)$ suppresses one-to-many and incorrect stitch assignments.

For a stitch $s = (e_1, e_2)$ with boundary-edge lengths ℓ_1 and ℓ_2, we define its seam-length imbalance as

$$\rho(s) = \left| \log \frac{\max(\ell_1, \varepsilon_\ell)}{\max(\ell_2, \varepsilon_\ell)} \right|, \tag{24}$$

where ε_ℓ is a small constant for numerical stability.

After mapping predicted stitches to reference identities, let $\mathcal{M}_i^{\text{stitch}}$ denote the matched stitch pairs. The reference-relative stitch-length reward is

$$r_i^{\text{stitch-length}} = \begin{cases} 1, & \widehat{S}_i = S^\star = 0, \\ \dfrac{\sum_{(\widehat{s}, s^\star) \in \mathcal{M}_i^{\text{stitch}}} \exp\left(-\frac{|\rho(\widehat{s}) - \rho(s^\star)|}{\tau_{\text{seam}}}\right)}{\max(\widehat{S}_i, S^\star, 1)}, & \text{otherwise.} \end{cases} \tag{25}$$

Unmatched stitches contribute zero through the denominator.

The final stitch-aware reward is

$$R_i^{\text{stitch}} = \frac{1}{2}\left(r_i^{\text{stitch-count}} + r_i^{\text{stitch-length}}\right). \tag{26}$$

C.3 GEOMETRY-AWARE REWARD

For each matched panel pair $(p, p^\star)$, we evaluate complementary size and corner consistency.

Size Distance. Let $\widehat{A}_{i,p}$ and $A^\star_{p^\star}$ denote panel areas, and let $\widehat{L}_{i,p}$ and $L^\star_{p^\star}$ denote panel perimeters. We define

$$d_{p,p^\star}^{\text{size}} = \frac{1}{2}\left(\left|\log \frac{\max(\widehat{A}_{i,p}, \varepsilon)}{\max(A^\star_{p^\star}, \varepsilon)}\right| + \left|\log \frac{\max(\widehat{L}_{i,p}, \varepsilon)}{\max(L^\star_{p^\star}, \varepsilon)}\right|\right). \tag{27}$$

Corner Distance. Let $\mathcal{M}_{p,p^\star}^{\text{edge}}$ denote the edge correspondence between a matched panel pair, and let θ denote the turning angle between adjacent boundary tangents. The corner distance is

$$d_{p,p^\star}^{\text{corner}} = \frac{\sum_{(e,e^\star) \in \mathcal{M}_{p,p^\star}^{\text{edge}}} |\widehat{\theta}_{i,p,e} - \theta^\star_{p^\star,e^\star}|}{|\mathcal{M}_{p,p^\star}^{\text{edge}}|}. \tag{28}$$

Table 3: Ablation study on sampling temperature.

Sampling Temp.	Panel Acc.↑	Edge Acc.↑	Stitch Acc.↑	Curve Type Acc.↑	Vertex $L_2\downarrow$	Transl. $L_2\downarrow$	Rot. $L_2\downarrow$	Curve Ctrl. $L_2\downarrow$	Cubic Ctrl. $L_2\downarrow$	Arc Ctrl. $L_2\downarrow$
0.1	93.63%	87.12%	83.58%	84.16%	3.7924	2.5357	0.0055	0.0865	0.1385	0.0429
0.3	92.94%	86.99%	83.89%	84.37%	3.7674	2.6409	0.0052	0.0892	0.1435	0.0425
0.5	92.60%	86.93%	83.64%	84.38%	3.7840	2.6381	0.0051	0.0850	0.1437	0.0424
0.7	93.37%	87.08%	83.66%	84.13%	3.7775	2.7715	0.0051	0.0896	0.1539	0.0424
0.9	93.52%	87.22%	83.87%	84.25%	3.7976	2.7274	0.0053	0.0871	0.1379	0.0421
1.0	93.16%	**87.32%**	83.72%	84.06%	3.7714	2.7539	0.0051	0.0850	0.1448	0.0425
1.1	93.55%	87.18%	83.69%	84.07%	3.7757	2.7098	0.0053	0.0875	0.1486	0.0425

Each panel-level distance is converted into a reward:

$$r_i^m = \frac{\sum_{(p,p^\star)\in\mathcal{M}_i^{\text{panel}}} \exp\left(-d_{p,p^\star}^m/\tau_m\right)}{D_i^{\text{panel}}}, \qquad m \in \{\text{size}, \text{corner}\}. \tag{29}$$

Non-finite distances contribute zero. The size term constrains panel area and perimeter, while the corner term preserves local boundary structure.

Mesh Quality. For completeness, the normalized mean-ratio triangle quality used in the main text is

$$q_{i,f}^{(v)} = \text{clip}\left(\frac{4\sqrt{3}A_{i,f}^{(v)}}{\sum_{k=1}^{3}(l_{i,f,k}^{(v)})^2}, 0, 1\right). \tag{30}$$

For the available mesh views $\mathcal{V}_i$, the mesh-quality reward is

$$R_i^{\text{mesh}} = \begin{cases} \min\limits_{v\in\mathcal{V}_i} Q_{0.05}\left(\{q_{i,f}^{(v)}\}_{f\in\mathcal{F}_i^{(v)}}\right), & \text{simulation succeeds,} \\ 0, & \text{otherwise.} \end{cases} \tag{31}$$

The fifth percentile emphasizes lower-tail triangle quality and is less sensitive to isolated numerical outliers than a strict minimum.

Finally,

$$R_i^{\text{geom}} = \omega_{\text{size}} r_i^{\text{size}} + \omega_{\text{corner}} r_i^{\text{corner}} + \omega_{\text{mesh}} R_i^{\text{mesh}}, \tag{32}$$

where $\omega_k \geq 0$ and $\sum_k \omega_k = 1$.

C.4 Reward-Scale Calibration

The exponential distance rewards use independently calibrated scales. For

$$m \in \{\text{seam}, \text{size}, \text{corner}\}, \tag{33}$$

we collect the finite positive distances

$$\mathcal{D}_m^+ = \left\{d \in \mathcal{D}_m : d > 10^{-12},\ d < \infty\right\}. \tag{34}$$

The scale is defined as

$$\tau_m = \max\left(\frac{\gamma\ \text{median}(\mathcal{D}_m^+)}{\log 2}, 10^{-8}\right), \tag{35}$$

where γ is a shared scale multiplier. With $\gamma = 1$, the median positive distance is mapped to a reward of 0.5.

D More Results

D.1 Additional Hyperparameter Sensitivity Analysis

We provide detailed topology and geometry metrics for the sampling temperature and GRPO group-size studies. These results complement the simulation success rates reported in Fig. **??**. In each experiment, we vary only one hyperparameter while keeping the remaining training and evaluation settings fixed.

Table 4: Ablation study on group size.

Group Size	Panel Acc.↑	Edge Acc.↑	Stitch Acc.↑	Curve Type Acc.↑	Vertex L_2 ↓	Transl. L_2 ↓	Rot. L_2 ↓	Curve Ctrl. L_2 ↓	Cubic Ctrl. L_2 ↓	Arc Ctrl. L_2 ↓
2	91.29%	85.27%	83.90%	84.19%	3.7912	2.7914	0.0053	0.0891	0.1432	0.0408
4	93.12%	87.12%	83.76%	84.22%	3.7789	2.7880	0.0056	0.0902	0.1538	0.0428
6	93.33%	87.27%	83.67%	84.04%	3.7878	2.6709	0.0050	0.0858	0.1319	0.0409
8	93.52%	87.22%	83.87%	84.25%	3.7976	2.7274	0.0053	0.0871	0.1379	0.0421
10	92.86%	85.89%	82.37%	82.69%	4.1900	2.7376	0.0058	0.0976	0.1419	0.0456
12	92.98%	85.14%	81.85%	82.07%	4.1583	2.6936	0.0056	0.0901	0.1436	0.0434

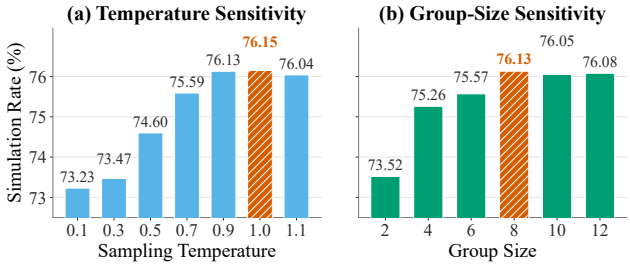

Figure 5: Influence of sampling temperature and group size on simulation success rate.

Sampling temperature. Table 3 reports the detailed results for different sampling temperatures. The topology and geometric fidelity metrics remain relatively stable over $T \in [0.1, 1.0]$, and no single temperature dominates every metric. Lower temperatures slightly favor panel accuracy and translation error, whereas moderate-to-high temperatures achieve competitive edge, stitch, and curve-control performance. Combined with the simulation results in Fig. 5, which improve from 73.23% at $T = 0.1$ to 76.15% at $T = 1.0$ and slightly decrease at $T = 1.1$, these results indicate that increased sampling diversity improves construction success without substantially compromising pattern fidelity.

Group size. Table 4 reports the detailed results for different GRPO group sizes. Increasing the group size from 2 to 6–8 generally improves structural accuracy and simulation success. In contrast, larger groups ($G = 10$ and $G = 12$) provide no additional simulation benefit and noticeably reduce edge, stitch, and curve-type accuracy while increasing several geometric errors. Together with the highest simulation success rate obtained at $G = 8$, these results support $G = 8$ as the best trade-off between performance and computational cost.



Figure 6: Qualitative cases of image-based generation.

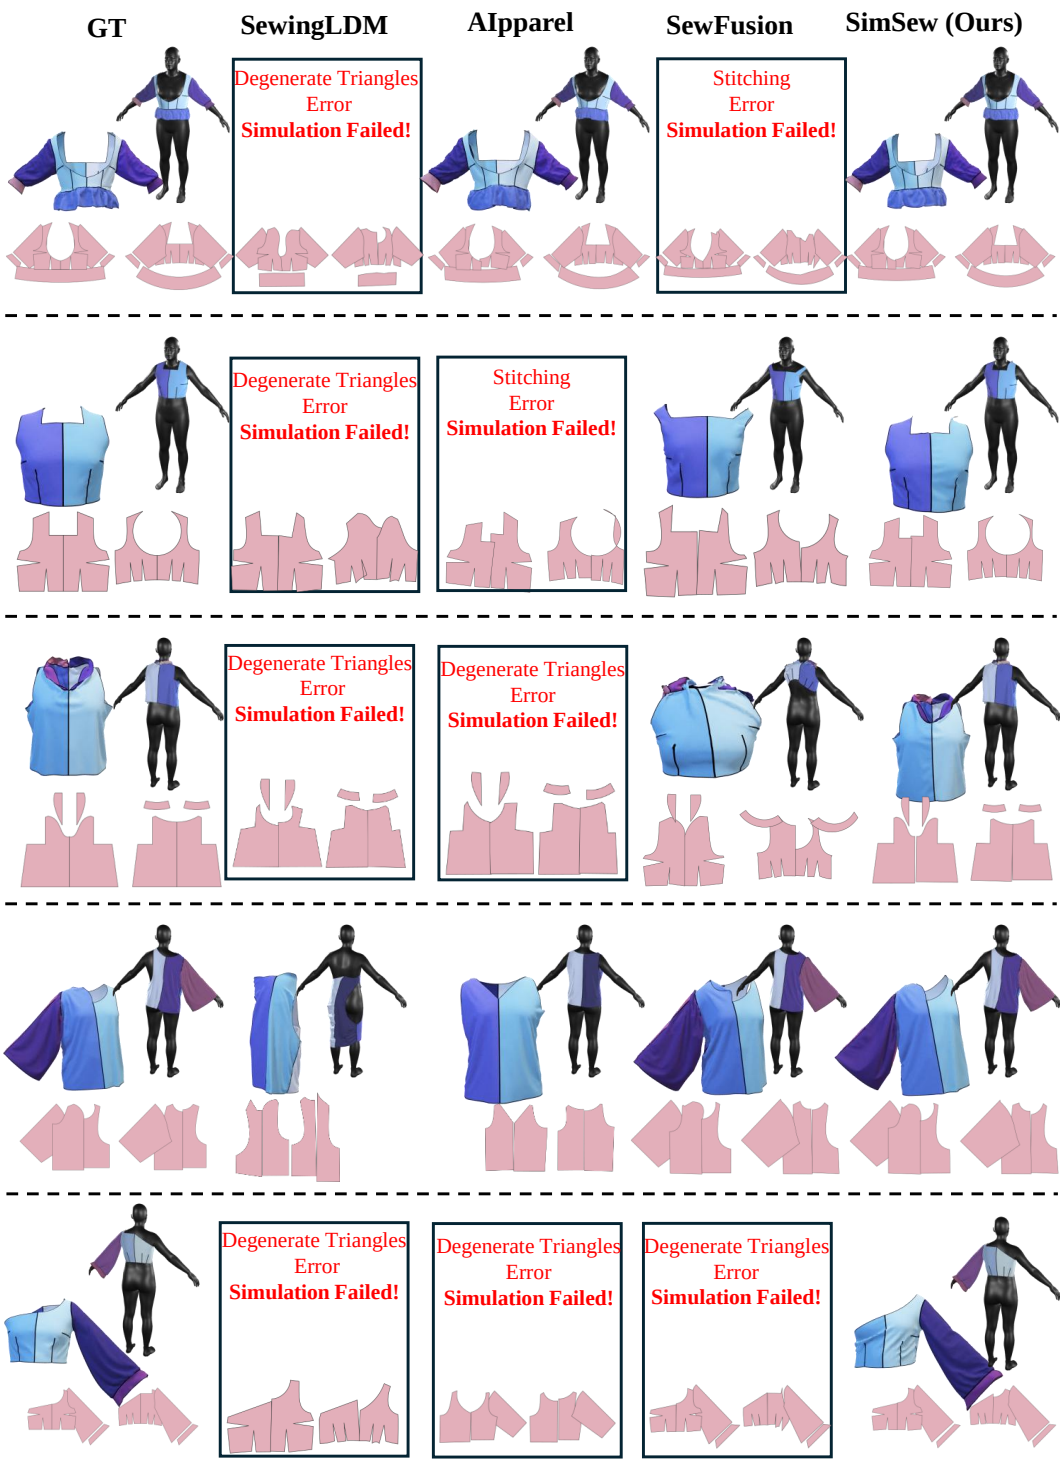

Figure 7: Qualitative cases of image-text-based generation. For clarity, text prompts are not shown.

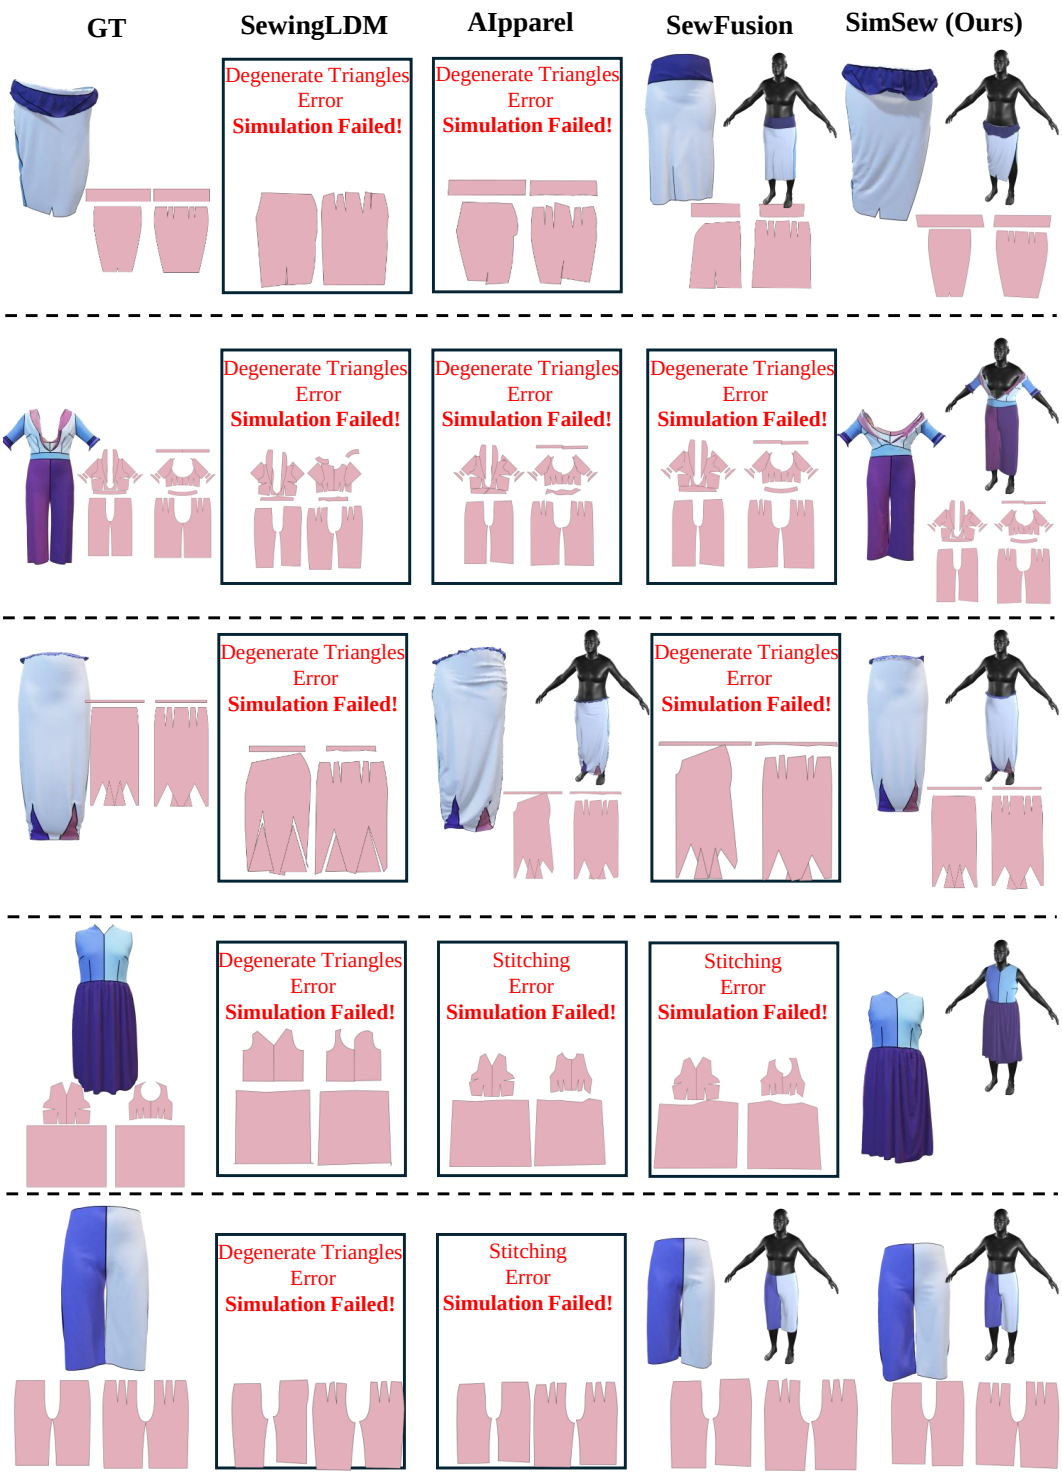

Figure 8: Qualitative cases of text-based generation. For clarity, text prompts are not shown.