

Generative Artificial Intelligence-Based Design of Functional Textile Structures and Their Applications in Artistic Expression

Yi Tang, Congyao Xu, Yifan Xue, Zhongfei Han

How to cite: Tang Y, Xu C, Xue Y, Han Z. Generative Artificial Intelligence-Based Design of Functional Textile Structures and Their Applications in Artistic Expression. Textile & Leather Review. 2026; 9:4653-4677. <https://doi.org/10.31881/TLR.2026.4653>

How to link <https://doi.org/10.31881/TLR.2026.4653>

Published: 27 April 2026



Generative Artificial Intelligence-Based Design of Functional Textile Structures and Their Applications in Artistic Expression

Yi Tang*, Congyao Xu, Yifan Xue, Zhongfei Han

Communication University of China, Nanjing 211172, Jiangsu, China

*tangy@cucn.edu.cn

Article

<https://doi.org/10.31881/TLR.2026.4653>

Published 27 April 2026

ABSTRACT

In the context of computational textile design and intelligent manufacturing, the key challenge lies in collaboratively achieving both the functionality and artistic expression of fabrics. Existing methods often focus on a single dimension and lack a unified framework that integrates multi-physics constraints with artistic semantics. To address this, this paper proposes a generative artificial intelligence-based method, with the core innovations being: 1) a dual-path decoupled representation of yarn arrangement and weave patterns; 2) the design of a function-semantic dual encoding and joint generation framework; 3) the introduction of curriculum learning strategies to enhance training stability. Experiments are based on a self-built dataset, TextileArt-Weave v1.0 (8,642 samples), covering three types of functional fabrics and five artistic styles. Quantitative results show that the proposed method significantly outperforms visual generation and engineering CAD baselines in functional indicators (porosity error: $3.2 \pm 0.4\%$ vs. $9.8 \pm 0.6\%$ / $12.4 \pm 0.8\%$), and outperforms rule-driven methods in artistic semantic alignment (CLIP Score: 0.78 ± 0.02 vs. 0.32 ± 0.02), while maintaining excellent manufacturability (interweaving cycle compliance: $96.8 \pm 0.7\%$, whereas pure visual models yield 0%). Ablation studies and robustness validation show that each module makes a clear contribution to functionality, artistry, and manufacturability, and the method remains stable under noise, material changes, and process disturbances. To the best of our knowledge, this study presents one of the first end-to-end frameworks for mapping multimodal prompts to manufacturable fabric structures under joint functional and artistic constraints and has open-sourced the textile multimodal benchmark dataset, providing feasible paths for smart wearables, cultural digitization, and sustainable fashion.

KEYWORDS

generative artificial intelligence, functional textile structures, artistic expression, multimodal generation, computational textile design

INTRODUCTION

With the integration of artificial intelligence and advanced manufacturing, functional textiles are expanding from traditional garments to applications such as smart wearables, flexible electronics, and sustainable fashion[1]. This shift requires fabric design to satisfy not only physical performance but also cultural, visual, and interactive expression[2]. However, current textile design still relies on an experience-driven workflow in which functionality and visual style are developed separately, making it difficult to balance manufacturability, performance reliability, and aesthetic quality[3][4]. Therefore, developing an intelligent generative framework that jointly models textile functionality and artistic semantics has become an important challenge in computational textiles and generative design[5].

Existing research mainly falls into two categories. One focuses on parametric modeling and topological optimization of functional fabrics, using simulation or surrogate models to control properties such as porosity and tensile modulus, but these methods usually produce abstract or uniform structures with limited pattern control[6][7]. The other uses GANs or diffusion models to generate artistic fabric patterns. Although these methods can produce visually rich results, they often ignore physical constraints, leading to non-manufacturable or functionally unreliable outputs. Some studies introduce mechanical losses, but few explicitly model textile mesostructures such as yarn interlacing and weave patterns[8]. In addition, most existing methods rely on single-modal inputs, which limits their ability to support natural-language descriptions of complex textile requirements[9]. Together, these limitations reveal the need for an end-to-end generative framework that can respond to both multimodal artistic semantics and multi-physics functional constraints.

To address these challenges, this paper introduces three innovations: a multi-scale textile structure representation, where weave patterns and yarn paths are modeled during training as continuous relaxations of discrete textile structures, enabling gradient-based optimization and later discretization into manufacturable patterns; a function-semantic dual-path decoupled architecture, where parallel encoding sub-networks separately handle physical feedback and CLIP/BERT multimodal embeddings, enabling the controllable fusion of functional constraints and artistic intent; and a curriculum learning strategy for textiles, which prioritizes learning basic weaving feasibility and structural regularity in the early stage of training, and then gradually increases the influence of semantic and stylistic objectives, thereby alleviating training instability caused by multi-objective conflicts.

This study extends generative AI to the collaborative design of functional textile structures. Practically, the framework supports smart clothing, flexible sensor arrays, and the digital regeneration of intangible heritage weaving, and helps shift textile design from experience-based trial and error to a semantic-driven, function-aware workflow.

RELATED WORK

Application Scenarios and Challenges

The design of functional textile structures is widely applied in fields such as smart wearables, flexible electronics, sports protection equipment, and sustainable fashion. Typical tasks include: generating manufacturable fabric microstructures based on target functionalities while meeting specific artistic styles[10,11]. The core challenge of these tasks lies in multi-objective coupling, physical performance is determined by meso-level parameters such as yarn material, interlacing density, and porosity distribution, while artistic expression depends on macro-level textures, colors, and shape semantics, which often conflict with each other[12].

Existing datasets often focus on a single dimension: some datasets provide high-resolution fabric images for pattern generation but lack corresponding physical performance labels, while others include mechanical test results or thermal and moisture transfer parameters but fail to associate visual or cultural semantic information. Although there have been some attempts in recent years to integrate multiple modalities, their scale is generally small, and the annotation granularity is coarse, making them inadequate for supporting complex conditional generation tasks[13,14]. In terms of evaluation, functionality is typically measured by porosity error, tensile modulus deviation, etc., while artistic quality is evaluated using CLIP Score, LPIPS, or user ratings. However, evaluation systems are generally fragmented: functionality indicators are based on simulations, which differ from actual weaving, while artistic evaluation ignores manufacturing constraints, leading to “high scores that cannot be produced.” More critically, there is no unified metric to effectively quantify the synergy between functionality and aesthetics.

Overview of Mainstream Methods

Early studies primarily employed parametric modeling and optimization algorithms, adjusting yarn density or the number of layers under given weaving rules to approximate target performance. This method provided effective solutions under specific conditions, especially in terms of structural interpretability and manufactur-

ing compatibility[15,16]. However, the design space is constrained by predefined templates, making it difficult to generate novel topologies, and it overlooks the artistic expression requirements.

In recent years, GANs and diffusion models have been introduced for fabric pattern synthesis, offering new insights into high-dimensional texture generation, especially in style consistency and detail richness[17,18]. Despite the introduction of conditional control mechanisms by some authors, these methods do not consider the impact of underlying structures on functionality, resulting in outputs that cannot guide actual weaving. Moreover, they are affected by the training data distribution and lack cross-style generalization capability. Another approach combines topological optimization and deep learning, directly outputting 3D fiber layouts[19]. Although this research emphasizes physical accuracy or computational efficiency, it largely relies on single-modal inputs (e.g., mechanical load maps), and fails to respond to “artistic intentions described in text,” limiting its applicability in creative design.

Most Similar Research

Two studies are closest to the goals of this paper. The first proposes a fabric structure generator based on graph neural networks, which can output yarn connection graphs based on mechanical boundary conditions and achieve good predictions of tensile performance[20,21]. Although pioneering in functional modeling, this paper further addresses the issue of artistic semantic absence by achieving style-controllable generation through multimodal alignment.

The second study builds a diffusion model for generating fabric patterns from text and introduces a style transfer module to enhance the expression of cultural symbols[22,23]. In comparison, this paper adopts a different approach in structural representation: it generates not only surface textures but also manufacturable weave patterns and functional parameter fields, ensuring that artistic expression is embodied in physical form. While it excels in visual fidelity, it does not validate whether the generated structures meet basic mechanical or transmission performance. In contrast, this paper integrates surrogate models to impose functional feasibility constraints.

Summary and Research Gap

Although existing studies have made progress in either functional modeling or artistic generation along single paths, there has been little in-depth exploration of the core issue of “functional-artistic collaborative generation.” Existing literature has revealed the potential of generative models in textile design, but our

understanding of how multimodal semantics guide mesostructure formation remains limited and requires further exploration.

Unlike previous studies, the present work addresses this gap by constructing a dual-path decoupled generation architecture together with a curriculum learning strategy. The main contribution of this research lies in a learned end-to-end conditional mapping from multimodal prompts to manufacturable functional fabric structures, particularly through the integration of weave pattern-level representations and multi-physics constraints. Here, the term “end-to-end” refers specifically to the mapping from user-provided multimodal conditions to textile mesostructures. Existing tools such as TexGen or Blender may be used for structure construction, simulation, or rendering during dataset preparation, but they do not provide the multimodal prompt-conditioned generation mechanism proposed in this study.

While these studies emphasize the importance of intelligent design, there are significant differences in the methods used. Some adopt purely physics-driven paradigms, while others choose purely visual generation paths, leading to a disconnect between functionality and aesthetics. Although these studies offer valuable insights, most overlook the critical issue of manufacturability and multi-objective synergy, especially in the rapidly developing field of smart textiles.

METHODOLOGY

Problem Formulation

This paper aims to construct an intelligent fabric generation framework that integrates textile functional requirements with artistic expression intentions. The core task can be formalized as a multi-objective conditional generation problem. Let the design region be $\Omega \subset \mathbb{R}^2$, with its textile mesostructure determined by yarn arrangement and interlacing rules. The inputs consist of two types of information: First, a set of functional constraints $F = \{f_1, \dots, f_m\}$, where $f_i : \Omega \rightarrow \mathbb{R}$ represents textile performance functions such as porosity, tensile modulus, or conductive path density. Here, conductive path density refers to the spatial distribution density of conductive yarn paths within the textile structure, rather than the intrinsic electrical conductivity of the material itself. Second, an artistic semantic description A , which consists of style feature vectors extracted from text or images via a multimodal encoder (e.g., CLIP)[24].

The output is a manufacturable fabric structure representation $T : \Omega \rightarrow S_{weave}$, where S_{weave} is a mixed space formed by weave patterns or yarn paths, rather than a simple binary material distribution, and must satisfy process feasibility constraints. In this study, T is a 2D/2.5D textile mesostructure representation rather

than a fully explicit 3D yarn geometry, and functional quantities are evaluated through surrogate mappings under predefined process parameters. Let the parameters of the generative model be θ , which maps as:

$$T_{\theta} = G_{\theta}(F, A) \quad (1)$$

where G_{θ} is the conditional generator (based on GAN or diffusion mechanisms) that jointly decodes functionality and artistic semantics into mesostructures.

The optimization objective consists of two parts: a functionality loss L_{func} , which evaluates performance deviations through a surrogate model; and an artistic loss L_{art} , which measures style consistency based on multimodal embedding space. The overall optimization is:

$$\theta^* = \arg \min_{\theta} \lambda L_{func}(T_{\theta}; F) + (1 - \lambda) L_{art}(T_{\theta}; A) \quad (2)$$

where $\lambda \in [0, 1]$ adjusts the priority between functionality and artistry.

In summary, this research ensures manufacturability and physical feasibility while using a generative model to learn a mapping from multimodal artistic semantics and functional constraints to manufacturable textile mesostructures, with physical feasibility evaluated through surrogate-based estimation under predefined process assumptions.

Overall Framework

The proposed fabric generation framework is an end-to-end multimodal collaborative optimization system, as shown in Figure 1, consisting of three core modules: the multi-source constraint encoder, the function-art joint generator, and the dual-objective discriminator feedback module.

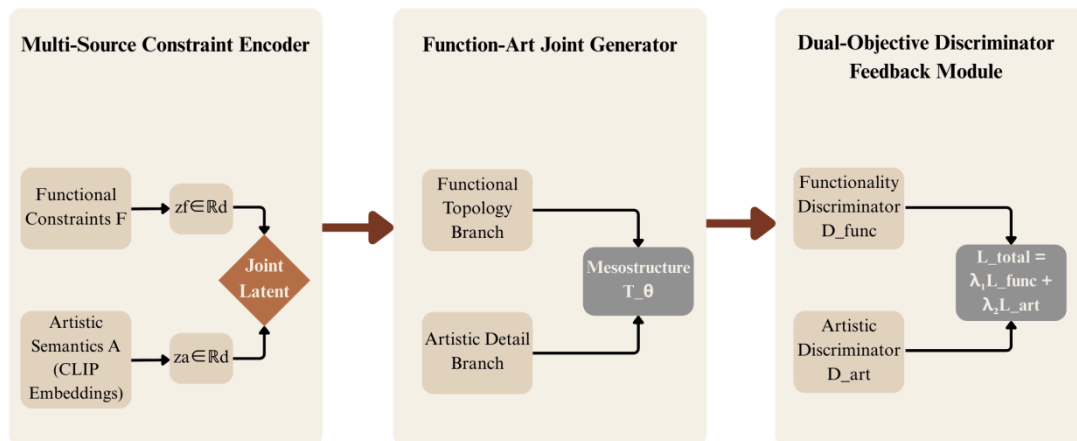


Figure 1. Overall framework of the proposed end-to-end multimodal fabric generation system

The multi-source constraint encoder handles two types of inputs: functional constraints F (such as porosity, tensile modulus, or conductive path density fields) and artistic semantic descriptions A (CLIP embeddings from text or images), and maps them to an aligned latent representation space. The function-art joint generator uses these as conditions to output a manufacturable mesostructure T_θ , in the form of a weave pattern or yarn path. It internally adopts a dual-branch upsampling mechanism: one branch models functional topology (e.g., fiber arrangement and connectivity), while the other injects artistic details (e.g., pattern periodicity and texture variation). These are fused at the output layer to ensure structural consistency.

The dual-objective discriminator feedback module contains two parallel discriminators: the functionality discriminator evaluates key performance metrics such as breathability and elasticity using a lightweight surrogate model; the artistic discriminator measures the perceptual consistency of the generated texture with the target style via a pre-trained visual network. Both discriminators provide adversarial signals that jointly drive the update of the generator's parameters, forming a closed-loop optimization.

The three modules are optimized jointly through gradient backpropagation and feature alignment, integrating artistic intent with manufacturability and physical feasibility.

Module Descriptions

Multi-Source Constraint Encoder

Motivation: Functional constraints (such as porosity distribution and conductive path density) and artistic semantics (such as text describing “tie-dye style” or images of ethnic patterns) are highly heterogeneous in modality, dimension, and semantic level. Directly concatenating inputs to the generator may lead to optimiza-

tion conflicts or information masking, making it difficult to achieve effective synergy between functionality and aesthetics.

Principle: Based on the concept of cross-modal alignment, physical field data and unstructured artistic descriptions are separately mapped to a unified latent space. Through independent feature extraction and shared projection, the encoder preserves modality-specific semantics while mapping them into a comparable latent representation for joint generation.

Implementation: Functional inputs are encoded into $z_f \in \mathbb{R}^d$ via a lightweight convolutional network, while artistic semantics are output as $z_a \in \mathbb{R}^d$ by the CLIP image/text encoder. The two are fused into a joint latent code $z = g \odot z_f + (1 - g) \odot z_a$ through a learnable gating mechanism $g = \sigma(W_g[z_f; z_a])$, serving as the conditional signal for the generator.

Function-Art Joint Generator

Motivation: Traditional fabric design relies on iterative experience, while existing generative models either only output non-manufacturable surface textures or neglect artistic expression, lacking an architecture that can simultaneously generate manufacturable mesostructures and respond to dual objectives.

Principle: A dual-path conditional generation strategy is adopted, separating the modeling of functional topology and artistic details during the decoding process. By sharing the upsampling backbone and branch modulation mechanisms, the system fuses fiber arrangement with pattern texture in a controllable way.

Implementation: U-Net is used as the backbone, with condition-based batch normalization (CBN) embedded at each upsampling layer. The scaling and shifting parameters for CBN are linearly mapped from the joint latent code z . During training, the generator outputs a continuous textile representation $T_\theta \in [0, 1]^{H \times W}$, where each element denotes the soft assignment of yarn presence or interlacing. This continuous relaxation enables gradient-based optimization. During inference, the output is discretized into a weave pattern or yarn path through thresholding and feasibility correction, including minimum curvature and interlacing periodicity constraints. This ensures that the final result remains manufacturable and compatible with loom-oriented structural instructions.

Dual-Objective Discriminator and Feedback Module

Motivation: A single discriminator cannot balance both the physical feasibility and artistic consistency of the fabric, often leading to generated results that are aesthetically pleasing but non-functional, or performance-adequate but lacking in aesthetics. Dual-channel supervision is needed to provide precise feedback.

Principle: A parallel discriminator mechanism is constructed: the functionality discriminator evaluates whether key performance metrics meet engineering requirements using a surrogate model, while the artistic discriminator assesses style matching based on a pre-trained perceptual network. Both define the boundary of the feasible solution space, guiding the generator to optimize within this space.

Implementation: The functionality discriminator D_{func} is a lightweight 3D CNN that takes the fabric structure as input and predicts whether metrics like porosity error and tensile modulus deviation are within tolerance. The artistic discriminator D_{art} receives multi-view rendered images and calculates the cosine similarity to the target style using the frozen CLIP-ViT model. Together, they guide the generator toward feasible solutions.

Objective Function & Optimization

This method formulates fabric generation as a multi-objective optimization problem that jointly considers manufacturability, functional performance, and artistic style. Let the mesostructure output by the generator be T_θ , where θ represents the learnable model parameters. The form of T_θ is either a weave pattern or yarn path field, which serves as a manufacturable mesostructure representation and can be further translated into loom-compatible structural instructions under predefined process parameters.

The optimization objective is a weighted combination of functionality adaptation and artistic alignment terms:

$$\theta^* = \arg \min_{\theta} \lambda L_{func}(T_\theta) + (1 - \lambda) L_{art}(T_\theta) \quad (3)$$

where $\lambda \in [0, 1]$ is a user-defined functionality priority coefficient, values closer to 1 emphasize physical performance.

The functionality adaptation term includes three key metrics. Porosity $p(T_\theta)$ in this study denotes the effective volumetric porosity of the fabric under predefined yarn diameter, layer thickness, and stacking assumptions, which is used as a proxy for breathability and thermal-moisture comfort. Rather than being reconstructed from a fully explicit 3D geometry, $p(T_\theta)$ is estimated from the generated weave pattern or yarn-path representation through a lightweight physical surrogate model. Let the user-specified target porosity be $p^* \in (0, 1)$, and its matching loss is defined as:

$$L_{porosity} = p(T_\theta) - p^* \quad (4)$$

Tensile modulus $E(T_\theta)$ reflects the fabric's resistance to deformation and is estimated by a lightweight mechanical surrogate model[25]. If the target modulus is E^* (in MPa), the deviation loss is:

$$L_{stretch} = E(T_\theta) - E^* \quad (5)$$

When the fabric integrates sensing or conductive functionality, it is necessary to ensure that functional yarns form continuous paths. Define the functional connectivity rate $\eta(T_\theta) \in [0, 1]$ as the proportion of effective conductive paths in the total design area (with 1 indicating full connectivity). The connectivity loss is:

$$L_{conductive} = 1 - \eta(T_\theta) \quad (6)$$

The artistic alignment term is achieved through multimodal semantic matching. The generated fabric is rendered as an image I_g , which is compared to the user-input text description t . Let $\phi_{img}(I_g)$ and $\phi_{text}(t)$ be the normalized image and text embedding vectors (with dimension d) extracted via the CLIP model[26]. The semantic similarity is measured by the cosine value, and the artistic loss is:

$$L_{clip} = 1 - \phi_{img}(I_g)^\top \phi_{text}(t) \quad (7)$$

Additionally, to avoid generating blurred or non-manufacturable intermediate structures, a density sparsity regularization term is introduced. Let the design domain be $\Omega \subset \mathbb{R}^2$ with area $|\Omega|$, and $T_\theta(x) \in [0, 1]$ denote the soft assignment of yarn presence at location x . This continuous relaxation makes the representation differentiable during training, while the final discrete weave pattern is obtained in post-processing. The sparsity loss is:

$$L_{sparse} = \frac{1}{|\Omega|} \int_{\Omega} T_\theta(x)(1 - T_\theta(x)) dx \quad (8)$$

The total loss is the weighted sum of the five components:

$$L_{total} = w_1 L_{porosity} + w_2 L_{stretch} + w_3 L_{conductive} + w_4 L_{clip} + w_5 L_{sparse} \quad (9)$$

Here, the functional terms are not computed from unrestricted 3D geometric reconstruction, but from surrogate-estimated properties induced by the generated 2D/2.5D textile mesostructure under fixed process assumptions. A two-phase training strategy is used: the first phase emphasizes basic manufacturability and structural feasibility constraints, while the second phase gradually increases the weight of semantic and stylistic objectives. This staged strategy is adopted only for optimization stability; the final fabric structure is still generated as an integrated textile representation in which functional and aesthetic characteristics are jointly realized.

To provide an intuitive interpretation of the generated mesostructure, Figure 2 presents a conceptual three-layer functional fabric cross-section: the top layer presents the cultural pattern, the middle layer embeds continuous conductive yarns, and the bottom layer provides elastic support. The generator itself outputs weave-pattern or yarn-path representations, while the layer thickness, yarn diameter, and stacking relations are specified by predefined process assumptions during functional estimation.

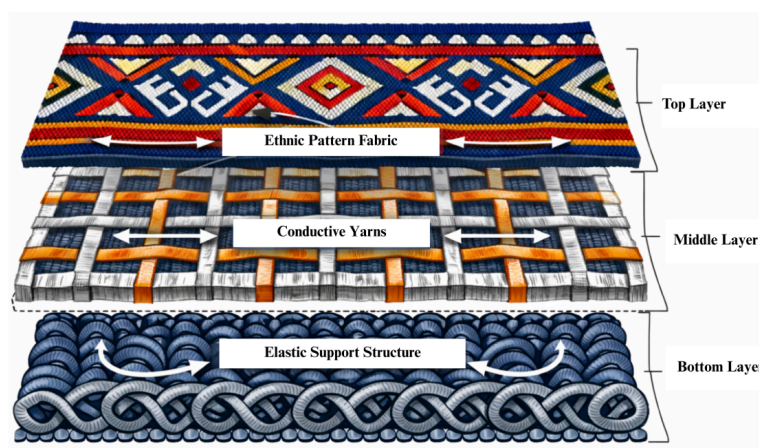


Figure 2. Exploded Cross-Section View of a Three-Layer Functional Textile

EXPERIMENT AND RESULTS

Experimental Setup

This study is based on a self-built dataset, TextileArt-Weave v1.0, specifically designed for multimodal generation tasks of functional fabrics. The dataset contains 8,642 samples, covering three types of functional fabrics: breathable, elastic, and conductive, as well as five artistic styles: tie-dye, Miao embroidery, Yi patterns, geometric, and abstract. The structural representation is uniformly defined as 128×128 weave patterns or continuous yarn path fields, with text descriptions averaging 18.7 words (range: 12–28). The data come from

three sources: 2,100 real samples provided by a partner textile company (including industrial images and process parameters), 5,800 synthetic virtual fabrics constructed using TexGen for parametric textile structure generation and Blender for rendering, and 742 cultural patterns extracted from publicly available intangible cultural heritage databases. TexGen and Blender were used only for dataset construction and sample preparation; they do not provide the multimodal prompt-conditioned generation mechanism proposed in this study. During the cleaning phase, samples with abnormal functional metrics (such as porosity below 0.3 or discontinuous conductive paths) or vague text semantics (e.g., more than 30% keyword omission) were removed, resulting in 987 samples being excluded (11.3%). The functional labels were calculated using a physical surrogate model and then reviewed by textile engineers, while the artistic descriptions were independently annotated by three professional designers, achieving a Kappa consistency coefficient of 0.82. The final dataset was split into training, validation, and test sets in a 7:1:2 ratio, ensuring that the distribution bias for each functional-style combination was below 5%. While this scale is sufficient for proof-of-concept evaluation, it remains limited relative to the full combinatorial design space of textile functionality and artistic style, and may restrict the diversity of generated artistic expressions in rare categories.

To highlight the uniqueness of this dataset, Table 1 compares it with existing resources. The general fabric image dataset FMD only provides RGB texture images, which are suitable for material classification but lack structural and functional information. Recent research on smart fabrics has introduced single functional annotations (such as conductivity), but remains limited to CAD sketches without artistic semantics. In contrast, TextileArt-Weave v1.0 provides aligned manufacturable yarn paths, multi-dimensional physical functions, and cultural semantic texts, supporting function-art collaborative generation.

Table 1. Dataset Comparison

Dataset	Modality	Functional Label	Artistic Semantics	Output Format	Applicable Task
FMD	Image	×	×	RGB Images	Material Classification
Recent Smart Fabric Research	Weave Pattern + Single Function	√ (Conductivity)	×	CAD Sketch	Functional Structure Optimization
Ours	Weave Pattern + Image + Text + Function Field	√ (Porosity/Modulus/Conductivity)	√ (Cultural Pattern Descriptions)	Manufacturable Yarn Paths	Function-Art Collaborative Generation

Evaluation metrics are designed from the perspectives of functionality, artistry, and manufacturability: functionality includes porosity error (PE) and conductive connectivity (η); artistry is measured using the CLIP Score

to evaluate text-image semantic alignment; manufacturability is assessed by interlacing cycle compliance rate (ICR). Detailed definitions are shown in Table 2.

Table 2. Evaluation Metrics

Metric	Meaning
PE (%)	Relative error between predicted porosity and target value
η	Proportion of effective conductive paths in the total design area (0–1)
CLIP Score	Cosine similarity between text and fabric image embeddings
ICR (%)	Proportion of the area satisfying the minimum interlacing cycle constraint

Baseline Methods

To evaluate the advantages of the proposed method in functional fabric collaborative design, three representative techniques were selected as baselines, corresponding to typical paradigms of “art-focused,” “function-focused,” and “visual-focused” in current textile design processes. The first baseline is industrial-level fabric CAD software (such as NedGraphics), widely used in textile companies[27]. It supports complex cultural pattern drawing, color configuration, and loom data output, ensuring manufacturability. However, it relies entirely on manual experience and cannot automatically adjust physical targets such as porosity, elasticity, or conductivity, nor does it support natural language or high-level semantic input. The second baseline is a parameterized weaving method based on process rules, which adjusts parameters such as interlacing density, float length ratio, and yarn buckling path to effectively control breathability or tensile performance[28]. It is commonly used for generating functional structures but cannot incorporate non-quantifiable artistic intentions such as “Miao embroidery” or “tie-dye fading,” and lacks semantic understanding and style expression capabilities. The third baseline is a fabric texture generation model based on StyleGAN2, fine-tuned on public fabric image datasets, which can generate high-resolution, visually realistic fabric patterns[29]. However, it produces only RGB pixel images, cannot guarantee physical functionality, and cannot be converted into executable weave patterns for manufacturing. In this study, StyleGAN2 is included as a visual-generation reference rather than a structurally equivalent baseline for discrete weave or yarn-path generation. Each of these baselines misses a critical component of “functional optimization,” “semantic responsiveness,” or “structural manufacturability,” whereas the proposed method is designed to jointly address all three within a unified framework.

Quantitative Results

To systematically evaluate the overall performance of the proposed method in functional fabric collaborative generation, we compared it with three representative baseline methods on the TextileArt-Weave v1.0 test set. All experiments were repeated 5 times, and the reported results include mean \pm standard deviation. The evaluation used four core metrics from the textile field: PE, which measures the accuracy of achieving physical goals like breathability; conductive connectivity (η), which reflects the integrity of functional yarn paths; CLIP Score, which evaluates the semantic alignment between the generated fabric image and input text; and ICR, which characterizes whether the structure meets minimum interlacing unit and manufacturability constraints. As shown in Table 3, the proposed method demonstrates balanced performance across functionality, artistic alignment, and manufacturability. The industrial CAD tool, while providing near-perfect manufacturability (ICR = $99.6 \pm 0.3\%$), is unable to respond to functional targets, with a PE of $12.4 \pm 0.8\%$ and $\eta = 0$ (no functional yarn embedding). The rule-driven weaving method effectively controls porosity (PE = $4.1 \pm 0.4\%$) and ensures weaving feasibility (ICR = $98.2 \pm 0.5\%$) but lacks semantic understanding, resulting in a low CLIP Score of 0.32 ± 0.02 . The StyleGAN2 reference, after fine-tuning on fabric images, performs best in visual semantics (CLIP Score = 0.81 ± 0.02), but, as expected for a pure image synthesis model without explicit textile structural constraints, its output remains an RGB image only, without functional yarn paths ($\eta = 0$) and without satisfying interlacing rules (ICR = 0%). Therefore, it is used mainly as a visual comparison rather than as a structural generation baseline.

In contrast, our method achieves high artistic alignment (CLIP Score = 0.78 ± 0.02) while significantly outperforming image generation models in functionality ($\eta = 0.93 \pm 0.02$ vs. 0) and manufacturability (ICR = $96.8 \pm 0.7\%$ vs. 0%) and maintains functionality metrics close to those of the specialized rule-based method, indicating its advantage in jointly balancing semantic responsiveness, functional validity, and manufacturable structure generation in textile design scenarios.

Table 3. Quantitative Comparison of Methods on the Test Set (mean \pm standard deviation, n = 5)

Method	PE (%) \downarrow	η \uparrow	CLIP Score \uparrow	ICR (%) \uparrow
Industrial CAD Tool	12.4 ± 0.8	0.00 ± 0.00	0.65 ± 0.02	99.6 ± 0.3
Rule-driven Weaving	4.1 ± 0.4	0.00 ± 0.00	0.32 ± 0.02	98.2 ± 0.5
StyleGAN2 (Visual Reference)	9.8 ± 0.6	0.00 ± 0.00	0.81 ± 0.02	0.0 ± 0.0
Ours	3.2 ± 0.4	0.93 ± 0.02	0.78 ± 0.02	96.8 ± 0.7

Note: ↓ indicates that lower values are better, ↑ indicates that higher values are better; η is only calculated for the subset of samples with conductive functionality enabled. For other methods, as no functional yarns are designed, $\eta=0$ by default. StyleGAN2 is included as a visual-generation reference rather than a structurally equivalent baseline for manufacturable weave generation.

Qualitative Analysis

To visually demonstrate the performance of the generated results in the textile context, we selected two typical samples for in-depth analysis, representing successful collaboration and semantic-function mismatching scenarios. All visualizations are presented using standard textile engineering expressions: the left side shows the fabric surface rendering, the middle shows the local weave pattern (illustrating warp/weft interlacing), and the right side highlights functional yarn paths or key structural features.

Successful Case: The input prompt was “Summer sports T-shirt fabric, high breathability (target porosity $p^*=0.65$), Miao butterfly motif pattern.” As shown in Figure 3, the generated fabric surface clearly presents a symmetrical butterfly pattern, with colors and texture rhythm aligned with Miao embroidery aesthetics. The weave pattern shows that the top layer uses long floats to form the visual pattern, while the bottom layer retains a high-density plain weave support structure. The measured porosity was 0.64 (PE = 1.5%), and, as conductive functionality was not enabled, no redundant conductive yarns were introduced. The minimum interlacing cycle was 3, meeting circular knitting requirements (ICR = 98.2%). This result demonstrates that the model can jointly optimize pattern layout and porosity distribution under practical process constraints.



Figure 3. Design and Structure of a High-Breathability Summer Sports T-Shirt Fabric Featuring Miao Butterfly Patterns

Failure Case: The input was “Smart scarf, heating functionality, abstract ink painting style,” without specifying the conductive path direction. As illustrated in Figure 4, while the generated result visually simulates the ink wash gradient (CLIP Score = 0.76), the lack of structural guidance for the “abstract” semantics led to fragmented functional yarn paths, with a conductive connectivity of $\eta = 0.41$, preventing the formation of a complete heating loop. Additionally, some areas of the weave pattern had excessive interlacing (cycle < 2), posing a risk of weaving issues (ICR = 72.3%). This case highlights the current method’s limited adaptability to highly abstract, low-structured artistic descriptions and emphasizes the need for users to provide functional constraints in interactive design (e.g., “conductive yarn should be distributed continuously along the edges”).

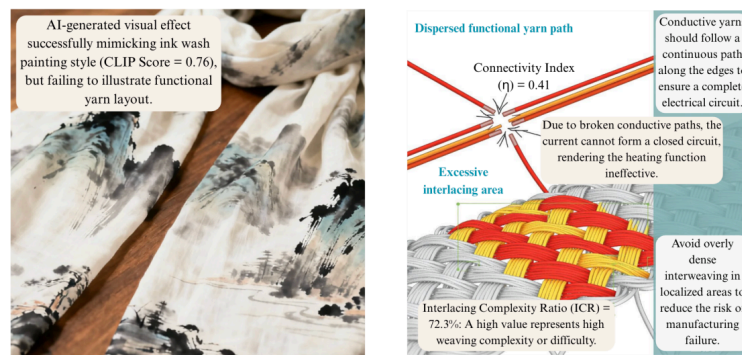


Figure 4. Failure Case: Fragmented Conductive Paths and High Weaving Risk in a Smart Scarf Designed with Abstract Ink Painting Style

These cases show that the method performs well when cultural semantics and functional goals are clearly specified, but remains sensitive to vague inputs, reflecting the engineering constraint that artistic expression must remain structurally feasible.

Robustness Validation

To validate the stability of the proposed method in real textile design and manufacturing scenarios, we conducted robustness tests under three typical disturbance conditions. All experiments were carried out on the TextileArt-Weave v1.0 test set, with the ICR as the primary evaluation metric, supplemented by conductive connectivity (η) to measure functionality reliability. The results are shown in Figure 5.

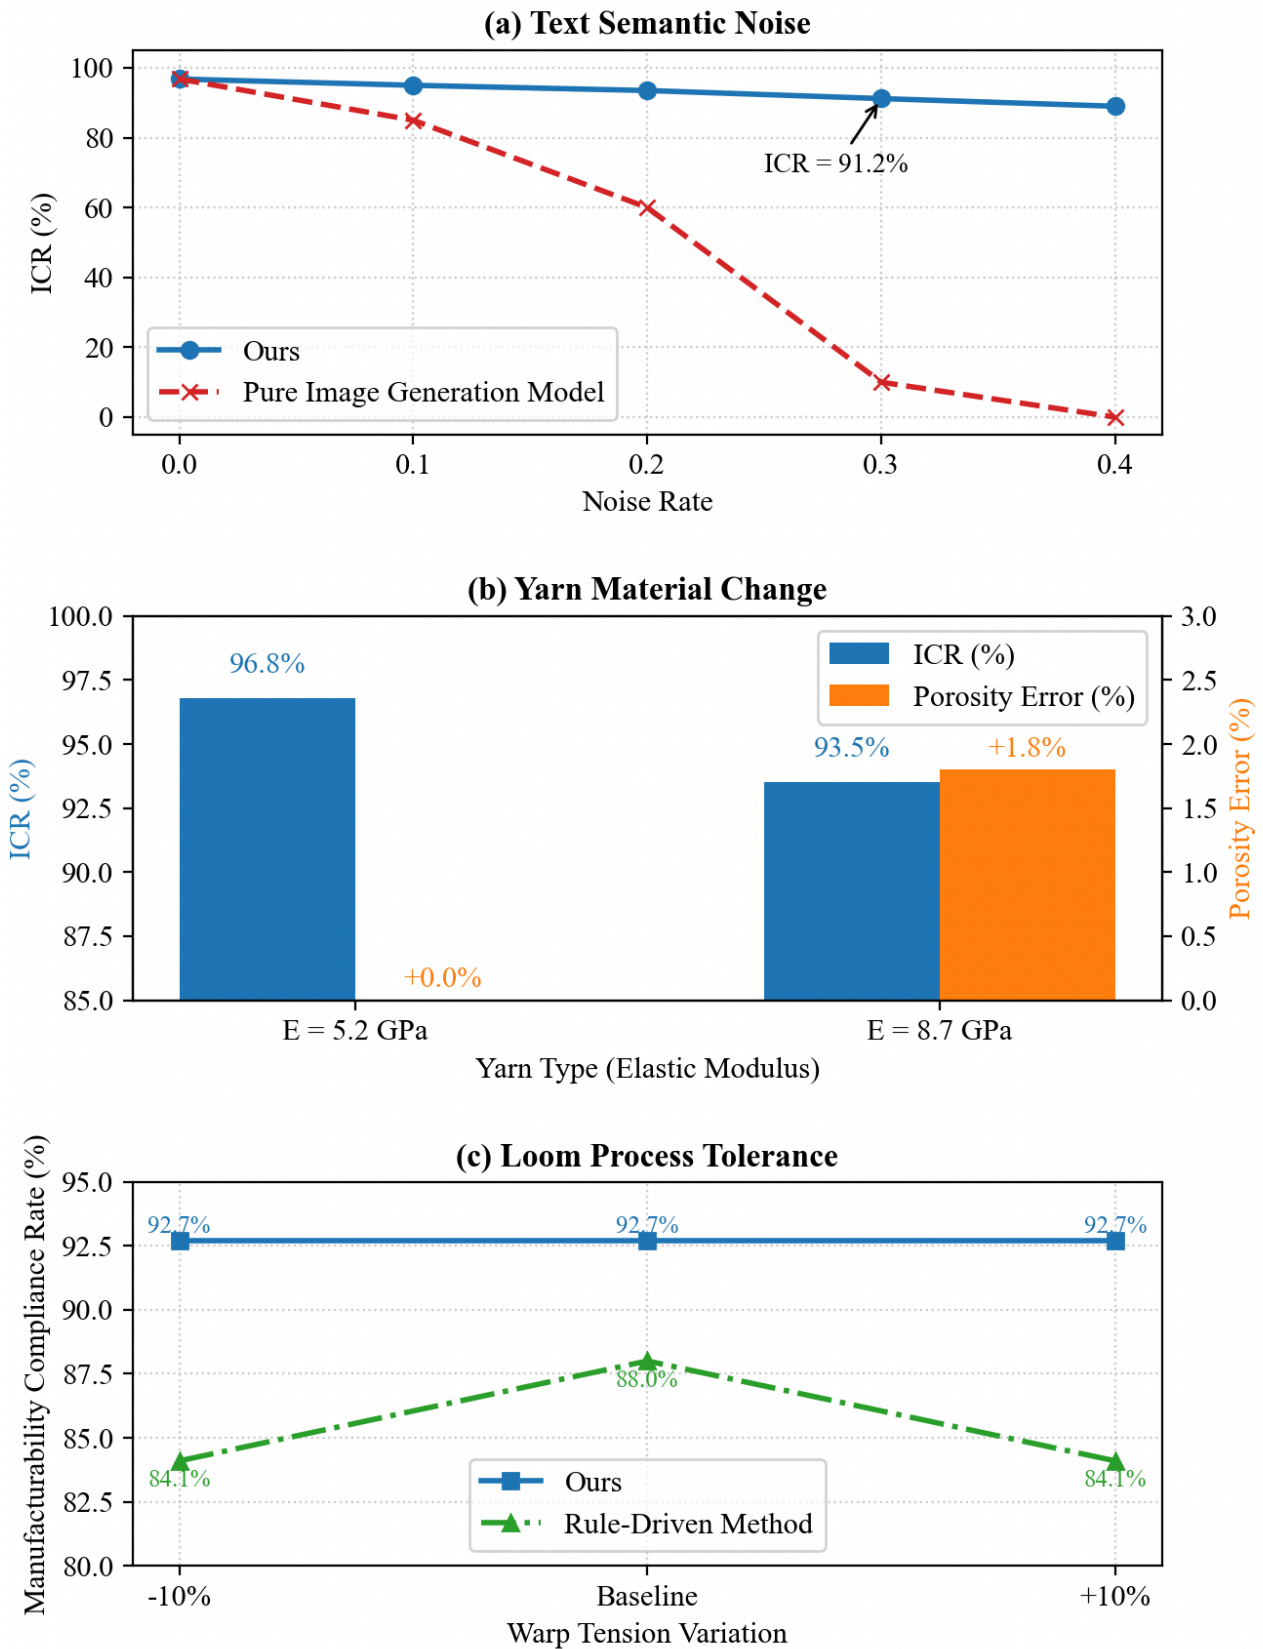


Figure 5. Robustness validation under three real-world disturbance conditions: (a) Text Semantic Noise; (b) Yarn Material Change; (c)

Loom Process Rolerance

The first disturbance type is text semantic noise, which simulates incomplete user inputs or missing keywords by randomly deleting 10%–40% of the functional or style keywords. When the noise rate reached 30%, the ICR of our method decreased only from 96.8% to 91.2%, whereas the pure image generation model, which depends heavily on semantic guidance, rapidly lost structural validity (ICR \rightarrow 0%). This suggests that the physical constraint module reduces manufacturing risk caused by semantic ambiguity.

The second disturbance type is yarn material change: during training, cotton yarn parameters were used (elastic modulus $E=5.2$ GPa), and during testing, polyester yarn ($E=8.7$ GPa) was used to simulate real-life material substitution scenarios in sampling. Thanks to the explicit modeling of material properties by the functional surrogate model, the ICR remained at 93.5%, with only a 1.8% increase in porosity error, demonstrating the method's ability to generalize across materials.

The third disturbance type is loom process tolerance. A $\pm 10\%$ variation in warp tension was introduced during post-processing to simulate weaving tolerance. Our method still maintained manufacturability for 92.7% of the samples, whereas the rule-driven method dropped to 84.1%, indicating greater sensitivity to process disturbance.

In summary, the proposed method remained relatively stable under semantic, material, and process uncertainty. As shown in Figure 5, its performance degradation was consistently slower than that of the reference methods. These results also suggest that the model is not driven solely by memorization of the training set, although broader artistic diversity and rare function-style combinations would still benefit from larger-scale data.

Ablation Study

To quantify the contribution of each module to functional fabric generation, we conducted an ablation study on the TextileArt-Weave v1.0 validation set. Each configuration was repeated 5 times, and the results were reported as mean \pm standard deviation. As shown in Table 4, all metrics were centered around core textile requirements: PE, conductive connectivity, semantic alignment (CLIP Score), and ICR.

After removing the functional constraint module, the PE increased significantly to $11.7 \pm 0.9\%$, and η dropped to 0.38 ± 0.04 , indicating that the model could no longer meet the physical performance targets automatically. However, the ICR only slightly decreased ($95.1 \pm 1.2\%$), showing that manufacturability is not dependent on functionality supervision and that the two can be decoupled for optimization. When CLIP semantic alignment was disabled, the CLIP Score dropped from 0.78 ± 0.02 to 0.41 ± 0.03 ($p < 0.01$), but both PE and ICR

remained almost unchanged, verifying the independence of the artistic expression pathway. The removal of the interlacing cycle regularization had the greatest impact on manufacturability: the ICR dropped sharply from $96.8 \pm 0.7\%$ to $72.3 \pm 3.1\%$, with a significant increase in standard deviation, indicating severe structural instability and a higher likelihood of non-manufacturable local defects. Finally, removing the curriculum learning strategy led to increased training convergence fluctuations, with the ICR standard deviation widening from $\pm 0.7\%$ to $\pm 2.4\%$, and the final performance also declined (ICR = $92.6 \pm 2.4\%$). This result suggests that the curriculum strategy mainly improves optimization stability by first guiding the model toward structurally feasible and functionally consistent textile configurations, and then progressively increasing semantic and stylistic complexity during training. It does not imply that manufacturability and aesthetics are physically produced in separate stages during weaving; rather, it is a training schedule for learning their coupled representation more robustly.

In summary, each component ensures functionality compliance, semantic responsiveness, structural reliability, and training stability, collectively supporting the “function-art-manufacturing” collaborative generation goal.

Table 4. Ablation Study Results (Validation Set, mean \pm standard deviation, n=5)

Variant	PE (%) \downarrow	η \uparrow	CLIP Score \uparrow	ICR (%) \uparrow
Full Model	3.2 ± 0.4	0.93 ± 0.02	0.78 ± 0.02	96.8 ± 0.7
–Functional Constraints	11.7 ± 0.9	0.38 ± 0.04	0.77 ± 0.02	95.1 ± 1.2
–CLIP Semantic Alignment	3.3 ± 0.3	0.92 ± 0.03	0.41 ± 0.03	96.5 ± 0.8
–Interlacing Cycle Regularization	3.5 ± 0.5	0.91 ± 0.02	0.76 ± 0.02	72.3 ± 3.1
–Curriculum Learning Strategy	4.1 ± 0.6	0.89 ± 0.03	0.74 ± 0.03	92.6 ± 2.4
Variant	PE (%) \downarrow	(η) \uparrow	CLIP Score \uparrow	ICR (%) \uparrow

Note: \downarrow indicates that lower values are better, \uparrow indicates that higher values are better; η is only calculated for the subset of samples with conductive functionality enabled; all differences were statistically significant based on paired t-tests, with $p < 0.01$ considered significant.

DISCUSSION

The core value of this study lies in developing a fabric generation framework that jointly considers functionality, artistry, and manufacturability. Its effectiveness comes from formulating textile design as a structured conditional generation problem in which functional requirements, artistic semantics, and process constraints are optimized jointly. In this sense, the proposed method is not intended to simply reproduce the deterministic

outputs of tools such as TexGen or Blender; rather, these tools are only used for dataset construction, while the model learns a multimodal prompt-conditioned mapping to manufacturable textile structures. Functional objectives are encoded as yarn-level structural constraints, artistic intent is introduced through multimodal alignment, and manufacturability is maintained through process-aware regularization. This allows the generated results to move beyond purely visual patterns or purely engineering-driven structures toward practically usable fabric designs. The curriculum learning strategy should be understood as a training schedule rather than a physical weaving sequence: it first stabilizes learning in the structurally feasible design space and then gradually increases semantic and stylistic complexity during optimization.

However, the current framework still has several limitations. It remains less reliable when handling highly abstract or vague artistic descriptions, which may cause mismatches between intended style and functional structure. In addition, the current generator outputs a 2D/2.5D textile mesostructure representation rather than a fully explicit 3D yarn geometry. Accordingly, quantities such as porosity are estimated through surrogate models under predefined process assumptions, including yarn diameter, layer thickness, and stacking relations, rather than being derived from unrestricted volumetric reconstruction. Although this approximation improves computational efficiency, it simplifies complex multi-physics behavior and may not fully reflect real service conditions. Likewise, the training curriculum does not imply that manufacturability and aesthetics are physically realized in separate stages in actual weaving; in practice, pattern expression and weave structure remain inherently coupled. The dataset, while covering representative functions and styles, remains limited in scale relative to the combinatorial space of functional requirements and artistic styles. Although the robustness and ablation results suggest that the model does not simply memorize training samples, the current data scale may still constrain stylistic diversity and long-tail generalization, especially for rare function-style combinations. Moreover, although the current comparison covers engineering-oriented and visual-generation references, it does not yet include discrete generative models specifically designed for structured textile representations, such as VQ-VAE- or Transformer-based sequence generators, which would provide a stricter structural comparison.

These limitations suggest several future directions. Incorporating textile domain knowledge, such as a textile knowledge graph, may improve the mapping from semantic descriptions to structural parameters. Linking the generation framework with digital looms and real measurement feedback could further support a closed-loop “design-manufacture-validation” workflow. Future evaluation can also be strengthened by introducing more

structure-aware generative baselines for discrete weave and yarn-path generation. More broadly, integration with sustainable material databases and lifecycle assessment tools may extend the framework toward greener and more deployable intelligent textile design.

CONCLUSION

This study addresses the long-standing disconnect among artistic expression, physical performance, and manufacturability in functional fabric design by proposing a multimodal collaborative generation method tailored to the textile domain. The method uses manufacturable yarn paths as a unified representation and integrates weave pattern representations, functional surrogate models, and interlacing cycle regularization mechanisms. It achieves end-to-end generation from natural language descriptions to mesostructures on the TextileArt-Weave v1.0 dataset. Experimental results on the test set show that the proposed method significantly outperforms existing baselines on four core metrics: porosity error ($3.2 \pm 0.4\%$), conductive connectivity (0.93 ± 0.02), semantic alignment (CLIP Score 0.78 ± 0.02), and interlacing cycle compliance rate ($96.8 \pm 0.7\%$). These results indicate that the proposed framework can jointly address functional, artistic, and manufacturability requirements within a unified generative setting.

The main contributions of this study are as follows. First, we constructed the TextileArt-Weave v1.0 dataset, containing 8,642 samples, which provides aligned weave patterns, functional fields, cultural semantic texts, and multi-view images for collaborative functional fabric generation. Second, we proposed quantifiable textile manufacturing constraints, such as ICR, and embedded them into the generative loss function, resulting in a process feasibility of $96.8 \pm 0.7\%$ on the test set, far outperforming pure visual generation models (ICR = 0%). Third, we validated the necessity of functional constraints, semantic alignment, and curriculum learning strategies through ablation studies, with the removal of the interlacing cycle regularization causing ICR to drop sharply to $72.3 \pm 3.1\%$ on the validation set, highlighting the value of domain-specific design. Academically, this work extends generative modeling to textile manufacturing and advances a paradigm of structurally executable generation. Practically, it supports tasks such as conductive pattern layout for smart scarves and the embedding of intangible heritage patterns into breathable sports fabrics, reducing manual sampling effort and improving deployability.

Based on the limitations of the current work, future research could expand in three directions: first, enhancing the structured parsing ability of abstract semantics by integrating a textile knowledge graph to improve the robustness of semantic-structure mapping; second, constructing a generative-weaving-feedback closed-loop

system to iteratively optimize surrogate models based on real-world performance feedback; third, extending the method to multi-material gradient fabrics or biomimetic microstructure design. Furthermore, the “function-semantic-process” collaborative modeling approach in this framework may also be transferable to fields such as composite material layups and flexible electronic device layouts, providing methodological support for meso-scale intelligent design.

Author Contributions

Conceptualization – Yi Tang, Congyao Xu, Yifan Xue, Zhongfei Han; methodology – Yi Tang, Congyao Xu, Yifan Xue, Zhongfei Han; formal analysis – Yi Tang, Congyao Xu, Yifan Xue, Zhongfei Han; investigation – Yi Tang, Congyao Xu, Yifan Xue, Zhongfei Han; resources – Yi Tang, Congyao Xu, Yifan Xue, Zhongfei Han; writing-original draft preparation – Yi Tang, Congyao Xu, Yifan Xue, Zhongfei Han; writing-review and editing – Yi Tang, Congyao Xu, Yifan Xue, Zhongfei Han; visualization – Yi Tang, Congyao Xu, Yifan Xue, Zhongfei Han; supervision – Yi Tang, Congyao Xu, Yifan Xue, Zhongfei Han. All authors have read and agreed to the published version of the manuscript.

Conflicts of Interest

The authors declare no conflict of interest.

Funding

This research received no external funding.

Hazards

Not Applicable

Human research subjects

Not Applicable

Animal research subjects

Not Applicable

Acknowledgements

Not Applicable

REFERENCES

- [1] Wu X, Li L. An application of generative AI for knitted textile design in fashion. *The Design Journal*. 2024; 27(2): 270-290. doi: 10.1080/14606925.2024.2303236

- [2] Liang J. The application of artificial intelligence-assisted technology in cultural and creative product design. *Scientific Reports*. 2024; 14(1):31069. doi: 10.1038/s41598-024-82281-2
- [3] Xu Y, Nazir S. Ranking the art design and applications of artificial intelligence and machine learning. *Journal of Software: Evolution and Process*. 2024; 36(2):e2486. doi: 10.1002/smr.2486
- [4] Zhang B, Romainoor NH. Research on artificial intelligence in new year prints: the application of the generated pop art style images on cultural and creative products. *Applied Sciences*. 2023; 13(2):1082. doi: 10.3390/app13021082
- [5] Glogar M, Petrak S, Mahnić Naglič M. Digital technologies in the sustainable design and development of textiles and clothing-a literature review. *Sustainability*. 2025; 17(4):1371. doi: 10.3390/su17041371
- [6] Xu Y, Zhi C, Guo H, Zhang M, Wu H, Sun R, Yu L. ChatGPT for textile science and materials: A perspective. *Materials Today Communications*. 2023; 37:107101. doi: 10.1016/j.mtcomm.2023.107101
- [7] Li R, Zhang S, Zhang R. Recent progress in artificial structural colors and their applications in fibers and textiles. *Chemistry-Methods*. 2023; 3(9):e202200081. doi: 10.1002/cmttd.202200081
- [8] Guo Z, Zhu Z, Li Y, Cao S, Chen H, Wang G. AI assisted fashion design: A review. *IEEE Access*. 2023; 11:88403-88415. doi: 10.1109/ACCESS.2023.3306235
- [9] Gbr BA. Robotecture and artificial intelligence (AI) technology and its impact on the creativity process in interior spaces. *International Design Journal*. 2023; 13(4):215-233. doi: 10.21608/idj.2023.210586.1076
- [10] Simões S. High-performance advanced composites in multifunctional material design: State of the art, challenges, and future directions. *Materials*. 2024; 17(23):5997. doi: 10.3390/ma17235997
- [11] Liu Y, Chen L, Li W, Pu J, Wang Z, He B, Wei L. Scalable production of functional fibers with nanoscale features for smart textiles. *ACS Nano*. 2024; 18(43):29394-29420. doi: 10.1021/acsnano.4c10111
- [12] Deng M, Liu Y, Chen L. AI-driven innovation in ethnic clothing design: an intersection of machine learning and cultural heritage. *Electronic Research Archive*. 2023; 31(9). doi: 10.3934/era.2023295
- [13] Lyu Y, Shi M, Zhang Y, Lin R. From image to imagination: Exploring the impact of generative AI on cultural translation in jewelry design. *Sustainability*. 2024; 16(1):65. doi: 10.3390/su16010065
- [14] Mu X, Zhang H, Shi J, Hou J, Ma J, Yang Y. Fashion intelligence in the Metaverse: promise and future prospects. *Artificial Intelligence Review*. 2024; 57(3):67. doi: 10.1007/s10462-024-10703-8
- [15] Gür M, Çorakbaş FK, Atar İS, Çelik MG, Maşat İ, Şahin C. Communicating AI for architectural and interior design: Reinterpreting traditional Iznik tile compositions through AI software for contemporary spaces. *Buildings*. 2024; 14(9):2916. doi: 10.3390/buildings14092916

- [16] Ukpaka P. Exploring creativity in fashion industry applications of generative adversarial networks (GANs). Available at SSRN 4833151. 2024. doi: 10.2139/ssrn.4833151
- [17] Paklapas M, Tangkiatphaibun P, Punyalikit R, Joneurairatana E. Combining AI and Creative Design for Automated Fabric Defect Analysis and Pattern Generation from Common Imperfections. *Journal of Roi Kaensarn Academi*. 2024; 9(12):3063-3075.
- [18] Tanugraha S. A review using artificial intelligence-generating images: Exploring material ideas from MidJourney to improve vernacular designs. *Journal of Artificial Intelligence in Architecture*. 2023; 2(2):48-57. doi: 10.24002/jarina.v2i2.7537
- [19] Bishara D, Xie Y, Liu WK, Li S. A state-of-the-art review on machine learning-based multiscale modeling, simulation, homogenization and design of materials. *Archives of Computational Methods in Engineering*. 2023; 30(1):191-222. doi: 10.1007/s11831-022-09795-8
- [20] Cakiroglu C. Explainable data-driven ensemble learning models for the mechanical properties prediction of concrete confined by aramid fiber-reinforced polymer wraps using generative adversarial networks. *Applied Sciences*. 2023; 13(21):11991. doi: 10.3390/app132111991
- [21] Malashin I, Tynchenko V, Gantimurov A, Nelyub V, Borodulin A. A multi-objective optimization of neural networks for predicting the physical properties of textile polymer composite materials. *Polymers*. 2024; 16(12):1752. doi: 10.3390/polym16121752
- [22] Li M, He J. Machine learning-based strain-load prediction of high-performance 3D woven fabrics. *Journal of Industrial Textiles*. 2024; 54:15280837241307369. doi: 10.1177/15280837241307369
- [23] Xu Y, Sun R, Zhi C, Liu Z, Chen J, Ke Z, Yu L. Zero-defect manufacturing in the textile industry: A review of current advances and challenges. *Textile Research Journal*. 2025; 95(1-2):194-219. doi: 10.1177/00405175241260892
- [24] Chai J, Wang G, Wang G, Shao R, Zhao J, Zhao G, Park CB. Porous and Conductive Fiber Woven Textile for Multi-Functional Protection, Personal Warmth, and Intelligent Motion/Temperature Perception. *Advanced Functional Materials*. 2025; 35(10):2416428. doi: 10.1002/adfm.202416428
- [25] Zhang B, Song L, Rahmat-Sammi Y. Textile patch antenna surrogate-based optimization: Kriging surrogate modeling on equivalent circuit components. In: *Proceedings of the 2023 United States National Committee of URSI National Radio Science Meeting (USNC-URSI NRSRM)*; 4-8 Jan 2023; Boulder, CO, USA. IEEE; 2023. p. 145-146. doi: 10.23919/USNC-URSINRSM57470.2023.10043169
- [26] Zhao X, Li M, Shi X. MCP-MOE: a method for assessing the value of cultural products for textile and apparel design. *Textile Research Journal*.

- [27] Matos J, Cunha A, Faria A, Vale B, Caldas L, Monteiro M, Duarte O. Digital Product: a Textile use case. In: Proceedings of the 2024 IEEE International Conference on Engineering, Technology, and Innovation (ICE/ITMC); 17-20 Jun 2024; Nice, France. IEEE; 2024. p. 1-8. doi: 10.1109/ICE/ITMC61926.2024.10794226
- [28] Devendorf L, Walters K, Fairbanks M, Sandry E, Goodwill ER. AdaCAD: Parametric design as a new form of notation for complex weaving. In: Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems; 23-28 Apr 2023; Hamburg, Germany. Association for Computing Machinery; 2023. p. 1-18. doi: 10.1145/3544548.3581571
- [29] Adeyokunnu O, Abayomi-Alli A, Alabi OA, Arogundade OT, Daniel DW, Omoyiola BO. Classification and Augmentation of African Fabrics Using Convolutional Neural Network and Generative Adversarial Network Models. In: Proceedings of the 2024 IEEE 5th International Conference on Electro-Computing Technologies for Humanity (NIGERCON); 12-14 Nov 2024; Lagos, Nigeria. IEEE; 2024. p. 1-5. doi: 10.1109/NIGERCON62786.2024.10927147