Bio-Fusion Tectonics

The operation of AIGC in generative design and robotic additive manufacturing processes

Chengbi Duan[1], Dasong Wang[2], Hao Hua[3]
[1,2,3]School of Architecture, Southeast University, Nanjing, China
[1,2,3]{220230206|dasong.wang|whitegreen}@seu.edu.cn

This paper proposes a novel generative design approach that synthesize AI generated biomimetic pattern with robotic additive manufacturing constraints for developing innovative architectural tectonics. The design approach is established against the critical challenge in the field of AI-assisted architectural formation: the disjunction between early-stage design intentions and the feasibility of physical construction. Although Generative Artificial Intelligence tools offer abundant visual inspiration during conceptual design phases, they often lack embedded understanding of construction logic, resulting in formally expressive yet unbuildable outcomes. The proposed methodology integrates diffusion models, LoRA fine-tuning, and Control-Net to guide the generation of biomimetic grayscale textures, which are subsequently transformed into three-dimensional panelized components optimized for robotic additive fabrication. Using the "Morphusion: Composite Sprawling" (Figure 1) project as a design-research demonstrator, this study validates a complete workflow of "AI Generation—Parameter Optimization—Robotic Construction." The workflow incorporates human-dominant design intent, algorithmic control, and robotic execution, thereby enhancing both structural integrity and surface articulation in the fabricated elements. By introducing the generative design procedure, the paper expands the boundaries of constructive morphology and providing a dual foundation of theory and methodology for architectural construction in the era of spatial intelligence.

Keywords: *Architectural Tectonics, Generative Design, AIGC, Diffusion Model, Robotic Additive Manufacturing.*

INTRODUCTION

In recent years, paradigm shifts in computational design methodologies, alongside continuous advances in computational power, have significantly facilitated the integration of generative design strategies with robotic construction technologies. This convergence is profoundly reshaping architectural tectonic manifestation, materiality, and constructability. Generative design, in particular, is redefining the power dynamics between design and construction, as the authority of form generation gradually shifts from human designers to algorithmic systems and data-driven models.

Within rule-based generative design approaches, the logic of form generation typically relies on predefined parameters, algorithmic workflows, or geometric strategies. These algorithmic methods have demonstrated remarkable potential in producing complex

Figure 1
Morphusion: Composite Sprawling is a biomimetic architectural installation that combines AI-generated design, robotic 3D printing, and CNC metal fabrication, showcasing a growth-inspired spatial logic through human-machine collaboration and material-system integration

morphologies, often characterized by highly articulated formal detail and emergent geometric logic. For example, studies based on multi-agent systems have revealed the evolutionary potential of biomimetic structures (Snooks, 2022). However, despite their expressive capacity at the geometric level, such rule-driven systems remain limited in their responsiveness to ambiguous design intentions or nuanced aesthetic considerations.

In contrast, the rapidly evolving trajectory of AIGC represents a qualitatively different generative paradigm. AIGC is fundamentally defined by a descriptive, intention-to-output mapping process driven by textual or visual prompts (Del Campo, Carlson and Manninger, 2021). It proves particularly powerful in early-stage design, where inspiration and visual ideation are crucial. Diffusion models, when guided by semantic prompts, offer more fluid and intuitive forms of input, liberating the conceptual space of form-making from the deterministic logic of conventional algorithms. In current practice, AI is often employed as a source of visual inspiration in 2D image-based tasks, such as architectural façade design (Ma and Zheng, 2024), spatial layout sketching, or the digital reconstruction of traditional ornamental textures (Zhong et al., 2025). Yet these strategies tend to decouple design generation from construction processes, with fabrication logic treated as a downstream rationalization, rarely participating in the formative stages of design (Wu et al., 2021).

Meanwhile, robotic additive manufacturing has made significant progress in enabling high-resolution, geometrically controlled production of large-scale architectural components. Notably, the work of Benjamin Dillenburger and his team has yielded substantial outcomes in non-planar 3D printing toolpaths (Mitropoulou, Bernhard and Dillenburger, 2025) and shape variation techniques for self-supporting cantilevered structures (Huang et al., 2024). These contributions primarily focus on enhancing precision in path generation, material control, and structural optimization. However, fabrication is still treated as a post-design rational process, with geometries often predetermined and disconnected from the generative logic of design intention.

This study positions itself at the intersection of these two trajectories and proposes an integrated "design-to-construction" framework

powered by AIGC. The proposed method synthesizes grayscale image generation via diffusion models, LoRA-based fine-tuning, and ControlNet guidance mechanisms to produce biomimetic grayscale images embedded with constructible logic. At the generation stage, these images already encode fabrication parameters such as overhang angles, path density, and directional attributes. Subsequently, a grayscale-to-height mapping algorithm translates these images into 3D components optimized for robotic additive manufacturing, achieving a continuous transformation from design intention to constructible geometry.

This methodology is demonstrated through the "Morphusion: Composite Sprawling" project. In this case, semantic prompt-driven generation of venation-inspired biomorphic forms is combined with grayscale optimization algorithms rooted in construction logic, resulting in panelized components fabricated through robotic extrusion. The project illustrates a complete pipeline—from AI-guided generation and parametric optimization to robotic fabrication—exemplifying a full-chain realization of "semantic control–parameter tuning–material execution."

At the theoretical level, this paper proposes a construction-oriented model of intention mapping. Rather than treating design semantics and construction feasibility as trade-offs resolved in later phases, the model repositions them as co-productive generative drivers from the very outset of the design process. This coupling preserves the semantic richness of formal exploration while simultaneously addressing the geometric discipline required for physical feasibility. As such, the proposed framework introduces a continuous paradigm of architectural production that synthesizes abstraction and constraint, intention and execution, human and machine.

METHODOLOGY

This methodology, titled as "Bio-fusion", presents a biomimetic morphology generation and construction optimization methodology based on the Low-Rank Adaptation (LoRA) fine-tuning model, with a focus on the design and construction processes of the Morphusion project(Figure 2). The method utilizes generative

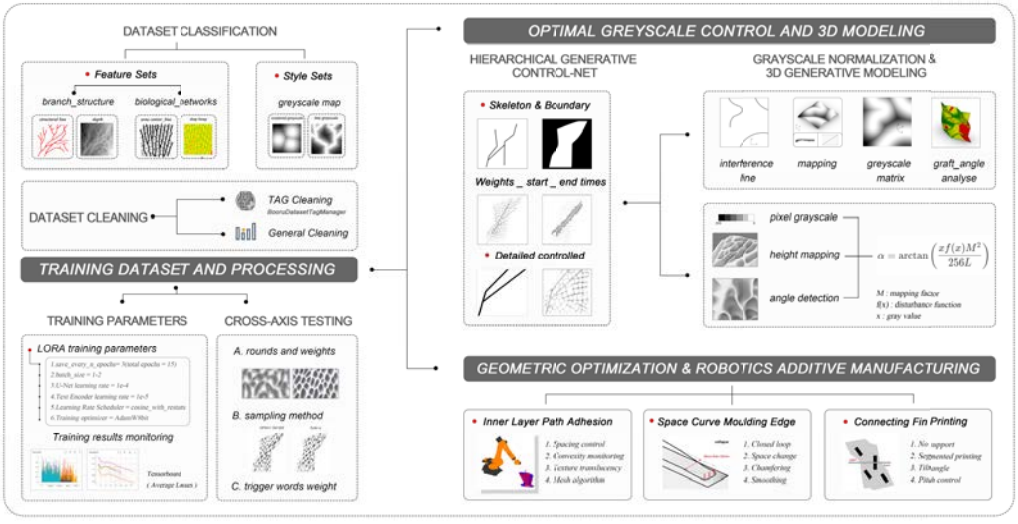

Figure 2
A methodological framework for bionic morphogenesis and construction process optimization based on Low-Rank Adaptation (LoRA) fine-tuning model.

artificial intelligence (Generative AI) to analyze the fractal geometric characteristics of vine structures and trains a fine-tuning model with hyperparameter topological geometry control capabilities and additive manufacturing parameter adaptability. The aim is to enhance both the design strategies and manufacturing accuracy of complex textured 3D printed components.

In addition to focusing on technical parameter optimization, the Morphusion experimental system is conceived as a testing platform to examine how generative semantics and construction logic can be integrated into spatially expressive and structurally tense architectural forms. The spatial logic is organized as a nested network of curvilinear, branching elements that progressively differentiate, forming a configuration characterized by suspended facades and spatial derivations. In terms of panel texture and structural strategy, the components embed biomimetic vein-like reinforcement ribs, which not only enhance the structural performance of individual panels but also contribute to a design language marked by tensile coherence and geometric order. While each panel's textural detail is uniquely generated by the trained model, they collectively maintain a unified construction logic and a consistent hierarchy of primary and secondary structural frameworks, all within the physical constraints of robotic additive manufacturing. This synergistic articulation of spatial logic and construction grammar demonstrates how AI-driven design

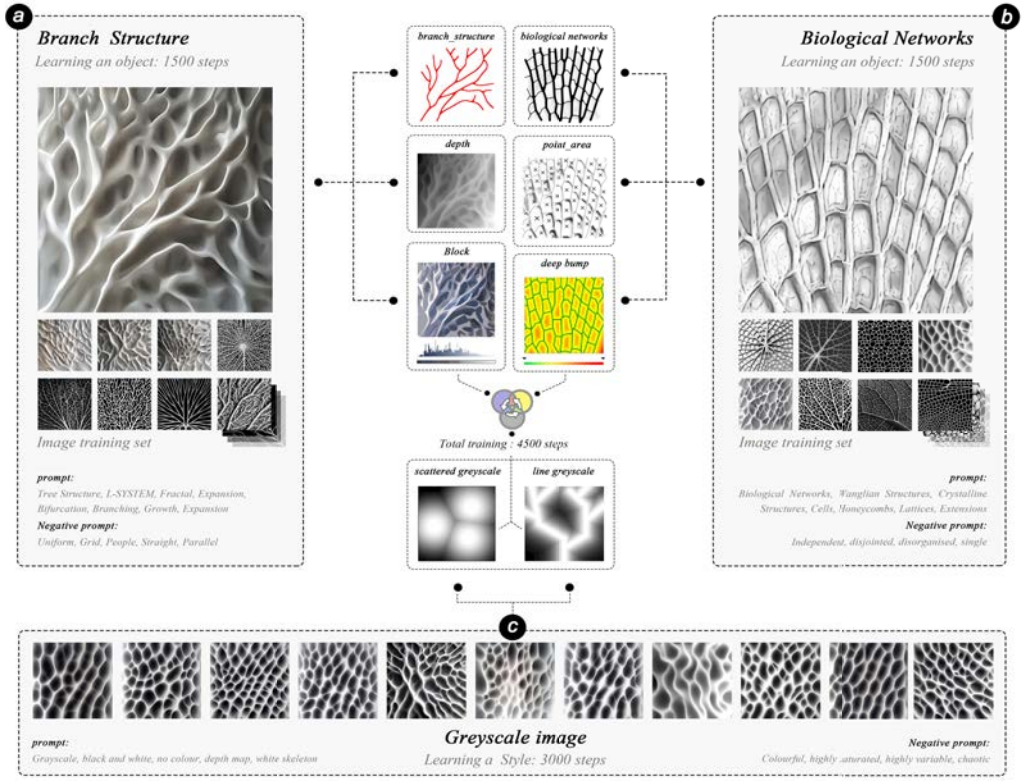

Figure 3
Dataset construction and style training: (a) Branch Structure; (b) Biological Networks; (c) Greyscale Image Training

workflows can embed the designer's logic and generative creativity into the fabrication process.

Training-set cleaning & synthesis

This section introduces the construction of a multi-level feature dataset, grayscale style constraints, and the implementation of data cleaning and label management for LoRA model training. The aim is to enhance the AI model's capacity for biomimetic morphology learning, controllable generation, and adaptability to additive manufacturing. In the Morphusion project, vine-like geometric features are deconstructed to establish a unified design language, ensuring that generated 2D grayscale images can be effectively mapped to 3D structures. The training dataset includes two feature sets and one style set, focused on learning the geometric growth patterns of vines and ensuring compatibility with 3D fabrication workflows.

1. Feature Sets

The feature sets simulate vine growth characteristics through two data types:

- Branch Structure (Figure 3a): Captures divergent growth and branching hierarchies, emphasizing fractal logic and morphological variation.

- Biological Networks (Figure 3b): Represents interconnections between vine branches by simulating stable, enclosed network forms with convex boundary conditions to reinforce adaptive connectivity.

2. Style Sets

The style set constrains output formats to ensure grayscale images can be converted into 3D models for additive manufacturing:

- Greyscale Style (Figure 3c): Uses white for raised skeletons and black for recessed areas, ensuring smooth gradients for 3D transitions and avoiding contrast-induced defects. Includes two variations (Scattered and Linear Greyscale) to support diverse printing strategies.

In the data cleaning and label management phase, the Booru Dataset Tag Manager is used to optimize images and manage trigger word labels by removing noise and redundant color. Dataset groups align with model trigger words, with standardized descriptions derived from feature sets. Trigger words are weighted at over three times normal descriptors to strengthen generation control. To avoid overfitting, unique labels tied to individual images are excluded during training to improve generalization.

LoRA model training with hyperparameter optimization

This study fine-tunes Stable Diffusion 1.5 using 512×512 pixel images. A layered learning rate (1e-4 for U-Net, 1e-5 for text encoder) is applied with cosine annealing restarts and the AdamW8bit optimizer to balance adaptability and memory efficiency. Training runs for up to 15 iterations, saving checkpoints every 3 iterations, with a batch size of 1–2 to suit an 8GB VRAM environment. Small-batch settings introduce gradient noise to improve generalization and morphology generation. Training progress is monitored via TensorBoard to track loss fluctuations and convergence trends (Figure 4).

Figure 4
LoRA hyperparameter training and Tensorboard monitoring

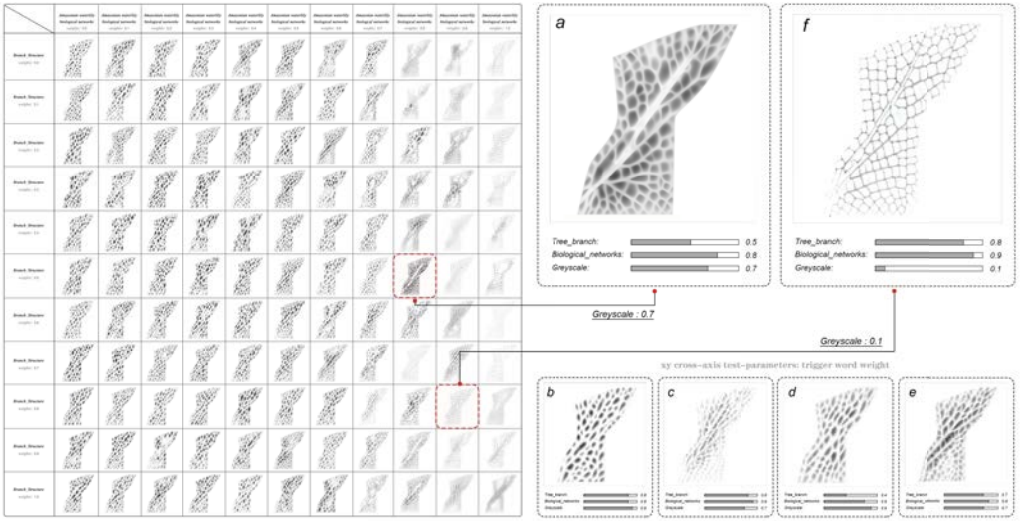

Figure 5
Cross-axis parameter testing 03: effects of prompt weights (branch structure, biological networks, and greyscale) on morphology generation

Cross-Axis Parameter Testing 01 – LoRA Iterations and Weighting: Low LoRA iterations (0–6) or weights below 0.5 lead to instability and detail loss, while overly high weights (0.9–1.0) produce sharp textures but cause local over-sharpening and grayscale loss, compromising 3D convertibility. Optimal plant vine morphology with stable grayscale and good 3D potential is achieved at 9–15 iterations and weights between 0.7–0.9.

Cross-Axis Parameter Testing 02 – Sampling Methods: This test compares sampling methods based on geometric detail, structural stability, and grayscale smoothness. DDIM and PLMS show significant detail loss and uneven transitions, proving unsuitable. Euler series offers faster computation but produces blurred textures and weak detail. DPM++ 2M and DPM++ 3M SDE yield consistently clearer details and smoother grayscale gradients.

Cross-Axis Parameter Testing 03 – The Impact of Trigger Word Weight on Generated Morphology(Figure 5)**:** Adjusting the weights of feature and style set trigger words reveals that combining Branch Structure and Biological Networks produces morphologies that effectively merge branching and network features, closely mimicking natural vine patterns. Greyscale weight significantly influences output style: lower values (0.1–0.3) yield line-drawing-like images (Figure 5F), ideal for contour analysis; higher values (0.7–0.9) create smoother gradients with enhanced relief and 3D depth (Figure 5A, B, E), better suited for grayscale-based modeling.

Hierarchical generative details via control-net

This section investigates ControlNet's hierarchical control over vine-inspired textures. By tuning weight parameters and guidance timing, the method balances structural precision with generative flexibility. In Morphusion, the Lineart module constrains the primary skeleton to match the steel framework while preserving expressive detail. Combined with the LoRA model, testing focuses on generating skeletons and networks, with grayscale weight set to zero to emphasize structural clarity in line-drawing form.

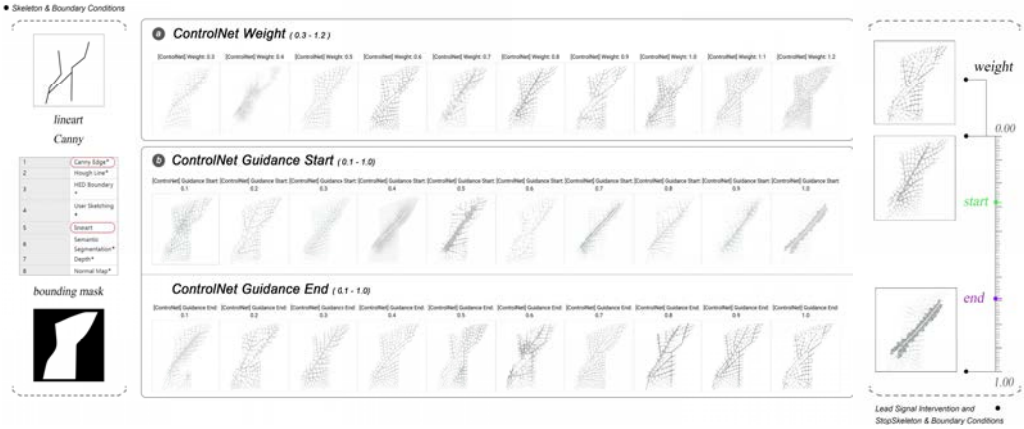

Figure 6
The Impact of ControlNet weight and guidance signal timing on shape generation

As ControlNet weight increases, the model increasingly relies on line-drawing inputs, improving morphological accuracy. Weights of 0.3–0.5 yield slightly blurred results; 0.6–1.0 offers clear structural definition; weights above 1.1 lead to overly rigid forms, limiting flexibility. In guidance timing tests, early-stage intervention secures core structure, mid-stage refines details, and late-stage release enhances generative freedom. Adjusting these parameters balances detail control with stylistic consistency(Figure 6).

Combining the LoRA trained model with ControlNet parameters (weight 1.1; guidance applied and withdrawn between stages 0.1–0.9), detailed generation tests were conducted on vine-like geometries (Figure 7). The process begins with preserving a red Y-shaped primary skeleton—symbolizing human-centered control—while ControlNet imposes structural constraints. A blue secondary structure then branches outward from the primary frame, forming its spatial support and extension. Finally, a green cellular ring network forms between the primary and secondary structures, enhancing the overall naturalism and complexity. This test confirms the LoRA model's effectiveness in generating vine-like biomimetic textures and identifies key ControlNet parameters, offering practical guidance for the controlled generation of diverse texture schemes.

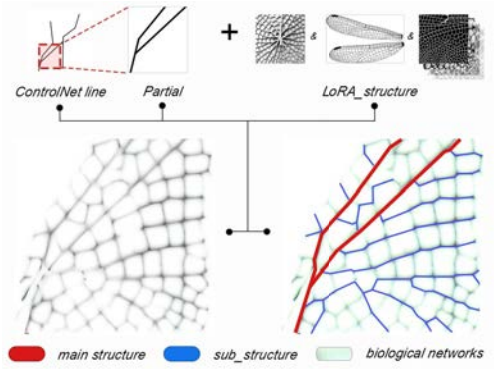

Figure 7
ControlNet for tendril biomimetic detail structure generation test

Grayscale normalization & 3d generative modeling

This section introduces a grayscale-based optimization method for 3D model generation, converting grayscale images into additive-manufacturing-ready geometries via height field mapping (Figure 8). The method refines draft angles and supports robotic arm printing. Grayscale pixel heights are used to form point clouds, from which Delaunay triangulation generates the mesh. A Bézier curve perturbation further enhances surface smoothness and detail, enabling accurate transformation from 2D biomimetic textures to 3D printable forms.

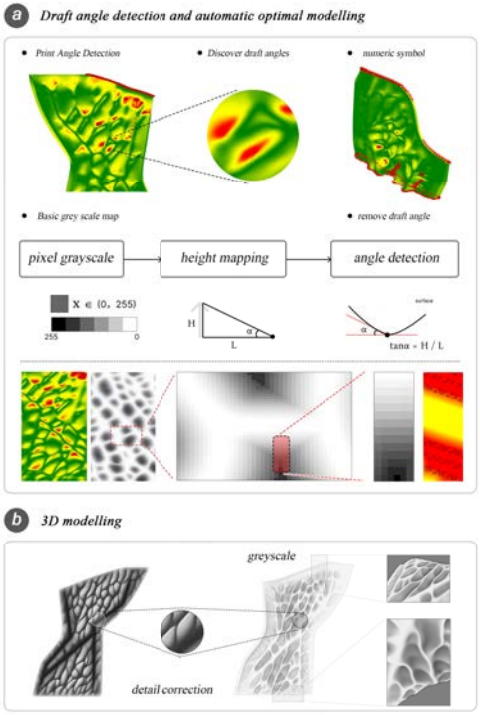

the Delaunay triangulation is applied to generate a mesh structure. Additionally, a Bézier curve perturbation function is employed to enhance surface smoothness and detail expression, enabling precise conversion from 2D biomimetic textures to 3D printed structures.

Pixel grayscale variation determines the model's draft angle. Grayscale values (x) range from 0 (black) to 255 (white) and are scaled using a mapping factor (M) to control the amplitude of morphological fluctuations. A Bézier-based

$$a = arctan\left(\frac{\pi f(x) M^2}{256L}\right) \quad (1)$$

perturbation function, replacing linear grayscale mapping, introduces smooth chamfered transitions in surface height. These height values, combined with the horizontal pixel spacing (L), are used to calculate the overhang printing angle (α). When α exceeds 40 degrees, adaptive inverse mapping is applied based on empirical parameters. This optimization process ensures a seamless transition from digital modeling to physical construction, bridging the qualitative 2D biomimetic morphology generation with the quantitative 3D printing parameters, thereby guaranteeing the print quality and manufacturability of the generative AI results.

Figure 8
Grayscale value optimization and additive manufacturing angle adjustment

Geometric optimization & robotics additive manufacturing

This study utilizes a KUKA KR210 R2700-2 six-axis robotic arm equipped with a 6 mm particle fusion extrusion nozzle to print PETG (K2012) granules. The material is melted via a three-stage temperature control at 230 °C and extruded in 8 mm-wide layers. The maximum draft angle is limited to under 42°, with all physical parameters

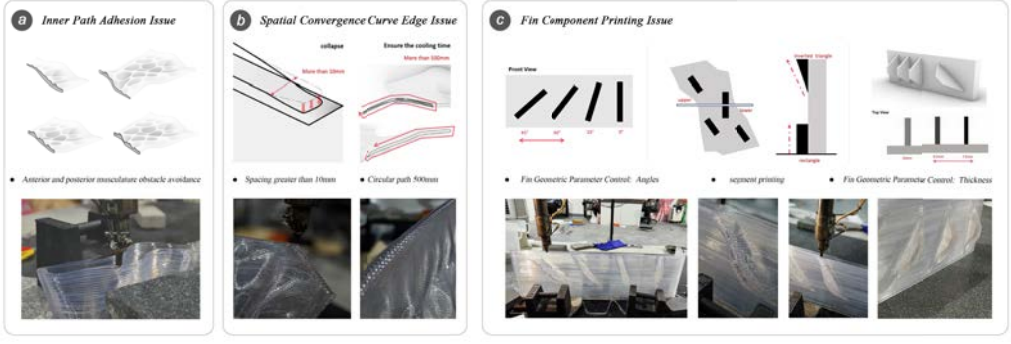

Figure 9
Detail-aware geometric optimization for fabrication: (a) optimizing inner path spacing to prevent adhesion; (b) circular path design to address spatial convergence and ensure cooling; (c) geometric control of fin component printing, focusing on angle and thickness

calibrated accordingly. To address the complexity of printing intricate textures, a series of topology optimization strategies are applied(Figure 9): (1) path spacing detection and mesh pulling in Grasshopper resolve material adhesion caused by overly dense inner textures; (2) an adaptive path planning algorithm adjusts spacing and print speed at overhanging convergence edges to prevent collapse; (3) for fin-like connectors, panel segmentation and horizontal baseline supports mitigate unsupported overhangs, ensuring print fidelity within a 7.5–8.5 mm thickness range.

After geometric optimization, 44 customized panel components are generated and converted into SRC-format robotic motion instructions. Using Rhino and Grasshopper, along with the KUKA|prc control plugin, the printing process is simulated. By tracking six-axis motion parameters and conducting collision detection, the workflow ensures mechanical safety during formal robotic fabrication.

RESULTS & DISCUSSION

Building on the proposed methodology, the Morphusion project realized a complete workflow from AI-driven generative design to physical fabrication, producing 44 complex panels with varied textures and boundary conditions. Each panel's primary skeleton was parametrically optimized to align with the steel framework, exhibiting vine-like undulations. The installation merges glacial transparency with biomimetic growth logic, extending from floor to facade to create a spatial experience of dynamic tension (Figure 10).

Figure 10
Spatial effects of the completed Morphusion project

A dual-layer style control mechanism guided the grayscale image generation. The first layer, trained via LoRA, ensured stylistic and geometric consistency; the second allowed designer-led aesthetic adjustments. Four designers employed the same workflow to generate the 44 panels. Results revealed that, even within a unified system, individual decisions shaped distinctive morphological outcomes—demonstrating "controlled variation." This tension between control and freedom reinforces designer agency, distinguishing AI-assisted workflows from conventional algorithmic uniformity (Figure 11).

Architecturally, while the study integrates design and fabrication, limitations remain. The model's generalization to heterogeneous materials and multi-scale components needs further testing. Robotic fabrication still faces challenges such as path deviations, thermal distortion, and delayed material response. Future work should incorporate structural feedback and behavioral recognition to improve robustness and scalability.

Bio-Fusion Tectonics serves as a research-through-design artifact, integrating parameter control, model response, and fabrication constraints within a human–machine co-design framework. As a prototype, it facilitates reflection on the evolving role of design authorship in AI-assisted processes. Its use of pretrained models reduces operational complexity, supporting iterative development of complex geometries. The project's value lies in advancing methodological inquiry into integrated design-to-construction workflows in architecture.

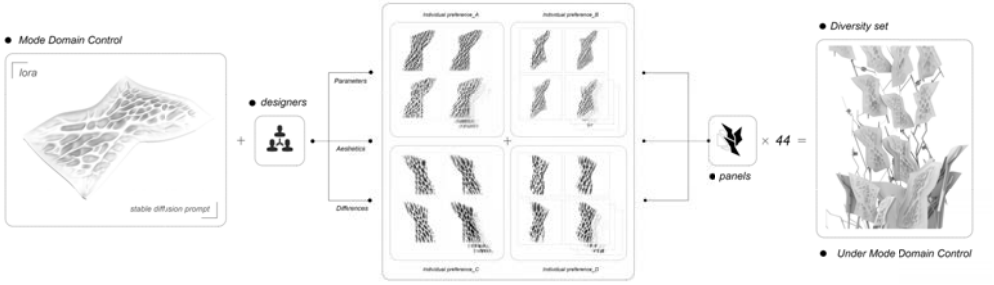

Figure 11
The two levels of control, pattern domain control and individual adjustment, work in concert for the Morphusion project

CONCLUSION

This study presents a generative AI-based method for biomimetic morphology generation and additive manufacturing optimization, integrating digital design with physical construction. Utilizing latent diffusion models, LoRA fine-tuning, and ControlNet, the workflow transforms style intentions into 2D grayscale images and subsequently into 3D geometric models. It achieves efficient cross-modal synthesis, expressive texture control, and addresses draft angle constraints in fabrication. Validating through an experimental fabrication project, it demonstrates a paradigm shift in design generation, emphasizing designer agency within an AI framework. The proposed "fixed process + variable intent" model supports scalability across design scenarios. Ultimately, this research contributes to a human–machine collaborative design approach that merges generative abstraction with manufacturing logic, offering an innovative pathway for computationally integrated architectural design.

REFERENCES

Snooks, R. (2022) 'Behavioral tectonics: agentBody prototypes and the compression of tectonics', *Architectural Intelligence*, 1(1), p. 9. https://doi.org/10.1007/s44223-022-00007-8

Del Campo, M., Carlson, A. and Manninger, S. (2021) 'Towards Hallucinating Machines - Designing with Computational Vision', *International Journal of Architectural Computing*, 19(1), pp. 88–103. Available at: https://doi.org/10.1177/1478077120963366

Ma, H. and Zheng, H. (2024) 'Text Semantics to Image Generation: A Method of Building Facades Design Base on Stable Diffusion Model', in C. Yan et al. (eds) *Phygital Intelligence*. Singapore: Springer Nature Singapore (Computational Design and Robotic Fabrication), pp. 24–34. Available at: https://doi.org/10.1007/978-981-99-8405-3_3

Zhong, X. *et al.* (2025) 'Image inpainting using diffusion models to restore eaves tile patterns in Chinese heritage buildings', *Automation in Construction*, 171, p. 105997. https://doi.org/10.1016/j.autcon.2025.105997

Wu, M. *et al.* (2021) 'Unraveling the capabilities that enable digital transformation: A data-driven methodology and the case of artificial intelligence', *Advanced Engineering Informatics*, 50, p. 101368. Available at: https://doi.org/10.1016/j.aei.2021.101368

Mitropoulou, I., Bernhard, M. and Dillenburger, B. (2025) 'Curvature-informed paths for shell 3D printing', *Automation in Construction*, 171, p. 105988. Available at: https://doi.org/10.1016/j.autcon.2025.105988

Huang, S. *et al.* (2024) 'Self-supporting lamellae: Shape variation methods for the 3D concrete printing of large overhang structures', *Additive Manufacturing*, 91, p. 104329. https://doi.org/10.1016/j.addma.2024.104329