ГОДИШНИК НА УНИВЕРСИТЕТА ПО АРХИТЕКТУРА, СТРОИТЕЛСТВО И ГЕОДЕЗИЯ СОФИЯ

Том 59 2026 Брой 3
Volume Issue

ANNUAL OF THE UNIVERSITY OF ARCHITECTURE, CIVIL ENGINEERING AND GEODESY SOFIA

DOI: 10.71167/uaceg.2026.590316

Received: 12.01.2026
Accepted: 26.02.2026

PROMPT-TO-FABRICATION: A PROPOSED MULTIMODAL PROCESS FOR THE AUTOMATED FABRICATION OF FAÇADE ELEMENTS

R. Shishkov[1], T. Nguyen[2], Z. Vugreshek[3]

Keywords: *generative design, diffusion models, text-to-3D, digital fabrication*

ABSTRACT

Architects increasingly use generative AI to produce novel façade concepts quickly, yet translating AI imagery into buildable components remains difficult. This paper proposes a prompt-to-fabricate workflow that couples text-to-image diffusion, AI-based 2D-to-3D mesh generation, and algorithmic geometry processing in Rhino/Grasshopper to convert visual intent into fabrication-ready panel geometry. The pipeline uses controllable image synthesis to explore design variations, then reconstructs a selected concept as a detailed mesh and rationalizes it through cleanup, topology control, thickness enforcement, and panelization. Physical production is achieved through a hybrid method in which CNC-milled XPS foam molds are used to cast or spray fiber-reinforced concrete such as GFRC. Preliminary results indicate that the approach can generate intricate, customized concrete façade elements with relatively rapid turnaround. The main advantages are fast concept iteration, continuity from generative design to fabrication constraints, and the ability to realize complex forms without fully manual modeling. Key limitations include partial controllability of AI outputs, uncertainty in mesh fidelity, the effort required for rationalization, and constraints imposed by materials and fabrication processes. The work demonstrates an integrated multimodal bridge from prompting to physical prototypes and motivates future research in performance-driven generative workflows and more automated construction pipelines.

[1] Radul Shishkov, Chief Assist. Prof. Dr. Arch., Dept. “Interior and Architectural Design”, Faculty of Architecture, UACEG, 1 H. Smirnenski Blvd., Sofia 1046, e-mail: rshishkov_far@uacg.bg
[2] Tung Nguyen, Arch. PhD student, Dept. “Interior and Architectural Design”, Faculty of Architecture, UACEG, 1 H. Smirnenski Blvd., Sofia 1046
[3] Zvonko Vugreshek, Arch. PhD student, Dept. “Interior and Architectural Design”, Faculty of Architecture, UACEG, 1 H. Smirnenski Blvd., Sofia 1046

1. Introduction

Recent advances in generative AI have introduced new possibilities for architectural design, particularly in the early concept phase. *Diffusion models* such as Stable Diffusion and Flux [1] can synthesize vivid architectural images from text prompts, enabling designers to quickly visualize ideas. Since 2022, there has been an "exponential emergence" of AI-generated architectural visuals, reflecting the growing interest in leveraging these tools for creative exploration. However, these 2D images – while rich in style and atmosphere – are typically disconnected from the 3D geometric data needed for fabrication. Simultaneously, the field of digital fabrication in architecture has matured, with technologies like robotic concrete printing, CNC mold-making, and advanced composites allowing construction of forms once deemed impractical. Notably, researchers [2] have produced bespoke concrete columns and panels using methods such as *smart dynamic casting* (robotically controlled slip-forming), *support-assisted 3D printing* of fresh concrete (e.g. Bartlett's salt-supported "Fossilized" concrete prints), and 3D-printed formworks that are later filled with high-performance concrete. Despite these advances, integrating AI-driven design generation with fabrication processes is still an open challenge. There is a research gap in connecting the prompt-based generation of architectural concepts to an automated "print" of physical components.

This paper addresses that gap by defining and testing a prompt-to-fabrication workflow that translates a text prompt into a fabrication-ready façade panel geometry and a physical prototype. The research objective is to demonstrate, at façade-panel scale, how diffusion-based ideation, AI image-to-3D reconstruction, and algorithmic rationalization can be coupled with a pragmatic hybrid fabrication route. The scientific contribution is threefold: (1) a multimodal, end-to-end pipeline that keeps continuity between visual intent and manufacturable geometry, (2) a rationalization strategy in Rhino/Grasshopper that operationalizes fabrication constraints (e.g., thickness, topology cleanup, panelization) on AI-derived meshes, and (3) a validated fabrication proof using CNC-milled XPS molds and GFRC casting/spraying to close the loop from prompt to physical artefact.

The approach builds upon several streams of prior work. In generative design, architects and researchers have begun exploring how diffusion models can aid the design. For instance, Cao et al. [3] integrated Stable Diffusion into conceptual design workflows and discussed its potential and limitations in practice. Other studies have applied AI specifically to façades – e.g. Kuang *et al.* [4] combined Stable Diffusion with rule-based expert systems (using *ControlNet* and fine-tuned models) to generate historically inspired façade designs in urban renewal. These efforts highlight the promise of AI in producing design options that balance creativity with context-specific rules. On the fabrication side, there is extensive research on producing free-form architectural elements. Techniques like additive manufacturing of concrete (3D concrete printing) allow tool-less fabrication of non-standard forms, eliminating traditional molds. However, 3D printing concrete has its own constraints (resolution limits, reinforcement challenges) and is not yet a plug-and-play solution for all scenarios. Conversely, conventional casting of complex shapes typically requires custom molds, which historically have been expensive and wasteful for unique pieces. For example, free-form GRC (glass fiber reinforced concrete) panels often demand one-off double-curvature molds, which can account for *40 – 80 % of production cost*. Researchers have thus sought more economical approaches, such as *mold-less fabrication*. Castañeda *et al.* [5] demonstrated a hybrid method where a 3D-printed plastic mesh acted as formwork and internal reinforcement for a GRC panel onto which concrete was sprayed. Their prototypes showed that eliminating traditional molds can drastically reduce cost and material waste.

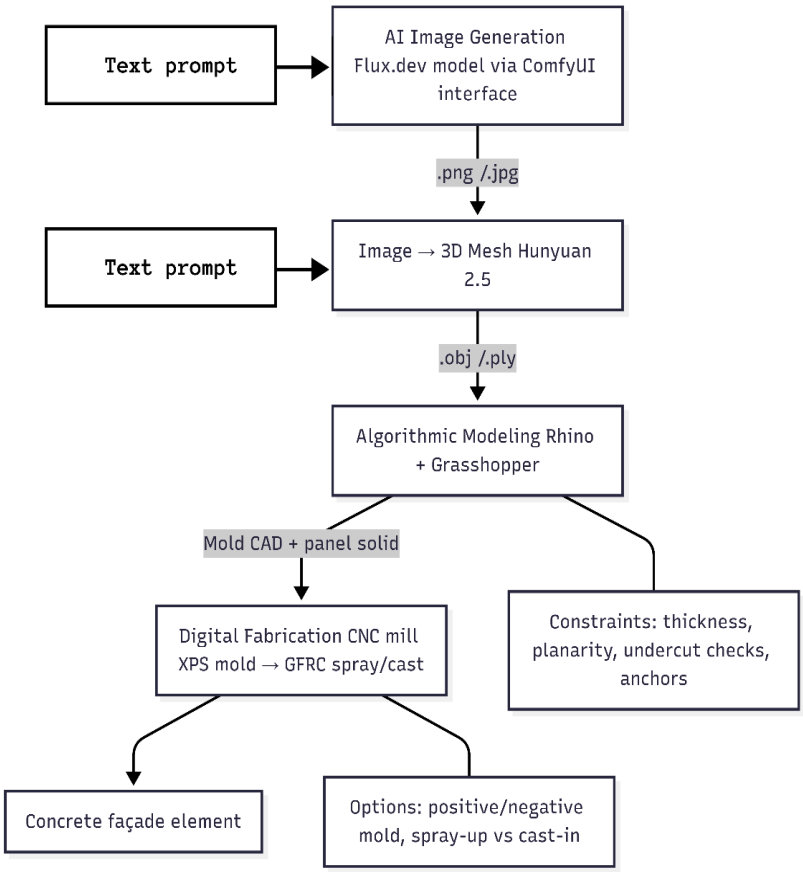


Figure 1. Mind map of our proposed workflow

These examples convey a few conclusions. In generative design, diffusion models are effective at producing stylistic and compositional variety, yet they typically lack explicit encoding of architectural semantics and constraints, which limits reliability when elements must satisfy dimensional, assembly, or code-driven requirements. Even when rule-based guidance is introduced, outputs often remain weakly constrained in terms of repeatable modules, tolerances, and thickness logic, which shifts the burden downstream to manual correction. In image-to-3D reconstruction, current models can generate detailed meshes rapidly, but the results often contain non-manifold regions, ambiguous depth interpretation, and topology artifacts that are incompatible with standard CAD/CAM preparation. On the fabrication side, concrete additive manufacturing enables tool-less production of non-standard forms, yet resolution, reinforcement integration, and process robustness remain limiting factors for consistent architectural cladding production. Conversely, high-fidelity casting of complex double-curvature panels often depends on bespoke molds, which introduces cost and waste penalties, especially for one-off components, despite ongoing research in mold-less or hybrid formwork methods. The core methodological gap, therefore, is not the absence of powerful tools, but the lack of an integrated procedure that converts generative outputs into controlled, watertight, and fabrication-feasible panel geometry with a tractable production method. Learning from these contexts, our proposed workflow uses a hybrid strategy: instead of expensive rigid molds or fully additive printing, we employ CNC-milled XPS foam forms combined with sprayed fiber-concrete, capitalizing on the precision of

subtractive fabrication and the flexibility of thin-shell GRC. Such trials have already been made by the Chair of Digital Design Methods at the Brandenburg Institute of Technology (BTU) in Cottbus, Germany as part of their Cyber Physical Facades Project [6]. The entire process is digitally driven, a workflow that is not new on its own and has been experimented upon for at least the last 15 years. But since the rise of Large Language models and their generative capabilities, we could now be linking AI-generated images to algorithmic model processing and finally to robotic fabrication to create a seamless pipeline from design intent to production. Finally, we reflect on the broader implications of prompt-to-fabrication workflows for architectural prototyping and scalable production, and outline directions for future research such as performance-driven generative design and more automated construction pipelines. Figure 1 shows generalization of our approach.

First, a diffusion model generates candidate façade designs from text prompts (Stage 1). Second, a 3D generation network converts the chosen 2D image into a volumetric mesh (Stage 2). Third, the mesh is imported into CAD (e.g. Rhino) for geometric cleanup and rationalization (Stage 3). Fourth, the finalized design is fabricated using XPS foam molds and GFRC casting (Stage 4). In the following sections, we detail each stage of the methodology and then discuss the results of initial experiments. We evaluate the benefits (e.g. rapid design iteration, high form complexity, integrated design-to-fabrication data flow) and challenges (e.g. accuracy of AI outputs, need for manual adjustments, material and sustainability considerations). Finally, we reflect on the broader implications of prompt-to-print fabrication for sustainable architecture, and outline directions for future research such as performance-driven AI design and fully automated construction workflows.

2. Methodology: Four-Stage Prompt-to-Fabrication Workflow

To realize the vision of automated AI-to-fabrication, we developed a workflow comprising four sequential stages. Each stage involves distinct tools and outputs, as outlined in Table 1. In this section, we describe the procedure and rationale for each stage in detail.

Table 1 presents a summary of the proposed four-stage *prompt-to-fabrication* workflow for façade element design and fabrication.

Table 1. Processes description

Stage	Description & Key Tasks	Tools / Methods	Output
1. Generative Design (2D)	Generate façade concept images using AI diffusion models based on text prompts (with optional image controls). Adjust style, composition, and articulation through prompt engineering and control parameters.	Set of photorealistic or concept 2D façade images reflecting desired styles and features.	Set of photorealistic or concept 2D façade images reflecting desired styles and features.
2. 2D-to-3D Conversion	Convert a selected 2D image into a 3D representation. Infer depth/geometry from the image and produce a three-dimensional mesh with texture. May involve neural implicit fields or diffusion-based 3D prediction.	Hunyuan3D 2.5 (image-conditioned 3D diffusion model); alternative pipelines (e.g. CLIP-guided NeRF or DreamFusion/Magic3D approaches); mesh reconstruction algorithms.	3D mesh model of the façade concept (e.g. OBJ/PLY), with basic texture/color information, approximating the design's geometry.

3. Mesh Rationalization	Import the AI-generated mesh into CAD environment. Clean and refine geometry: smooth noise, close holes, simplify topology. Segmentation for panels, embed attachment details. Prepare final manufacturing geometry and mold model.	Rhinoceros 3D + Grasshopper (parametric modeling scripts); mesh editing tools (QuadRemesh for topology optimization, Weaverbird for smoothing); custom scripts for thickness offset and mold extraction.	Fabrication-ready geometry – e.g. a watertight surface or solid of the panel, with proper thickness and details. This includes the mold shape (negative or positive) for CNC milling.
4. Hybrid Fabrication	Fabricate the physical façade element using a combination of CNC-milling and concrete casting: Mill the mold from XPS foam according to the prepared geometry. Then either spray or cast fiber-reinforced concrete onto the mold. Cure and finish the element as needed for the context.	CNC Mill (Standalone/Robotic Arm) to carve the XPS mold; Fiber-Reinforced Concrete (FRC) mixing (e.g. AR glass fiber GFRC mix); Spray equipment for GFRC or manual casting; optional reinforcement mesh or inserts placed during casting for unique features. This process is done either by hand or with a robotic arm.	Physical concrete façade panel, capturing the intricate geometry from the design. Fiber reinforcement yields a thin, strong shell. The panel can be demolded and is ready for installation or further finishing.

2.1. Stage 1: AI Generation of Façade Design Concepts

In the first stage, we use AI-based image diffusion models to produce a variety of façade and interior design concepts in 2D. The goal is to rapidly explore stylistic variations and complex articulations that would be time-consuming to conceive manually. Recent developments have made diffusion models both accessible and specifically tuned for architectural imagery. In this work, we leverage Flux.dev model – a latent diffusion model – through the ComfyUI interface with a specialized model called Flux. Flux is a custom-trained diffusion model focusing on architectural images (Black Forest Labs, 2025). By using Flux within ComfyUI, we benefit from both architectural domain knowledge (via the model) and a node-based workflow that allows fine-grained control over the image generation process (Fig. 2).

ComfyUI is a modular, visual programming interface for generative models. Instead of a simple text-box prompt, it enables the user to construct a graph of nodes representing steps like prompt inputs, diffusion samplers, and image post-processors. This flexibility is crucial for architectural use: we can introduce controlling nodes to steer the generation. For example, we incorporate ControlNet modules that generate and then use auxiliary inputs (edge drawings and normal maps) to constrain the diffusion output. In practice, we generated initial façade images by providing a base text prompt describing the desired design based on a particular style. We then added a normal map guide – essentially a coarse depth hint – to ensure the diffusion model produces consistent shading suggesting protrusions and recesses at intended locations (windows, columns, etc.). We also experimented with edge detection maps from reference sketches to impose certain geometric rhythms onto the AI output. Furthermore, we use generated images and rerun the main design in different geometric styles. Through these controls, the designer can achieve a balance between *surprising AI creativity* and *intentional design direction*. The ComfyUI pipeline allows chaining these controls: for instance, one can feed an initial AI-generated images back into the pipeline with a *ControlNet Normal* node to refine its depth consistency in a second pass.

The outputs of Stage 1 are high-resolution (1024×1024) images of candidate façade designs. Dozens of variations can be generated within minutes by altering prompts or random seeds. Some curated examples of these AI-generated façade concepts are illustrated below. We found that the Flux model tends to produce results with convincing architectural detail – e.g. realistic material textures, ornamentation, and lighting – often *indistinguishable from hand-crafted renderings*. This aligns with observations by Hansen (2024), who noted that ComfyUI paired with stable diffusion models allows architects to experiment locally with "an exceptional node-based interface" and achieve highly customized outputs. Indeed, compared to out-of-the-box tools like Midjourney (which operate as black-boxes in the cloud) [7], our local workflow offers full control over the model and parameters. The trade-off is a higher technical complexity and GPU computation requirement, but this is a reasonable cost for research and advanced design studios. *Limitations:* It should be acknowledged that diffusion models *lack explicit knowledge of engineering constraints* – they might generate forms that are beautiful but structurally implausible or violating brief requirements. In our experiments, some AI façade images included extremely thin finials or floating elements that would be challenging to fabricate. At this concept stage, we tolerate such issues, focusing on creative exploration. However, the next stages of the workflow must address these gaps by adding realism and constraints to the design. For better consistency in the experimental design, we generate all the façades in frontal view and defined by square boundary. This allows for a better comparison between the iterations, as well as helps to put restriction on potential geometry. We also experimented with using a generated image as input to algorithm and changing the geometry via different prompts. Results showed consistency in the general approach and concept and opportunities for design iterations. Figure 2 shows one cycle of generations of that design experiment.

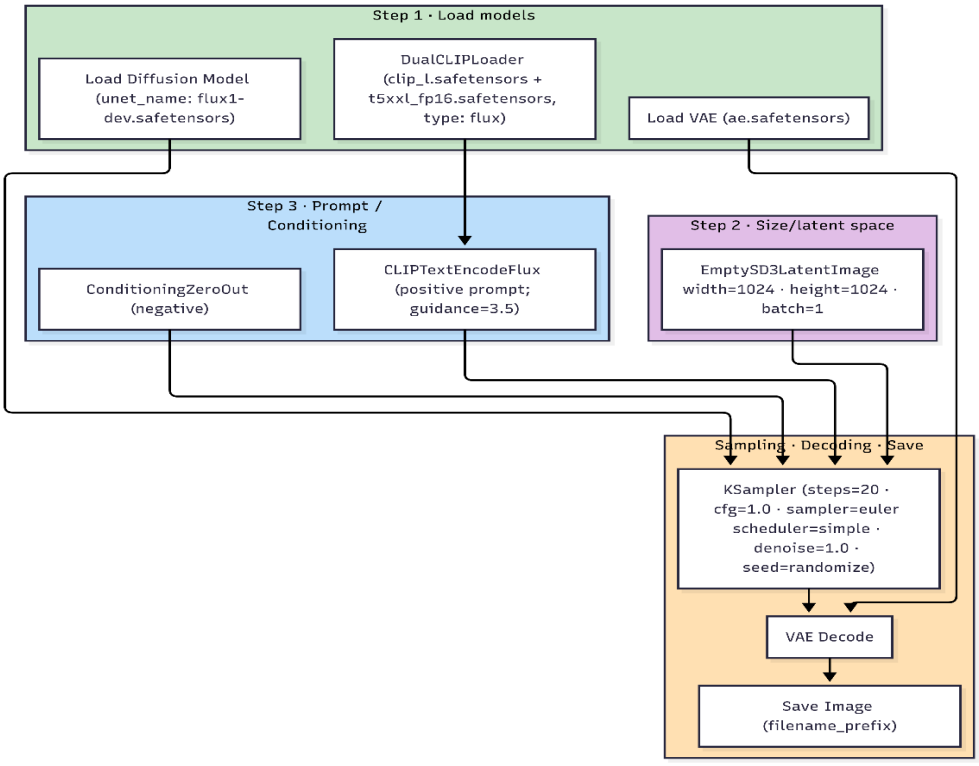


Figure 2. Text to image generation workflow

Figure 3 shows experimental results from this algorithm.



Figure 3. Experimental results of AI generated facades design, ComfyUI/Flux

Figure 4 shows design iterations on established design.



a) *b)* *c)* *d)* *e)*

Figure 4. Geometry style iteration:
a) input image; b) softer curvilinear style; c) mixed styles; d) hard edges style; e) ultra rigid style

Same text-to-image workflow was tested on an interior wall design task. Some results of these iterations are shown in Figure 5.



Figure 5. Text-to-image, design iterations of interior, wall panels, ComfyUI/Flux

2.2. Stage 2: Converting 2D Images to 3D Meshes

After selecting a promising 2D concept from Stage 1, the next step is to derive a three-dimensional geometry that represents this design. This is a non-trivial step: one must *"infer"* the depth and volume from a single image. Traditionally, architects would interpret an elevation or rendering and manually model a 3D version using CAD software. Here, we aim for automation using AI-based 2D-to-3D techniques. To enhance geometric inference, we employ Google Gemini 2.5 Flash Image as an intermediate step. When provided with a single façade image and an instruction to output orthographic and axonometric projections, the model generates consistent top, front, right, left, and axonometric views. While these views are not CAD-precise, they reduce ambiguity by eliminating perspective distortion and providing multiple vantage points of the same concept. In practice, these additional projections help constrain the interpretation of hidden geometry and proportions, offering more reliable reference inputs for downstream reconstruction.

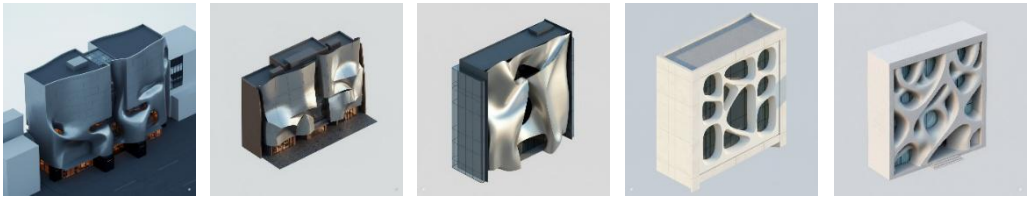


Figure 6. Results from Google Gemini 2.5 flash: axonometric views

We utilize a state-of-the-art model, Hunyuan3D 2.5 [8], which is a recently released 3D diffusion model by Tencent that can produce textured 3D assets conditioned on input images. Hunyuan3D 2.5 employs a two-stage approach: first generating a coarse geometry (as a point cloud or implicit field) and then refining it into a detailed mesh with high-quality textures. Notably, it introduces a 10-billion-parameter "LATTICE" model that significantly improves shape accuracy and surface smoothness, achieving a close alignment between the generated 3D model and the reference image. We found this advantageous for façade applications, as architectural forms often require clean planar regions and sharp edges, which the model handles better than previous versions.

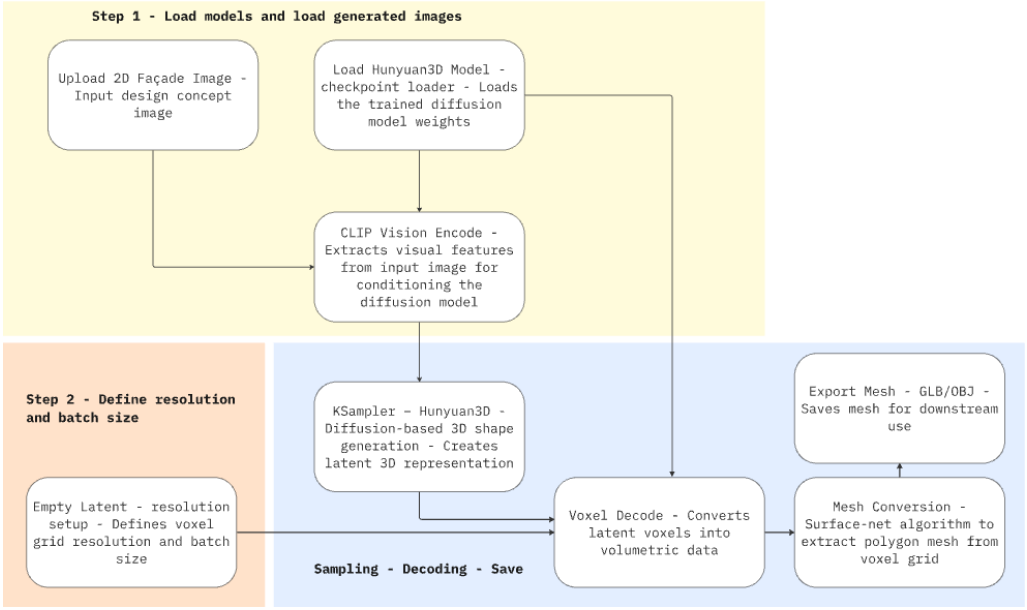


Figure 7. Mind map of stage 2 of our proposed process

Our process in Stage 2 (Fig. 7) begins by feeding the chosen façade image into the Hunyuan3D 2.5 pipeline. The model outputs a triangle mesh (OBJ format) with an associated texture map. In our test runs, the raw output mesh typically had on the order of 50 – 100 k faces for a facade piece roughly 3 – 5 meters in width. The mesh captures major protrusions (such as balconies, window frames, cornices) and recesses (deep window openings, etc.) present in the 2D image, translating them into a volumetric form. This is achieved through the model's training on large datasets of images with known 3D counterparts, enabling it to predict plausible depth even from a single view. Recent research corroborates this capability: for instance, Magic123 [9] demonstrated that optimizing a Neural Radiance Field (NeRF) conditioned on an image, then refining it with a mesh-based diffusion prior, can yield high-quality textured 3D meshes from "in-the-wild" images. Our approach is conceptually similar, as Hunyuan3D's diffusion transformer leverages both 2D and 3D priors to construct geometry.

It is important to discuss the *implicit neural field* aspect here. Methods like NeRF represent scenes as continuous volumetric fields that can be rendered from any angle. We did consider using NeRF-based pipelines (e.g. feeding our image into a CLIP-guided NeRF or the DreamFusion approach). DreamFusion by Poole et al. [10] famously allowed text-to-3D by optimizing a NeRF with guidance from a pre-trained diffusion model. Magic3D [11] improved

on this by a coarse-to-fine strategy, producing higher-resolution meshes than DreamFusion. In preliminary trials, we found direct NeRF optimization from a single image can recover basic depth but struggled with fine architectural detail and tended to produce noisy surfaces. The diffusion-based mesh predictors like Hunyuan3D gave superior results in terms of completeness of geometry (e.g. it produces a closed mesh with reasonable thickness, whereas NeRF requires an additional step to extract a mesh via marching cubes). Therefore, we proceeded with the Hunyuan3D output as the base geometry for the façade element.

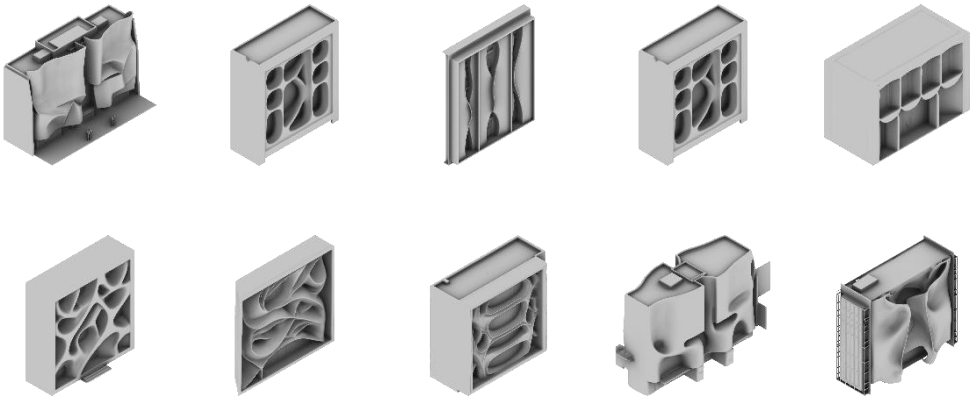


Figure 8. Results from AI generated façade designs from Hunyuan 3D imported in Rhinoceros 8

Even so, the Stage 2 output is not final. The AI-generated mesh, while visually aligned to the concept, often needs geometric refinement before fabrication. Common issues include *non-uniform tessellation* (clusters of very dense triangles in detailed areas and sparse facets elsewhere), *minor holes or self-intersections* from AI artifacts, and *lack of flat backs* (the model might only generate the front façade and have an arbitrarily closed backside). These are addressed in Stage 3. Nonetheless, it is remarkable that current AI can produce a starting 3D model that captures the essence of a conceptual sketch. For instance, one of our test images showed a curvilinear perforated screen; the Hunyuan3D mesh successfully extruded the perforations and gave them depth, creating a lattice-like panel with thickness, which provided an excellent point of departure for detailed modeling. This demonstrates how diffusion-based "mesh prediction" is narrowing the gap between a purely visual concept and a tangible design model – indeed, Hunyuan3D 2.5 reportedly *"closes the gap between generated and handcrafted 3D shapes"* in quality. With Stage 2, we harness that capability to move one step closer to fabrication.

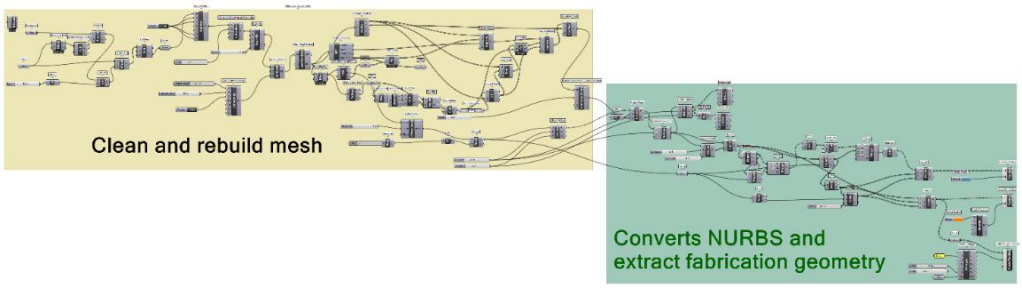


Figure 9. Grasshopper algorithm

2.3. Stage 3: Mesh Processing and Rationalization in Grasshopper

Stage 3 is a critical bridge between the *"free-form" AI geometry* and the *"constraint-aware" fabrication model*. Here we import the Stage 2 mesh into Rhino 3D and use Grasshopper, Rhino's algorithmic modeling environment (Fig. 9), to rationalize the design. The term *rationalization* in architectural geometry refers to modifying a design to satisfy practical constraints – whether they be structural, material, or fabrication-related. In our workflow, the key rationalization tasks include cleaning the mesh topology, penalization and preparing mold geometry (Fig. 10).

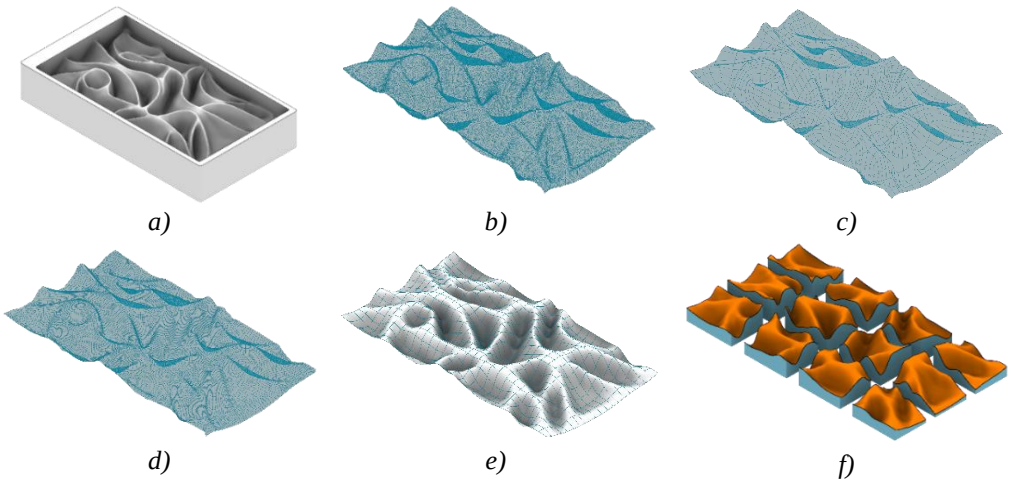


Figure 10. Stages of algorithmic modeling:
a) mesh from Hunyuan; b) extracted mesh; c) Quad remesh; d) Catmull-clark smoothened mesh; e) NURBS surface; f) Fabrication molds

Mesh Cleanup: We first address any mesh pathologies present in the AI-generated model. Using Grasshopper's mesh tools, we remove redundant faces and spikes. A common step is applying *QuadRemesh*, which retopologizes an arbitrary mesh into a more uniform grid of quadrilaterals. This simplifies subsequent operations and yields a cleaner topology for applying panelization or offsets. As noted by forum practitioners, an initial *uniform remeshing* followed by *dynamic relaxation* can greatly improve a mesh for fabrication. We use the *QuadRemesh* with relatively small number of faces (around 200) and adaptive value for curvature at 50 %. This provides clean topology, that would follow with density the curvature of the mesh. After that we smooth the mesh additionally with *Catmull-Clark* algorithm, that provides well refined clean mesh. Next step is to transform it into NURBS surface. We do that by projecting the points of a flat surface onto the mesh and then rebuilding the surface from the projected points. The result is a continuously untrimmed surface, that is easily modelled and trimmed in next steps. This works particularly well for our fabrication proposal, since the NURBS geometry is easily editable. Mold Modeling: With clean solid panel geometry, we derive the mold shape needed for casting. We have two options: a negative mold or a positive mold depending on the fabrication approach. For a *cast-in mold* approach, we would mill a negative cavity of the panel into foam. For a *spray-over* approach, we mill a positive foam piece that exactly mimics the panel's exterior, and we will spray concrete on it. In our implementation, we favored the spray method (common in GFRC production), so we treat the panel's front face as the mold's surface. We create a mold model by simply taking the panel surface, projecting it to a flat plane, lofting the two contours

and then combining all the surfaces into a solid closed Brep. The divisions in the two directions for the panels are set a parameter. Grasshopper's parametric nature made it easy to adjust mold dimensions.

By the end of Stage 3, we have a fully rationalized digital model consisting of a concrete panel of uniform thickness that has been adjusted for fabrication and a corresponding foam mold geometry. The transformations in this stage ensure that, when we proceed to create the piece, there will be no geometric surprises – everything will fit the constraints of our chosen fabrication method and material. It is worth noting that this stage still requires significant expertise and judgement. Automated rationalization is an ongoing research problem, so we combined algorithmic tools with manual tuning in our process. As Pantazis and Gerber [12] observed, an effective generative design system often requires an evaluation feedback loop. In future iterations, we plan to integrate structural or fabrication feedback, such as stress analysis or mold cost estimation, into the Grasshopper script to automatically adjust the geometry. For now, Stage 3 outcomes are vetted by designers. Once they are satisfactory, we move to production. The algorithm we created for mesh rationalization and penalization is generic and robust, but it cannot be automatically applied to any AI design since topologies and curvatures can vary greatly. Nevertheless, the successful outcome of this stage for one of the facade designs proves the applicability of diffusion AI models in the design process. Bridging the gap between AI design and computational design is an important step that opens a wide range of possibilities for further research and development. Future work could include testing the algorithm on more AI-generated meshes and creating different approaches for different topologies.

2.4. Stage 4: Hybrid Digital Fabrication (Foam Molds + Fiber-Reinforced Concrete)

In the final stage, the digital becomes physical. We propose fabrication of the façade elements using a hybrid method combining subtractive manufacturing (milling a foam mold) and additive casting (spraying or pouring concrete). This approach capitalizes on the strengths of each method. CNC milling can quickly create precise, free-form molds, while concrete casting yields robust, fire-safe building components with complex shapes. Crucially, using fiber-reinforced concrete (FRC) specifically glass fiber-reinforced concrete (GFRC) enables us to produce thin, lightweight panels that accurately capture the mold's details. Recent advancements in GFRC formulations continue to enhance durability and weather resistance, establishing it as a favored choice for architectural applications.

Mold Fabrication: We use XPS (extruded polystyrene) foam blocks to make our molds. XPS is inexpensive and easy to mill, and its fine cell structure captures good detail. Based on the mold geometry from Stage 3, we generate CNC toolpaths. We use a 3-axis CNC router for prototyping. If the mold has regions that cannot be accessed from one direction, we either mill the mold in multiple parts or use a router with more available axes or a robot arm. This approach is advantageous because it can produce complex dies and molds in a single setup. Shorter, more rigid tools can be used to improve precision and surface finish. This reduces the need for subsequent hand finishing of the mold.

For example, a concave recess under an overhang was addressed by splitting the mold into two interlocking foam pieces. The foam is usually oversized, and locator pins are included to ensure the pieces align properly. We mill with a combination of rough cuts, which use a larger tool to remove bulk, and finish passes, which use a small ball-nose bit to create fine grooves. The result is a foam mold that is the inverse of the panel's front surface (for a negative mold) or identical to the panel's front surface (for a positive mold). The quality of the milled foam surface is generally sufficient for concrete. However, we sometimes seal it with a thin coat of plaster or

resin to prevent concrete moisture from soaking in and to achieve a smoother finish. However, for GFRC spray-up, minimal sealing is sufficient since contact time with the wet mix is short, and modern polymer admixtures can improve the concrete's interface with the mold surface.

Notably, using foam as formwork is a sacrificial technique; the foam will likely be destroyed or discarded after casting. Although this results in waste, it is a deliberate trade-off to avoid building heavy, reusable molds for one-time use. To address sustainability concerns, the industry is shifting towards low-GWP (global warming potential) XPS foam, which can reduce embodied carbon by over 90 % and often incorporates recycled materials. The literature recounts many instances of foam formwork in free-form construction. For example, the curved forms of the Bus Stop in Hoofddorp were cut from polystyrene and coated for concrete casting. In our context, each unique panel requires a unique foam mold. This is acceptable for bespoke elements or small batches. For larger quantities, one might invest in a more durable mold. Looking forward, emerging methods like mineral foam 3D printing (F3DP) suggest a future of creating waste-free, stay-in-place formwork that provides additional thermal functionality.

Concrete Spraying/Casting: With the mold ready, we cast the final part using GFRC. We chose GFRC due to its high strength-to-weight ratio and ability to form thin, intricate shapes. We prepared a mixture of fine sand, cement, a polymer admixture, water, and alkali-resistant (AR) glass fibers. To ensure long-term durability, it is crucial that the fibers comply with established standards such as ASTM C1666/C1666M. This standard specifies a minimum zirconia content that resists alkali attack from the cement matrix. The fiber content is typically 3 – 4 % by weight of the mixture. There are two main process variants:

- **Spray-up method:** This is the most common method for GFRC panels. First, we spray a thin "face coat", a fiber-free mortar, onto the mold using a specialized mist coat gun. This captures detail and avoids fiber exposure to the surface. Immediately after, we spray the fiber-rich GFRC backer mix. Using a pneumatic spray gun, we deposit and build up the mix to the desired thickness (often ~10 – 15 mm). We ensure even coverage, especially in corners. For a negative mold, we spray into the cavity, and for a positive mold, we spray onto the exterior of the foam. Although manual spraying is highly flexible, recent advancements include automated robotic concrete spraying (ARCS), in which robotic arms follow digital toolpaths to spray GFRC with high precision. This enables consistent thickness across complex geometries.

- **Cast-in mold method:** As an alternative, we cast self-compacting, fiber-reinforced concrete into a foam mold. We used this method for a panel with a solid cross-section (approximately 50 mm thick). We poured and lightly vibrate the concrete (though excessive vibration can cause fibers to settle). This method is like traditional precasting, except it uses a disposable mold. This method works well for simpler shapes. However, for very intricate thin sections, spraying is preferable to ensure fiber distribution and avoid trapped air.

During the casting process, we can embed any necessary reinforcements, hardware, or special features, whether aesthetic or not. For example, we placed a light steel wire mesh in one panel for additional reinforcement and cast bolt anchors into the back for mounting. These elements were planned in Stage 3, and the mold was modified locally to accommodate them.

Curing and Demolding: After casting, the element is kept damp to ensure proper curing. GFRC, particularly when formulated with modern polymer admixtures, gains strength quickly. Two to three days are sufficient for handling pieces. We then demold by separating or breaking away the foam. In our spray-up cases on positive foam, the foam bonded to the concrete; therefore, we peeled it off in chunks and used a solvent on the residual bits. In cases involving a

negative mold, we were able to lift the panel out of the cavity because we had tapered the mold slightly in Grasshopper, giving it a "draft angle". Inevitably, most of the foam cannot be reused. However, the concrete panel emerges with the intended form and a smooth surface finish. Minor hand finishing, such as sanding the edges and patching small air voids, may be necessary, but no major rework is needed if the process is well-controlled.

In summary, Stage 4 involves producing a tangible artifact to validate the design. This process uses a hybrid digital fabrication approach. CNC provides the formwork and concrete casting provides material realization. This combination of methods joins the precision of digital models with the proven performance of concrete as a material. Standardized, alkali-resistant fiber reinforcement ensures that complex, slender forms attain the necessary strength and durability for building use. This process closes the loop from a conceptual AI prompt to a physical, architectural-quality component.

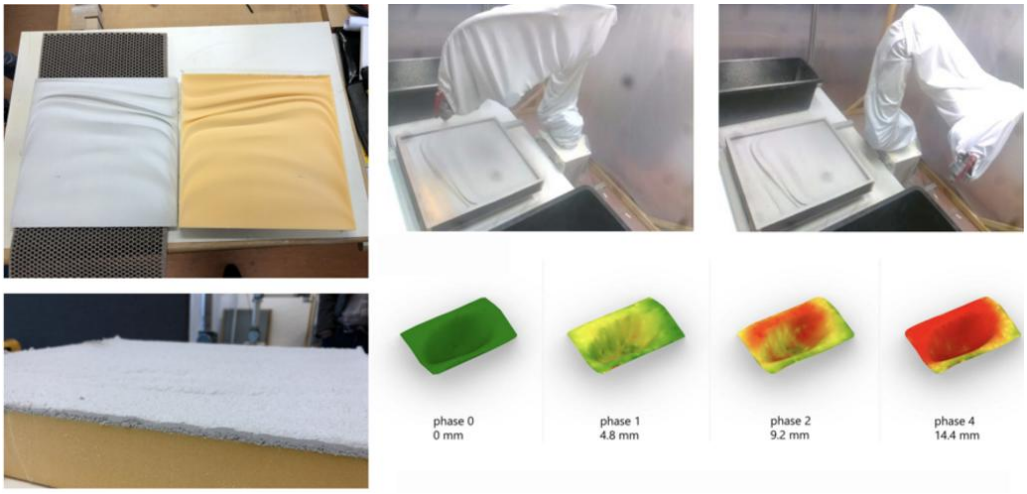


Figure 11. The XPS milled and robotically sprayed molds from the project "Cyber Physical Facades" by the Brandenburg Institute of Technology (BTU) in Cottbus, Germany [13]

3. Results

Using the proposed workflow, we produced several prototype façade designs, which allow us to evaluate the outcomes, benefits, and challenges of the approach. Overall, the *prompt-to-fabricate* process was successful in translating abstract AI-generated ideas into physical architectural components. We discuss the results in terms of design fidelity, fabrication quality, efficiency, and areas for improvement.

Design Fidelity: One key metric is how closely the final product reflects the original AI-generated design intent. In our trials, we observed that major design features conceived in Stage 1 were preserved through Stage 4. This indicates a high degree of fidelity; the pipeline did not introduce significant deviations from the design vision. Minor differences occurred in extremely small details, and mesh openings were ignored in the mesh processing stage. Additionally, certain dramatic overhangs generated by the AI had to be toned down for structural reasons. Despite these adjustments, the essence of the AI design – its style and language – remained intact. This result is encouraging because it suggests that, by combining advanced AI models with careful rationalization, we can maintain creative designs without the heavy manual reinterpretation often needed when translating 2D sketches to 3D.

Efficiency and Workflow Assessment: The time from concept to final product was significantly compressed using our pipeline. Generating and selecting an AI concept took on the order of hours (including multiple tries), converting to 3D and cleaning the mesh was a few hours of computational and manual work, and fabricating the panel took a few days (dominated by CNC milling time and concrete curing). In a conventional process, an equivalent level of exploration and fabrication (imagine an architect hand-sketching dozens of ideas, then making a detailed Rhino model, then producing a mold through industrial means) could take weeks or more. Thus, a clear benefit is the acceleration of the iterative loop. An architect can physically test forms earlier in the design process, leading to a more informed design development. This speed does not come without caveats: some stages required specialized knowledge (e.g. fine-tuning a diffusion model, or writing Grasshopper scripts), meaning the pipeline is efficient only for those equipped with such digital skills. One can foresee, however, that as these AI and CAD tools become more user-friendly or packaged, the barrier to entry will lower. There is also a learning curve; our first panel was essentially a "learning prototype", and subsequent ones were much faster as we reused Grasshopper definitions and had a better sense of what the AI output quality would be.

Challenges and Limitations: Through the experimentation, several challenges became evident:

- **Controlling AI Outputs:** While Stage 1 yielded many inspiring designs, controlling functional aspects in the AI image was difficult. For example, ensuring that a window opening appears exactly where needed or that a pattern repeats with a certain module size is not straightforward with a text prompt alone. Additional guiding images (via ControlNet) helped, but there is a need for AI models that understand architectural semantics (like knowing what a "window" is in terms of solid-void requirements). When those were lacking, we had to enforce such things later in Grasshopper, effectively altering the design from the original image. This indicates that a more constrained generative approach or a human-in-the-loop refinement might be necessary for truly architectural use-cases.

- **Mesh Conversion Quality:** The AI mesh of Stage 2, while impressive, sometimes had ambiguous or non-manifold geometry. Especially in regions where the 2D image was unclear (e.g. a dark shadowy area), the AI would fill in with a "blobby" geometry or double surfaces. Cleaning this required manual attention. In one case, an open balcony in the image was interpreted as a solid volume by the model, and we had to manually hollow it out in Rhino. More generally, the system showed a tendency to resolve smooth, fluid-like forms more convincingly than angular or orthogonal geometries, which were often distorted or misrepresented. Such issues suggest that current image-to-3D models, though advanced, are not infallible and benefit from human correction. Research is rapidly evolving here – newer models or multi-view approaches (using more images or known constraints) could alleviate this in future iterations.

- **Rationalization Effort:** Stage 3 proved to be the most labor-intensive in practice. Automating it fully is challenging. We relied on Grasshopper scripts, but these often-needed tweaking per design. Each façade geometry had its own quirks (one might need more smoothing; another needed a different offset strategy). In a sense, we replaced the manual 3D modeling of a design with manual problem-solving on an AI mesh. The expertise required is different but still significant. This is a limitation to address if such workflows are to become mainstream. Embedding

more intelligence in Stage 2 (so the output is fabrication-ready, separated in different parts) is one approach; another is improving Stage 3 tools (e.g. more robust algorithms for automatic thickening, feature detection, etc.). Despite the effort, the rationalization was doable for relatively modest-sized panels (on the order of 1 – 2 m). Scaling to very large façade sections might compound these issues.

- **Fabrication Constraints:** Our hybrid method demonstrated efficacy at the panel scale; however, practical limitations must be considered. In the case of larger elements, it may be necessary to employ sectional foam molds due to limitations imposed by the machine and the foam block size. This approach introduces seams that must be meticulously managed during the digital design phase to ensure optimal alignment. Additionally, while GFRC is ideal for thin cladding that can support its own weight and resist wind loads, it is not intended for primary structural applications. If load-bearing panels are desired for the structural facade, additional reinforcement measures must be implemented. Such measures may include the use of rebar and the incorporation of thicker sections, which, in turn, introduces complexity to the casting process. Consequently, the present workflow is optimally suited for non-structural façade cladding or screen elements. Another constraint was the need for skilled labor in the spraying of concrete. Although the involvement of a human operator introduces slight variability in our case, the field of robotic fabrication is rapidly addressing this challenge. Cognitive robotic platforms for concreting are being developed by advanced systems such as Automated Robotic Concrete Spraying (ARCS) and research initiatives such as the RoBétArmé project. These systems utilize real-time three-dimensional scanning and feedback loops to regulate the trajectory of the spray and the thickness of the material, thereby striving to ensure greater consistency and quality than that of manual application. While the prospect of full autonomy remains encumbered by present challenges, these advancements suggest a future in which robotic automation may be capable of surmounting the limitations of skilled labor regarding the large-scale production of complex concrete forms.

- **Integration with Design Process:** From a workflow standpoint, integrating AI design with traditional architectural design processes remains an open question. The AI may produce designs that do not meet building codes or client requirements, necessitating iterations in which the architect guides the AI. Although our process was linear, design is iterative in practice. For example, one might fabricate a panel, evaluate it, and then adjust the design. Our pipeline supports iteration. For example, one could adjust the Grasshopper constraints and re-fabricate. However, incorporating higher-level feedback, such as “make this panel cheaper” or “lighter”, would require linking with analysis tools or cost estimation. A performance-driven design could be considered for future work.

- **Preliminary Outcomes:** To summarize, our results validated the core premise: it is possible to generate, materialize, and evaluate unconventional façade designs rapidly using AI, algorithmic modeling and digital fabrication. Our figures showcase the preliminary results, but they only represent a small part of the large design space that could be explored through generative AI. Due to the limitations of this paper, only a few of our tests are shown. The results highlight the benefits of multimodality: text prompts provide creative breadth, the image-to-3D step accelerates modeling, and CNC+GFRC provides physical proof – all within a relatively short loop.

At the same time, the identified challenges form a roadmap for improvement. For example, to reduce rationalization time, one could develop a library of Grasshopper definitions that automatically detect and resolve common mesh issues. To improve AI control, one could train diffusion models on paired data of facade images and 3D models so that the model inherently knows how to generate more "modellable" images. This is an area of active research. Integrating evaluation metrics (structural and environmental) into the generative loop is another promising approach, ensuring from the beginning that designs are beautiful, functional, and sustainable. Encouragingly, frameworks are being proposed that could be merged with our pipeline, such as multi-criteria agent systems. The AI could generate options, which would then be filtered or adjusted by performance-oriented algorithms before fabrication.

Overall, the prompt-to-print workflow demonstrates a novel synergy between AI and digital fabrication. It capitalizes on the strength of AI in idea generation and the precision of CAD/CAM in realization. The results so far are compelling, showing that architectural elements of great complexity and custom character can be achieved faster and potentially at lower cost than with conventional means. The approach fosters creativity (by freeing designers to explore wild ideas with AI support) while also grounding those ideas in making (by forcing consideration of fabrication early on). This synergy is key to advancing sustainable design-build practices – by iterating quickly and physically, architects and engineers can refine designs for optimal performance and resource use.

4. Conclusion

This paper presented "Prompt-to-fabrication", a multimodal workflow that links AI-driven design with automated fabrication to produce custom façade elements. The proposed four-stage process – spanning AI image generation, 3D reconstruction, parametric rationalization, and hybrid concrete fabrication – offers a proof-of-concept for how future architectural design might seamlessly integrate creativity and making. Through our implementation and prototypes, we demonstrated that a text prompt can evolve into a real concrete panel in a matter of days, opening up new possibilities for rapid prototyping and mass customization in architecture. The approach leverages cutting-edge tools: by using diffusion models architects can tap into a rich latent space of architectural styles and quickly visualize ideas; by employing AI 3D generation (Hunyuan3D, etc.), those visions can be automatically extrapolated into workable geometry; by applying algorithmic design tools (Grasshopper), the often messy AI output can be tamed into buildable form; and finally, by adopting a hybrid fabrication strategy (CNC-milled foam + GFRC), the leap from digital model to physical artifact is made efficient and feasible for one-off designs.

The novelty of this workflow lies in its end-to-end integration. Unlike traditional separations between design and construction, prompt-to-fabrication collapses these phases, suggesting a future where architects might produce their ideas as easily as they sketch them. This has profound implications. Architects can explore a far greater diversity of forms without worrying from the outset about how to draw or model them – the AI provides a fluent brainstorming partner, and the fabrication pipeline ensures that even highly complex forms can be realized. As AI takes a more significant role in form-finding, architects' expertise will shift towards curation and orchestration – guiding the AI with the right prompts, setting the constraints for rationalization, and deciding on material and fabrication strategies that align with project goals. This workflow exemplifies that shift: the designer defines objectives and constraints, and the AI+fabrication system generates solutions within that framework. In doing so, it maintains the architect's authorship of the process while utilizing the AI as a powerful co-author.

Of course, this study is an initial step. We focused on a façade panel scale, but scaling up to entire building facades or structures will bring additional challenges (and likely require breaking designs into modular components, integrating with structural frames, etc.). Future research should investigate the integration of structural analysis and optimization directly into the workflow. For example, generative AI could propose forms which are then immediately evaluated for wind load or thermal performance; the feedback could refine the prompt. Another extension is towards fully automated fabrication: while we used human spraying, one could envision a robotic arm that mills foam and then robotically sprays concrete or even directly prints concrete, all guided by the digital model – essentially a robot construction pipeline triggered by an AI design. This aligns with the concept of "Digital Twin" fabrication where the digital model is directly realized by robotic means, minimizing translation errors.

In conclusion, the prompt-to-fabrication workflow demonstrates a focused integration of AI-driven ideation and fabrication-aware digital processing. It leverages diffusion models for rapid exploration, image-to-3D reconstruction for accelerated geometry acquisition, and CAD-based rationalization to enforce manufacturability before physical production. The results indicate that façade elements with high geometric differentiation can be prototyped faster than fully manual modeling-to-mold workflows. This shifts the architect's effort from exhaustive modeling to constraint definition, curation, and verification – an operational change that is essential if generative systems are to contribute to reproducible, buildable outcomes rather than only image-level novelty.

Acknowledgment

The authors would like to thank colleagues and peers for their constructive feedback during the development of this work. We also acknowledge the open-source communities behind ComfyUI, Grasshopper, and related AI model developers whose tools made this research possible.

REFERENCES

1. *Black Forest Labs*. FLUX.1 [dev] – model card. 2025, retrieved from https://huggingface.co/black-forest-labs/FLUX.1-dev.

2. *Anton, A., Yoo, A., Bedarf, P., Reiter, L., Wangler, T., Dillenburger, B.* Vertical Modulations: Computational design for concrete 3D printed columns. In Proc. ACADIA 2019: Ubiquity and Autonomy (pp. 596 – 605), Austin, TX: ACADIA.

3. *Cao, Y., Abdul Aziz, A., Mohd Arshard, WNR.* Stable diffusion in architectural design: Closing doors or opening new horizons? // International Journal of Architectural Computing. 2025, 23(2):339 – 357, doi:10.1177/14780771241270257.

4. *Kuang, Z., Zhang, J., Li, Y., Fukuda, T. et al*. Preserving architectural heritage in urban renewal: a stable diffusion model framework for automated historical facade generation. npj Heritage Science, 2025, 13:256.nature.com.

5. *Castañeda Vergara, E., Lauret Aguirregabiria, B., Lirola Pérez, J. M., Ovando Vacarezza, G.* Innovative free-form glass fiber reinforced concrete (GRC) panel with 3D printed internal grid. Revista de la Construcción, 2018, 17(3):479 – 490, doi: https://doi.org/10.7764/RDLC.16.3.479.

6. *Vukorep, I., Zimmermann, G., Sablotny, T.* Robot-Controlled Fabrication of Sprayed Concrete Elements as a Cyber-Physical-System. 2020, 10.1007/978-3-030-49916-7_94.

7. *Tan, L., Luhrs, M.* Using Generative AI Midjourney to enhance divergent and convergent thinking in an architect's creative design process. The Design Journal, 2024, 27, doi: 10.1080/14606925.2024.2353479.

8. *Lai, Z., Zhao, Y., Liu, H. et al*. Hunyuan3D 2.5: Towards High-Fidelity 3D Assets Generation with Ultimate Details. 2025, arXiv:2506.16504 [cs.CV].arxiv.orgarxiv.org.

9. *Qian, X., Yang, C., Feng, C., Zhang, Y., Li, H.* (2024). Magic123: One image to high-quality 3D object generation using both 2D and 3D diffusion priors. arXiv:2306.17843.arxiv.org

10. *Poole, B., Jain, A., Barron, J. T., Mildenhall, B., Liu, C.* DreamFusion: Text-to-3D using 2D diffusion. 2023, arXiv:2209.14988.arxiv.org.

11. *Lin, C. H., Gao, J., Tang, L., Takikawa, T., Zeng, X., Huang, X., Kreis, K., Fidler, S., Liu, M.-Y., Lin, T. Y.* Magic3D: High-resolution text-to-3D content creation. arXiv. 2023, https://arxiv.org/abs/2211.10440.

12. *Pantazis, E., Gerber, D.* A framework for generating and evaluating façade designs using a multi-agent system approach. // International Journal of Architectural Computing, 2018, 16(4): 251 – 268.

13. https://www.b-tu.de/fg-digitales-entwerfen/forschung/projekte, visited on 10.01.2026.

ОТ ПРОМПТ ДО ПРОИЗВОДСТВО: ПРЕДЛОЖЕНИЕ ЗА МУЛТИМОДАЛЕН ПОДХОД ЗА АВТОМАТИЗИРАНЕ НА ПРОИЗВОДСТВО НА ФАСАДНИ ЕЛЕМЕНТИ

Р. Шишков[1], Т. Нгуен[2], З. Вугрешек[3]

Ключови думи: *генеративен дизайн, дифузионни модели, текст към 3D, дигитално фабрикиране*

РЕЗЮМЕ

Архитектите все по-често използват генеративен изкуствен интелект за бързо създаване на нови фасадни концепции, но превръщането на AI изображения в изпълними строителни елементи остава трудно. Настоящата статия предлага работен процес от тип prompt-to-fabricate, който свързва text-to-image дифузионни модели, AI-базирано 2D-to-3D генериране на мрежи и алгоритмична геометрична обработка в Rhino/Grasshopper, за да преведе визуалното намерение в геометрия, готова за производство. Процесът използва контролируема синтеза на изображения за изследване на варианти, след което реконструира избраната концепция като детайлен 3D меш и я рационализира чрез почистване, контрол на топология, осигуряване на дебелина и панелизация. Физическото

[1] Радул Шишков, гл. ас. д-р арх., кат. „Интериор и дизайн за архитектурата“, УАСГ, бул. „Хр. Смирненски“ № 1, 1046 София, e-mail: rshishkov_far@uacg.bg

[2] Тунг Нгуен, докторант, арх., кат. „Интериор и дизайн за архитектурата“, УАСГ, бул. „Хр. Смирненски“ № 1, 1046 София

[3] Звонко Вугрешек, докторант, арх., кат. „Интериор и дизайн за архитектурата“, УАСГ, бул. „Хр. Смирненски“ № 1, 1046 София

изработване се реализира чрез хибриден метод, при който CNC-фрезовани калъпи от XPS пяна се използват за отливане или напръскване на фиброармиран бетон, като GFRC. Първоначалните резултати показват, че подходът може да генерира сложни, персонализирани бетонни фасадни елементи в сравнително кратки срокове. Основните предимства са бърза итерация на концепции, непрекъснатост между генеративния дизайн и производствените ограничения, и възможност за реализация на комплексни форми без изцяло ръчно моделиране. Ключови ограничения са частичната контролируемост на AI изходни данни, несигурността в точността на мешовете, усилието за рационализация, ограниченията на материалите и производственият процес. Работата демонстрира интегриран мултимодален мост от промпт към физически прототипи и очертава посоки за бъдещи изследвания в енеративния дизайн и по-автоматизирани строителни процеси.