Prompt-to-Geometry: Voxel-Based Architectural Generation via Cellular Automata and LLMs

Andrei Vince, Aidan Oh & Prof. Sabri Gokmen (Advisor)

Summer Research Scholars Program • UNC Charlotte • May–July 2025

ABSTRACT & PROBLEM

This research introduces a novel framework for generating 3D architectural geometry from natural language. By integrating Large Language Models (LLMs), Cellular Automata (CA), and conditioned Diffusion Models, our system automates the translation of abstract concepts into tangible voxel models and high-fidelity renderings.

Current Challenge: Traditional architectural design workflows require extensive manual modeling and iteration, creating barriers between conceptual ideas and 3D visualization. Existing generative tools lack the precision and control needed for architectural applications.

KEY CONTRIBUTIONS

Novel Integration: First system combining LLMs with 3D cellular automata for architectural generation.

Semantic Parsing: Automated extraction of architectural metadata from natural language.

Procedural Generation: Scalable voxel-based geometry creation with emergent properties.

Visual Consistency: ControlNet-conditioned rendering maintains geometric fidelity.

Production Ready: Runpods cloud backend enables real-time architectural ideation.

TECHNICAL VALIDATION

Key Performance Metrics:

- **Generation Time:** ~17 seconds per model (A40 GPU)
- **Test Configurations:** 500+ unique models generated
- **Rendering Quality:** 1024×1024 resolution
- **Voxel Resolution:** 64×64×64 grid
- **Project Duration:** 12 Weeks
- **Pipeline Stages:** 4

WORKFLOW

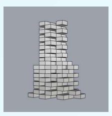

1. CELLULAR AUTOMATA GENERATION

Custom 3D CA algorithm generates voxel-based geometry from initial parameters and rules.

→

2. PHOTOREALISTIC RENDERING

SDXL with ControlNet generates realistic images maintaining spatial and material fidelity.

→

3. PROMPT EXTRACTION

GPT-4 interprets design and generates descriptive prompts for documentation.

Example prompt output:

"A rectilinear glass building with three visible floors; the ground floor is clearly taller approximately double the height of the upper floors creating strong volumetric distinction. Facade is composed of..."

→

4. DATABASE INTEGRATION

All geometry, images, and prompts stored for retrieval and analysis.

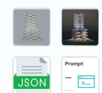

RESULTS & DEMONSTRATIONS

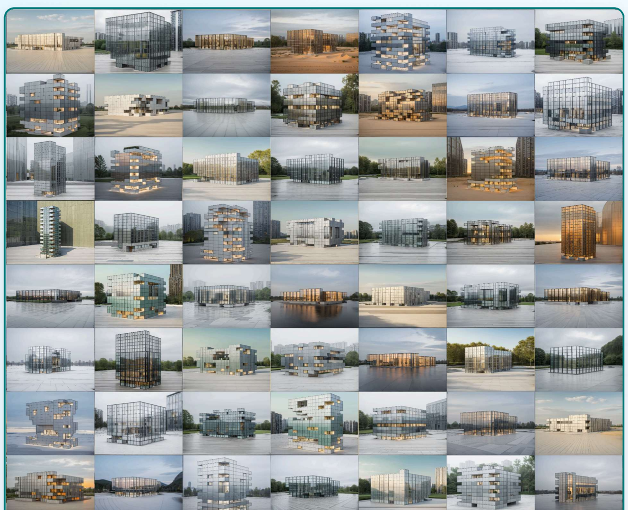

Design Variations: A grid showcasing the wide variety of architectural forms and styles generated by the system from different prompts, demonstrating the model's flexibility and creative potential.

Key Achievements

- **Geometric Accuracy:** Maintains spatial relationships from voxel to render
- **Material Consistency:** Realistic textures and lighting conditions
- **Scalable Output:** From single buildings to urban blocks
- **Rapid Iteration:** 17-second generation enables real-time design
- **AI Integration:** Seamless LLM-to-geometry workflow

SYSTEM PERFORMANCE

The complete pipeline demonstrates robust performance across diverse architectural typologies, successfully generating over 500 unique building configurations during testing phases.

CORE INNOVATION

The core novelty lies in using procedural voxel algorithms, guided by AI-inferred design parameters, to rapidly accelerate early-stage architectural ideation, bridging the gap between linguistic creativity and geometric complexity.

Scan to visit repository

TOOLS & STACK

- **CAD & Scripting:** Rhino, Grasshopper
- **Voxel Generation:** Python (voxel_ca.py)
- **Language Processing:** OpenAI API (GPT-4)
- **Image Synthesis:** SDXL + ControlNet
- **Cloud Compute:** Runpods (GPU A40)
- **Version Control:** GitHub

ACKNOWLEDGEMENTS

This research was supported by the UNC Charlotte Office of Undergraduate Research through the Summer Research Scholars Program. Special thanks to Prof. Sabri Gokmen for guidance and the School of Architecture for computational resources.

Our weekly mentor check-ins have supported the progress we've made toward the larger goal of building this tool and dataset pipeline that demonstrates the creative and technical integration of LLMs into early-stage architectural workflows.

FUTURE WORK

- **Urban Scale:** Extend to multi-building urban layouts with contextual relationships
- **Interactive Design:** Web-based interface for real-time parameter adjustment
- **Constraint Integration:** Incorporate building codes, FAR, and structural requirements
- **Dataset Expansion:** Scale training data for fine-tuned architectural language models
- **Multi-Modal Input:** Support for sketch-based and image-based prompts