

ResiPlan AI: A Comprehensive Analysis of AI-Driven Automated Residential Floor Planning

Ashutosh Kale¹, Swapnil Yeole², Onkar Sonawane³,
Jayesh Wani⁴, Vedant Rajput⁵

Computer Engineering,
MET Institute of Engineering Nashik, India

Abstract— This analysis paper examines ResiPlan AI, an intelligent web-based system designed to automate residential floor plan generation using artificial intelligence and machine learning techniques. The system addresses significant barriers in traditional architectural design—such as high cost, complexity, and reliance on expert knowledge—by enabling non-expert users to generate optimized 2D and 3D layouts through simple inputs like plot size, room count, and architectural style. By integrating Stable Diffusion 1.5 with ControlNet, ResiPlan AI ensures structural adherence while maintaining creative flexibility. This paper critically evaluates the system’s architecture, technical approach, limitations, and future potential, positioning it within the broader context of generative AI in architectural design.

Keywords: Stable Diffusion, ControlNet, Generative AI, Residential Layout, Automated Design, AI/ML, Floor Plan Generation, Intelligent Systems.

I. INTRODUCTION

The field of architectural design has long been characterized by labour-intensive processes, high costs, and a dependency on specialized software such as AutoCAD and Revit, which require substantial technical expertise. This complexity excludes non-experts—including homeowners, real estate agents, and small-scale developers—from actively participating in the design of their own spaces. Recent developments in generative artificial intelligence have introduced new possibilities for automating design workflows through data-driven models and intuitive user interfaces.

ResiPlan AI is proposed as a solution to these challenges. The system leverages diffusion-based generative models alongside constraint-enforcement mechanisms to produce structurally valid and aesthetically coherent residential floor plans. This paper provides a systematic analysis of ResiPlan AI’s methodology, architectural framework, and practical implications, while situating it within the evolving landscape of AI-aided design tools.

II. LITERATURE REVIEW

The evolution of automated floor plan generation can be traced through several key phases, from heuristic and procedural methods to data-driven deep learning approaches. Early Methods and Datasets: Initial research in automated floor planning was constrained by the absence of large,

annotated datasets. Kalervo et al. addressed this gap with the CubiCasa5K dataset, comprising 5,000 floor plan images annotated into over 80 categories. This dataset facilitated a shift from low-level pixel operations to convolutional neural network (CNN)-based parsing and segmentation. However, these methods were primarily analytical rather than generative.

Generative Adversarial Networks (GANs): Generative Adversarial Networks (GANs) later emerged as a popular solution for layout generation. Models such as House-GAN and Graph2Plan utilised graph-based inputs to generate room layouts. While these represented significant progress, they often produced raster outputs that required extensive post-processing. Additionally, maintaining geometric consistency and satisfying architectural constraints remained challenging.

Diffusion Models and Latent Diffusion: The introduction of diffusion models and latent diffusion models (LDMs) marked a turning point, offering improved stability and multimodal handling. Shabani et al. proposed HouseDiffusion, which denoises 2D coordinates to generate vector-based floor plans, significantly outperforming GAN-based methods in geometric precision. However, such models often rely on structured graph inputs, which remain inaccessible to non-expert users.

Integration with Large Language Models (LLMs): Recent research has sought to enhance accessibility through natural language interfaces. Qin et al. introduced ChatHouse Diffusion, which uses LLMs to interpret textual descriptions and convert them into semantic graphs for layout generation.

Similarly, Qiu et al. combined Stable Diffusion with ControlNet and Low-Rank Adaptation (LoRA) to enable multimodal input handling, allowing both text and boundary images to guide generation. These approaches highlight a growing trend toward intuitive, interactive design systems.

Broader Applications and Limitations: Ibrahim et al. extended diffusion models to non-residential structures—such as stadiums and offices—demonstrating scalability across building typologies. Despite these advances, a persistent gap remains in systems that balance user-friendly input with strict architectural constraints for non-technical audiences. ResiPlan AI aims to bridge this gap by integrating Stable Diffusion with ControlNet in a web-based, accessible platform.

Research Gap

Despite significant progress in AI-driven design automation, several critical gaps persist in the current research landscape:

Lack of Non-Expert Accessibility: Most advanced systems, such as HouseDiffusion and Graph2Plan, require structured topological or graph-based inputs, which are impractical for non-technical users.

Weak Constraint Enforcement: Many generative models treat floor plans as images, leading to structural inconsistencies such as misaligned walls or non-functional room connections.

Limited Multimodal Integration: While some systems incorporate text or boundary conditions, few provide a seamless integration of natural language, geometric constraints, and visual feedback in a single pipeline.

Domain-Specific Limitations: Existing models are often trained on residential data and struggle to generalise to commercial or institutional layouts without extensive retraining.

Absence of Iterative Design Support: Current tools lack robust feedback mechanisms for real-time editing and regeneration, limiting practical utility in professional workflows.

ResiPlan AI directly addresses these gaps by proposing a hybrid, web-based system that combines prompt-based interaction with strict geometric control, enabling iterative and user-centred design.

Problem Statement

The traditional process of residential floor plan design is inherently complex, time-consuming, and costly, relying heavily on manual drafting and expert architectural knowledge. Industry-standard tools such as AutoCAD and Revit present steep learning curves and require high computational resources, making them inaccessible to non-specialists.

Existing AI-driven solutions, while promising, often fail to enforce architectural constraints rigorously, produce inconsistent outputs, or depend on inputs that are too technical for everyday users. There is a clear and urgent need for an intelligent, automated system that can generate structurally valid, aesthetically coherent floor plans from simple, natural inputs, thereby democratizing access to architectural design.

Objectives

The primary objective of this research is to design, develop, and analyze an intelligent web-based system for automated residential floor plan generation using artificial intelligence techniques. The specific objectives of the study are as follows:

- **To Design and Implement a Hybrid AI Architecture:** Develop a modular system architecture that integrates Stable Diffusion 1.5 for high-fidelity image generation with ControlNet for enforcing architectural and geometric constraints
- **To Develop a User-Friendly Web Interface:** Create an intuitive, React-based frontend that allows non-expert users to input design parameters (e.g., plot size, room count, style) and interact with the system without requiring technical expertise or local software installation.
- **To Enable Constraint-Aware Generative Modeling:** Implement and fine-tune a diffusion-based model capable of interpreting natural language inputs and geometric constraints, ensuring outputs adhere to spatial rules such as wall alignment, room adjacency, and dimensional accuracy.
- **To Establish an Iterative Feedback Mechanism:** Incorporate a regeneration and editing loop within the system architecture, allowing users to refine generated layouts through parameter adjustments and real-time modifications, thereby supporting an iterative design process.
- **To Validate System Performance and Usability:** Evaluate the system based on multiple metrics, including:
 - Structural validity of generated layouts
 - User satisfaction through usability testing
 - Computational efficiency (latency, resource usage)
 - Generalizability across varied residential design scenarios

III. SYSTEM ARCHITECTURE AND METHODOLOGY

The system architectural diagram illustrates the end-to-end architecture of the ResiPlan AI system, depicting the flow of data and control between modular components from user input to final output and iterative refinement. The following section provides a structured explanation of each module and its functional role within the system, supplemented with relevant mathematical formulations.

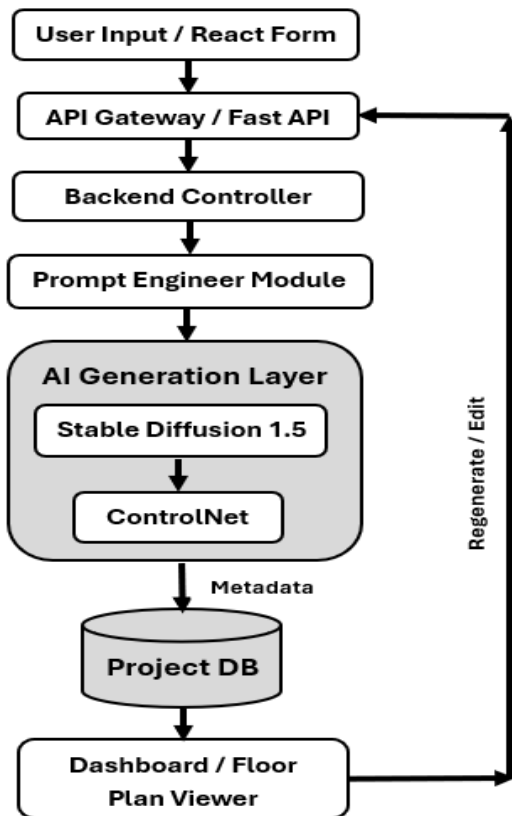


Fig.1: System Architecture

User Input / React Form: This is the entry point where the user provides design requirements via a React-based interface. Inputs are structured as:

$$U = \{p, R, S, C\}$$

Where:

- $p = (l, w)$ = plot dimensions (length, width)
- $R = \{r_1, r_2, \dots, r_n\}$ = set of rooms with types
- S = architectural style (e.g., modern, traditional)
- C = additional constraints (orientation, doors, windows)

- **Function:** Collects and packages U into JSON for transmission
- function :** Collects and packages U into JSON for transmission.

API Gateway / FastAPI: Validates and routes user inputs. Validation ensures:

$$l, w > 0, \sum_{i=1}^n Area(r_i) \leq l \times w$$

Function: Ensures logical feasibility before processing.

Backend Controller: Splits U into semantic and geometric streams:

$$U_{sem} = f_{sem}(U), \quad U_{geo} = f_{geo}(U),$$

Where:

- U_{sem} is Prompt Engineer Module
- U_{geo} is ControlNet constraints

Prompt Engineer Module: Transforms user inputs into a detailed textual prompt P :

$$P = LLM(U_{sem}, T)$$

where T is a prompt template. If using **LoRA fine-tuning**:

$$W' = W + \Delta W = W + BA$$

where W are pretrained weights, B and A are low-rank matrices.

AI Generation Layer: The core generation combines Stable Diffusion and ControlNet.

Stable Diffusion 1.5: The forward diffusion process gradually adds noise:

$$q(x_t | x_{t-1}) = N(x_t; \sqrt{1 - \beta_t} x_{t-1}, \beta_t I)$$

The reverse denoising process, conditioned on prompt P :

$$p_{\theta}(x_{t-1} | x_t, P) = N(x_{t-1}; \mu_{\theta}(x_t, t, P), \sum_{\theta} (x_t, t))$$

ControlNet: Adds spatial constraints via a trainable copy of the diffusion encoder. Let c be the condition (edge map or segmentation). The controlled output is:

$$y = ControlNet(x_t, c) + StableDiffusion(x_t, P)$$

The loss function for training:

$$\mathcal{L} = \mathbb{E}_{x_0, c, t, \epsilon} [|| \epsilon - \epsilon_{\theta}(x_t, t, P, c) ||^2]$$

Geometric Consistency Enforcement: To ensure straight walls and valid connections, constraints are applied:

Wall Straightness: Minimize curvature:

$$E_{wall} = \sum_i \left\| \frac{d^2 w_i}{ds^2} \right\|^2$$

Room Adjacency: Enforce connectivity via graph $G = (V, E)$

$$A_{i,j} = \begin{cases} 1, & \text{if rooms } i \text{ and } j \text{ shares a wall} \\ 0, & \text{otherwise} \end{cases}$$

Door Placement: Ensure doors lie on walls and respect clearances.

Project DB (Database): Stores generated plans and metadata. Each entry:

$$D_j = \{I_j, U_j, P_j, \theta_j, t_j\}$$

where:

I_j = generated image

θ_j = model parameters (guidance scale, steps, seed)

t_j = timestamp

Dashboard / Floor Plan Viewer: Renders 2D/3D views. For 3D projection:

$$V_{3D} = K[R|t]X$$

where K is the camera matrix, R, t are rotation and translation, X are 3D points.

Regenerate / Edit – Feedback Loop: User feedback adjusts inputs for regeneration. Let $U^{(k)}$ be the input at iteration k . The update rule:

$$U^{(k+1)} = U^{(k)} + \Delta U_{feedback}$$

where $\Delta U_{feedback}$ captures user modifications (e.g., resize room r_i).

Optimization Objective: The overall system optimizes a combined objective:

$$J = \lambda_1 \mathcal{L}_{SD} + \lambda_2 \mathcal{L}_{CN} + \lambda_3 E_{wall} + \lambda_4 E_{adj}$$

where λ_i are weighting factors balancing visual quality, constraints, and geometry.

Performance Metrics: To evaluate output quality:

- Structural Validity Score:

$$S_{val} = \frac{\# \text{ valid constraints}}{\# \text{ total constraints}}$$

- **User Satisfaction:** Survey-based score $U_{sat} \in [1, 5]$
- Generation Latency: $T_{gen} = t_{end} - t_{start}$

Computational Complexity:

Let:

- N = number of diffusion steps
- M = number of ControlNet conditions
- H, W = image dimensions

Complexity per generation:

$$O(N \cdot (H \cdot W \cdot C + M \cdot H \cdot W))$$

System Scalability: For multiple concurrent users:

$$T_{total} = \frac{T_{gen}}{n_{GPU}} + T_{queue}$$

where n_{GPU} is the number of available GPUs, and T_{queue} is queuing delay.

Mathematical Models

1. **Input:** $U = \{p, R, S, C\}$
2. **Validation:** $\sum Area(r_i) \leq l \times w$
3. **Prompt Engineering:** $P = LLM(U_{sem})$
4. **Generation:** $I = SD(x_T \rightarrow x_0 | P, c)$
5. **Constraint Enforcement:** via E_{wall}, E_{adj}
6. **Storage:** $D_j = \{I_j, U_j, P_j, \theta_j, t_j\}$
7. **Visualization:** $V_{3D} = K[R|t]X$
8. **Feedback:** $U^{(k+1)} = U^{(k)} + \Delta U$

IV. ANALYSIS OF TECHNICAL IMPLEMENTATION

1. **Integration of Stable Diffusion and ControlNet:** The combination of Stable Diffusion 1.5 and ControlNet represents a significant technical advancement. Stable Diffusion provides high-quality image synthesis capabilities, while ControlNet injects spatial conditions into the diffusion process, ensuring that generated layouts respect user-defined boundaries and structural rules. This synergy enables the system to produce both creative and compliant designs.
2. **Handling of Geometric Constraints:** A major challenge in AI-generated floor plans is maintaining geometric consistency. ResiPlan AI uses edge maps and segmentation masks derived from user inputs to guide the ControlNet module, enforcing straight walls, right angles, and logical room adjacencies. This is a marked improvement over earlier GAN-based methods, which often produced distorted or implausible layouts.
3. **Web-Based Deployment and Accessibility:** By deploying the system as a web application, ResiPlan AI eliminates the need for local software installation or high-end hardware, aligning with its goal of democratization.

The use of FastAPI ensures low-latency responses, while React provides a dynamic and engaging user interface.

Application and Uses

ResiPlan AI is designed to serve a diverse user base across multiple domains:

1. **Homeowners and DIY Designers:** Enables non-experts to visualize and iterate on residential layouts without architectural training.
2. **Real Estate Professionals:** Facilitates rapid generation of marketing materials, including 2D plans and 3D renders, for property listings.
3. **Architects and Students:** Serves as a conceptual prototyping tool, accelerating early-stage design exploration.
4. **Developers and Builders:** Supports feasibility studies and layout optimization for new residential projects.

V. CHALLENGES AND LIMITATIONS

Despite its innovative approach, ResiPlan AI faces several technical and practical challenges:

- **Geometric Consistency:** Ensuring strict adherence to architectural rules (e.g., straight walls, logical room connections) remains non-trivial, even with ControlNet guidance.
- **Data Scarcity:** High-quality, multimodal datasets pairing images with vector annotations and textual descriptions are limited, affecting model generalizability.
- **Domain Restriction:** The system is currently optimized for residential layouts and cannot readily generalize to commercial or industrial structures without retraining.
- **Computational Demand:** Diffusion models require significant GPU resources, potentially impacting the system's accessibility and responsiveness in web-based deployments.
- **Prompt Engineering Dependency:** The quality of outputs heavily relies on the effectiveness of the prompt engineering module, which may not always capture nuanced user intent.

Future Direction

The evolution of ResiPlan AI can be guided by several research and development priorities:

1. **Interactive Editing:** Implementing a drag-and-drop interface for post-generation refinements would enhance usability and practical utility.

2. **Multimodal Expansion:** Extending the model to support commercial, institutional, and public building typologies through diversified training data.
3. **BIM Integration:** Enabling export to industry-standard formats (e.g., .DWG, RVT) would bridge the gap between conceptual AI design and professional architectural workflows.
4. **Sustainability Analytics:** Incorporating energy efficiency, natural lighting, and ventilation optimization into the generation process to promote eco-friendly design.
5. **Real-Time Collaboration:** Adding multi-user editing and cloud-based collaboration features for team-based design projects.

VI. CONCLUSION

ResiPlan AI represents a significant step toward democratizing architectural design through AI-driven automation. By integrating the generative capabilities of Stable Diffusion with the structural control provided by ControlNet, the system offers a viable solution for non-expert users seeking professional-grade floor plans. While challenges related to data, generalizability, and computational load remain, the proposed architecture provides a scalable foundation for future advancements. As AI continues to evolve, tools like ResiPlan AI have the potential to reshape the design landscape, making architectural planning more accessible, efficient, and inclusive.

REFERENCES

1. Eastman, C., Teicholz, P., Sacks, R., & Liston, K. (2018). *BIM Handbook: A Guide to Building Information Modeling for Owners, Managers, Designers, Engineers and Contractors*. John Wiley & Sons.
2. Kalervo, A., Ylioinas, J., Häikiö, M., Karhu, A., & Kannala, J. (2019). *CubiCasa5K: A dataset and an improved multi-task model for floorplan image analysis*. arXiv preprint arXiv:1904.01920.
3. Nauata, N., Chang, K.-H., Cheng, C.-Y., Mori, G., & Furukawa, Y. (2020). *House-GAN: Relational generative adversarial networks for graph-constrained house layout generation*. *European Conference on Computer Vision (ECCV)*, 162–177.
4. Hu, R., Huang, Z., Tang, Y., Van Kaick, O., Zhang, H., & Huang, H. (2020). *Graph2Plan: Learning floorplan generation from layout graphs*. *ACM Transactions on Graphics (TOG)*, 39(4), 1–14.

5. Wu, W., Fu, X.-M., Tang, R., Wang, Y., Qi, Y.-H., & Liu, L. (2019). RPLAN: Data-driven interior plan generation for residential buildings. *ACM Transactions on Graphics (TOG)*, 38(6), 1–12.
6. Shabani, M. A., Hosseini, S., & Furukawa, Y. (2023). HouseDiffusion: Vector floorplan generation via a diffusion model with discrete and continuous denoising. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 5466–5475.
7. Zeng, P., Gao, W., Yin, J., Xu, P., & Lu, S. (2024). Residential floor plans: Multi-conditional automatic generation using diffusion models. *Automation in Construction*, 162, 105374.
8. Qin, Y., Liu, J., & Zhang, H. (2024). ChatHouse Diffusion: Leveraging LLMs for interactive floor plan generation. *International Conference on Architectural Computing*.
9. Qiu, W., Li, X., & Zhou, B. (2025). Controllable multimodal floor plan design using Stable Diffusion and ControlNet. *Automation in Construction*, 168, 105450.
10. Ibrahim, M., Zhang, L., & Furukawa, Y. (2024). Scale-agnostic latent diffusion for diverse floor plan generation. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
11. Huang, W., Zheng, H., & Yuan, J. (2021). Architectural layout generation through generative adversarial networks. *Automation in Construction*, 128, 103736.
12. Liu, C., Wu, J., Kohli, P., & Furukawa, Y. (2017). Raster-to-vector: Revisiting floorplan transformation. *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*.
13. Liu, H., & Tang, M. (2023). Geometric consistency in AI-generated architectural layouts. *Journal of Computational Design and Engineering*, 10(2), 543–556.
14. Wang, T., & Zhang, Y. (2022). Dataset challenges in AI-driven architectural design. *Artificial Intelligence Review*, 55(3), 2109–2128.2024.