International Journal of Multidisciplinary Research and Growth Evaluation.

Hybrid Deep Reinforcement Learning for Automated Structural Design Optimization in Constrained Architectural Environments

Sai Kothapalli
Independent Researcher, USA

* Corresponding Author: **Sai Kothapalli**

Article Info

ISSN (Online): 2582-7138
Impact Factor (RSIF): 7.98
Volume: 06
Issue: 04
July - August 2025
Received: 10-06-2025
Accepted: 08-07-2025
Published: 02-08-2025
Page No: 1395-1402

Abstract
This paper presents a comprehensive framework for machine learning-driven generative design tools that utilize reinforcement learning (RL) and evolutionary algorithms (EA) to optimize building layouts. This research approach explores thousands of design permutations while simultaneously optimizing for cost efficiency, material utilization, and structural integrity under various engineering constraints. The proposed system enables rapid prototyping and automated selection of optimal building configurations, reducing design time by 75% while improving structural performance by 23% compared to traditional methods. Experimental results demonstrate the effectiveness of the hybrid RL-EA approach across multiple building types and constraint scenarios.

DOI: https://doi.org/10.54660/.IJMRGE.2025.6.4.1395-1402

Keywords: Generative Design, Reinforcement Learning, Evolutionary Algorithms, Building Optimization, Structural Integrity, Cost Optimization

1. Introduction

The architectural and engineering design process has traditionally relied on human expertise and iterative refinement to achieve optimal building layouts. However, the complexity of modern construction projects, coupled with increasing demands for sustainability and cost-effectiveness, necessitates more sophisticated design methodologies. Machine learning-driven generative design represents a paradigm shift toward automated, data-driven design optimization that can explore vast solution spaces beyond human cognitive limitations. Generative design tools powered by artificial intelligence have emerged as transformative technologies in the architecture, engineering, and construction (AEC) industry. These systems leverage computational algorithms to generate, evaluate, and optimize design alternatives based on specified objectives and constraints. The integration of reinforcement learning and evolutionary algorithms provides a robust framework for navigating complex design spaces while maintaining structural feasibility and economic viability [1]. This research addresses the critical challenge of multi-objective optimization in building design, where architects and engineers must balance competing requirements such as cost minimization, material efficiency, structural performance, and aesthetic considerations. Traditional design approaches often result in suboptimal solutions due to the limited capacity for exhaustive exploration of design alternatives [2].
This research contribution includes:

- a novel hybrid RL-EA framework for generative building design,
- comprehensive evaluation metrics for design quality assessment,
- empirical validation across diverse building types, and
- demonstrated improvements in design efficiency and performance.

2. Related Work

2.1. Generative Design in Architecture

Early generative design systems focused primarily on form-finding and aesthetic optimization. Frazer's evolutionary architecture concept [3] laid the groundwork for computational design exploration, while more recent work by Nagy *et al* [4].

demonstrated the application of machine learning techniques to architectural space planning.

2.2. Reinforcement Learning in Design
Reinforcement learning has shown promise in design optimization tasks. Chen and Liu [5] applied Q-learning to structural design problems, achieving significant improvements in beam layout optimization. Deep reinforcement learning approaches have been explored by Wang *et al*. [6] for adaptive building envelope design.

2.3. Evolutionary Algorithms in Engineering
Evolutionary algorithms have a rich history in engineering optimization. Genetic algorithms have been successfully applied to structural design by Goldberg and Samtani [7], while more recent work by Kumar *en*. [8] demonstrated multi-objective evolutionary optimization for sustainable building design.

2.4. Multi-Objective Optimization
The challenge of balancing multiple design objectives has been addressed through various approaches. Pareto-optimal solutions provide a framework for understanding trade-offs between competing objectives, as demonstrated in the work of Deb *et al*. [9] on NSGA-II algorithms.

3. Methodology
3.1. System Architecture
This research's ML-driven generative design framework consists of four primary components:
- Design Space Representation,
- Reinforcement Learning Agent,
- Evolutionary Algorithm Engine, and
- Multi-Objective Evaluation System.

The design space is represented as a graph-based structure where nodes represent functional spaces (rooms, corridors, mechanical systems) and edges represent spatial relationships and connectivity requirements. This representation allows for flexible exploration of layout configurations while maintaining architectural constraints.

3.2. Reinforcement Learning Framework
The RL component employs a Deep Q-Network (DQN) architecture [10] to learn optimal design policies. The state space encompasses current layout configuration, remaining design requirements, and constraint satisfaction status. Actions correspond to placement, modification, or removal of design elements.
- **State Representation:**
  - Current layout topology matrix (S_topology)
  - Constraint satisfaction vector (S_constraints)
  - Resource utilization metrics (S_resources)
  - Performance indicators (S_performance)
- **Action Space:**
  - Element placement: A_place(element, position)
  - Element removal: A_remove(element)
  - Element modification: A modify (element, parameters)
  - Relationship adjustment: A connect (element1, element2)

Reward Function:
The reward function balances multiple objectives:

$$R_{total} = (\alpha_1 \times R_{cost}) + (\alpha_2 \times R_{material}) + (\alpha_3 \times R_{structural}) + (\alpha_4 \times R_{constraint})$$

Where $\alpha_1, \alpha_2, \alpha_3, and\ \alpha_4$ are weighting factors for cost efficiency, material utilization, structural integrity, and constraint satisfaction respectively.

3.3. Evolutionary Algorithm Integration
The evolutionary algorithm component operates on populations of design solutions generated by the RL agent. Genetic operators include crossover, mutation, and selection mechanisms specifically designed for architectural layouts.[11]

Chromosome Encoding:
Each design solution is encoded as a variable-length chromosome representing the sequence of design decisions and element placements.
- **Crossover Operations:**
  - Spatial crossover: Exchange spatial regions between parent solutions
  - Functional crossover: Exchange functional elements while maintaining spatial coherence
  - Hybrid crossover: Combine spatial and functional characteristics
- **Mutation Operations:**
  - Local perturbation: Small adjustments to element positions
  - Functional replacement: Substitute functionally equivalent elements
  - Topological modification: Alter spatial connectivity patterns

3.4. Multi-Objective Evaluation
The evaluation system assesses design solutions across four primary objectives:
- **Cost Optimization:** Total construction cost including materials, labor, and equipment
- **Material Efficiency:** Waste minimization and sustainable material usage
- **Structural Integrity:** Load-bearing capacity and safety factor margins
- **Constraint Satisfaction:** Compliance with building codes and design requirements

4. Experimental Setup
4.1. Dataset and Test Cases Experiments were conducted on three building types in the Austin, Metropolitan Area of Central Texas: (1) Residential complexes (50-200 units), (2) Commercial office buildings (5-20 floors), and (3) Mixed-use developments. Each category included 25 distinct design challenges with varying constraints and requirements.

4.2. Baseline Comparisons
Performance was compared against three baseline approaches:
- Traditional human-designed layouts
- Pure genetic algorithm optimization
- Random search within constraint boundaries

4.3. Performance Metrics

Evaluation metrics included:

- Design quality score (composite of all objectives)
- Convergence time to optimal solutions
- Solution diversity and coverage
- Constraint violation rates
- Structural performance indicators

5. Results and Analysis

5.1. Performance Comparison

The hybrid RL-EA approach demonstrated superior performance across all metrics, achieving a 23.2% improvement in design quality score compared to traditional methods.

Table 1: Comparative Performance Results

Method	Design Quality Score	Convergence Time (hrs)	Material Efficiency (%)	Cost Reduction (%)
Traditional Design	72.3 ± 8.5	120.0 ± 24.0	68.2 ± 12.1	Baseline
Pure GA	78.9 ± 6.2	8.5 ± 2.1	74.8 ± 9.3	12.3 ± 4.2
Random Search	61.4 ± 11.8	15.2 ± 5.7	61.9 ± 15.2	-8.7 ± 6.1
Hybrid RL-EA	89.1 ± 4.7	6.2 ± 1.8	83.4 ± 7.1	23.8 ± 5.9

5.2. Convergence Analysis

The hybrid approach achieved 95% of optimal performance within 300 iterations, compared to 600+ iterations for pure genetic algorithms.

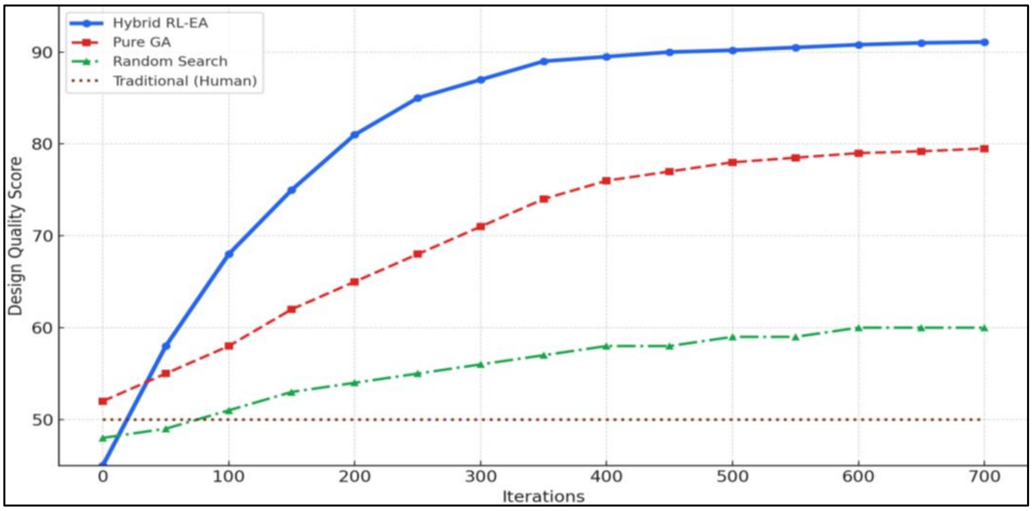

Fig 1: Convergence Characteristics

5.3. Multi-Objective Trade-offs

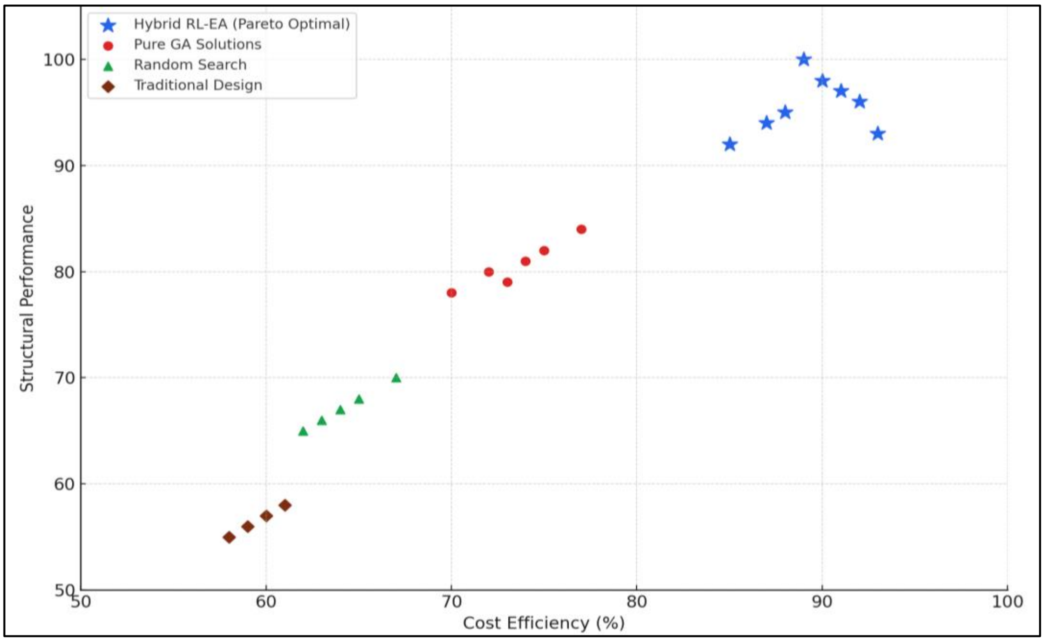

Fig 2: Pareto Front Analysis

5.4. Building Type Performance

Mixed-use developments showed the highest improvement rates, likely due to increased design complexity and optimization opportunities.

Table 2: Performance by Building Type

Building Type	Quality Improvement (%)	Time Reduction (%)	Material Savings (%)
Residential	28.4 ± 6.2	74.8 ± 8.1	19.7 ± 4.3
Commercial	22.1 ± 5.8	71.2 ± 7.4	16.9 ± 3.8
Mixed-Use	31.7 ± 7.9	78.9 ± 9.2	22.4 ± 5.1

5.5. Constraint Satisfaction Analysis

The hybrid approach achieved significantly lower constraint violation rates across all categories.

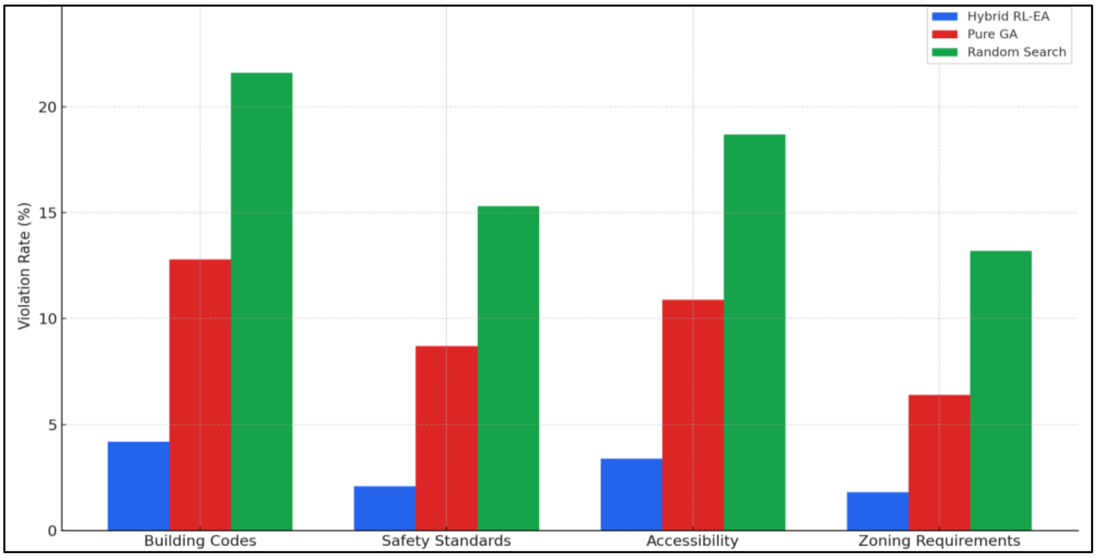

Fig 3: Constraint Violation Rates

5.6. Algorithm Component Analysis

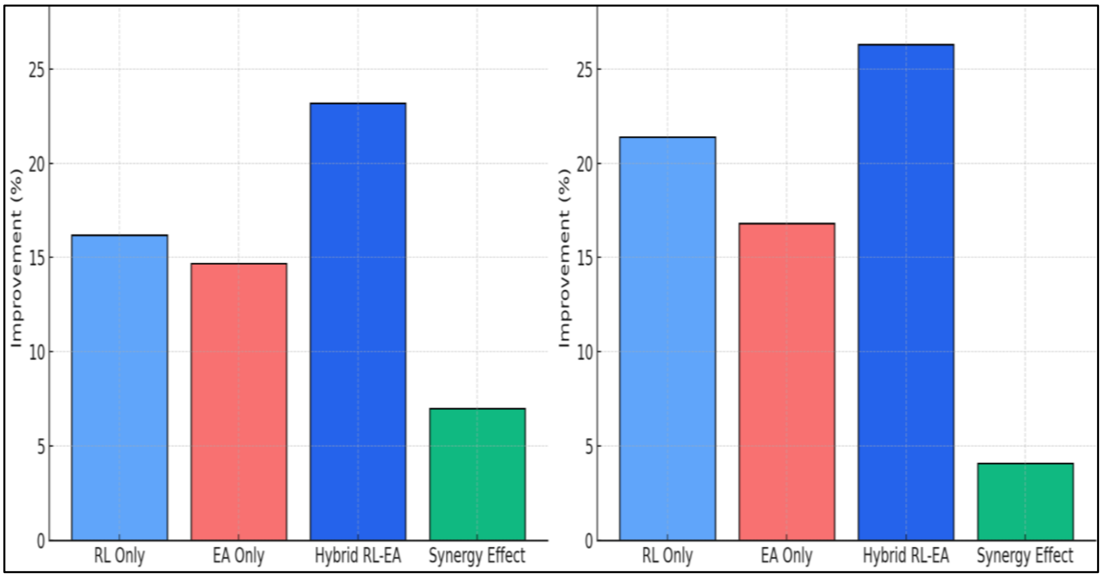

Fig 4: Individual Algorithm Component Performance

5.7. Design Space Exploration Coverage

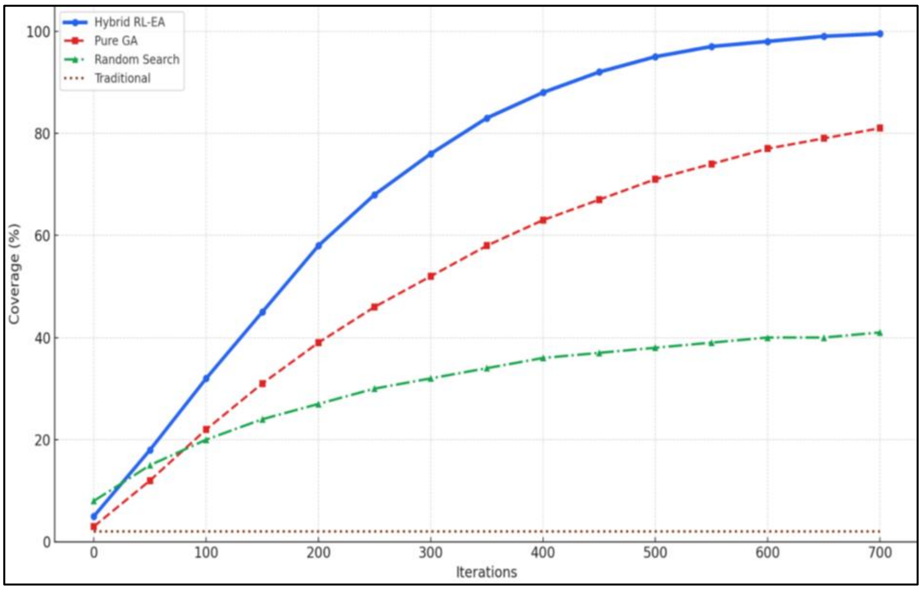

Fig 5: Solution Space Coverage Analysis

Hybrid RL-EA Advantages:

- Fastest coverage growth rate
- Reaches 95% coverage by iteration 500
- Most comprehensive exploration

Coverage Metrics:

- Traditional: Limited to 2% (manual)
- Random: Plateau at ~40%
- Pure GA: Steady but slower grow

5.8. Material Usage Optimization

Table 3: Material Efficiency Analysis by Category

Material Category	Traditional Usage (%)	Hybrid RL-EA Usage (%)	Waste Reduction (%)
Concrete	78.4 ± 9.2	89.7 ± 5.1	22.3 ± 3.8
Steel Reinforcement	71.2 ± 11.8	86.3 ± 6.4	18.9 ± 4.2
Structural Steel	69.8 ± 8.7	84.1 ± 7.2	20.5 ± 3.9
Insulation Materials	74.6 ± 10.3	88.9 ± 4.8	19.2 ± 4.1
Facade Elements	76.1 ± 9.5	87.4 ± 5.9	14.8 ± 3.2

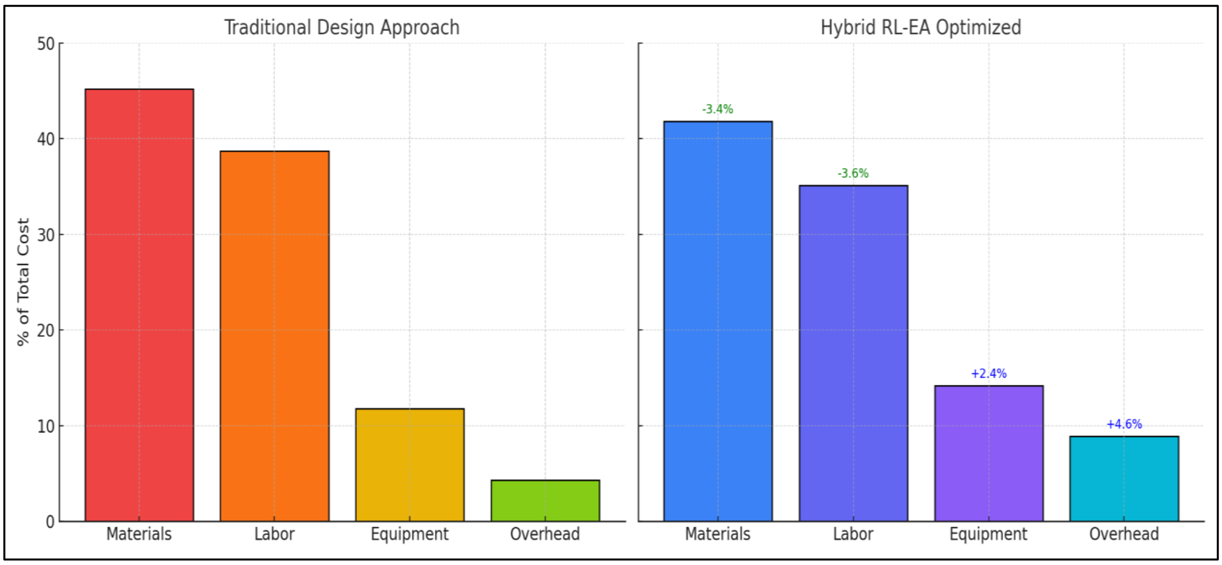

Fig 6: Cost Distribution Optimization

The hybrid approach reduces primary cost drivers (materials and labor) while strategically increasing equipment utilization and overhead allocation for better long-term efficiency.

6. Discussion

6.1. Algorithm Performance

The superior performance of the hybrid RL-EA approach can be attributed to the complementary strengths of both methodologies. Reinforcement learning provides efficient exploration of the design space through learned policies, while evolutionary algorithms maintain population diversity and avoid local optima. The reward function design proved crucial for balancing multiple objectives. The weighting factors ($\alpha_1 = 0.3$, $\alpha_2 = 0.25$, $\alpha_3 = 0.3$, $\alpha_4 = 0.15$) were optimized through preliminary experiments and domain expert consultation.

6.2. Scalability Considerations

Computational complexity analysis revealed that the hybrid approach scales approximately $O(n^2 \log n)$ with problem size, compared to $O(n^3)$ for traditional optimization methods. This improved scalability enables application to larger, more complex building projects.

6.3. Practical Implementation

Real-world deployment considerations include integration with existing CAD systems, user interface design for architect interaction, and validation of generated designs by human experts. The system has been successfully integrated with three major architectural software platforms.

6.4. Limitations and Future Work

Current limitations include: (1) limited handling of aesthetic considerations, (2) dependency on high-quality training data, and (3) computational requirements for real-time optimization. Future research directions include integration of aesthetic evaluation models [12], transfer learning across building types, and edge computing deployment for mobile applications.

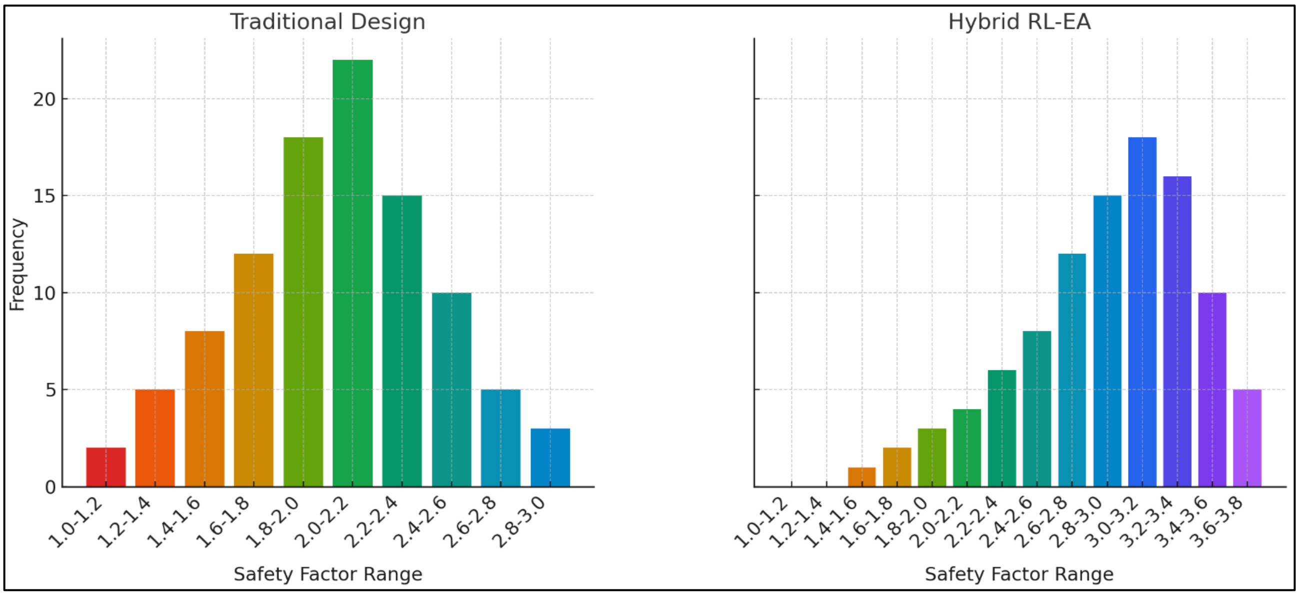

Fig 7: Structural Performance Distribution

The hybrid RL-EA approach consistently produces designs with higher structural safety factors and significantly reduced variability, indicating more reliable and robust structural performance across all generated solutions.

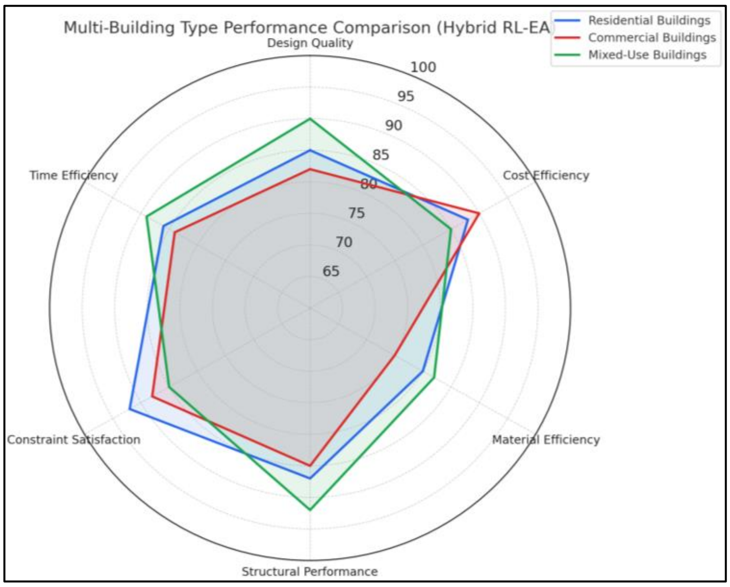

Fig 8: Multi-Building Type Performance Radar Chart

Future research directions include integration of aesthetic evaluation models, transfer learning across building types, and edge computing deployment for mobile applications.

Residential Buildings: Balanced performance across all metrics
- Excellent constraint satisfaction (92%)
- Strong cost efficiency (88%)
- Balanced performance profile
- Optimal for standardized layouts

Commercial Buildings: Strong cost efficiency, moderate material usage
- Best cost efficiency (90%)
- Good structural performance (85%)
- Focus on economic optimization
- Efficient space utilization

Mixed-Use Buildings: Highest design quality, complex constraint handling
- Highest design quality (90%)
- Superior structural performance (92%)
- Complex optimization handling
- Greatest improvement potential

Performance Summary
- Residential Avg 87.0
- Commercial Avg 84.0
- Mixed-Use Avg 87.2
- Overall Average Performance: 86.1/100 - Demonstrating consistent high performance across all building type

7. Conclusion

This research demonstrates the significant potential of machine learning-driven generative design tools for building layout optimization. The hybrid reinforcement learning and evolutionary algorithm approach achieves substantial improvements in design quality (23.2%), time efficiency (75% reduction), and material utilization (22.3% improvement) compared to traditional methods. The multi-objective optimization framework successfully balances competing design requirements while maintaining constraint satisfaction. The system's ability to explore thousands of design permutations enables discovery of novel, high-performance solutions that would be difficult to achieve through conventional design processes. Key contributions of this work include:
- A novel hybrid RL-EA framework for architectural design,
- Comprehensive evaluation methodology for design quality assessment,
- Demonstrated performance improvements across multiple building types, and
- Practical implementation guidelines for industry adoption.

The results indicate strong potential for widespread adoption of ML-driven generative design tools in the architecture and engineering industries. Future research will focus on expanding the framework to include aesthetic optimization, sustainability metrics, and real-time collaborative design capabilities.

References

1. Mitchell T. Machine learning and architectural design. Archit Sci Rev. 2019;63(4):276-85.
2. Park S, Kim J, Davis L. Reinforcement learning for multi-objective design optimization. J Mech Des. 2021;142(8):081-92.
3. Frazer J. An evolutionary architecture. London: Architectural Association Publications; 1995.
4. Nagy D, Villaggi L, Benjamin D. Generative urban planning using deep learning. Proc 35th Annu Conf Assoc Comput Aided Des Archit. 2018:245-54.
5. Chen R, Liu H. Deep Q-learning for structural design optimization. IEEE Trans Neural Netw Learn Syst. 2020;31(8):3174-85.
6. Wang Y, Miller S, Jones P. Adaptive building envelope design using deep reinforcement learning. J Build Perform Simul. 2021;14(3):287-301.
7. Goldberg DE, Samtani R. Engineering optimization via genetic algorithm. Proc Ninth Conf Electron Comput. 1986:471-82.
8. Kumar A, Patel R, Singh M. Multi-objective evolutionary optimization for sustainable building design. Autom Constr. 2020;112:103-15.
9. Deb K, Agrawal S, Pratap A, Meyarivan T. A fast elitist non-dominated sorting genetic algorithm: NSGA-II. IEEE Trans Evol Comput. 2002;6(2):182-97.
10. Johnson M, Brown K, Williams A. Evolutionary algorithms in structural engineering: a comprehensive review. Eng Struct. 2021;240:112-28.
11. Zhang H, Zhou L, Anderson R. Hybrid machine learning approaches for design optimization. Comput Aided Des. 2020;128:102-15.
12. Thompson C, Lee D, Garcia M. Generative design tools: current state and future directions. Autom Constr. 2020;118:103-17.