WWW.IJRPR.COM

International Journal of Research Publication and Reviews

Journal Homepage: www.ijrpr.com ISSN 2582-7421

HAVEN: A Review of An AI assisted House Planning and Construction Cost Estimation System

Prof. Heena Kachhela, Bhavesh Hatwar, Bhuvan Bhonde, Vansh Giradkar, Chetan Kolhe, Akansha Kawale

Priyadarshini College of Engineering, Nagpur

ABSTRACT

Computational methods are increasingly being explored to support different aspects of residential building planning. This paper reviews twenty-two studies covering architectural layout generation, multi-criteria decision-support methods, and data-driven construction cost analysis. The review compares the principal techniques, reported outcomes, evaluation settings, and limitations described in the selected studies. In architectural design, graph-based optimization and deep generative models have expanded the ways in which spatial layouts can be represented, generated, and reconstructed. Decision-support research offers established methods for handling competing criteria, while cost-analysis studies report growing use of machine-learning models, including ensemble and neural-network approaches. Despite these developments, recurring challenges include geometric inconsistencies, dependence on structured input, limited diversity in datasets, and variation in model performance across evaluation settings. The paper consolidates these findings and highlights general research needs related to usability, reliability, data quality, and external validation in computational approaches to residential planning.

Index Terms — computational architecture, architectural layout generation, generative models, decision-support methods, machine learning, built-environment analysis

I. INTRODUCTION

Residential building planning involves the organization of interior spaces and the consideration of functional, technical, and economic constraints. Computational research has addressed these concerns through a range of methods, including graph-based and generative models for spatial layout, multi-criteria techniques for structured evaluation, and statistical or machine-learning approaches for construction-related analysis. Each area has developed its own representations, datasets, and evaluation practices. Understanding the strengths and limitations of these approaches is important when assessing their suitability for practical built-environment applications.

This review examines twenty-two studies relevant to computational methods in residential building planning and related built-environment tasks. It summarizes the methodological contributions and reported findings, considers the limitations identified by the original authors or evident from their evaluations, and compares the approaches within their respective research areas. Particular attention is given to recent developments, the data and evaluation settings used, and recurring issues affecting practical applicability. Where two studies report the same underlying method at different stages of publication (conference and extended journal form), only the original version is treated as a distinct contribution to avoid duplicate representation of a single body of work.

The remainder of this paper is organized as follows. Section II reviews the literature in three thematic groups: generative and graph-based layout design, multi-criteria decision-support methods, and machine-learning-based construction analysis. Section III presents a comparative summary of all reviewed studies. Section IV discusses recurring research challenges. Section V outlines general directions for future research, and Section VI concludes the paper.

II. LITERATURE REVIEW

A. Generative and Graph-Based Architectural Layout Design

Yu et al. [1] conducted a systematic review of twenty-nine studies on architectural design collaboration within immersive virtual environments, organizing the literature around three perspectives: the use of such environments within the architecture-engineering-construction industry generally, their role in supporting remote collaboration, and their specific application to collaborative design tasks. The review found that, despite growing adoption, collaborative tools available at the time supported only simplified sketching and modelling operations; parametric design capability was typically supplied through external plugins rather than integrated natively, and higher-order analyses such as cost estimation, land-use assessment, and building-performance

simulation were largely unsupported within the collaborative environment itself. This finding is significant for any system that aims to combine design generation with downstream evaluation, since it suggests that such integration has not been a design priority in prior collaborative tools.

Sepúlveda et al. [5] addressed layout generation more directly, developing a Unity-based plugin in which a user specifies rooms and door connections as a node-edge graph together with size and aspect-ratio constraints. A combination of stochastic and steepest-ascent hill-climbing search then converts this graph into an optimized, navigable three-dimensional blueprint, which the user may further edit. While the approach demonstrates that non-expert users can generate structurally valid layouts through simple graph specification, the authors evaluated it on only ten test cases, restricted output to tile-based (axis-aligned) geometry, and noted that the stochastic search component, while computationally efficient, does not guarantee a high-quality result on every run.

A more recent generative approach, MaskPLAN [6], addresses a limitation implicit in most prior graph-to-layout systems: the assumption that the user provides a complete specification before generation begins. MaskPLAN instead employs a graph-structured dynamic masked autoencoder together with five attribute-specific transformer modules, allowing a layout to be generated from a partial specification — for instance, a single room's size and adjacency preference — while the model infers the remaining attributes. Because the transformer modules are trained jointly, the model captures dependencies between attributes such as room size, location, and connectivity rather than treating them independently, which the authors identify as a shortcoming of earlier user-guided generators.

House-GAN++ [7] takes an adversarial approach to the same problem, combining a graph-constrained relational generative adversarial network with a conditional generative adversarial network so that a layout produced in one generation cycle becomes a constraint for the next, allowing progressive refinement toward a result comparable in quality to those produced by professional architects. The evaluation, however, was conducted entirely offline against pre-collected datasets; the authors do not report testing the system in a live, iterative session with a human designer, and separately observe that the model's door-placement behaviour does not always correspond to the order in which a human designer would place doors during actual design practice.

Takeuchi et al. [10] propose a single-stage alternative to the multi-stage pipelines typical of GAN-based layout generation. Their model applies a diffusion process parameterized by beta distributions — chosen because wall-junction coordinates are continuous while their connections are discrete — to generate both simultaneously in one pass. The resulting model, at approximately five million parameters, is over an order of magnitude smaller than comparable multi-stage systems exceeding one hundred million parameters, while achieving comparable generation quality as measured by Fréchet Inception Distance. The authors report, however, that 36.6 percent of generated samples contain geometrically invalid crossing walls, that room-type classification errors persist, and — notably for usability — that the model requires the number of wall junctions to be specified in advance, a parameter an end user is unlikely to know.

ActFloor-GAN [11] incorporates a human-centric signal into the generation process by predicting room types conditioned on a human-activity map, a heatmap-style representation of anticipated occupant movement within the space, using a reformulated cycle-consistency constraint during adversarial training. Architect evaluators rated the resulting layouts as professionally credible and topologically sound. The authors report two specific weaknesses: the model confuses room types with similar visual and geometric signatures, such as kitchens and bathrooms, or among bedroom sub-types, and the model's internal decision process offers limited interpretability. Evaluation was also confined to residential building types.

Beyond generation, LayoutGMN [12] addresses the related problem of measuring structural similarity between two layouts, offering a graph-based means of assessing correspondence between floor-plan structures. The authors train a graph matching network with a triplet loss on weakly supervised intersection-over-union labels, producing a similarity metric that better matches human similarity judgments than raw geometric overlap or prior baseline methods. The approach nonetheless depends on substantial labelled supervision and does not model hierarchical structure within complex layouts, limiting its demonstrated applicability to relatively simple layout comparisons.

Two further studies extend graph-based methods to layout-adjacent problems. An automated semantics and topology extraction method [9] recovers room identity and connectivity directly from raster (scanned image) floor plans by detecting geometric primitives and reconstructing structure through mixed-integer programming, though at the cost of extensive manual parameter tuning and limited ability to infer irregular structural geometry. Simonsen et al. [17] instead work from CAD-derived graphs rather than raster images, training a graph neural network to classify graph nodes as doors or non-doors and introducing a new dataset for this task; the method substantially outperforms a Faster R-CNN raster-based baseline, particularly under strict localization requirements, though accuracy falls to approximately 75 percent under the strictest threshold tested, and the method is limited to door detection rather than full architectural-element recognition.

Finally, Xu et al. [13] apply reinforcement learning to an architecturally adjacent problem — integrated-circuit floorplanning — formulating component placement as a Markov decision process and combining a graph convolutional network with a reinforcement-learning agent to outperform heuristic placement methods on area and wirelength objectives, with the trained agent shown to transfer to differently sized problem instances without full retraining. The authors note that scalability remains limited on very large problem instances, a limitation attributable to the growth of the state-space representation. Although the application domain differs from residential design, the combination of graph-based state representation with reinforcement learning, together with the demonstrated transferability across problem sizes, offers a methodological pattern applicable to spatial layout problems more broadly.

Two further studies extend the layout-generation problem to alternative input and output modalities. Wei and Li [18] address text-based floorplan generation, in which a user supplies a natural-language description rather than a graph or bubble diagram. A preliminary comparison established that graph input outperforms text input on its own, motivating a dual-stage framework in which a graph generator, pre-trained on script-generated descriptions

and fine-tuned with a large-language-model-based alternative, augments the textual input with a serialized graph before a sequence-to-sequence floorplan generator produces the output. The augmented approach improved bounding-box and boundary accuracy over text alone, but the authors report that both graph generators struggled specifically at extracting spatial relations between rooms, and that excluding relation information from even ground-truth graphs reduced performance to a level comparable with generated graphs — indicating that relation extraction, rather than room identification, is the binding constraint on text-based generation.

Zhang et al. [19] address a complementary problem: reconstructing a standardized floor plan from a hand-drawn sketch rather than generating one from a specification. Their framework defines twelve categories of structure descriptors (door and window elements, structural points, and direction indicators), recognizes them with a YOLOX-based detector, and reconstructs standardized geometry through descriptor connection, coordinate adjustment, and door and window restoration sub-algorithms. The method substantially outperformed sketch-synthesis baselines on line and angle reconstruction accuracy and operated within an average of 0.474 seconds per plan, meeting real-time requirements for the robot-navigation application the authors target. The framework is, however, explicitly limited to common indoor floor-plan structures, and its reconstruction accuracy for basic structural points, while high overall, remained lower than its accuracy for doors and windows, indicating that fine-grained line-and-junction recognition remains the harder sub-problem even within a purpose-built descriptor scheme.

B. Multi-Criteria Decision-Support Methods

A second body of literature addresses the evaluation of alternatives when multiple, potentially conflicting criteria are involved. The reviewed examples focus mainly on supplier and inventory decisions, and examine methods for assigning criterion weights and assessing the stability of resulting rankings. A study introducing the Least-to-Most Range (LMR) weighting method [2] proposes a criteria-weighting technique using two complementary rating scales, applied to a green-supplier-selection problem, and demonstrates lower inconsistency than the widely used Analytic Hierarchy Process and Best-Worst Method when many criteria must be weighted simultaneously. The method's validation, however, is limited to a single numerical case study, and it operates only on fixed, crisp inputs without accommodating uncertainty in the underlying data.

A more comprehensive framework [3] combines subjective methods (AHP and Fuzzy AHP), an objective method (Entropy weighting), and hybrid combinations of the two, followed by TOPSIS ranking, applied across seven criteria and twenty-one sub-criteria for supplier selection. Through twenty-one sensitivity-analysis scenarios, the study finds that hybrid subjective-objective weighting schemes produce the most stable rankings under parameter perturbation, an important finding wherever the reliability of ranking results matters. The study's scope, however, remains fixed to a particular criteria set and supplier pool, and the subjective weighting components remain susceptible to evaluator bias.

A third study [16] applies stochastic optimization, rather than classical MCDM weighting, to a related planning problem: allocating replenishment orders across multiple pre-selected suppliers under uncertain lead times. Comparing order-splitting and order-release-flexibility strategies, the authors find that combining both minimizes expected total cost more effectively than either strategy alone. The model, however, considers cost as the sole optimization criterion, excluding quality and environmental factors, and does not incorporate transportation-mode selection.

C. Data-Driven Construction Cost Analysis

Tajziyehchi [4] examines cost-growth prediction using a dataset of 139 construction projects and 281 candidate features, applying LASSO regularization to reduce the feature set to sixteen significant predictors before comparing statistical regression against artificial neural networks; a Random Forest model achieved the highest coefficient of determination among the models tested. The dataset is drawn from a single geographic region, raising questions about generalizability to other markets, and features missing in more than seventy percent of records were excluded outright, a decision that risks discarding informative, if incomplete, predictors.

An earlier study by Du et al. [8] applies a genetic algorithm to optimize the initial weight configuration of a backpropagation neural network, improving both convergence speed and prediction accuracy relative to a standard backpropagation baseline, validated on real engineering cost-estimation cases. The source material does not report an explicit limitations discussion, though the underlying weakness the method addresses — the slow convergence and inconsistent accuracy of unoptimized backpropagation networks — is only partially resolved by genetic weight initialization, as later ensemble and gradient-boosted approaches in this review demonstrate.

A hybrid stacking approach [14] combines Random Forest, XGBoost, and Logistic Regression to classify construction cost into low, medium, and high bands. Notably, the stacked ensemble achieved 89.5 percent accuracy, marginally below the 91 percent achieved by a standalone Random Forest model — an outcome the authors report without fully resolving, and one that carries a broader methodological implication: ensemble stacking should not be assumed superior to a well-tuned individual model without direct empirical comparison.

A further study [15] trains a backpropagation neural network on eleven cost-impact indicators across 144 prefabricated-construction projects, comparing the Adam, stochastic gradient descent, and Adadelta optimization algorithms; Adam produced the most accurate and stable results, with a maximum reported error of 3.89 percent. The authors observe, however, that certain atypical projects produced substantially higher error than the average case, indicating that the model's reliability degrades for projects with unusual characteristics, and that the training data is limited to a single construction methodology.

Three further studies extend the cost-estimation cluster with more recent architectures and a systems-level perspective. R and Kanaga [20] propose a two-stage hybrid stacked ensemble that combines Random Forest, XGBoost, LightGBM, and a ReLU-activated multi-layer perceptron as base learners, whose out-of-fold predictions feed a specification-tier-specific Ridge meta-learner, trained and evaluated on approximately 100,000 Indian residential

construction records. The model achieved a mean absolute percentage error of 4.86 percent and an R-squared of 0.9868, outperforming every individual baseline and prior published benchmarks the authors compare against; an ablation study confirmed that both the tier-specific meta-learner split and the inclusion of the weaker MLP base learner each contributed measurable accuracy gains, illustrating that stacking benefits depend on architectural choices such as stratification rather than being automatic. The system further incorporates an adaptive retraining mechanism triggered jointly by the volume of expert-verified actuals and validation error drift, addressing post-deployment model staleness. The authors note that the dataset is geographically concentrated to two Indian metropolitan markets and that the model currently produces point estimates without calibrated prediction intervals.

Xu et al. [21] propose a hybrid GRU-BPNN architecture that couples a Gated Recurrent Unit for temporal feature extraction with a Backpropagation Neural Network for non-linear regression, trained end-to-end with a single joint loss rather than the loosely coupled two-stage training common in prior GRU-BPNN combinations. Evaluated with mean squared error, root mean squared error, mean absolute error, and mean absolute percentage error, the enhanced model achieved a MAPE of 2.10 percent, outperforming a plain BPNN, a loosely coupled GRU-BPNN, and RNN, LSTM, and Bi-GRU baselines trained on the same data; an ablation study attributed the improvement jointly to the GRU's temporal feature extraction and the end-to-end optimization strategy. The study is comparatively narrow in scope, evaluating on a single dataset without the multi-region or multi-construction-type validation called for elsewhere in this review.

Jin [22] takes a systems-engineering rather than a predictive-modelling approach, presenting a BIM-integrated construction cost control platform architected around a five-layer client-server model with a data-mining layer for classification and association-rule analysis of cost information. Rather than predicting a single project cost figure, the system tracks contract, planned, and actual cost against a five-level deviation-alarm scheme across the construction timeline, enabling real-time material demand tracking and early-warning notification to project managers. This orientation toward continuous monitoring and multi-party data integration, rather than point-in-time estimation, is a complementary capability largely absent from the predictive cost-estimation studies reviewed elsewhere in this paper, though the system as described predates the gradient-boosted and deep-learning approaches now dominant in the field and relies on comparatively dated infrastructure choices.

III. COMPARISON AND ANALYSIS

Table I summarizes the focus area, core technique, principal reported outcome, and principal limitation of each reviewed study, providing a consolidated basis for the discussion of broader research challenges in Section IV. Read across the table, a clear methodological progression is visible within the layout-generation cluster: early work relies on constrained local-search optimization over user-specified graphs [5], adversarial refinement then improves output fidelity toward professional quality [7], and more recent studies move to masked-generative [6] and diffusion-based [10] models capable of operating on partial or compact specifications with substantially fewer parameters. Accuracy and geometric validity, however, do not improve monotonically alongside this progression — the diffusion model [10] reports a notable rate of geometrically invalid output, indicating that parameter efficiency and generation fidelity can remain in tension.

The decision-support studies cover a smaller but methodologically consistent set of approaches. All three [2], [3], [16] treat cost as the dominant or sole criterion, and the two weighting studies [2], [3] report improved ranking stability from hybrid subjective-objective schemes, although only [3] tests this through formal sensitivity analysis. The cost-analysis studies span classical regression, genetic-algorithm-tuned neural networks, gradient-boosted ensembles, recurrent-network hybrids, and optimizer comparisons. Tree-based ensemble methods (Random Forest, XGBoost, LightGBM) match or exceed neural-network baselines in several reported evaluations [4], [14], [20]. However, direct comparisons of stacking with individual models produce different findings: [14] reports that stacking underperforms a standalone Random Forest, while [20] reports improved performance from a tier-specific stacked model. These results emphasize that model performance depends on dataset characteristics and implementation choices. Studies [18] and [19] also indicate that recovering spatial relationships and fine structural details remains challenging even when broad room or element identification is successful. Finally, [22] focuses on continuous cost monitoring, whereas [4], [8], [14], [15], [20], and [21] primarily address prediction. These studies therefore represent distinct but related perspectives within construction cost management.

TABLE I

Comparative Summary of Reviewed Studies

Ref.	Focus Area	Core Technique	Reported Outcome	Principal Limitation
[1]	Design collaboration	Systematic literature review (29 studies)	Maps three research angles in ImVE-based collaboration	Tools offer only shallow sketching/modelling; parametric design not built-in
[2]	Criteria weighting	LMR weighting method	Lower inconsistency than AHP/BWM	Validated on a single numerical case only

Ref.	Focus Area	Core Technique	Reported Outcome	Principal Limitation
[3]	Supplier selection	AHP, Fuzzy AHP, Entropy, TOPSIS, sensitivity analysis	Hybrid weighting most stable across 21 scenarios	Fixed criteria set; subjective bias remains
[4]	Cost prediction	LASSO feature selection + ANN/regression/RF	Random Forest gave the best R^2	Single-region dataset (Alberta); high-missing features dropped
[5]	Layout generation	Graph input + hill-climbing optimization	Automated, editable 3D blueprint from a bubble diagram	Only 10 test layouts; tile-format only
[6]	Layout generation	Masked autoencoder + 5 transformers (MaskPLAN)	Generates full layouts from partial input	Prior models could not handle incomplete specifications
[7]	Layout refinement	Graph-constrained relational GAN + conditional GAN	Near architect-quality iterative layouts	Tested offline only; door placement counter-intuitive
[8]	Cost prediction	GA-optimized BP neural network	Improved convergence and accuracy over plain BP	No explicit limitations reported; slow baseline convergence
[9]	Plan digitization	Geometric-primitive detection + MIP	Extracts semantics/topology from raster floor plans	Heavy manual parameter tuning; limited structural inference
[10]	Layout generation	Single-stage graph beta-diffusion model	5M-parameter model rivals 100M+ multi-stage GANs	36.6% samples show crossing edges; junction count must be pre-specified
[11]	Layout generation	Activity-guided adversarial network (cycle-consistency)	Architect-rated professional-quality, human-centric plans	Confuses visually similar room types; low interpretability
[12]	Layout similarity	Graph Matching Network + triplet loss	Outperforms IoU-based baselines at matching human judgment	Needs strong supervision; no hierarchical representation
[13]	Chip floorplanning	GCN + reinforcement learning (MDP formulation)	Better area/wirelength than heuristics; transfers to new sizes	Scalability limited on very large problem instances
[14]	Cost classification	Stacked Random Forest + XGBoost + Logistic Regression	89.5% stacked vs. 91% standalone Random Forest	Ensemble stacking did not outperform the best single model
[15]	Cost prediction	BP neural network with Adam/SGD/Adadelta	Adam optimizer most accurate and stable (144 projects)	Large error on atypical projects; single construction type

Ref.	Focus Area	Core Technique	Reported Outcome	Principal Limitation
[16]	Supplier/inventory planning	Stochastic lead-time optimization (order splitting + flexibility)	Combined strategy minimizes expected cost	Cost-only model; no quality/environmental criteria
[17]	Element detection	Graph Neural Network on CAD-derived graphs	Outperforms Faster R-CNN, especially at strict IoU	~75% accuracy at IoU > 0.95; doors only, no other elements
[18]	Text-to-layout generation	Dual-stage graph-augmented Seq2Seq (T5)	Graph augmentation improves text-only floorplan generation	Relation extraction accuracy far below room extraction accuracy
[19]	Sketch reconstruction	YOLOX descriptor detection + geometric reconstruction	0.932 line / 0.922 angle accuracy; ~0.47s per plan	Limited to common indoor layouts; point recognition weaker than door/window
[20]	Cost prediction	Stacked RF+XGBoost+LightGBM+MLP, tier-specific Ridge	MAPE 4.86%, R^2 0.9868 on ~100k Indian records	Geographically concentrated data; no prediction intervals
[21]	Cost prediction	End-to-end joint GRU + BPNN	MAPE 2.10%, best among RNN/LSTM/Bi-GRU baselines	Single dataset; no multi-region validation
[22]	Cost control systems	BIM + data mining, 5-layer client-server architecture	Real-time multi-party cost tracking with 5-level deviation alarms	Dated infrastructure; monitoring-oriented, not predictive

IV. RESEARCH CHALLENGES

Across the reviewed studies, several recurring challenges emerge. These are summarized below at the level of methods, data, and evaluation practices.

First, layout-generation methods continue to face issues related to geometric validity and structural consistency. Crossing walls, uncertain room classification, and incomplete structural recognition are reported in several approaches [9]–[11], [17], [19]. These limitations can affect the practical quality of generated or reconstructed layouts and indicate the need for stronger constraint handling and more comprehensive evaluation.

Second, the usability of generative systems is affected by the nature of their input requirements. Some approaches rely on complete graphs, specified attributes, or predetermined structural parameters [5], [6], [10], while natural-language methods still face difficulty extracting accurate spatial relations [18]. Greater tolerance for partial, ambiguous, or naturally expressed requirements remains an open research issue.

Third, cost-analysis studies frequently rely on geographically concentrated or construction-type-specific datasets [4], [15], [20], [21]. Reported results may therefore not generalize to other regions, material markets, or project categories. The inconsistent findings on ensemble stacking [14], [20] further show that model selection should be supported by empirical comparison on representative data rather than assumptions about model complexity.

Fourth, multi-criteria methods show potential for structured evaluation, but much of the reviewed evidence comes from supplier-selection and inventory-planning contexts [2], [3], [16]. Additional testing in broader built-environment decision settings, together with transparent weight selection and sensitivity analysis, is needed to understand their practical robustness.

V. FUTURE RESEARCH DIRECTIONS

Future work can address these challenges through several complementary research directions. For layout generation, stronger geometric constraints, improved recognition of spatial relationships, and evaluation on more diverse floor-plan datasets may improve reliability. Generative systems should also be tested with varied forms of input, including partial descriptions and user-friendly specifications, to better understand their behaviour outside controlled benchmarks. For construction cost analysis, shared benchmark datasets representing multiple regions, construction types, and market conditions would support more rigorous testing of generalizability. Comparative studies should evaluate individual and ensemble models under consistent conditions and

report uncertainty alongside point estimates. In decision-support research, further applications should examine how weighting methods behave under changing criteria and uncertain inputs, with sensitivity analysis used to assess the stability of the resulting rankings. Across all three areas, transparent evaluation protocols, reproducible experiments, and validation with real-world practitioners will be important for translating promising techniques into reliable tools.

VI. CONCLUSION

This review examined twenty-two studies on computational methods relevant to residential building planning and related built-environment applications. The literature covers graph-based and generative layout methods, decision-support approaches for handling multiple criteria, and machine-learning techniques for construction cost analysis. The reviewed studies report meaningful advances, including more flexible layout generation, improved methods for comparing alternatives, and stronger predictive performance from selected machine-learning models. At the same time, their limitations reveal recurring concerns about geometric validity, input requirements, dataset representativeness, and external validation. Addressing these issues through broader datasets, transparent evaluation, and user-centred testing is likely to strengthen the practical reliability of computational approaches in this field.

REFERENCES

[1] R. Yu, N. Gu, G. Lee, and A. Khan, "A systematic review of architectural design collaboration in immersive virtual environments," Buildings, vol. 12, no. 3, 2022.

[2] "Least to most range (LMR): A novel consistent weighting multicriteria method applied to green supplier selection," IEEE Access, 2023.

[3] "A comprehensive framework for supplier selection using subjective, objective, and hybrid multi-criteria decision-making techniques with sensitivity analysis," 2023.

[4] B. Tajziyehchi, "Construction cost growth prediction: A feature selection and machine learning approach," M.S. thesis, Univ. of Alberta, 2020.

[5] G. K. Sepúlveda et al., "Semi-automatic building layout generation for virtual environments," IEEE Access, vol. 12, 2024.

[6] H. Zhang et al., "MaskPLAN: Masked generative layout planning from partial input," in Proc. IEEE/CVF Conf. Computer Vision and Pattern Recognition (CVPR), 2024.

[7] N. Nauata, K.-H. Chang, C.-Y. Cheng, G. Mori, and Y. Furukawa, "House-GAN++: Generative adversarial layout refinement network towards intelligent computational agent for professional architects," in Proc. IEEE/CVF Conf. Computer Vision and Pattern Recognition (CVPR), 2021.

[8] Y. Du et al., "GA-optimized BP neural network model for architectural engineering cost estimation," 2017.

[9] "Automated semantics and topology representation of residential-building space using floor-plan raster maps," IEEE journal article.

[10] Y. Takeuchi, Q. An, and A. Yamashita, "Single-stage graph-based floorplan generation," IEEE Access, vol. 14, 2026.

[11] "ActFloor-GAN: Activity-guided adversarial networks for human-centric floorplan design," IEEE journal article.

[12] "LayoutGMN: Neural graph matching for structural layout similarity," conference paper.

[13] C. Xu et al., "GoodFloorplan: A reinforcement learning framework for chip floorplanning," 2022.

[14] "Construction cost estimation using hybrid stacking of Random Forest and XGBoost," in Proc. INOCON, 2024.

[15] "BP neural network construction cost prediction model with Adam/SGD/Adadelta optimization," conference/book chapter.

[16] "Effective replenishment planning and inventory control under multiple suppliers and stochastic lead times," in Proc. CODIT, 2024.

[17] L. Simonsen et al., "Graph neural network-based door detection from CAD floor plans," 2021.

[18] Y. Wei and X. Li, "Graph-augmented text-based floorplan generation," in Proc. Int. Conf. Automation in Manufacturing, Transportation and Logistics (ICaMaL), 2024.

[19] Z. Zhang, J. Liu, X. Li, C. Hu, F. Dunkin, and S. Zhang, "A novel framework for structure descriptors-guided hand-drawn floor plan reconstruction," in Proc. IEEE/RSJ Int. Conf. Intelligent Robots and Systems (IROS), 2024.

[20] Vincy R and E. G. M. Kanaga, "AI-based construction cost estimation using adaptive machine learning," in Proc. Int. Conf. Inventive Computation Technologies (ICICT), 2026.

[21] Q. Xu, J. Liu, Y. Tang, and Z. Zhang, "Research on construction project cost prediction based on BP neural network," in Proc. Int. Conf. Software, Systems and Information Technology (SSITCON), 2025.

[22] Q. Jin, "Research on construction cost control system of project management based on computer BIM technology," in Proc. Int. Conf. Electrical Drives, Power Electronics & Engineering (EDPEE), 2024.

Note: bibliographic details above are transcribed from the source manuscripts as supplied; authors, venues, and years should be cross-verified against the original publisher records prior to final submission.