AI-Assisted Urban Density Planning Through Integrated Environmental and Structural Simulation

Kirstie M. Lacey
Department of Computer Science, University of Houston, Houston, TX, USA.
kirstie1987@uh.edu

Tirza van der Laarse
Department of Computer Science, Colorado State University, Fort Collins, CO, USA.
tirza.vanderlaarse57@colostate.edu

Nakul Bawa
School of Information Technology, University of Cincinnati, Cincinnati, OH, USA.
nakulb@uc.edu

Abstract

Urban density planning must reconcile land-use intensity with solar access, daylight, outdoor thermal conditions, and structural feasibility. Existing workflows often evaluate these dimensions sequentially or through simplified indicators, which obscures trade-offs and makes early-stage design search opaque. This manuscript develops a methodological framework for AI-assisted density planning that integrates environmental and structural simulation through spatially explicit decision variables, surrogate learning, constrained multi-objective optimization, and uncertainty-aware decision rules. The framework represents an urban block as a vector of height, footprint, spacing, and program variables; couples environmental and structural simulators through a common geometry and scenario interface; and replaces repeated high-fidelity calls with Gaussian-process surrogates trained through active learning. A constrained Pareto formulation links environmental penalties, structural demand and risk, construction cost, and solar availability. Robustness is introduced through expectation and conditional value-at-risk terms, while a transparent utility rule supports planner selection among nondominated plans. The contribution is not a trained model or a reporting of measured performance, but a coherent formulation that specifies inputs, simulation assumptions, learning architecture, constraints, and failure modes. The framework clarifies how AI can accelerate design-space exploration without displacing engineering checks or environmental evidence. It provides a basis for reproducible density studies, comparable surrogate diagnostics, and defensible early-stage planning decisions.

Keywords

urban density; surrogate modeling; multi-objective optimization; building performance simulation; structural safety; solar access

1. Introduction

Urban density is not a single design choice but a bundle of coupled decisions about height, footprint, spacing, land use, and public-space allocation. These decisions affect travel behavior, service viability, solar exposure, daylight, ventilation potential, outdoor thermal comfort, and structural demand. Empirical synthesis of the built environment and travel literature shows that density, diversity, and design variables are associated with travel

outcomes, although effect sizes vary with context and measurement [1]. A city-systems perspective further treats urban form as an emergent property of interacting decisions rather than a static container for activity [2]. Density planning therefore requires more than a scalar floor-area ratio; it requires a representation that preserves the spatial relationships through which environmental and structural performance are produced.

Environmental performance is especially sensitive to urban geometry. Solar and daylight availability depend on street width, orientation, building height, and mutual shading, and simplified indicators can misrepresent the distribution of accessible radiation across façades and public spaces [3]. Urban morphology also conditions radiation availability in ways that cannot be reduced to average density without losing the spatial variation that matters for passive design and outdoor comfort [4]. Structural performance introduces another layer of coupling: building height and footprint influence lateral-force demand, foundation loads, material use, and constructability, while spacing and adjacency affect shared structural systems and site constraints. Building information modeling provides a common digital representation for geometry, materials, and performance attributes [5], and simulation-based optimization has been used to search building design alternatives under performance objectives [6]. Yet early-stage density planning still lacks a coherent formulation that connects fine-grained environmental simulation, structural checks, AI-assisted surrogate search, and uncertainty-aware selection.

This article addresses that gap by developing an integrated methodological framework for AI-assisted urban density planning. The central contribution is a surrogate-assisted, constrained multi-objective formulation in which urban density variables are evaluated jointly through environmental and structural simulation, learned response surfaces, and robust decision rules. The framework does not assume that a single density optimum exists. Instead, it organizes the search for nondominated plans and makes explicit the assumptions under which surrogate predictions can support planning judgments. The argument proceeds from problem formulation to simulation coupling, surrogate learning, multi-objective optimization, uncertainty treatment, and technical limitations.

2. Problem Formulation and Knowledge Gap

The planning problem can be represented through a spatially explicit decision vector for a set of urban blocks. Let each block i have a height variable, a footprint or coverage variable, a spacing or setback variable, and a program-mix variable. The decision vector collects these quantities across N blocks, with feasibility constraints imposed by zoning, land ownership, daylight protection, fire access, and infrastructure capacity. The representation must preserve enough geometric detail for simulation while remaining tractable for optimization.

$$x = \{h_i, \lambda_i, g_i, \mu_i\}_\{i = 1\}\^\{N\}$$

Here, x is the complete density plan, h_i is building height, λ_i is footprint coverage, g_i is spacing or setback, and μ_i is the program mix assigned to block i. The vector is not merely a list of parcel attributes; it defines a three-dimensional urban geometry from which environmental and structural indicators are computed. The choice of these variables is deliberate. Height and coverage control floor-area intensity, spacing controls shading and access, and program mix shapes occupancy schedules and internal loads. A scalar density target cannot represent these distinctions, and a purely geometric metric cannot indicate whether a plan is environmentally or structurally feasible.

The integrated evaluation maps the decision vector into environmental and structural response vectors. Environmental responses may include annual solar irradiation on roofs and façades, daylight autonomy proxies, sky-view factor, outdoor mean radiant temperature indicators, and natural ventilation potential. Structural responses may include lateral drift, gravity-load utilization, foundation demand, material intensity, and a qualitative robustness indicator. The simulators are treated as distinct but coupled mappings because they consume a common geometry and scenario definition while producing different classes of evidence.

$$e = \mathrm{E}(x; \theta_E), q = Q(x; \theta_Q)$$

In this expression, e is the environmental response vector, q is the structural response vector, E and Q are the respective simulation mappings, and θ_E and θ_Q collect scenario assumptions such as weather data, material properties, occupancy schedules, soil conditions, and code-based load factors. The coupling is not numerical in the sense of a monolithic solver; it is representational. A shared geometry and scenario layer ensures that the same plan is evaluated under consistent orientation, terrain, and boundary conditions. This matters because environmental and structural conclusions can diverge if each analysis uses a different abstraction of the same urban block.

The knowledge gap is that most density-planning tools either optimize simplified indicators without simulation or run detailed simulations only for a small set of manually selected alternatives. Simplified indicators are useful for screening but can miss nonlinear shading and structural demand effects. Detailed simulation can capture those effects but is too costly for broad design-space exploration when every candidate requires multiple environmental and structural runs. AI-assisted planning is most defensible when it is framed as a surrogate-accelerated search over simulation-consistent candidates, with explicit constraints and uncertainty, rather than as a replacement for engineering or environmental analysis.

3. Methodological Framework

The proposed framework has four interacting components: a parametric urban geometry generator, a simulation-coupling layer, a surrogate learning layer, and a constrained multi-objective optimizer with decision support. The geometry generator converts the decision vector into building massing, envelope surfaces, street canyons, and parcel-level attributes. The coupling layer assigns consistent weather, orientation, material, occupancy, and load assumptions to both environmental and structural analyses. The learning layer approximates selected simulator outputs to reduce computational cost during search. The optimizer searches for nondominated plans and identifies plans that satisfy hard constraints.

Surrogate modeling is appropriate because simulation calls are expensive and the design space is nonconvex. Gaussian-process regression is a useful default when response surfaces are smooth but highly nonlinear and when uncertainty estimates are needed for active learning [7]. Bayesian optimization provides a principled way to select simulation points that balance predicted performance against uncertainty [8]. In this framework, surrogate learning is not presented as a validated replacement for simulation. It is a screening and search device whose predictions are periodically checked against high-fidelity runs, with discrepancies feeding back into the training set and model structure.

$$\hat{y_m}(x) = \mu_m + r(x)^T R^{-1}(y_m - \mu_m 1)$$

Here, $\hat{y}_m(x)$ is the predicted value of simulator output m at plan x, μ_m is the estimated mean response, r(x) is the vector of correlations between x and the observed design points, R is the correlation matrix among observed points, y_m is the vector of observed simulator outputs, and 1 is a vector of ones. The equation describes a standard Gaussian-process predictor. Its value in this framework is not only a point estimate; the conditional variance provides a local measure of surrogate uncertainty that can guide additional simulation runs. A plan with promising predicted performance but high predictive variance is a candidate for direct evaluation rather than immediate acceptance.

The learning layer can be organized as an active-learning loop. An initial design of experiments is generated to cover the feasible region, including boundary cases such as maximum height under solar-access constraints and minimum spacing under structural or fire-access constraints. High-fidelity simulations populate the first training set. The surrogate is then fitted jointly or separately for environmental and structural outputs. An acquisition rule selects candidates that are likely to improve the Pareto front, violate constraints, or reduce uncertainty in regions where decisions are sensitive. This loop continues until a computational budget or a stability criterion is reached. The framework does not specify a fixed budget because it depends on simulator cost and planning scale; instead, it requires that the stopping rule be documented and that surrogate error be reported by output type and design region.

The simulation-coupling layer must resolve several representational issues. Geometry must be watertight enough for radiation and structural analysis. Envelope surfaces must carry material and glazing properties. Street canyons and open spaces must be represented as part of the environmental domain, not as empty gaps between buildings. Structural models may use simplified frame or equivalent lateral-force representations at early stages, but the level of simplification must be consistent with the decision being made. If the planning question concerns massing and density rather than detailed member sizing, a structural feasibility surrogate can be based on code-level demand ratios and drift indicators rather than full nonlinear analysis. This choice reduces computational burden while preserving the constraints that matter for early-stage planning.

4. Simulation Coupling and Surrogate Learning

The integrated search is formulated as a constrained multi-objective problem. Environmental penalties, structural demand or risk, construction cost, and solar availability are treated as separate objectives because they are not reducible to a single monetary metric without contestable assumptions. Multi-objective evolutionary algorithms are well suited to nonconvex problems with mixed continuous and discrete variables, and elitist nondominated sorting has become a standard approach for approximating Pareto fronts [9]. Performance assessment of such algorithms must consider convergence, diversity, and coverage rather than a single best value [10]. In density planning, diversity is not a cosmetic property; it gives planners alternative plans with different trade-off profiles.

Sensitivity analysis is necessary because the ranking of plans can change when weather, material, occupancy, or load assumptions vary. Global sensitivity analysis can identify which input assumptions dominate output variance and which interactions matter [11]. Bayesian calibration provides a formal route for updating simulator assumptions against observed data when such data exist, while preserving a distinction between calibration parameters and model discrepancy [12]. In the absence of site-specific observations, the same formalism can be used to propagate scenario uncertainty and to avoid presenting a single deterministic ranking as definitive.

Spatial data infrastructure affects both model inputs and decision credibility. Parcel boundaries, building footprints, street networks, terrain, and land-use attributes are often assembled from multiple sources with different accuracy and update cycles. Volunteered geographic information and other non-authoritative sources can fill gaps, but their provenance and positional uncertainty must be recorded [13]. Urban building energy modeling illustrates the difficulty of scaling detailed simulation to districts: archetypes, geometry simplification, and occupancy assumptions strongly condition results [14]. The framework therefore treats input provenance as part of the model, not as metadata added after optimization.

The objective vector combines environmental and structural criteria with cost and solar access. A generic formulation is:

$$\min_x F(x) = \left(C_e(x), C_q(x), C_c(x), -A_s(x)\right)$$

Here, C_e(x) is an environmental penalty that aggregates undesirable deviations such as excessive shading of public space, poor daylight potential, or heat-related exposure; C_q(x) is a structural demand or risk penalty derived from utilization, drift, and foundation indicators; C_c(x) is a construction-cost proxy that may include material intensity and site work; and A_s(x) is an aggregate solar-access benefit. The negative sign converts solar availability into a minimization objective. This formulation is deliberately modular: a planning authority may replace the cost proxy with a different fiscal indicator or add a public-space quality term, but the underlying trade-off structure remains explicit.

Hard constraints are separated from objectives because not every desirable property should be traded against cost or solar gain. Structural safety, fire access, minimum daylight protection, and zoning limits are constraints rather than soft preferences. A general constraint set can be written as:

$$g_j(x) \le 0, h_k(x) = 0, x \in X$$

In this expression, g_j are inequality constraints such as maximum drift or minimum spacing, h_k are equality constraints such as fixed public-space area in a redevelopment parcel, and X is the bounded feasible design space. The distinction between constraints and objectives is a design decision with ethical and legal consequences. A plan that violates a structural safety constraint should not enter the Pareto set merely because it performs well on solar access or cost.

Uncertainty enters through scenario variables θ, which may include weather years, material properties, occupancy schedules, soil parameters, and load factors. A robust formulation minimizes a combination of expected performance and a risk measure:

Here, E_θ denotes expectation over scenario variables, CVaR_α is the conditional value at risk at confidence level α, and κ is a nonnegative weight that controls risk aversion. The conditional value at risk captures the average loss in the worst tail of the performance distribution, which is relevant when structural demand or heat exposure has asymmetric consequences. This expression is a proposed decision formulation, not a calibrated model. Its use requires specifying the scenario distribution and the risk weight through a transparent planning process rather than hiding those choices inside an optimizer.

Decision support must also accommodate multiple actors with different preferences. A single Pareto front can be large, and planners may need to select among plans that differ in cost, solar access, structural intensity, and spatial equity. A utility-based selection rule can be applied after nondominated plans are generated:

$$x* = argmin_\{x \in X\}\, U(F(x);\, w), w \in W$$

Here, U is a transparent utility or scoring function, w is a vector of preference weights, and W is the admissible weight set. The rule does not replace the Pareto front; it makes the final selection explicit and revisable. If weights are uncertain, the framework can report the stability of x across W rather than presenting a single winner. This is important because density decisions often involve distributional effects that are not captured by aggregate objectives alone.

A recent generative optimization study of building blocks for density, solar, and structural performance demonstrates the relevance of integrated block-level search and multi-criteria evaluation [17]. The framework proposed here differs in emphasis: it treats simulation consistency, surrogate uncertainty, hard constraints, and scenario robustness as first-class elements of the planning problem. Generative search can produce many plausible massing configurations, but planning credibility depends on how those configurations are evaluated and how surrogate error is controlled. Decision-support tools for early-stage zero-energy building design show that performance feedback is most useful when it is embedded in the design workflow rather than delivered as a post hoc score [18]. The integrated density framework follows that logic at the block and district scale.

5. Multi-Objective Density Optimization and Decision Logic

The optimizer operates on a population of candidate plans. Each candidate is decoded into geometry, checked against hard constraints, and evaluated by surrogates. Promising or uncertain candidates are sent to high-fidelity simulation. The resulting observations update the surrogate models. Nondominated sorting and crowding-distance mechanisms can be used to maintain a diverse approximation of the Pareto front [9]. Performance indicators should be reported alongside the front, including hypervolume or coverage where appropriate, so that the search is not judged solely by the best single plan [10]. Because surrogate error varies across the design space, the optimizer should maintain separate error estimates for environmental and structural outputs rather than a single aggregate accuracy figure.

The objective formulation has several technical consequences. First, solar availability and environmental penalty may conflict: dense massing can improve solar access on some roofs while shading streets and lower façades. Second, structural demand scales nonlinearly with height and slenderness, so a small increase in density can produce a disproportionate increase in material intensity. Third, cost proxies can reward configurations that are environmentally poor if externalities are omitted. The multi-objective formulation keeps these tensions visible. It does not assume that a compensatory monetary value can resolve them. Instead, it produces a set of plans whose trade-offs can be inspected, debated, and constrained by policy.

The surrogate ensemble must be diagnosed by output and region. A global correlation coefficient can hide poor performance in the high-density region where structural nonlinearity is strongest. Error metrics should therefore be reported separately for interpolation and extrapolation regimes, for constrained and unconstrained candidates, and for each objective component. Where surrogate error exceeds a threshold in a decision-relevant region, direct simulation should be mandatory. This is especially important for structural outputs because a surrogate that is adequate for ranking may be inadequate for safety verification. The framework treats surrogate predictions as search guidance, not as a substitute for code-compliant structural analysis.

The decision logic should also expose scenario dependence. If the ranking of plans changes across weather years or load assumptions, the final recommendation should identify that instability. Robust optimization can reduce sensitivity, but it can also produce conservative plans that underuse density capacity. The risk weight κ in the robust objective is therefore a policy variable as much as a mathematical one. A low κ favors expected performance; a high κ favors protection against adverse scenarios. Documenting this choice is more honest than presenting a single robust optimum as if it were value-neutral.

6. Technical Discussion

The main strength of the framework is that it aligns AI-assisted search with the evidentiary structure of urban planning. Machine learning is used where it is most defensible: to approximate expensive simulation mappings, to guide sampling, and to explore a constrained design space. It is not used to certify structural safety or to replace environmental measurement. This division of labor matters because density decisions have long time horizons and distributional consequences. A plan that performs well under a surrogate but poorly under detailed simulation can create costly redesign or, worse, unsafe or unlivable conditions.

Several failure modes require attention. Surrogate extrapolation is the first. If the optimizer proposes plans outside the convex hull of the training data, Gaussian-process variance may rise, but the acquisition function may still favor such plans if predicted performance is attractive. A trust-region or penalty on extrapolation distance can mitigate this risk. The second failure mode is geometry simplification. Environmental simulations are sensitive to street-canyon representation and surface orientation, while structural models are sensitive to connectivity and load paths. A shared geometry layer reduces inconsistency, but it does not eliminate the need for domain-specific model checks. The third failure mode is objective aggregation. If environmental penalties combine dissimilar effects such as daylight loss and heat exposure without clear weights, the resulting scalar can be misleading. The Pareto formulation reduces this risk by keeping objectives separate until decision support.

Co-simulation and model interoperability are practical requirements for integrated analysis. Building energy and control systems are often modeled in different tools, and co-simulation standards can support coupled execution without forcing a single monolithic environment [20]. Simulation engines differ in their capabilities, numerical assumptions, and output conventions [21]. The framework therefore requires a mapping layer that records which outputs come from which simulator, under which assumptions, and with what known limitations. This is not bureaucratic overhead; it is what makes cross-project comparison possible.

Urban climate and outdoor comfort add another scale of analysis. Street design and urban canopy geometry influence airflow, radiation exchange, and temperature within the urban canopy layer [22]. Density plans that improve solar access on roofs may reduce street-level ventilation or increase radiant loading. The integrated framework should therefore include canopy-level indicators when the planning question concerns public space or heat exposure, and should state clearly when the analysis is limited to building-scale performance. The level of model fidelity should follow the decision: a massing study may use simplified canopy indicators, while an environmental impact assessment may require coupled microclimate simulation.

Institutional use of the framework depends on transparent assumptions and reproducible reporting. Scenario definitions, constraint sets, surrogate training data, error diagnostics, and Pareto-front summaries should be documented in a form that allows independent review. The framework does not prescribe a single software stack, but it does require that the optimizer not conceal constraint violations or surrogate uncertainty. This is consistent with the broader lesson from building performance simulation: model credibility depends on fitness for purpose, not on computational sophistication alone [19]. AI assistance is valuable when it makes the design space and trade-offs more legible, not when it turns a complex decision into an opaque score.

7. Conclusion

This article has developed a methodological framework for AI-assisted urban density planning through integrated environmental and structural simulation. The framework represents density as a spatially explicit decision vector that defines height, coverage, spacing, and program mix across urban blocks. It couples environmental and structural simulators through a shared geometry and scenario layer, uses Gaussian-process surrogates and active learning to reduce simulation burden, and formulates density search as a constrained multi-objective problem with explicit solar, environmental, structural, cost, and public-space criteria. Robustness is addressed through scenario expectations and conditional value-at-risk terms, while a utility-based selection rule supports transparent choice among nondominated plans. The required generative block optimization reference is positioned within this integrated search logic, and the broader literature on decision support, co-simulation, simulation capabilities, and urban canopy climate informs the framework's boundary conditions. The contribution is a coherent formulation rather than an empirical claim of performance. Its value lies in making assumptions, constraints, surrogate uncertainty, and trade-offs explicit so that AI-assisted density planning can be scrutinized as a technical and planning practice rather than accepted as an automated output. Future applications should report simulation provenance, surrogate diagnostics by region, constraint satisfaction, and the sensitivity of selected plans to scenario and preference weights. Such reporting would allow the framework to be tested, compared, and refined across urban contexts without overstating what any single model can establish.

References

1. Ewing, R., & Cervero, R. (2010). Travel and the built environment: A meta-analysis. Journal of the American Planning Association, 76(3), 265–294.
2. Batty, M. (2013). The new science of cities. MIT Press.
3. Compagnon, R. (2004). Solar and daylight availability in the urban fabric. Energy and Buildings, 36(4), 321–328.
4. Robinson, D. (2006). Urban morphology and indicators of radiation availability. Solar Energy, 80(12), 1643–1648.
5. Sacks, R., Eastman, C., Lee, G., & Teicholz, P. (2018). BIM handbook: A guide to building information modeling for owners, designers, engineers, contractors, and facility managers (3rd ed.). Wiley.
6. Nguyen, A. T., Reiter, S., & Rigo, P. (2014). A review on simulation-based optimization methods applied to building performance analysis. Applied Energy, 113, 1043–1058.

7. Forrester, A. I. J., Sóbester, A., & Keane, A. J. (2008). Engineering design via surrogate modelling: A practical guide. Wiley.

8. Shahriari, B., Swersky, K., Wang, Z., Adams, R. P., & de Freitas, N. (2016). Taking the human out of the loop: A review of Bayesian optimization. Proceedings of the IEEE, 104(1), 148–175.

9. Deb, K., Pratap, A., Agarwal, S., & Meyarivan, T. (2002). A fast and elitist multiobjective genetic algorithm: NSGA-II. IEEE Transactions on Evolutionary Computation, 6(2), 182–197.

10. Zitzler, E., Thiele, L., Laumanns, M., Fonseca, C. M., & da Fonseca, V. G. (2003). Performance assessment of multiobjective optimizers: An analysis of issues and review. IEEE Transactions on Evolutionary Computation, 7(2), 117–132.

11. Saltelli, A., Ratto, M., Andres, T., Campolongo, F., Cariboni, J., Gatelli, D., Saisana, M., & Tarantola, S. (2008). Global sensitivity analysis: The primer. Wiley.

12. Kennedy, M. C., & O'Hagan, A. (2001). Bayesian calibration of computer models. Journal of the Royal Statistical Society: Series B (Statistical Methodology), 63(3), 425–464.

13. Goodchild, M. F. (2007). Citizens as sensors: The world of volunteered geography. GeoJournal, 69(4), 211–221.

14. Reinhart, C. F., & Cerezo Davila, C. (2016). Urban building energy modeling – A review of a nascent field. Building and Environment, 97, 196–202.

15. Kusiak, A., Tang, F., & Xu, G. (2011). Multi-objective optimization of HVAC system with an evolutionary computation algorithm. Energy, 36(5), 2440–2449.

16. Hong, T., Chou, S. K., & Bong, T. Y. (2000). Building simulation: An overview of developments and information sources. Building and Environment, 35(4), 347–361.

17. Tabasi, S. F., Luo, D., Rafizadeh, H. R., Le Nguyen, K., & Banihashemi, S. (2025). Generative optimization of building blocks for density, solar and structural performance. Journal of Building Engineering, 111, 113307.

18. Attia, S., Gratia, E., De Herde, A., & Hensen, J. L. M. (2012). Simulation-based decision support tool for early stages of zero-energy building design. Energy and Buildings, 49, 2–15.

19. Hensen, J. L. M., & Lamberts, R. (Eds.). (2011). Building performance simulation for design and operation. Spon Press.

20. Wetter, M. (2011). Co-simulation of building energy and control systems with the Functional Mock-up Interface standard. Journal of Building Performance Simulation, 4(2), 103–116.

21. Crawley, D. B., Hand, J. W., Kummert, M., & Griffith, B. T. (2008). Contrasting the capabilities of building energy performance simulation programs. Building and Environment, 43(4), 661–673.

22. Oke, T. R. (1988). Street design and urban canopy layer climate. Energy and Buildings, 11(1–3), 103–113.