AI-Assisted Design of Robotic Healthcare Facility Construction Systems for Improving Operational Efficiency and Infection Control

Michal Kleczar
Department of Computer Science, George Mason University, Fairfax, VA, USA.
michal2002@gmu.edu

Maria Pilar Canet
Department of Computer Science and Engineering, University at Buffalo, Buffalo, NY, USA.
contactmariapilar@buffalo.edu

Maks Gabryel
School of Electrical Engineering and Computer Science, Oregon State University, Corvallis, OR, USA.
maksgabryel@oregonstate.edu

Abstract

Healthcare facility construction and renovation create coupled operational and infection-control problems: airborne and surface contamination must be contained while schedules, material flows, and robotic tasks are coordinated in constrained clinical environments. Existing construction-robotics research has advanced task automation and on-site autonomy, and healthcare design research has established environmental controls for infection prevention, but these bodies of knowledge are rarely combined in a single design model. This manuscript develops an AI-assisted co-design framework for robotic healthcare facility construction systems that treats infection control as a dynamic constraint on robotic planning rather than as a post hoc environmental add-on. The framework represents the construction state through a building-information-model graph, robot fleet states, environmental fields, contamination estimates, and human activity; it couples task allocation and sequencing with a physics-informed airflow and infection-risk surrogate, then solves a multi-objective problem over completion time, infection risk, energy use, and disruption under uncertainty. The central methodological contribution is the formulation of robot-oriented construction design as a constrained, partially observable decision problem in which ventilation zoning, material logistics, and robot paths are optimized jointly. Analytical assessment shows why efficiency gains from robot density and parallel task execution can be offset by contaminant resuspension, pressure imbalance, and schedule fragility when infection-control constraints are ignored. The framework provides a defensible basis for designing robotic construction systems that are both operationally efficient and compatible with healthcare infection-prevention requirements, while identifying calibration, validation, and human-oversight conditions necessary for practical use.

Keywords

healthcare construction, construction robotics, infection control, physics-informed machine learning, digital twin, multi-objective optimization

1. Introduction

Healthcare-associated infections and construction-related environmental contamination remain persistent risks in hospital environments, particularly when new construction or renovation occurs near occupied clinical areas [1], [2]. Environmental infection-control guidance emphasizes containment, pressure relationships, ventilation, filtration, and surface cleanability, but these controls are usually specified after the construction method and logistics have been chosen [2]. Evidence-based healthcare design further links layout, airflow, visibility, and material flows to safety and operational performance, yet the design process often separates clinical planning from construction execution [3]. This separation is consequential for robotic construction because robots alter not only labor productivity but also dust generation, movement patterns, temporal occupancy, and the spatial distribution of work.

Airborne transmission and contaminant dispersion depend on ventilation, airflow direction, source strength, and exposure duration [4]. Studies of hospital wards show that ventilation configuration can produce materially different contaminant distributions even when nominal air-change rates are similar [5]. Standards for healthcare ventilation specify pressure cascades, outdoor-air rates, filtration, and zoning requirements that constrain where construction activity can occur and how temporary barriers must be maintained [6]. A robotic construction system that ignores these constraints may improve schedule performance while creating unacceptable contamination pathways or disrupting clinical pressure relationships.

Construction robotics has developed from single-task machines toward mobile, sensor-rich systems capable of on-site autonomy [7], [8]. Research on on-site autonomous construction emphasizes unstructured environments, localization, task uncertainty, and safety [9], while collective robotic construction explores decentralized coordination and scalability [10]. These advances provide the technological basis for robotic healthcare construction, but they do not specify how robot design and task planning should respond to infection-control requirements. The transition from controlled factory settings to construction sites changes assumptions about geometry, tolerances, temporary services, and environmental variability [14]. In healthcare settings, those general site uncertainties are compounded by clinical occupancy, negative-pressure barriers, and the need for verifiable cleanliness.

Building information modeling and digital twins offer a shared representation for geometry, equipment, schedules, and sensor data [11], [12]. Physics-informed machine learning can combine mechanistic airflow and contaminant-transport relationships with data-driven approximations, which is important when high-fidelity simulation is too slow for repeated planning queries [13]. The unresolved problem is how to embed these capabilities in a design framework that jointly selects robotic construction methods, task sequences, and environmental controls. Treating infection control as a constraint on robotic planning, rather than as an independent engineering package, changes the optimization landscape and exposes trade-offs that are invisible in separate design workflows.

This manuscript makes one coherent contribution: it formulates AI-assisted design of robotic healthcare facility construction systems as a constrained, partially observable decision problem that couples robot task allocation, material logistics, airflow zoning, and infection-risk estimation. The framework is methodological and analytical; it does not report a field deployment or a calibrated empirical evaluation. Instead, it specifies the state representation, scheduling relationships, contamination surrogate, risk model, multi-objective objective, and robust decision rule needed to reason about operational efficiency and infection control together. The analysis identifies conditions under which robot density, parallel execution, and

aggressive task sequencing improve performance, and conditions under which they increase contamination risk or schedule fragility. It also clarifies the calibration, validation, and human-oversight requirements that any practical implementation would need to satisfy.

2. Design Requirements and Problem Formulation

A robotic healthcare construction system must satisfy three interacting classes of requirements. The first is clinical infection control: temporary barriers must maintain pressure differentials, exhaust air must be filtered or directed away from occupied zones, surfaces and tools must be cleanable, and material movement must not create uncontrolled openings [2], [6]. The second is construction performance: tasks have precedence, resource constraints, robot reachability, material delivery limits, and inspection points. The third is operational continuity: construction activity may need to be scheduled around clinical shifts, sterile procedures, or patient transport routes. These requirements are not independent. A shorter robotic task sequence may concentrate dust generation in a critical period, while a conservative containment plan may idle robots and increase project duration.

The design problem is to choose the robot fleet, task assignment, task timing, path plans, and environmental-control actions that minimize a vector of operational and infection-control costs subject to safety and feasibility constraints. The decision variables include binary assignments of tasks to robots, start times, path and speed profiles, material delivery schedules, and control settings for temporary ventilation and barriers. The constraints include task precedence, robot capability and reachability, collision avoidance, pressure-zone limits, maximum allowable contaminant concentration, and inspection release rules. A useful framework must represent these variables explicitly because informal or purely rule-based coordination cannot reveal when an efficiency improvement trades against infection risk.

The problem is partially observable in practice. Robot positions and task states can be measured, but contaminant concentrations, airflow patterns, and barrier integrity are known only through sparse sensors and models. Clinical activity also changes unpredictably. The design framework therefore needs a state representation that separates known building geometry from estimated environmental and contamination fields. It also needs a decision rule that remains safe under model error. These requirements motivate the state-based formulation and robust optimization approach developed in the next section.

3. AI-Assisted Co-Design Framework

3.1 Digital construction state and robot-oriented design

The framework represents the construction state at planning time t as a tuple containing the building information graph, robot fleet state, environmental field, contamination estimate, and human activity state. The state is written as

$$s_t = (G, R, E_t, C_t, H_t)$$

where G is the building information graph with spaces, barriers, service zones, and precedence relations; R contains robot poses, capabilities, payload states, and availability; E_t includes pressure, airflow, temperature, and particulate measurements or estimates; C_t is the estimated contaminant field; and H_t encodes clinical and construction personnel activity. This representation is deliberately heterogeneous because the design problem couples discrete task logic with continuous environmental fields. A purely geometric model cannot represent pressure cascades or contamination, while a purely environmental model cannot represent robot reachability or task precedence.

The transition from factory automation to on-site construction requires robot design assumptions to be revisited: site geometry is less controlled, temporary services vary, and tolerances must be negotiated with existing structures [14]. In healthcare construction, robot-oriented design also has to account for barrier penetrations, cleanability of robot surfaces, and the need to move materials without holding doors open or disrupting pressure zones. The state representation supports these concerns by making barriers, pressure zones, and human activity first-class variables rather than background context.

3.2 Task allocation, sequencing, and learning

Given a set of construction tasks I and robots R, the schedule must respect precedence relationships. For every precedence arc from task j to task i, the start of i cannot precede the finish of j plus any required separation time. This is expressed as

$$T_i_start \geq T_j_finish + \tau_ji \; for \; all \; (j,i) \in P$$

where T_i_start and T_j_finish are task start and finish times, τ_ji is the minimum lag required by inspection, curing, cleaning, or contamination decay, and P is the precedence set. The completion time of the construction plan is the latest task finish:

$$T_makespan = max_i \, (T_i_start + \sum_r \, d_ir \, x_ir)$$

where d_ir is the duration of task i when performed by robot r and x_ir is a binary assignment variable equal to one when task i is assigned to robot r. The duration is not a fixed constant in practice; it depends on travel distance, material availability, obstacle density, and task complexity. Deep learning can approximate these duration and feasibility relationships from geometry, robot state, and prior task records, but the resulting model must be treated as a planning prior rather than as a guarantee [15]. Reinforcement learning is appropriate for adaptive sequencing when the state space is large and task outcomes are stochastic, provided that safety constraints are enforced by a separate shielding or feasibility layer [16].

A task-level objective can combine assignment cost, makespan, and due-date penalty:

$$J_task = \sum_i \sum_r \, c_ir \, x_ir + \lambda_T \, T_makespan + \lambda_D \sum_i \, max(0, T_i_start - T_i_due)$$

where c_ir is the direct cost of assigning task i to robot r, λ_T and λ_D are weighting coefficients, and T_i_due is the target completion time for task i. This expression is not presented as an empirically fitted objective; it defines the trade-off structure that the design framework must expose. The coefficients are policy choices that should be set with construction managers and infection-prevention staff because they determine whether schedule speed or disruption avoidance dominates.

3.3 Airflow and infection-risk surrogate

Infection control requires a model of contaminant transport. For a conservative contaminant with concentration C, airflow velocity v, turbulent diffusivity D, inactivation rate κ, and source term S, the transport relationship can be written as

$$\partial C/\partial t + v \cdot \nabla C = \nabla \cdot (D \nabla C) - \kappa C + S$$

where the left side represents advection and accumulation, the first right-hand term represents diffusion, the second represents inactivation or settling, and S represents generation from construction activity, material handling, or clinical sources. This relationship is a physical basis for a surrogate model that predicts how robot tasks, barrier states, and ventilation settings affect contaminant concentration in adjacent zones. Physics-informed

learning can approximate the solution operator while retaining the governing structure, which is preferable to a purely black-box model when extrapolating to new barrier configurations [13]. Sparse identification methods can also be used to discover reduced-order relationships from sensor streams when the full geometry is not resolvable [17].

For a zone z, a Wells-Riley-type infection-risk surrogate can be written as

$$P_z = 1 - \exp\left(-I_z q p \frac{t_z}{Q_z}\right)$$

where I_z is the number of infectious sources in the zone, q is the quanta generation rate, p is the breathing rate of susceptible occupants, t_z is exposure duration, and Q_z is the clean-air flow rate. This expression is a planning surrogate, not a clinical diagnosis. Its value is that it makes exposure duration and clean-air flow explicit in the construction plan. A robot task that increases t_z or reduces Q_z by disturbing a barrier has a direct risk consequence that can be compared with its schedule benefit. Airborne transmission can be reduced by ventilation, filtration, and source control, but the effectiveness depends on flow organization and occupancy [18]. Healthcare construction design also has practical implications for containment, cleaning, and commissioning that must be represented in the planning model [19].

3.4 Multi-objective optimization and uncertainty

The co-design problem combines schedule, infection risk, energy, and disruption. A normalized scalarization is

$$F = w_T\, T_makespan/T_ref + w_P\, P_max/P_ref + w_E\, E_energy/E_ref + w_D\, D_disruption/D_ref$$

where T_ref, P_ref, E_ref, and D_ref are reference scales, P_max is the maximum zone infection-risk surrogate, E_energy is robot and temporary ventilation energy, D_disruption is a measure of clinical disruption, and w_T, w_P, w_E, and w_D are nonnegative weights. The optimization is subject to precedence, reachability, collision, pressure-zone, and risk constraints. This formulation is useful because it exposes the Pareto structure of the design problem: reducing T_makespan by increasing robot density may increase P_max through resuspension and barrier traffic, while reducing P_max by extending containment may increase T_makespan and energy use. Digital twins provide the integration layer for updating the state, running surrogate queries, and comparing design alternatives [20].

Because task durations, source strengths, airflow patterns, and barrier performance are uncertain, a robust decision rule is more appropriate than optimizing a single nominal estimate:

$$\min_{x,u} \max_{\theta \in \Theta} F(x, u; \theta)$$

where x and u denote task and control decisions, θ is an uncertain parameter vector, and Θ is the admissible uncertainty set. This rule selects a plan whose worst-case weighted cost over the uncertainty set is minimized. In practice, Θ should be constructed from sensor calibration ranges, known robot performance variability, and clinically justified bounds on source strength. The robust formulation prevents the optimizer from exploiting model regions where confidence is low, but it can be conservative. Sensitivity analysis should therefore report how the selected plan changes as Θ expands, so that decision makers can distinguish genuine trade-offs from model-derived caution.

4. Analytical Assessment and Design Trade-offs

The first trade-off concerns robot density and parallel execution. Increasing the number of robots can reduce T_makespan by parallelizing independent tasks, but it also increases traffic, material handling, and barrier crossings. In healthcare construction, each barrier crossing is a potential pressure disturbance and a source of contamination spread. The schedule benefit is therefore conditional on the containment design and task grouping. A design that assigns robots to enclosed work zones with dedicated material airlocks can exploit parallelism more safely than a design that requires repeated movement through occupied corridors. The optimization should penalize barrier openings and pressure-zone violations directly, rather than treating them as scheduling afterthoughts.

The second trade-off concerns ventilation and energy. Higher clean-air flow reduces the infection-risk surrogate for a given source strength and exposure duration, but it increases energy use and may alter pressure relationships if temporary systems are unbalanced. The framework should therefore optimize ventilation settings jointly with task timing. A short, high-generation task might be scheduled when the zone is unoccupied and ventilation can be temporarily increased without affecting clinical areas. A longer task with lower generation might be acceptable under standard containment if the pressure cascade is verified. The risk equation makes exposure duration explicit, which allows the planner to compare these options rather than relying on fixed rules.

The third trade-off concerns model fidelity and decision speed. High-fidelity computational fluid dynamics can represent airflow and contaminant transport in detail, but it is too slow for repeated optimization over many task sequences. A physics-informed surrogate can provide fast queries, but its error must be characterized in the regions that matter for decisions. A useful design practice is to reserve high-fidelity simulation for verification of the final candidate plans and for generating training data near critical barriers and airlocks. The surrogate should be calibrated against sensor measurements during commissioning and updated as barriers are moved. Without this calibration loop, the optimizer may produce plans that are safe only in the model.

The fourth trade-off concerns robustness and operational flexibility. A deterministic plan may achieve a shorter nominal makespan by tightly sequencing tasks and minimizing buffer time. However, if a robot is delayed, a material delivery is late, or a pressure reading drifts, the plan may require emergency barrier openings or rushed work. The robust formulation reduces this fragility by selecting plans that perform acceptably across a set of plausible disturbances. The cost is a longer nominal schedule or a more conservative containment configuration. This cost is justified when infection-control violations are unacceptable, but the stringency of the uncertainty set should be set through clinical risk assessment rather than by default.

Failure mechanisms also need explicit treatment. Contaminant resuspension can occur when robots move across dusty surfaces or when negative-pressure barriers are opened repeatedly. Pressure imbalance can arise when temporary exhaust is undersized or when doors are held open for material transfer. Robot downtime can create idle zones where ventilation and containment are maintained without productive work, increasing energy use. Sensor drift can bias the risk surrogate and cause the optimizer to underestimate exposure. Human oversight is therefore not an external add-on; it is a necessary part of the control architecture. Infection-prevention staff should review barrier plans, pressure-zone assignments, and

exception handling, while construction managers should review schedule feasibility and material logistics.

5. Implementation Architecture and Boundary Conditions

A practical implementation would use a layered architecture. The building information model supplies geometry, service zones, and task precedence. A digital twin maintains the current state using robot telemetry, environmental sensors, and schedule updates. A surrogate layer estimates contaminant transport and infection risk for candidate plans. An optimization layer solves the multi-objective or robust problem and returns task assignments, timings, paths, and control settings. A safety layer checks the proposed plan against hard constraints before execution and can veto or modify actions that would violate pressure or contamination limits. The safety layer should be rule-based and independently verifiable, while the learning components operate as advisory or bounded planning tools.

The framework has clear boundary conditions. It requires a reliable building information model with sufficient spatial and service-zone detail; otherwise, robot reachability and barrier placement cannot be evaluated. It requires sensor calibration for pressure, airflow, and particulate measures; otherwise, the risk surrogate is unanchored. It requires a task taxonomy that distinguishes dust-generating, wet, and low-generation activities, because source strength is task-dependent. It also requires clinical input on acceptable disruption and exposure thresholds. These are not minor implementation details; they determine whether the optimization is meaningful. The framework should therefore be introduced incrementally, beginning with planning support and human-approved task sequences, then moving to supervised execution and, only after validation, to higher levels of autonomy.

The approach is preferable to separate design workflows because it exposes couplings that are otherwise hidden. A construction planner may select a robot fleet and sequence for speed, and an infection-control engineer may then add containment measures that increase cost and delay. Co-design can identify a different task grouping, a different barrier layout, or a different robot configuration that achieves comparable speed with lower risk. The approach is also preferable to purely rule-based automation because it can evaluate many combinations and quantify trade-offs. However, it does not replace engineering judgment or clinical oversight. Its value depends on the quality of the state representation, the transparency of the objective weights, and the rigor of the safety layer.

6. Conclusion

This manuscript developed an AI-assisted design framework for robotic healthcare facility construction systems that treats operational efficiency and infection control as coupled objectives. The framework represents construction state through a building information graph, robot fleet states, environmental fields, contamination estimates, and human activity; it formulates task sequencing and assignment with precedence and completion-time relationships; it embeds a contaminant-transport relationship and an infection-risk surrogate; and it solves a multi-objective or robust optimization problem over schedule, risk, energy, and disruption. The central argument is that robotic construction in healthcare settings cannot be designed as a faster version of conventional construction. Barrier crossings, dust generation, pressure-zone integrity, and exposure duration are affected by robot density, task order, and material logistics, so these variables must be optimized together.

The analytical assessment identified several defensible trade-offs. Parallel robot execution can improve schedule performance but increases contamination and barrier-management risk

unless work is organized into contained zones with controlled material transfer. Higher ventilation reduces risk but increases energy use and can disturb pressure relationships if temporary systems are unbalanced. Physics-informed surrogates enable fast planning but require calibration and verification near critical barriers. Robust optimization reduces schedule fragility and safety violations but can be conservative. These trade-offs are not resolved by a single algorithm; they require explicit objective weights, uncertainty bounds, and human oversight. The framework provides a structured basis for that decision process and a foundation for future validation in controlled construction mock-ups and clinically supervised planning studies.

References

1. World Health Organization. (2009). Natural ventilation for infection control in health-care settings. World Health Organization.

2. Centers for Disease Control and Prevention. (2003). Guidelines for environmental infection control in health-care facilities. MMWR Recommendations and Reports, 52(RR-10), 1–42.

3. Ulrich, R. S., Zimring, C., Zhu, X., DuBose, J., Seo, H.-B., Choi, Y.-S., Quan, X., & Joseph, A. (2008). A review of the research literature on evidence-based healthcare design. HERD: Health Environments Research & Design Journal, 1(3), 61–125.

4. Li, Y., Leung, G. M., Tang, J. W., Yang, X., Chao, C. Y. H., Lin, J. Z., Lu, J. W., Nielsen, P. V., Niu, J., Qian, H., Sleigh, A. C., Su, H. J. J., Sundell, J., Wong, T. W., & Yuen, P. L. (2007). Role of ventilation in airborne transmission of infectious agents in the built environment—A multidisciplinary systematic review. Indoor Air, 17(1), 2–18.

5. Qian, H., Li, Y., Nielsen, P. V., Hyldgaard, C. E., Wong, T. W., & Chwang, A. T. Y. (2006). Dispersion of exhaled droplet nuclei in a two-bed hospital ward with three different ventilation systems. Indoor Air, 16(2), 111–128.

6. ASHRAE. (2017). Ventilation of health care facilities (ANSI/ASHRAE/ASHE Standard 170-2017). ASHRAE.

7. Bock, T., & Linner, T. (2016). Construction robots: Elementary technologies and single-task construction robots. Cambridge University Press.

8. Saidi, K. S., O'Brien, J. B., & Lytle, A. M. (2008). Robotics in construction. In B. Siciliano & O. Khatib (Eds.), Springer handbook of robotics (pp. 1079–1099). Springer.

9. Melenbrink, N., Werfel, J., & Menges, A. (2020). On-site autonomous construction robots: Towards unsupervised building. Automation in Construction, 119, 103312.

10. Petersen, K. H., Napp, N., Stuart-Smith, R., Rus, D., & Kovac, M. (2019). A review of collective robotic construction. Science Robotics, 4(28), eaau8479.

11. 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.

12. Tao, F., Zhang, H., Liu, A., & Nee, A. Y. C. (2019). Digital twin in industry: State-of-the-art. IEEE Transactions on Industrial Informatics, 15(4), 2405–2415.

13. Raissi, M., Perdikaris, P., & Karniadakis, G. E. (2019). Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. Journal of Computational Physics, 378, 686–707.

14. Luo, D., & Yu, L. (2021). From factory to site—designing for industrial robots used in on-site construction. In Automating cities: design, construction, operation and future impact (pp. 87-109). Singapore: Springer Singapore.

15. LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. Nature, 521(7553), 436–444.

16. Sutton, R. S., & Barto, A. G. (2018). Reinforcement learning: An introduction (2nd ed.). MIT Press.

17. Brunton, S. L., Proctor, J. L., & Kutz, J. N. (2016). Discovering governing equations from data by sparse identification of nonlinear dynamical systems. Proceedings of the National Academy of Sciences, 113(15), 3932–3937.

18. Morawska, L., Tang, J. W., Bahnfleth, W., Bluyssen, P. M., Boerstra, A., Buonanno, G., Cao, J., Dancer, S., Floto, A., Franchimon, F., Haworth, C., Hogeling, J., Isaxon, C., Jimenez, J. L., Kurnitski, J., Li, Y., Loomans, M., Marks, G., Marr, L. C., ... Yao, M. (2020). How can airborne transmission of COVID-19 indoors be minimised? Environment International, 142, 105832.

19. Bartley, J. M., Olmsted, R. N., & Haas, J. P. (2010). Current views of health care design and construction: Practical implications for safer, cleaner environments. American Journal of Infection Control, 38(5), S1–S12.

20. Jones, D., Snider, C., Nassehi, A., Yon, J., & Hicks, B. (2020). Characterising the digital twin: A systematic literature review. CIRP Journal of Manufacturing Science and Technology, 29, 36–52.