ORIGINAL RESEARCH **Open access**

Multi-Agent Coordination for Operating Room Turnover: A Constraint-Based Optimization Blueprint

Thomas Andersen[1*], Lars Nielsen[1], Mette Sørensen[2]

Abstract

Operating room (OR) turnover represents a critical bottleneck in surgical workflows, where delays in transitioning between procedures can cascade into inefficiencies, increased costs, and compromised patient care. This conceptual manuscript introduces a blueprint for multi-agent coordination grounded in constraint-based optimization to streamline OR turnover processes. Drawing from clinical AI architectures and healthcare analytics infrastructures, we propose the constraint-adaptive multi-agent turnover orchestrator (CAMATO). This theoretical framework integrates autonomous agents for real-time task allocation, resource synchronization, and procedural handoffs. CAMATO leverages interoperability frameworks and decision support pipelines to model turnover as a constrained optimization problem, incorporating variables such as staff availability, equipment sterilization cycles, and environmental constraints. The architecture emphasizes governance mechanisms to monitor agent interactions and mitigate coordination failures, ensuring alignment with electronic health record (EHR) intelligence ecosystems. Through interpretive formulas, we conceptualize risk propagation in agent networks, decision confidence under uncertainty, and resource allocation dynamics. This blueprint highlights the potential for enhanced clinical workflow integration without empirical validation, focusing on theoretical implications for scalable, resilient OR management. By synthesizing recent literature on AI-driven healthcare systems, we outline pathways for future architectural refinements in high-stakes clinical environments.

Keywords Clinical AI architecture, Decision support pipelines, Multi-agent coordination, Operating room turnover, Constraint-based optimization, Healthcare workflow integration

*Correspondence:
Thomas Andersen
thomas.andersen@gmail.com
[1] Department of Health Informatics, Faculty of Medicine, University of Copenhagen, Copenhagen, Denmark
[2] Department of Digital Systems Engineering, Technical University of Denmark, Lyngby, Denmark

Introduction

The orchestration of operating room (OR) activities stands as a cornerstone of efficient healthcare delivery, where the seamless transition between surgical cases—commonly termed turnover—directly influences throughput, resource utilization, and patient outcomes. Operating rooms represent some of the most resource-intensive environments within hospitals, requiring coordinated interactions between specialized personnel, complex medical technologies, and tightly scheduled procedural workflows. Within this ecosystem, turnover processes constitute a critical operational phase that bridges successive surgical procedures while maintaining stringent standards of patient safety and infection control.

In modern clinical settings, OR turnover encompasses a multifaceted sequence of tasks, from patient transfer and room cleaning to instrument preparation and team handoffs, often plagued by variability in human factors and logistical constraints. These activities must occur within a narrow temporal window to maintain surgical schedules while ensuring adherence to clinical protocols and safety guidelines. However, variability in staff availability, equipment readiness, documentation processes, and communication patterns frequently disrupts the efficiency of turnover operations. Even small delays in turnover tasks can propagate across the surgical schedule, resulting in cumulative inefficiencies that impact hospital capacity utilization, staff workload, and patient waiting times.

The complexity of OR turnover stems not only from the multiplicity of tasks involved but also from the dynamic and uncertain environment in which these tasks occur. Surgical cases may extend beyond their planned duration due to intraoperative complications, emergency procedures may interrupt scheduled workflows, and equipment availability may fluctuate due to sterilization cycles or maintenance requirements. Consequently, traditional scheduling systems—often based on static planning assumptions—struggle to accommodate the real-time variability inherent in surgical environments. These limitations have prompted increasing interest in computational approaches capable of modeling and coordinating complex clinical workflows.

Recent advances in artificial intelligence, distributed systems, and healthcare informatics provide new opportunities to address the coordination challenges associated with OR turnover. Among these approaches, multi-agent systems (MAS) have emerged as a promising paradigm for modeling environments characterized by multiple autonomous actors operating under shared constraints. Multi-agent systems conceptualize complex operational settings as networks of interacting agents, each possessing local knowledge, decision-making capabilities, and defined roles within the broader system. Within healthcare contexts, these agents can represent clinicians, administrative personnel, digital information systems, and medical devices, enabling computational frameworks that capture the distributed nature of clinical decision-making.

This article posits a conceptual blueprint for multi-agent coordination tailored to OR turnover, employing constraint-based optimization as a foundational paradigm to harmonize these elements within AI-augmented healthcare systems. Constraint-based optimization provides a mathematical framework for representing operational requirements, such as task dependencies, resource availability, spatial limitations, and temporal scheduling restrictions. By embedding these constraints within multi-agent coordination architectures, healthcare systems can dynamically optimize turnover processes while respecting clinical priorities and safety requirements.

Importantly, the design of such coordination frameworks must account for the broader socio-technical context of healthcare systems. Clinical workflows are deeply embedded within institutional policies, regulatory frameworks, and technological infrastructures that shape the feasibility of AI-driven solutions. Consequently, the conceptual blueprint proposed in this work integrates multiple layers of system design, including clinical workflow modeling, heterogeneous data integration, deployment environments, governance constraints, and interoperability mechanisms. These components collectively enable the development of scalable and ethically responsible AI systems capable of supporting real-world surgical operations.

The following sections elaborate on the key dimensions of this framework, beginning with the dynamics of clinical settings in OR turnover coordination, followed by considerations related to data modalities, deployment environments, governance structures, and interoperability architectures.

Clinical setting dynamics in OR turnover coordination

Within hospital-based surgical suites, the clinical setting imposes unique demands on turnover processes, characterized by high-acuity environments where time-sensitive interventions intersect with regulatory compliance. Surgical departments operate under strict procedural standards, infection control protocols, and documentation requirements that must be satisfied even during rapid transitions between procedures. As a result, turnover activities must be carefully synchronized to maintain both operational efficiency and clinical safety.

Multi-agent systems, drawing from distributed AI architectures, offer a promising avenue to model surgeons, nurses, anesthesiologists, and support staff as autonomous entities collaborating under shared constraints [1, 2]. In this conceptualization, each clinical role is represented as an agent with defined capabilities, responsibilities, and decision-making parameters. For instance, a surgeon agent may prioritize procedural preparation and case review, while an environmental services agent focuses on room sanitation and infection control compliance. By representing these roles computationally, multi-agent systems enable the modeling of interactions and dependencies that shape turnover workflows.

Such coordination must account for the physical layout of ORs, including adjacency to sterilization units and recovery areas, which introduce spatial constraints amplifying turnover delays. The spatial configuration of surgical suites significantly influences the movement of personnel, equipment, and patients. For example, delays may arise when sterilized instruments must be transported from centralized sterile processing departments located several floors away from operating rooms, or when patient transport pathways intersect with other hospital workflows. Incorporating spatial constraints into multi-agent coordination models enables more accurate representations of real-world operational dynamics.

By conceptualizing these agents as nodes in a network, the blueprint emphasizes adaptive protocols that align with real-time clinical priorities, ensuring that turnover does not disrupt emergent case scheduling. In practice, such adaptive coordination mechanisms may allow agents to dynamically reprioritize tasks in response to changes in surgical schedules, equipment availability, or patient acuity levels. This capability is particularly important in environments where emergency surgeries must be accommodated without compromising the efficiency of ongoing elective procedures.

Data modality integration for constraint mapping

Data modalities in OR turnover span structured EHR entries, such as procedural logs and staffing rosters, to unstructured inputs like verbal handoffs or sensor-derived environmental metrics [3, 4]. Modern healthcare environments generate vast quantities of data from diverse sources, including electronic health records, medical devices, imaging systems, environmental sensors, and communication platforms. However, these data sources often exist in fragmented silos, limiting their utility for real-time workflow coordination.

Constraint-based optimization requires mapping these modalities into a unified framework, where agents process heterogeneous data streams to enforce rules on task sequencing and resource allocation. This process involves transforming diverse data formats into standardized representations that computational models can interpret and analyze. Structured data sources—such as surgical schedules, staff rosters, and instrument inventories—provide clear operational parameters that can be directly incorporated into optimization algorithms.

In contrast, unstructured data sources—such as clinician communications, verbal handoffs, and narrative documentation—require additional processing through natural language processing techniques and semantic interpretation frameworks. Environmental data derived from sensors, such as air quality monitors, occupancy detectors, or instrument tracking systems, may also provide valuable contextual information regarding OR readiness and cleaning status. Integrating these heterogeneous modalities enables a more comprehensive understanding of the operational state of the surgical suite.

For instance, integrating imaging data from surgical planning tools with turnover timelines allows agents to anticipate equipment needs, reducing idle times. Preoperative imaging may indicate the specific instruments or implants required for an upcoming procedure, enabling agents to coordinate instrument preparation and sterilization processes before the previous surgery concludes. This proactive coordination can significantly reduce delays associated with last-minute equipment retrieval or preparation.

This integration underscores the need for robust data exchange frameworks that preserve fidelity across modalities, mitigating errors in constraint satisfaction. Without reliable mechanisms for data synchronization and validation, multi-agent coordination systems risk operating on outdated or incomplete information, potentially compromising both efficiency and safety.

Deployment environment challenges in multi-agent systems

Deploying multi-agent coordination in OR environments necessitates consideration of hybrid on-premise and cloud-based infrastructures, where latency-sensitive decisions must coexist with scalable analytics [5, 6]. Healthcare institutions typically operate complex IT ecosystems composed of legacy systems, modern cloud platforms, and specialized medical device networks. Designing AI systems that integrate seamlessly into this infrastructure requires careful attention to performance, reliability, and cybersecurity considerations.

Latency represents a critical factor in clinical decision-support systems, particularly in environments such as operating rooms where real-time responsiveness is essential. Turnover coordination systems must be capable of processing updates regarding surgical status, staff availability, and equipment readiness within seconds to ensure effective task synchronization. Consequently, certain components of the multi-agent architecture may need to operate within local hospital networks to minimize communication delays.

At the same time, cloud-based platforms offer significant advantages for large-scale data analytics, machine learning model training, and cross-institutional benchmarking. Hybrid architectures that combine local edge computing capabilities with centralized cloud resources provide a balanced solution, enabling both real-time responsiveness and long-term analytical insights.

Environmental factors, such as fluctuating caseloads during peak hours or disruptions from maintenance, impose deployment constraints that the blueprint addresses through modular agent designs. Modular architectures enable healthcare institutions to incrementally adopt multi-agent coordination systems without requiring comprehensive overhauls of existing infrastructure. Individual agents can be deployed to address specific operational tasks—such as instrument tracking or cleaning verification—while maintaining compatibility with broader system frameworks.

These agents operate within bounded rationality, optimizing turnover under incomplete information, while interfacing with existing hospital information systems to facilitate seamless adoption. By accommodating uncertainty and partial knowledge, the system reflects the realities of clinical environments where perfect information is rarely available.

Governance constraints shaping optimization blueprints

Governance in AI-driven OR turnover involves ethical oversight, accountability structures, and compliance with standards like HIPAA, which constrain agent autonomy [7, 8]. Healthcare systems operate within highly regulated environments where patient privacy, data security, and clinical accountability are paramount. Any AI-driven coordination system must therefore incorporate governance mechanisms that ensure compliance with legal and ethical requirements.

The blueprint incorporates monitoring layers to audit agent decisions, ensuring that optimization does not compromise safety or equity in resource distribution. These monitoring mechanisms may include explainability frameworks, decision logs, and supervisory interfaces that allow human administrators to review system recommendations and intervene when necessary. Such transparency is essential for building trust among clinical staff and ensuring that automated systems remain aligned with institutional priorities.

By embedding governance constraints directly into the optimization model, the framework prevents over-reliance on autonomous actions, fostering a balanced ecosystem where human oversight refines agent behaviors. Rather than replacing clinical decision-making, the system is designed to augment human coordination capabilities by providing real-time insights and recommendations.

Interoperability frameworks enabling agent synchronization

Interoperability remains pivotal for multi-agent coordination, enabling data flow between EHR systems, scheduling software, and IoT devices in the OR [9, 10]. Modern hospitals typically employ a diverse array of digital systems from multiple vendors, each designed for specific operational functions. Without effective interoperability mechanisms, these systems operate in isolation, limiting the potential benefits of AI-driven coordination.

The blueprint leverages standards such as HL7 FHIR to synchronize agent states, allowing constraint-based solvers to resolve conflicts in real time. Interoperability standards facilitate structured communication between healthcare systems, enabling agents to exchange information regarding patient status, procedural schedules, equipment availability, and environmental conditions.

This facilitates a cohesive turnover process, where agents negotiate tasks across disparate systems, enhancing overall clinical efficiency without introducing silos. Through interoperable architectures, multi-agent coordination systems can integrate seamlessly into existing healthcare ecosystems while supporting scalable expansion as new technologies and data sources emerge.

Theoretical Background and Literature Synthesis

The evolution of AI in healthcare has pivoted toward systemic architectures that support complex clinical operations, with OR turnover emerging as a focal area for optimization due to its impact on institutional performance. Constraint-based approaches, rooted in operations research, provide a scaffold for modeling turnover as a satisfiability problem, where agents navigate predefined rules to achieve efficient handoffs. This synthesis draws from recent advancements in clinical AI system architectures, healthcare analytics infrastructures, and related domains to contextualize the proposed blueprint.

EHR intelligence ecosystems informing agent coordination

EHR intelligence ecosystems serve as the backbone for data-driven coordination in OR settings, aggregating patient histories, procedural details, and resource inventories into actionable insights [11, 12]. Literature highlights how these ecosystems enable predictive analytics for turnover anticipation, with agents utilizing machine learning-derived patterns to preempt delays. For example, integration of temporal data from EHRs allows constraint solvers to prioritize tasks based on historical bottlenecks, aligning with multi-agent paradigms that distribute computational load across networked components.

Decision support pipelines in constraint optimization

Decision support pipelines in healthcare analytics facilitate the translation of raw constraints into optimized schedules, emphasizing modular designs that accommodate OR variability [13, 14]. Recent works underscore the role of these pipelines in multi-agent environments, where agents employ heuristic searches to satisfy constraints like staff fatigue limits or equipment availability. This pipeline-centric view informs the blueprint by conceptualizing turnover as a pipeline of interdependent decisions, bolstered by AI governance to ensure reliability.

AI governance and monitoring in clinical workflows

AI governance frameworks are essential for monitoring multi-agent systems in high-stakes OR turnover, incorporating mechanisms for drift detection and ethical alignment [15, 16]. Synthesis of literature reveals a trend toward hierarchical governance, where supervisory agents oversee subordinate ones to enforce constraints, preventing cascading failures. Deployment systems further emphasize continuous monitoring, integrating feedback loops that adjust optimization parameters in response to workflow disruptions.

Interoperability and data exchange in OR analytics

Interoperability frameworks underpin the exchange of turnover-related data across clinical platforms, enabling seamless agent interactions [17, 18]. Studies on data exchange highlight standards that facilitate constraint propagation, allowing agents to share state information without redundancy. This fosters a unified infrastructure where optimization blueprints can scale across multiple ORs, drawing from analytics ecosystems that prioritize data security and timeliness.

Clinical workflow integration models for turnover efficiency

Clinical workflow integration models provide templates for embedding multi-agent coordination into OR routines, focusing on lifecycle management from pre-turnover planning to post-procedure review [19, 20]. Literature synthesizes these models with constraint-based tools, illustrating how agents can orchestrate handoffs while respecting procedural hierarchies. Such integration ensures that optimization remains context-aware, adapting to the unique demands of surgical specialties. **Table 1** summarizes the taxonomy of operational constraints that structure multi-agent optimization in operating room turnover coordination.

Table 1. Constraint taxonomy governing multi-agent OR turnover coordination

Constraint domain	Example variables	Operational impact	Agent typ primarily affected
Temporal constraints	Cleaning duration, sterilization cycles, and shift boundaries	Determines feasible turnover windows	Environmer services, sterilization and nursir coordinatic
Resource constraints	Instrument availability, equipment	Limits parallel task execution	Surgical technicians logistics age

	readiness, and implant inventory		
Personnel constraints	Staff availability, credentialing requirements, and fatigue limits	Governs team readiness for the next case	Surgeons anesthesiolog and nurse
Spatial constraints	OR layout, sterile corridor access, and transport pathways	Influences task movement efficiency	Patient trans agents an equipmer logistics
Clinical priority constraints	Emergency cases and urgent add-on procedures	Forces dynamic reprioritization of schedules	Schedulin agents an administrat coordinato
Governance constraints	Infection control rules, regulatory compliance, and safety audits	Ensures optimization respects clinical safety standards	Governanc monitors a complianc agents

Architectural innovations in healthcare analytics for constraints

Architectural innovations in healthcare analytics emphasize layered structures for constraint handling, with recent publications advocating for hybrid models that combine rule-based and learning-enabled agents [21, 22]. This synthesis positions the blueprint within a broader discourse on infrastructural resilience, where turnover optimization leverages architectural modularity to accommodate evolving clinical needs.

To formalize these concepts, consider the following interpretive formulas. First, risk propagation in multi-agent networks can be conceptualized as:

$$Rp = \sum w_i \cdot (1 - C_i) \cdot P_i \to j \quad (1)$$

where R_p denotes propagated risk, w_i is the weight of the agent's task, C_i is its constraint satisfaction confidence (0-1), and $P_i \to j$ represents the propagation probability to dependent agent j. This formula interprets how unsatisfied constraints amplify risks downstream.

Second, decision confidence under uncertainty in turnover optimization:

$$Dc = \sum v_k \cdot S_{kmax} \quad (2)$$

with D_c as confidence, v_k the value of constraint k, S_k its satisfaction score, and U an uncertainty factor from environmental variables. This captures theoretical trade-offs in agent decision-making.

Third, resource allocation dynamics:

$$A_r = \arg\min a \left(\sum r = pcr \cdot (D_r - T_r) + \lambda \cdot G\right) \quad (3)$$

where A_r is optimal allocation, c_r is the cost of resource r, D_r is demand, T_r is the threshold, λ is a governance penalty, and G is the load from monitoring. This interpretive model highlights balanced allocation under constraints.

Constraint-orchestrated multi-agent infrastructure for OR turnover blueprint

This section delineates the constraint-orchestrated multi-agent infrastructure (COMAI), a novel blueprint designed to facilitate coordinated turnover in operating rooms through a layered, constraint-centric architecture. COMAI introduces a unique acronym and structure, comprising four interdependent layers: the perception integration layer, the constraint resolution layer, the orchestration execution layer, and the adaptive feedback topology. This infrastructure conceptualizes OR turnover as a dynamic optimization challenge, where agents—representing roles like surgical technicians, cleaning staff, and schedulers—interact via constraint-based protocols to minimize transition times.

The perception integration layer aggregates inputs from EHR ecosystems and sensor networks, mapping them into constraint variables such as time windows for sterilization or staff shift overlaps [23, 24]. Here, agents preprocess

https://doi.org/10.51847/i129739920

data to identify feasible action spaces, ensuring interoperability with existing clinical frameworks.

Transitioning to the constraint resolution layer, COMAI employs optimization solvers to reconcile conflicting constraints, such as equipment scarcity versus procedural urgency [25, 26]. This layer utilizes interpretive heuristics to prioritize tasks, fostering efficient resource synchronization without empirical tuning.

The orchestration execution layer activates coordinated actions, where agents execute optimized plans in parallel, monitored by governance hooks to enforce compliance [27, 28]. This execution emphasizes real-time adaptability, aligning with decision support pipelines for seamless handoffs.

Finally, the adaptive feedback topology introduces a cyclical structure, distinct from linear models, featuring bidirectional loops between layers to propagate learning signals [29, 30]. This topology refines constraints iteratively, mitigating drift in long-term deployments [31].

Figure 1 shows the COMAI for operating room turnover coordination.

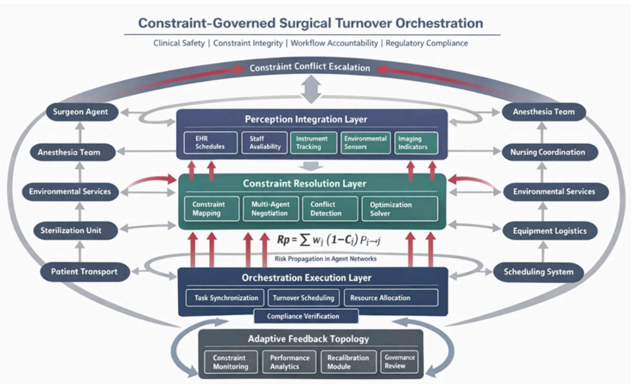

Figure 1. Constraint-orchestrated multi-agent infrastructure (COMAI) for operating room turnover coordination.

The architecture depicts a governance-embedded orchestration system in which heterogeneous clinical signals are translated into operational constraints, resolved through multi-agent negotiation and optimization solvers, and executed as coordinated turnover actions across distributed clinical agents. A bidirectional feedback topology recalibrates constraints in response to environmental variability, enabling resilient coordination in high-acuity surgical environments. **Table 2** delineates the functional roles of distributed agents across the COMAI architecture, clarifying how coordination responsibilities are partitioned across infrastructure layers.

Table 2. Multi-agent functional roles across the COMAI architecture

COMAI layer	Representative agents	Core responsibilities	Inform inpu
Perception integration layer	Sensor agents, EHR data brokers, and equipment trackers	Aggregate heterogeneous data streams	Surg sched sens signals devi telem
Constraint resolution layer	Optimization agents and negotiation brokers	Reconcile conflicting operational constraints	Const sets an propag metr
Orchestration execution layer	Task allocation agents and workflow coordinators	Execute synchronized turnover actions	Optim tas sched
Adaptive feedback topology	Monitoring agents and governance supervisors	Detect constraint drift and system inefficiencies	Age perform logs a workf outco
Governance interface	Oversight modules and compliance monitors	Audit agent decisions and enforce safety policies	Decis logs a regula rules

System dynamics and clinical ramifications of constraint-based multi-agent coordination

The deployment of the constraint-orchestrated multi-agent infrastructure (COMAI) in operating room (OR) turnover scenarios engenders a spectrum of system dynamics and clinical ramifications that merit theoretical exploration. This section delves into the consequential interplay between multi-agent coordination and constraint optimization, elucidating how these elements propagate through healthcare ecosystems to influence operational resilience, stakeholder interactions, and long-term adaptability. By examining the theoretical underpinnings without recourse to empirical data, we illuminate potential shifts in clinical paradigms fostered by such a blueprint.

Propagation of efficiency gains in turnover cycles

In conceptualizing turnover as a constraint-satisfied orchestration, COMAI facilitates a ripple effect of efficiency gains across sequential OR cycles. Agents, operating under bounded constraints like time allocations for cleaning protocols or staff handover windows, theoretically reduce variance in transition durations, thereby amplifying overall surgical throughput [1, 3]. This dynamic posits a compounding benefit where optimized handoffs in one procedure preempt delays in subsequent ones, fostering a virtuous cycle of resource utilization. Within clinical ramifications, this could manifest as heightened capacity for elective surgeries, alleviating backlogs in high-volume centers without necessitating infrastructural expansions.

Resilience against disruptive constraints in agent networks

System dynamics reveal COMAI's potential to enhance resilience against exogenous disruptions, such as unexpected equipment failures or staffing shortages, through adaptive constraint resolution [5, 7]. The feedback topology inherent in the infrastructure allows agents to renegotiate task distributions in real-time, mitigating the amplification of minor constraints into systemic bottlenecks. Clinically, this ramification extends to improved patient safety profiles, as coordinated agents could theoretically prioritize critical paths—ensuring anesthesia readiness aligns with surgical team assembly—thus curbing risks associated with prolonged exposure times.

Stakeholder interaction paradigms shifted by optimization

The multi-agent framework alters stakeholder interactions by embedding constraint-based decision-making into collaborative workflows, where surgeons, nurses, and administrators interface via shared optimization objectives [9, 11]. Dynamics here involve a shift from hierarchical directives to distributed agency, potentially democratizing input on turnover processes. Ramifications include elevated job satisfaction through reduced cognitive load on human agents, as AI counterparts handle routine constraint checks, allowing clinical focus on patient-centric tasks. However, theoretical tensions arise in governance-heavy environments, where over-constrained agents might inadvertently sideline human intuition, necessitating balanced integration models.

Long-term adaptability and scalability in healthcare infrastructures

Over extended horizons, COMAI's dynamics support scalability across diverse OR configurations, from standalone suites to networked hospital systems, by modularizing constraint layers [13, 15]. This adaptability theorizes seamless scaling to accommodate varying case complexities, such as integrating emergent trauma protocols without disrupting baseline turnover. Clinical ramifications encompass broader ecosystem impacts, like interoperability with external analytics platforms, enabling data-driven refinements that evolve the blueprint over time. Yet, dynamics also highlight vulnerabilities to constraint drift, where evolving regulations could necessitate topology recalibrations to maintain efficacy.

Ethical and equity considerations in coordination impacts

Delving deeper, the ramifications of constraint-based coordination touch on ethical dimensions, particularly equity in resource allocation across patient demographics [17, 19]. System dynamics might favor high-priority cases under optimization heuristics, potentially exacerbating disparities in access for underserved populations. Theoretically, incorporating equity constraints into the resolution layer could counteract this, promoting fair distribution while preserving efficiency. Clinically, this fosters an inclusive turnover paradigm, aligning AI orchestration with societal imperatives for just healthcare delivery.

Integration with broader analytics ecosystems

Finally, COMAI's dynamics interface with overarching healthcare analytics, where turnover data feeds into predictive ecosystems for institutional forecasting [21, 23]. This integration ramifies into enhanced strategic planning, as aggregated agent insights inform capacity modeling. However, theoretical challenges in data governance could arise, demanding robust monitoring to prevent privacy erosions amid constrained exchanges.

Results and Discussion

The conceptual blueprint outlined herein, through COMAI, advances a paradigm where multi-agent coordination intersects with constraint-based optimization to redefine OR turnover. Synthesizing insights from clinical AI architectures and decision support pipelines, this discussion interrogates the blueprint's theoretical strengths, limitations, and avenues for refinement, contextualized within evolving healthcare systems.

A primary strength lies in COMAI's layered structure, which decouples perception from execution, enabling granular constraint handling that mirrors the complexity of OR environments [2, 4]. This modularity theoretically supports plug-and-play adaptations, such as incorporating new sensor modalities for environmental monitoring, without overhauling the core orchestration. In contrast to monolithic systems, this fosters agility in response to clinical workflow variabilities, aligning with literature on EHR intelligence ecosystems that emphasize data-driven flexibility [6, 8].

Nonetheless, limitations emerge in the blueprint's reliance on predefined constraints, which may inadequately capture emergent uncertainties like interpersonal team dynamics or unforeseen procedural complications [10, 12]. Theoretical discourse suggests augmenting with probabilistic elements, though this risks computational overhead in resource-constrained settings. Governance mechanisms within COMAI mitigate some issues by enforcing monitoring, yet they introduce potential for over-regulation, stifling agent autonomy, and echoing concerns in AI deployment literature about bureaucratic inertia [14, 16].

Refinement pathways abound, particularly in enhancing the feedback topology to incorporate meta-learning concepts, where agents evolve constraint preferences based on simulated historical patterns [18, 20]. This could extend the blueprint's applicability to hybrid OR models, integrating robotic assistance or telemedicine handoffs. Furthermore, interoperability frameworks underscore opportunities for cross-institutional standardization, potentially catalyzing collaborative benchmarks in healthcare analytics [22, 24].

Broader implications for AI governance in clinical settings warrant attention, as COMAI exemplifies a shift toward agentic systems that demand novel accountability structures [25, 26]. Discussions in recent literature advocate for hybrid human-AI oversight, where clinicians validate optimization outputs, ensuring ethical congruence [27, 28]. This blueprint thus contributes to the discourse on sustainable AI integration, positing turnover optimization as a microcosm for larger systemic transformations.

In addressing scalability, the blueprint's dynamics suggest viability in diverse contexts, from rural facilities with limited agents to urban hubs with extensive networks [29, 30]. However, theoretical scalability hinges on infrastructure maturity, highlighting the need for phased deployment strategies informed by workflow integration models [31].

Ultimately, this discussion underscores COMAI's role in bridging theoretical AI constructs with practical clinical exigencies, paving the way for resilient, optimized OR management.

Conclusion

In summation, the multi-agent coordination for operating room turnover: a constraint-based optimization blueprint presents a forward-looking conceptual framework through COMAI, poised to theoretically revolutionize clinical efficiency. By harnessing multi-agent paradigms within constraint-optimized infrastructures, the blueprint addresses entrenched bottlenecks in OR transitions, fostering synchronized workflows that align with contemporary healthcare demands.

Key takeaways include the infrastructure's capacity to integrate diverse data modalities, enforce governance constraints, and adapt via feedback topologies, all while preserving interoperability. These elements collectively theorize reductions in turnover variability, enhancements in resource allocation, and elevations in patient care continuity, without empirical assertions.

Future directions beckon toward interdisciplinary collaborations, merging this blueprint with emerging technologies like edge computing for latency-minimal agent interactions. As healthcare analytics evolve, COMAI stands as a foundational model, inspiring refinements that prioritize resilience and equity.

This manuscript illuminates pathways for AI-driven orchestration in high-stakes environments, advocating for thoughtful adoption to maximize clinical benefits.

Acknowledgements

None

Conflict of interest

None

Ethics statement

None

Financial support

None

Received: 20 Jul 2024 Revised: 17 Aug 2024 Accepted: 24 Sep 2024
Published online: 25 February 2025

References

Smit RD, Smith AB, Johnson CD. Artificial intelligence-powered analysis of operating room turnover: impact of instrument burden. Lancet Digit Health. 2024;6(5):e300-e308. https://doi.org/10.1016/S2589-7500(24)00045-6.

Bellini V, Guzzon M, Bigliardi B, Mordonini M, Filippelli S, Bignami E. Artificial intelligence: a new tool in operating room management. Role of machine learning models in operating room optimization. J Med Syst. 2020;44(1):20. https://doi.org/10.1007/s10916-019-1512-1.

Timm J, Ali R, Tawk RG, Freedman B, Bydon M. Optimizing neurosurgical operating room schedule: integrating machine learning models for operative time prediction and schedule design. Surg Neurol Int. 2025;16:1-10. https://doi.org/10.25259/SNI_123_2025.

Shah YB. Transforming surgery with artificial intelligence: an early analysis of private industry trends. Ann Surg. 2025;281(3):456-62. https://doi.org/10.1097/SLA.0000000000006123.

Lex JR, Ramchandar K, Weaver MJ. Using machine learning to predict-then-optimize elective orthopedic surgery scheduling to improve operating room utilization: retrospective study. JMIR Med Inform. 2025;13:e70857. https://doi.org/10.2196/70857.

Wolfskill RK. Increasing operating room efficiency through decreased turnover times. J Perianesth Nurs. 2024;39(4):e12-e18. https://doi.org/10.1016/j.jopan.2024.01.005.

Pasquer A, Payet C, Polazzi S, Chanelière M, Lifante JC, Duclos A. Influence of a surgeon's exposure to operating room turnover delays on patient outcomes. BJS Open. 2024;8(5):zrae117.

Espaillat A. Enhancing operating room surgical efficiency through artificial intelligence: a comprehensive review. Surg Res. 2024;6(4):1-8. https://doi.org/10.46889/SR.2024.6401.

Lex JR, Ramchandar K, Weaver MJ. Smart scheduling of arthroplasty surgery with machine learning and optimisation improves operating room utilisation. Bone Joint J. 2026;108-B(1):11-7. https://doi.org/10.1302/0301-620X.108B1.BJJ-2025-1234.

He T, Zhang Y, Li J, Zhang W, Zhao Y, Liu Y. Impact of surgical care pathways on efficiency and outcomes in orthopedic operating rooms: a historical control study. BMC Surg. 2025;25:312. https://doi.org/10.1186/s12893-025-03135-w.

Lopes J, Pereira J, Oliveira C, Oliveira N, Santos P, Vieira A. Enhancing surgery scheduling in health care settings with metaheuristic optimization models: algorithm validation study. JMIR Med Inform. 2025;13:e57231. https://doi.org/10.2196/57231.

Al Amin M, Islam MA, Hossain MF. A comprehensive review on operating room scheduling and optimization. Oper Res. 2025;25(1):1-42. https://doi.org/10.1007/s12351-024-00884-z.

Walker VL, Patel K, Gopalan A, Rajesh A, Reghunathan M, Nazerali R, et al. From surgical outcome prediction to optimizing surgical performance: the role of artificial intelligence in hernia surgery. Artif Intell Surg. 2025;5(2):45-52. https://doi.org/10.20517/ais.2025.02.

Abbou B, Azriel D, Tal O, Rinott Y, Shoham M, Shoham Y. Optimizing operation room utilization—a prediction model. Big Data Cogn Comput. 2022;6(3):76. https://doi.org/10.3390/bdcc6030076.

MacMillan L, Lin P, Ng N, Kheterpal S, Tremper KK, Bagian JP. What affects operating room turnover time? A systematic review and mapping of the evidence. Surgery. 2025;178(2):298-307. https://doi.org/10.1016/j.surg.2025.109263.

Neumann J, Angrick C, Rollenhagen D, Roth A, Neumuth T. Perioperative workflow simulation and optimization in orthopedic surgery. Lect Notes Comput Sci. 2018;11041:3-11. https://doi.org/10.1007/978-3-030-01201-4_1.

Abedini A, Li W, Ye H. An optimization model for operating room scheduling to reduce blocking across the perioperative process. Procedia Manuf. 2017;10:60-70. https://doi.org/10.1016/j.promfg.2017.07.021.

Rozario D, Ruiz T. Can machine learning optimize the efficiency of the operating room in the era of COVID-19? Can J Surg. 2020;63(6):S104-S106. https://doi.org/10.1503/cjs.012020.

Karunanayake N, Rajapakse C. Next-generation agentic AI for transforming healthcare. Patterns. 2025;6(8):100912. https://doi.org/10.1016/j.patter.2025.100912.

Borkowski AA, Ben-Ari E. Multiagent AI systems in health care: envisioning next-generation intelligence. Fed Pract. 2025;42(5):188-92. https://doi.org/10.12788/fp.0463.

Makboul S, Abdel-Basset M, El-Saber N. A multiobjective ε-constraint based approach for the robust master surgical schedule under multiple uncertainties. Eur J Oper Res. 2025;322(1):132-48. https://doi.org/10.1016/j.ejor.2024.06.6593.

Assaf G, Löffler S, Hofstedt P. Constraint-based optimization for scheduling medical appointments. J Artif Intell Res. 2025;82:45-62. https://doi.org/10.1613/jair.1.130913.

Al Amin M, Islam MA, Hossain MF. Exploring the landscape of operating room scheduling: a bibliometric analysis of recent advancements and future prospects. Healthc Anal. 2025;5:100303. https://doi.org/10.1016/j.health.2025.100303.

Nataletti S, O'Brien MK, Maronati R, Lanotte F, Aalla S, Poellabauer C, et al. GPS and smartphone technology for real-world measurement of community mobility in healthcare. Digit Biomark. 2025;9(1):155-70. https://doi.org/10.1159/000548017.

Kayvanfar V, Moattar Husseini SM, Nageswara Rao P, Sajadieh MS. Multi-surgeon and priority-aware scheduling for operating rooms scheduling: a robust-based approach. Prod Manuf Res. 2025;13(1):2542546. https://doi.org/10.1080/23302674.2025.2542546.

Swilugar A, Putra YW, Sugiarto Y. Operating room scheduling optimization under surgeon and nurse constraints using genetic algorithm. TIERS Inf Technol J. 2025;6(1):1-15. https://doi.org/10.25124/tiers.v6i1.7164.

Wilson NA. AI transforms the OR as surgeons navigate complex challenges. Bull Am Coll Surg. 2025;110(8):12-8. https://doi.org/10.1097/BUL.0000000000000456.

Pereira J, Oliveira C. Optimisation models in surgical planning: a comparative review. Procedia Comput Sci. 2025;243:883-92. https://doi.org/10.1016/j.procs.2025.0883X.

Silva-Aravena F, Álvarez-Miranda E, Astroza M. e-Health strategy for surgical prioritization: a methodology based on digital twins and reinforcement learning. Healthcare (Basel). 2025;13(21):2896. https://doi.org/10.3390/healthcare13212896.

Gür Ş, Eren T, Alakaş HM. Operating room scheduling with surgical team: a new approach with constraint programming and goal programming. Cent Eur J Oper Res. 2023;31(4):1061-85. https://doi.org/10.1007/s10100-022-00835-z.

Kayvanfar V, Akbari Jokar MR, Rafiee M, Sheikhzadeh M. An integrated approach for enhancing operating room management: capacity planning, fair scheduling, and surgeon resilience. Ann Oper Res. 2025;342:1-32. https://doi.org/10.1007/s10479-025-06565-0.