Optimizing Resource Allocation for Crisis-Resilient Healthcare Robotics: An Integrated MLR, MDP, and Petri Net Approach

Ikram Dahamou[1]*, Ayoub Elbazzazi[2], Cherki Daoui[3]
Faculty of Science and Technology, Computer Engineering,
Sultan Moulay Slimane University, B.P. 523, Beni Mellal, Morocco[1,2]
Faculty of Science and Technology, Mathematics and Computer Science
Sultan Moulay Slimane University, B.P. 523, Beni Mellal, Morocco[3]

***Abstract*—Global crises, such as pandemics and climate-related disasters, place unprecedented strain on healthcare systems, exposing weaknesses in resource management and patient care. This study aims to address these challenges by developing an integrated computational framework for crisis-resilient healthcare robotics. We propose a unified approach that combines Multiple Linear Regression (MLR), Markov Decision Processes (MDPs), and Petri Nets. MLR is applied to predict the Average Length of Stay (ALOS) using patient and hospital data. These forecasts inform MDPs, which guide admission and triage decisions under uncertainty. Petri Nets are employed to model and validate patient flow and hospital workflows, ensuring feasibility and efficiency. Case studies, including ICU bed prioritization and disaster logistics, demonstrate that the proposed framework improves adaptability and resource utilization while supporting structured decision guidance. Simulation results highlight enhanced system efficiency, better patient prioritization, and reduced congestion during surge conditions. The integration of predictive analytics, probabilistic optimization, and workflow modeling provides a robust decision-support system for healthcare robotics in crisis scenarios. This interdisciplinary framework offers practical solutions for improving resilience, scalability, and patient outcomes, providing a structured foundation for enhancing resilience and coordination in healthcare systems facing future emergencies.**

Keywords*—Healthcare robotics; Markov Decision Processes; Petri Nets; Multiple Linear Regression; crisis management; ICU resource allocation

I. Introduction

Global healthcare systems are increasingly challenged by recurring crises such as pandemics, extreme weather events, and large-scale disasters [1], [2]. These events expose persistent weaknesses in continuity of care, inefficiencies in resource management, and risks to patient outcomes.

A. Motivation and Contribution Perspective

Rather than introducing new individual algorithms, this work focuses on how existing predictive, decision, and workflow models can be structurally combined to support crisis-aware hospital operations. The core contribution lies in a decision-oriented modeling chain (MLR $\rightarrow$ MDP $\rightarrow$ Petri Net), which constrains hospital dynamics to decision-feasible trajectories. This structured coupling enables tractable reachability analysis and formal verification of admission policies under severe resource constraints, which cannot be achieved using predictive or decision models alone.

To address this gap, we propose an integrated computational framework that combines:

- Multiple Linear Regression (MLR) for forecasting critical indicators such as the Average Length of Stay (ALOS) [5], [6],
- Markov Decision Processes (MDPs) for guiding triage and admission policies under uncertainty [7], [8],
- Petri Nets for validating hospital workflows and ensuring feasibility of robotic interventions under operational constraints [10], [11].

This unified approach integrates statistical forecasting, probabilistic decision-making, and workflow simulation into a single coherent system.

B. Main Contributions

The main contributions of this work are summarized as follows:

- We propose a decision-constrained modeling chain that links predictive analytics, sequential decision-making, and workflow verification, reducing the reachable state space of hospital systems during crises.
- We introduce an explainable approximate MDP-based admission policy, designed for environments where full transition probabilities cannot be reliably identified.
- We demonstrate how Petri Net reachability analysis can be used to formally verify the feasibility, boundedness, and safety of data-driven admission decisions.

C. Paper Organization

The remainder of this study is organized as follows:

- Section II reviews related work on healthcare robotics and computational modeling, discusses applications

*Corresponding author.

and current limitations of robotic systems in healthcare, presents the proposed improvements and analyzes resource management during global crises.

- Section III details the computational models (MLR, MDPs, Petri Nets).
- Section IV provides case analysis and intermediate results.
- Section V develops Petri Net modeling and analysis.
- Section VI defines the initial marking and reachability analysis.
- Section VII describes integration with robotic systems through an illustrative workflow and offers a comparative analysis of reachability approaches.
- Section VIII concludes the study with perspectives for future research.

II. Related Work

The integration of computational models in healthcare robotics has attracted increasing attention in recent years. Several studies have applied Markov Decision Processes (MDPs) to model clinical decision-making under uncertainty, particularly for triage optimization and ICU resource allocation [7], [8]. For example, Hauskrecht (2000) demonstrated the use of MDPs in managing treatment policies for chronic diseases. More recent works have explored reinforcement learning and optimization-based strategies for hospital resource allocation under pandemic-induced surge conditions [24], [25], [26], highlighting the importance of adaptive decision-making in crisis scenarios.

Petri Nets have been widely used to model complex hospital workflows, including patient flow in emergency departments and coordination of clinical tasks [10], [11], [17]. They provide powerful tools for identifying bottlenecks and ensuring safety constraints in distributed systems. Recent probabilistic and time-dependent Petri Net extensions have further enabled formal modeling of clinical pathways under uncertainty [27].

Multiple Linear Regression (MLR) has been applied for forecasting hospital metrics such as length of stay (LOS), readmission risks, and cost estimation [5]. However, these studies generally focus on prediction as a standalone task and do not embed regression outputs into real-time robotic or autonomous decision workflows.

In contrast to prior works that treat predictive modeling, decision optimization, and workflow analysis as separate components, the present framework integrates MLR, MDPs, and Petri Nets within a unified decision-support modeling chain. While existing literature addresses forecasting, optimization, or workflow verification independently, relatively few approaches structurally couple predictive analytics with sequential decision-making and formal system validation.

This methodological gap motivates the integrated approach proposed in this study, where predictive estimates inform admission control policies, and decision-constrained transitions are subsequently verified through Petri Net reachability analysis.

Petri Nets have been widely used to model complex hospital workflows, including patient flow in emergency departments and coordination of clinical tasks [10], [11]. They provide powerful tools for identifying bottlenecks and ensuring safety constraints in distributed systems. Recent probabilistic and time-dependent Petri Net extensions have further enabled formal modeling of clinical pathways under uncertainty [27]. Recent Petri Net models have also been applied to healthcare systems under COVID-19 surge conditions [18].

A. Main Applications of Healthcare Robotics

- Surgical Robots: Systems like *da Vinci* enable minimally invasive surgery with higher precision and reduced recovery time [15].
- Assistive Robots: Devices such as PARO or Robear help elderly and disabled patients with daily activities and emotional support [12].
- Diagnostic and Rehabilitation Robots: AI-based robots analyze medical images or assist physical recovery, improving outcomes [22], [23].
- Logistics Support: Autonomous systems (e.g., TUG robots) deliver drugs and supplies, increasing hospital efficiency [1].

B. Challenges in Crisis Contexts

Despite these advances, crises such as pandemics expose key weaknesses:

- Resource rigidity: Robots are often task-specific and struggle to adapt to new priorities [3], [4].
- Limited real-time adaptation: Pre-programmed workflows cannot manage sudden surges in patient volume [8].
- Poor integration: Robotic subsystems rarely connect seamlessly to hospital information systems or triage networks [10].
- High cost and scalability issues: Advanced units remain inaccessible to resource-limited hospitals [2].

C. Resource Management Under Global Crises

Pandemics and disasters further aggravate shortages of ICU beds, staff, and equipment [19], [5]. Traditional allocation approaches remain static, manually driven, and fragmented, lacking predictive and adaptive capacity [7], [2]. Integrating data-driven models such as Petri Nets, Markov Decision Processes, and Multiple Linear Regression can improve crisis resource allocation [20].

- Real-time decision-making and prioritization of tasks,
- Interoperability with hospital information systems,
- Scalability and coordination during resource surges.

This hybrid framework promotes proactive, self-adjusting robotic systems capable of supporting hospitals under both routine and emergency conditions. The constraints and challenges outlined in this section motivate the need for an integrated decision-support framework capable of anticipating resource usage, regulating admissions, and ensuring operational safety.

III. Computational Models for Crisis-Aware Optimization

This work relies on the integration of three complementary modeling paradigms—Multiple Linear Regression (MLR), Markov Decision Processes (MDPs), and Petri Nets—to support crisis-aware decision-making in hospital systems. Each model plays a distinct and clearly delimited role: prediction, decision guidance, and formal workflow validation.

A. Predictive Layer: Multiple Linear Regression for ALOS Estimation

Multiple Linear Regression (MLR) is used to forecast the Average Length of Stay (ALOS), a key operational indicator for hospital capacity planning [5]. Accurate ALOS prediction enables anticipation of future bed occupancy and resource commitment, which is critical during crisis situations such as pandemics or mass-casualty events [6].

1) Model formulation: The MLR model estimates ALOS as:

$$Y = \beta_0 + \sum_{i=1}^{n} \beta_i X_i + \epsilon \tag{1}$$

where, Y denotes ALOS, X_i represent patient- and admission-related variables, β_i are regression coefficients, and ϵ is the error term.

2) Why Multiple Linear Regression?: MLR is selected as a transparent baseline that supports interpretability and auditability in healthcare operations. Unlike complex nonlinear predictors, MLR allows direct inspection of coefficients, facilitates sensitivity analysis, and reduces integration friction when predictions are embedded into the downstream decision layer.

Although nonlinear models may achieve marginally higher predictive accuracy, the primary objective of this work is to ensure explainability and seamless coupling with the decision-theoretic layer. Therefore, MLR is adopted as a structurally compatible and interpretable predictive backbone.

3) Data and predictive performance: The model is trained on data derived from the MIMIC-III database [6], including demographic information, clinical severity, emergency admission rate, satisfaction score, and total hospital charges.

4) Preprocessing and reproducibility details: We used MIMIC-III structured variables consistent with ALOS modeling. Missing values were handled by: 1) removing features with $> 30\%$ missingness, and 2) median imputation for continuous variables and mode imputation for categorical variables. Continuous predictors were standardized using z-score normalization: $x' = (x - \mu)/\sigma$ computed on the training set only. Categorical variables were one-hot encoded.

Feature selection followed a two-step procedure: a) correlation screening and variance inflation factor (VIF) control to mitigate multicollinearity, and b) stepwise regression using AIC to retain parsimonious predictors. We report the final set of predictors in Table I.

TABLE I. Key Predictors Impacting Average Length of Stay (ALOS)

Variable	Coef.	Interpretation
Emergency_Rate	+0.043	Increases ALOS
Patient_Severity_Index	+0.037	Prolongs stay
Satisfaction_Score	-0.031	Reduces stay duration
Total_Hospital_Charges	+0.022	Associated with longer stay

We evaluated the model using 5-fold cross-validation and report the average values of R^2, RMSE, and MAE across folds. The coefficient of determination ($R^2 = 0.86$) indicates that 86% of the variability in ALOS is explained by the model. The RMSE of 0.38 days reflects limited quadratic prediction error, while the MAE of 0.48 days suggests that the average absolute deviation between predicted and observed ALOS remains below half a day, indicating stable and clinically acceptable predictive performance.

The feature correlations are illustrated in Fig. 1.

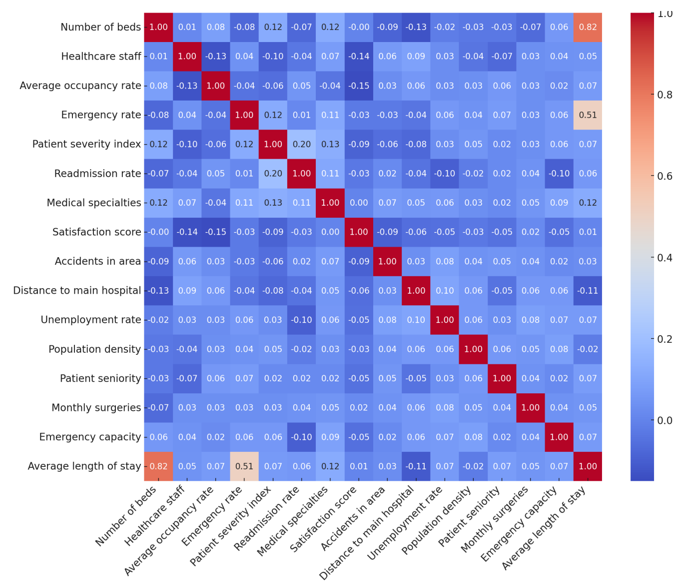

Fig. 1. Correlation matrix between input features and ALOS.

5) Operational role of MLR: Rather than acting as a decision-maker, MLR provides anticipatory information that feeds the downstream decision layer. Predicted ALOS values quantify expected resource occupation and are explicitly embedded into the MDP state representation.

B. Decision Layer: Crisis-Aware Admission Control via MDP

The Markov Decision Process (MDP) is not used to compute an optimal policy in the classical sense. Instead, it represents a constrained, explainable approximation of an optimal MDP policy, suitable for crisis contexts where transition probabilities are uncertain or rapidly evolving [7], [8].

1) State and action representation: Each MDP state captures the current operational condition of the hospital, including bed availability, ICU load, patient severity, and predicted ALOS from the MLR model. The action space is deliberately restricted to three clinically meaningful decisions: *admit*, *waitlist*, or *reject*.

2) Formal MDP definition and coupling with MLR: We model admission control as a Markov Decision Process (MDP) defined by the tuple $M = (S, A, P, R, \gamma)$, following the classical MDP formulation [9]. Each state $s_t \in \mathcal{S}$ summarizes the hospital operational condition at decision time t:

$$s_t = \big(b_t, o_t, \bar{a}_t, \hat{\ell}_t, \sigma_{\hat{\ell},t}, u_t\big) \tag{2}$$

where, b_t is the number of available beds, o_t is the ICU occupancy level, $\bar{a}_t$ denotes arrival pressure (e.g., emergency rate), $\hat{\ell}_t$ is the MLR-predicted ALOS for the incoming patient (or cohort), $\sigma_{\hat{\ell},t}$ is the prediction uncertainty (estimated from residual variance), and u_t denotes patient acuity/severity.

The action space is $\mathcal{A} = \{\mathit{Admit}, \mathit{Waitlist}, \mathit{Reject}\}$.

Transition dynamics are expressed as:

$$\mathcal{P}(s_{t+1} \mid s_t, a_t) \tag{3}$$

where, bed availability evolves according to admissions/discharges:

$$b_{t+1} = b_t - \mathbb{I}[a_t = \mathit{Admit}] + d_t \tag{4}$$

and d_t represents discharges, approximated using the predicted LOS distribution induced by $\hat{\ell}_t$.

The reward function encodes clinical utility and congestion penalties:

$$\begin{aligned}\mathcal{R}(s_t, a_t) = \alpha u_t \mathbb{I}[a_t = \mathit{Admit}] - \beta \mathbb{I}[b_t = 0 \wedge a_t = \mathit{Admit}] \\ - \gamma_1 \hat{\ell}_t \mathbb{I}[a_t = \mathit{Admit}] - \gamma_2 \sigma_{\hat{\ell},t} \mathbb{I}[a_t = \mathit{Admit}]\end{aligned} \tag{5}$$

where, $\alpha, \beta, \gamma_1, \gamma_2 \geq 0$ balance severity prioritization, overflow risk, expected resource occupation, and robustness to prediction uncertainty.

3) Decision logic and interpretation: The reward structure reflects institutional priorities rather than numerical optimality. Admissions associated with long predicted stays are penalized under high occupancy, while high-severity patients are prioritized. As a result, the derived policy acts as a decision-guidance mechanism that balances urgency and sustainability.

4) Illustrative admission decisions: Algorithm 1 presents a simplified rule-based representation of the MDP-derived logic, introduced for interpretability only.

Algorithm 1 MDP-Based Admission Policy

```
Require: CurrentBeds, TotalBeds,
         SeverityIndex, PredictedALOS
Ensure: AdmissionDecision
if CurrentBeds ≥ TotalBeds then
    AdmissionDecision ← Reject
else
    if SeverityIndex ≥ 8 then
        AdmissionDecision ← Admit
    else if PredictedALOS ¡ 5 then
        AdmissionDecision ← Admit
    else
        AdmissionDecision ← Waitlist
    end if
end if
return AdmissionDecision
```

Representative admission outcomes are reported in Table II, illustrating how severity and predicted ALOS jointly shape decisions under constrained capacity.

TABLE II. MDP DECISIONS BASED ON PRIORITY SCORE COMBINING SEVERITY INDEX AND PREDICTED ALOS.

Patient	Severity Index	Predicted ALOS	Priority Score	MDP Decision
P1	7	8.29	9.86	Waitlist
P2	4	5.98	5.01	Waitlist
P3	8	4.46	13.77	Admit
P4	5	6.97	6.51	Waitlist
P5	7	3.22	12.39	Admit
P6	3	8.90	1.55	Reject
P7	7	6.15	10.93	Admit
P8	8	5.77	13.12	Admit
P9	5	9.49	5.26	Waitlist
P10	4	8.09	3.95	Reject

C. Formal Validation Layer: Petri Net Modeling

Petri Nets are used as a formal verification and simulation layer to validate the feasibility and safety of MDP-driven decisions within hospital workflows [10], [11].

1) Reachability analysis: A Petri Net is defined as:

$$PN = (P, T, F, W, M_0)$$

and its reachability graph enumerates all possible system states reachable from the initial marking [13], [14]. Reachability analysis enables the verification of critical properties such as deadlock-freeness, boundedness, and guaranteed admission of high-priority patients. This positions Petri Nets as the scientific validation core of the proposed framework.

D. Methodological Complementarity

In summary:

- MLR anticipates future resource demand,
- The MDP provides explainable decision guidance under uncertainty,
- Petri Nets formally validate workflow consistency.

This layered architecture ensures predictive, decision-oriented, and formally verified crisis-aware optimization.

IV. Comparative Example Analysis of Decision-Aware and Direct Petri Net Modeling

This section presents a comparative example analysis between two hospital workflow modeling approaches under crisis conditions, highlighting the analytical benefits of introducing a decision layer prior to Petri Net construction.

A. Method A: Direct Petri Net Modeling

In the baseline approach, the hospital workflow is modeled directly as a Petri Net without explicit decision constraints. Patient arrivals, admissions, treatments, and discharges are represented as transitions, while hospital resources are modeled as bounded places.

The initial marking is defined as:

$$M_0 = (1, 0, 0, 0, 0, 0),$$

Indicating that one patient enters the system while all other places are initially empty. By firing all enabled transitions, a corresponding reachability graph is obtained. Although this approach accurately captures workflow dynamics, it generates a large and weakly constrained state space that may include clinically unrealistic configurations.

B. Method B: MLR → PDM → Petri Net Modeling (Proposed)

The proposed approach introduces a decision layer prior to Petri Net construction. Admission decisions are regulated by the Resource Logic Model (RLM) and the Process Decision Model (PDM), derived from the MDP framework and informed by predicted ALOS. Only policy-compliant decisions are translated into Petri Net transitions.The resulting decision-aware Petri Net structure is shown in Fig. 2.

In particular, transitions from waitlist to admission are enabled exclusively when priority and resource availability conditions are satisfied.

C. Reachability Graph Derived from the Decision-Aware Model

The reachability graph associated with the proposed modeling chain is shown in Fig. 3. Each node corresponds to a feasible system configuration, while edges represent authorized transition firings. Compared to direct Petri Net modeling, the reachability graph is significantly more compact and structured, containing only clinically admissible and resource-feasible states.

D. Comparative Discussion

The comparison highlights three major advantages of the proposed approach:

- Controlled State-Space Complexity: Decision constraints limit unnecessary reachability expansion.
- Decision Alignment: Petri Net execution remains consistent with MDP-based admission priorities.
- Formal Verifiability: Boundedness, liveness, and deadlock-freeness are easier to establish.

E. Interpretation and Practical Implications

This comparative example confirms that integrating decision intelligence prior to Petri Net modeling produces a more interpretable and analytically tractable system. The proposed modeling chain preserves workflow realism while preventing state-space explosion, making it particularly suitable for crisis-aware healthcare management.

F. Computational Complexity and Runtime Considerations

The predictive layer (MLR) performs inference in $O(pd)$ per time step, where p denotes the number of predictors and d the number of patients scored at that step. The decision layer relies on a compact discrete-state representation; under the proposed rule-based approximation, policy evaluation is $O(1)$ per patient decision.

Petri Net reachability analysis is exponential in the worst case. To mitigate state-space explosion, we constrain both the initial marking and the set of policy-compliant transitions, thereby limiting the number of reachable states.

We report runtimes measured on a standard workstation for: 1) MLR inference per patient, 2) decision computation per patient, and 3) reachability graph construction for the considered markings. Finally, we discuss scalability to multi-department settings by decomposing the global Petri Net into interacting subnetworks and evaluating departments in parallel.

V. Petri Net Modeling and Analysis

Fig. 4 presents a Petri Net model designed to represent patient flow and hospital resource management under global crisis conditions. The model follows the classical Petri Net formalism, which has been widely applied to healthcare process modeling [13], [14], [16].

A. Places and Transitions

Places represent critical system states and hospital resources, including Emergency Capacity, Emergency Rate, Personnel, Beds, Operations, and Satisfaction. Each place captures the availability or status of a resource or patient group at a given time.

Transitions correspond to operational events that modify the system state, such as Emergency Visits, Admission, Treatment, Operations, and Evaluation. Transition firings consume and generate tokens, which represent patients or resources flowing through the hospital system.

The model dynamics explicitly encode causal dependencies:

- Admission depends on emergency capacity, emergency rate, population density, patient age, accident zone proximity, and distance to the hospital, reflecting a multi-criteria admission decision process.
- Treatment requires the simultaneous availability of beds and medical personnel.
- Patients requiring surgery pass through the Operations transition, mobilizing specialized resources.

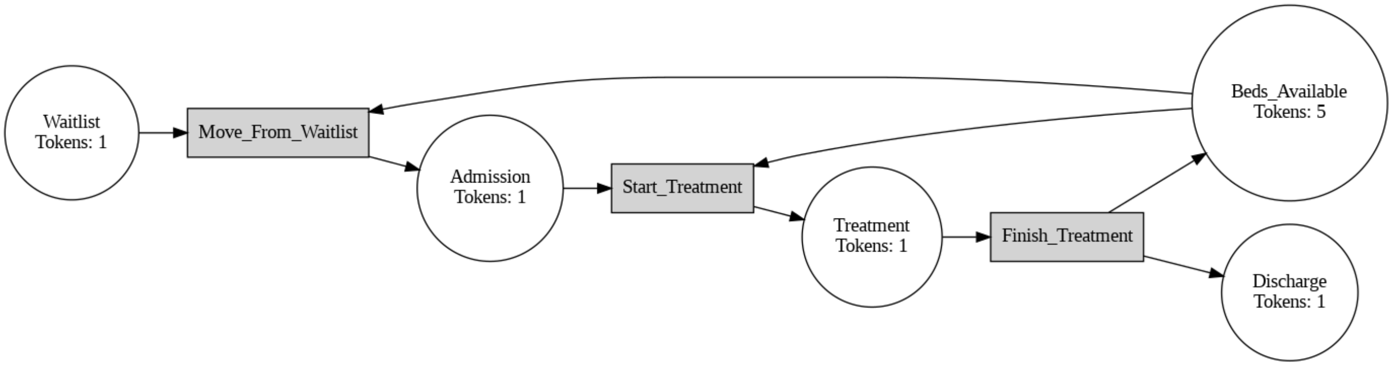

Fig. 2. Petri Net representation of patient flow derived from PDM and MDP policy.

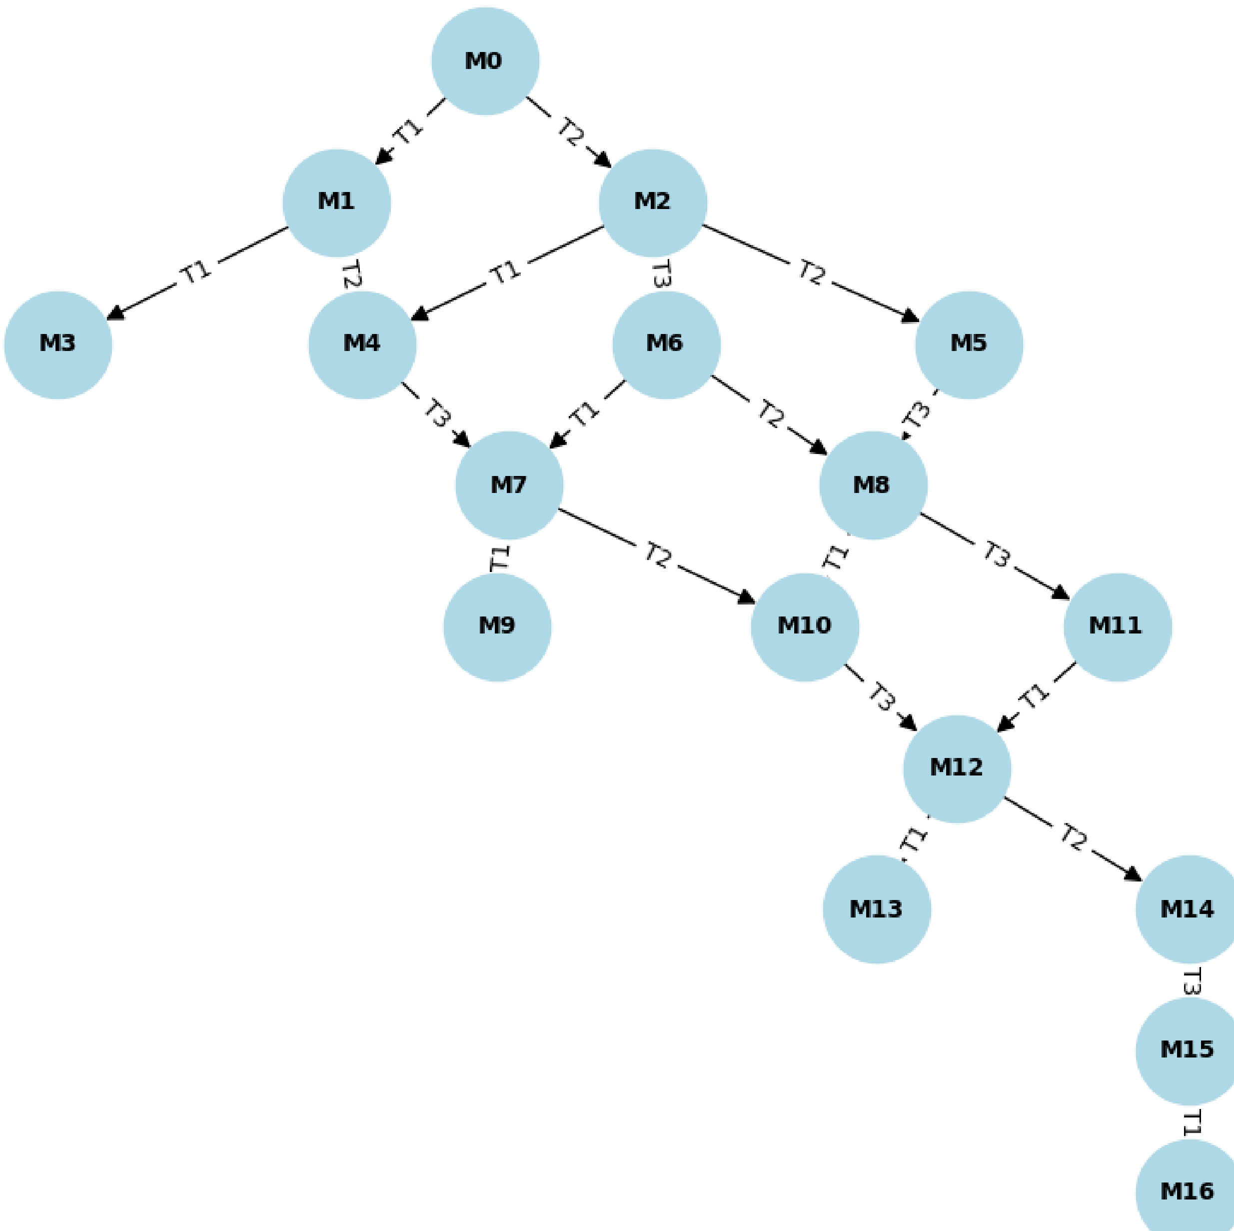

Fig. 3. Reachability graph derived from MLR → PDM → Petri Net.

- A subset of treated patients may undergo Readmission, affecting system congestion and satisfaction.
- The Evaluation transition aggregates outcomes such as occupancy, treatment completion, and readmissions to assess system performance.
- The Satisfaction place reflects patient-perceived quality of care.

B. Functional Support of the Petri Net Model

The Petri Net model supports:

- analysis of patient and resource flows,
- identification of bottlenecks and capacity limits,
- formal verification of boundedness, liveness, and deadlock-freeness,
- evaluation of prioritization and allocation strategies.

When integrated with decision and optimization models (e.g., MDPs and learning-based policies), the Petri Net provides a formal validation layer, ensuring that decision-driven hospital operations remain safe and feasible during crisis scenarios.

VI. Initial Marking and Reachability Graph Analysis

A. Initial Marking M_0

The initial marking M_0 defines the starting configuration of the hospital system. It specifies the initial distribution of tokens across places, each token representing a patient or an available resource.

$$M_0 = [1, 2, 0, 2, 0, 0, 1, 0, 1, 1, 1, 1, 1, 0, 1].$$

Each component corresponds to a specific place in the Petri Net. For instance:

- Place P1 contains one token, representing a patient ready for admission.
- Place P2 contains two tokens, representing two available beds.

This initial configuration represents a controlled and realistic hospital state, suitable for detailed reachability analysis.

B. Reachability Graph Analysis

Starting from M_0, all enabled transitions are recursively fired to generate the reachability graph, shown in Fig. 5. Each node represents a reachable marking, while edges correspond to transition firings.

To avoid combinatorial explosion, the initial marking intentionally includes a limited number of patient tokens. This ensures that all reachable states can be exhaustively analyzed while preserving graph interpretability.

The resulting reachability graph:

- captures all admissible system evolutions,
- illustrates resource consumption and release cycles,
- reaches terminal states under defined operational constraints,
- reflects realistic hospital dynamics.

These results confirm that the Petri Net model is bounded, live, and operationally valid.

1) Operational KPIs from petri net simulation: Beyond structural properties (boundedness, liveness, deadlock-freeness), we quantify: i) throughput time, ii) resource utilization, and iii) bottleneck intensity.

Let N_D be the number of completed discharges over horizon T. The throughput is:

$$\text{Throughput} = \frac{N_D}{T}. \quad (6)$$

For a resource place P_r (e.g., beds or personnel) with capacity C_r, utilization is:

$$\text{Utilization}(P_r) = \frac{1}{T}\int_0^T \frac{m_r(t)}{C_r}\,dt, \quad (7)$$

where, $m_r(t)$ is the marking of P_r at time t. A bottleneck is identified when utilization remains above a threshold (e.g., > 0.85) while queue places (e.g., waitlist) accumulate tokens.

VII. Integration with Robotic Systems

A. Illustrative Scenario

To illustrate the interaction between Multiple Linear Regression (MLR), Markov Decision Processes (MDPs), and Petri Nets (PNs), we consider a hospital facing a COVID-19 surge with limited capacity (50 beds). The objective is to optimize patient admission and resource utilization in real-time.

1) Step 1: Forecasting with MLR: The MLR model predicts the Average Length of Stay (ALOS) using patient-specific features such as age, severity, and admission type. Predicted ALOS values (e.g., 7, 4, and 2 days) estimate future bed occupancy and inform downstream decision-making.

2) Step 2: Decision-making with MDP: The MDP incorporates predicted ALOS together with the current hospital state (available beds, patient severity) to derive an admission policy. The reward function balances patient survival, throughput, and bed utilization, ensuring prioritization of critical cases while maintaining operational flow.

3) Step 3: Simulation and validation with petri nets: Decisions produced by the MDP are injected into the Petri Net model, which simulates patient and resource transitions. Places represent treatment states and resources, transitions model admission, treatment, and discharge events, and tokens represent patient flow. The Petri Net verifies that MDP-driven decisions preserve boundedness and avoid deadlock. Similar Petri Net models have been used to coordinate human–robot collaboration in surgical environments [21].

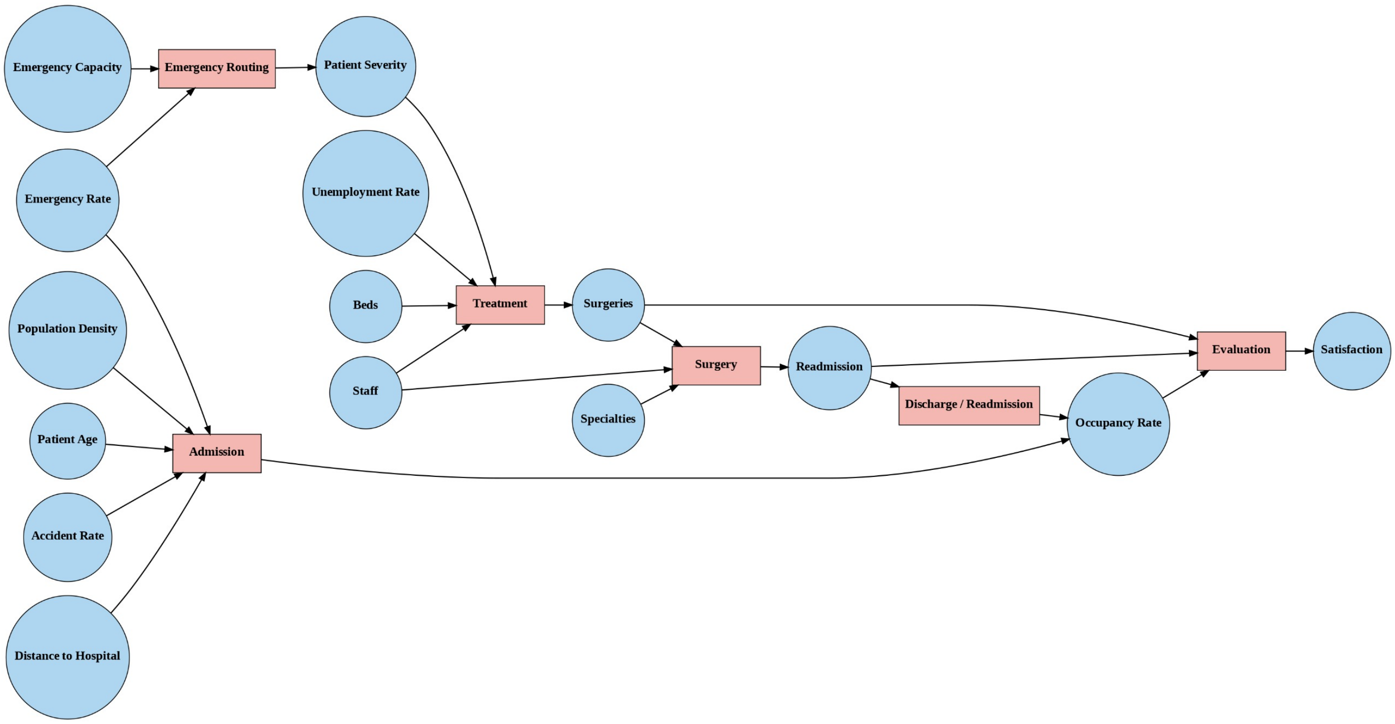

Fig. 4. Petri Net model with extended operational states and resource dependencies.

4) Integrated workflow: This pipeline combines prediction, decision optimization, and formal validation to support intelligent robotic systems for healthcare crisis management.

$$\textbf{MLR} \rightarrow \textbf{MDP} \rightarrow \textbf{Petri Net} \rightarrow \textbf{Simulation}$$

B. Priority-Based Decision Mechanism

Based on MLR predictions, patients are ranked using a Priority Score:

$$\text{Priority Score} = (\alpha \cdot \text{Severity Index}) - (\gamma \cdot \text{Predicted ALOS}),$$

where α and γ control the trade-off between clinical urgency and expected resource occupation.

Admission decisions follow:

- Admit if the Priority Score exceeds threshold θ_1,
- Waitlist if the score lies between θ_2 and θ_1,
- Reject if the score falls below θ_2.

This mechanism ensures fairness, efficiency, and transparency, enabling adaptive and explainable decision-making suitable for robotic healthcare systems.

C. Stress-Testing Prediction Errors and Policy Robustness

To evaluate robustness, we perturb the MLR outputs by additive noise: $\tilde{\ell} = \hat{\ell} + \epsilon$, $\epsilon \sim \mathcal{N}(0, \sigma^2)$, and measure policy stability as the percentage of actions unchanged under perturbations. We report: 1) decision flip rate, 2) overflow events, and 3) average congestion (waitlist tokens) as functions of σ. This analysis quantifies how forecasting uncertainty propagates through admission decisions and affects hospital dynamics.

VIII. Conclusion

This study presented an integrated computational framework combining MLR, MDPs, and Petri Nets to optimize healthcare robotics in crisis scenarios. The proposed system demonstrates strong predictive performance, supports real-time decision-making under uncertainty, and validates hospital workflows through formal modeling.

A. Key Contributions

The main contributions of this work are summarized as follows:

- Development of a unified MLR–MDP–Petri Net framework tailored to crisis-resilient healthcare robotics and hospital resource allocation.
- Integration of predictive analytics (ALOS forecasting) into robotic decision workflows, improving adaptability, bed turnover, and coordination during surge conditions.
- Validation of the framework through case studies such as ICU bed prioritization and crisis logistics, demonstrating improvements in responsiveness, efficiency, and patient outcomes.
- Comparative analysis of reachability graphs, showing that the RLM $\rightarrow$ PDM $\rightarrow$ Petri Net chain provides

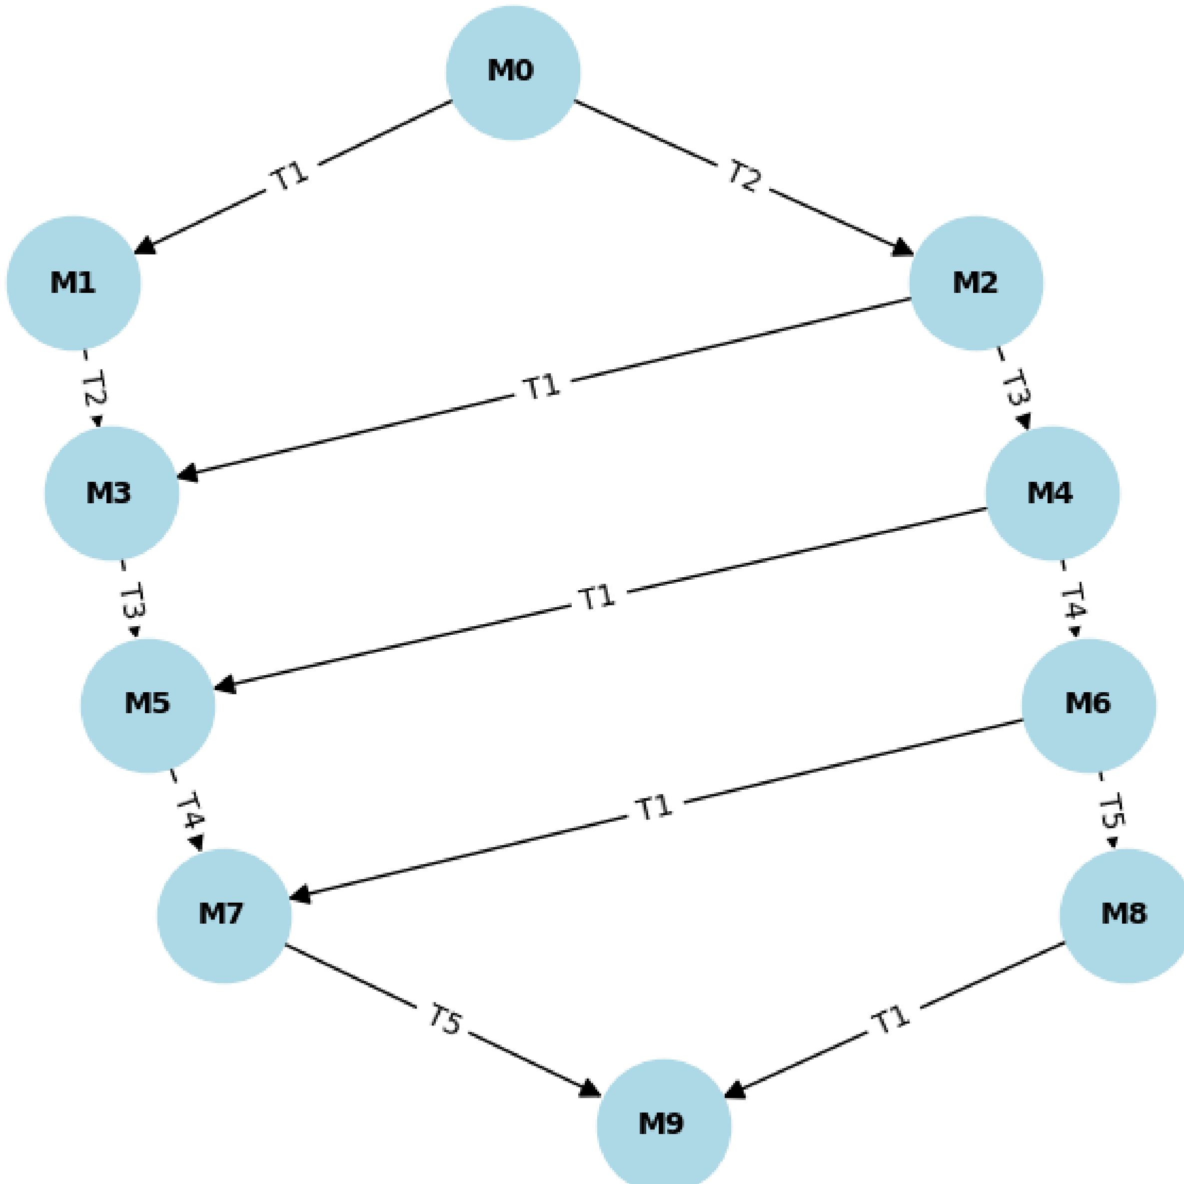

Fig. 5. Reachability graph derived from the initial marking M0, illustrating possible hospital system evolutions.

optimized decision paths while preserving analytical tractability.

Beyond these contributions, this study highlights that reachability graphs are a powerful tool to verify essential system properties such as liveness and boundedness. They provide transparency in analyzing distributed systems, hospital workflows, and robotic decision protocols under stress conditions.

B. Future Perspectives

An important direction for future work is the integration of MDPs with Deep Reinforcement Learning (DRL). Combining these techniques would allow healthcare robots to:

- Learn continuously from real-time sensor data and improve decision-making through experience.
- Adapt dynamically to unpredictable and rapidly evolving clinical environments.
- Optimize multi-agent coordination, where multiple autonomous robots collaborate to enhance hospital workflow efficiency.

This evolution in computational modeling contributes to the advancement of structured decision-support frameworks for healthcare robotics. By coherently integrating predictive modeling, decision optimization, and formal workflow validation, the proposed approach illustrates how interdisciplinary modeling techniques can enhance hospital adaptability under stress conditions. Although further empirical and large-scale validation is required, such integrated methodologies provide a foundation for improving resilience and coordination in healthcare systems facing future crisis scenarios.

Recent digital twin frameworks further support the integration of predictive and operational models in healthcare systems [28].

REFERENCES

[1] G.-Z. Yang, B. J. Nelson, R. R. Murphy, H. Choset, H. Christensen, S. H. Collins *et al.*, "Combating COVID-19—The role of robotics in managing public health and infectious diseases," *Science Robotics*, vol. 5, no. 40, p. eabb5589, 2020.

[2] A. Venkatesh and G. D. Curfman, "Preparing for the next pandemic—The role of hospital design," *New England Journal of Medicine*, vol. 382, p. e87, 2020.

[3] G.-Z. Yang, J. Cambias, K. Cleary, E. Daimler, J. Drake, P. E. Dupont *et al.*, "Medical robotics—Regulatory, ethical, and legal considerations for increasing levels of autonomy," *Science Robotics*, vol. 2, no. 4, p. eaam8638, 2017.

[4] R. Ferrer and I. Martín-Loeches, "Robotics in critical care: Toward a new era of hospital automation," *Intensive Care Medicine*, vol. 47, pp. 1161–1163, 2021.

[5] H. M. Krumholz, Y. Wang, J. A. Mattera *et al.*, "An administrative claims model suitable for profiling hospital performance," *Circulation*, vol. 113, no. 13, pp. 1693–1701, 2006.

[6] A. E. W. Johnson, T. J. Pollard, L. Shen *et al.*, "MIMIC-III, a freely accessible critical care database," *Scientific Data*, vol. 3, no. 1, pp. 1–9, 2016.

[7] M. Hauskrecht, "Value-function approximations for partially observable Markov decision processes," *Journal of Artificial Intelligence Research*, vol. 13, pp. 33–94, 2000.

[8] S. M. Shortreed, E. Laber, D. J. Lizotte, T. S. Stroup, J. Pineau, and S. A. Murphy, "Informing sequential clinical decision-making through reinforcement learning," *Medical Decision Making*, vol. 31, no. 3, pp. 353–361, 2011.

[9] M. L. Puterman, *Markov Decision Processes: Discrete Stochastic Dynamic Programming*, Hoboken, NJ, USA: Wiley, 2014.

[10] W. M. P. van der Aalst, "The application of Petri nets to workflow management," *Journal of Circuits, Systems and Computers*, vol. 8, no. 1, pp. 21–66, 1998.

[11] B. F. van Dongen, W. M. P. van der Aalst, and H. M. Verbeek, "Verification of EPCs using Petri nets," in *Business Process Management*, Berlin, Germany: Springer, 2005, pp. 172–190.

[12] M. Martins and L. Santos, "Assistive robotics for independent living: A review," *Robotics*, vol. 10, no. 3, p. 112, 2021.

[13] T. Murata, "Petri nets: Properties, analysis and applications," *Proceedings of the IEEE*, vol. 77, no. 4, pp. 541–580, 1989.

[14] W. Reisig, *Understanding Petri Nets: Modeling Techniques, Analysis Methods, Case Studies*, Berlin, Germany: Springer, 2012.

[15] D. M. Herron and M. Marohn, "The da Vinci surgical system," *Surgical Endoscopy*, vol. 15, no. 11, pp. 1345–1350, 2001.

[16] N. Salimi and E. Seidgar, "Modeling and analysis of healthcare processes using Petri nets," *Journal of Biomedical Informatics*, vol. 82, pp. 96–109, 2018.

[17] L. Wang, Y. Chen, and T. Zhang, "Petri net modeling for healthcare workflow optimization: A systematic review," *Healthcare*, vol. 10, no. 2, p. 365, 2022.

[18] W. Zhang, L. Xu, and M. Li, "Petri net-based modeling of healthcare systems under COVID-19 outbreak," *IEEE Access*, vol. 8, pp. 146383–146394, 2020.

[19] World Health Organization, "Ethical considerations in developing a public health response to pandemic influenza," WHO Technical Report Series, no. WHO/CDS/EPR/GIP/2007.2, 2020.

[20] S. Kumar and A. Singh, "Petri net simulation of critical resource allocation during COVID-19," *Simulation Modelling Practice and Theory*, vol. 107, p. 102213, 2021.

[21] B. Farahani and G.-Z. Yang, "Petri net models for coordinating human–robot collaboration in surgery," in *Proc. IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2020, pp. 2254–2260.

[22] A. Esteva, A. Robicquet, B. Ramsundar, V. Kuleshov, M. DePristo, K. Chou, C. Cui, G. S. Corrado, S. Thrun, and J. Dean, "A guide to deep learning in healthcare," *Nature Medicine*, vol. 25, pp. 24–29, 2019.

[23] B. H. Dobkin and P. W. Duncan, "Clinical trials of rehabilitation robotics with Lokomat," *Neurorehabilitation and Neural Repair*, vol. 23, no. 6, pp. 600–608, 2009.

[24] K. Zong and C. Luo, "Reinforcement learning based framework for COVID-19 resource allocation," *Computers & Industrial Engineering*, vol. 167, p. 107960, 2022, doi: 10.1016/j.cie.2022.107960.

[25] F. Soroush, B. Nabilou, A. Faramarzi, and H. Yusefzadeh, "A study of the evacuation and allocation of hospital beds during the Covid-19 epidemic: A case study in Iran," *BMC Health Services Research*, vol. 22, p. 864, 2022, doi: 10.1186/s12913-022-08286-7.

[26] R. Maleki, M. Taghizadeh-Yazdi, R. Ghasemi, *et al.*, "A Hybrid Mathematical-Simulation Approach to Hospital Beds Capacity Optimization for COVID-19 Pandemic Conditions," *Operations Research Forum*, vol. 5, p. 107, 2024, doi: 10.1007/s43069-024-00389-7.

[27] T. Le Moigne, C. Mahulea, and G. Faraut, "Optimizing Clinical Pathways: A Probabilistic Time Petri Net Approach," *IFAC-PapersOnLine*, vol. 58, no. 1, pp. 96–101, 2024, doi: 10.1016/j.ifacol.2024.07.017.

[28] G. Pellegrino, M. Gervasi, M. Angelelli, A. Corallo, *et al.*, "A Conceptual Framework for Digital Twin in Healthcare: Evidence from a Systematic Meta-Review," *Information Systems Frontiers*, vol. 27, pp. 7–32, 2024, doi: 10.1007/s10796-024-10536-4.