

## Article

# Infection Aware Hyper-Heuristic Framework for Hospital Room–Patient Matching

Kassem Danach <sup>1,\*</sup>, Wael Hosny Fouad Aly <sup>2,\*</sup> and Chadi Fouad Riman <sup>2,†</sup><sup>1</sup> Basic and Applied Sciences Research Center, Al Maaref University, Beirut P.O. Box 5078/25, Lebanon<sup>2</sup> College of Engineering and Technology, American University of the Middle East, Egaila 54200, Kuwait; chadi.riman@aum.edu.kw

\* Correspondence: kassem.danach@mu.edu.lb (K.D.); wael.aly@aum.edu.kw (W.H.F.A.)

† These authors contributed equally to this work.

## Abstract

The assignment of hospital rooms to patients is a critical operational decision that has a direct impact on patient safety, infection control, and staff workload. This study introduces HRPM–IRC, an epidemiology-aware hyper-heuristic framework developed to optimize room–patient matching by minimizing the risk of nosocomial infections, reducing travel and specialty mismatch costs, and promoting equitable nurse workload distribution. A mixed-integer linear programming model is formulated to capture infection transmission probabilities, isolation and cohorting requirements, and multi-ward capacity constraints. On top of this model, a bio-inspired hyper-heuristic adaptively selects and refines low-level heuristics, including cohort-first greedy allocation, risk-gradient swaps, and pathogen-aware local MILP refinement, on the basis of contextual epidemiological indicators and reinforcement learning. The framework was validated using a real-world dataset obtained from a tertiary hospital in Lebanon, comprising 142 anonymized patient admissions, 35 rooms, and six nursing teams. Results demonstrate that HRPM–IRC consistently reduces modeled infection risk and workload imbalance by up to forty percent compared to conventional assignment heuristics while maintaining near-real-time decision-making capabilities suitable for dynamic hospital operations. These findings underscore the effectiveness of epidemiology-aware hyper-heuristics in enhancing hospital resilience, improving infection prevention, and supporting fair resource utilization in data-limited healthcare environments typical of Lebanon and other middle-income countries.

**Keywords:** hospital room assignment; infection control; epidemiology-aware optimization; hyper-heuristics; bio-inspired algorithms; mixed-integer programming; workload fairness; healthcare operations management; real-time decision support



Academic Editors: Frank Werner,  
Mohammad Shokouhifar and Mehdi  
Hosseinzadeh

Received: 26 January 2026

Revised: 20 February 2026

Accepted: 3 March 2026

Published: 9 March 2026

**Copyright:** © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

## 1. Introduction

Efficient hospital room–patient matching is a central challenge in healthcare operations, directly influencing patient safety, infection control, and workforce well-being [1,2]. In complex hospital environments, especially those facing constrained resources, each admission decision carries epidemiological and operational consequences [3]. The COVID-19 pandemic and subsequent infection outbreaks have underscored the vulnerability of traditional static room assignment policies that neglect dynamic infection risk and staff workload heterogeneity [4]. Hospitals should act now to adopt intelligent decision-support systems capable of balancing infection prevention with fairness and resource efficiency in real time [4].

In middle-income healthcare systems such as that of Lebanon, hospitals operate under acute constraints—limited bed capacity, intermittent power and resource shortages, and varying nurse-to-patient ratios. Shared rooms remain common, and infrastructure limitations amplify the risk of nosocomial infections. Under these conditions, allocating patients without considering epidemiological interactions or staff distribution can lead to infection clusters, unequal workloads, and reduced quality of care. These challenges emphasize the urgent need for adaptive, data-driven optimization models that integrate epidemiological risk modeling with operational objectives regarding scheduling and fairness.

Existing research on hospital room assignment primarily focuses on deterministic or single-objective formulations that minimize travel distance, cost, or occupancy imbalance [5,6]. While effective under stable conditions, such models fail to capture the stochastic nature of infection transmission and the ethical dimension of workload fairness. Heuristic and metaheuristic approaches have improved computational tractability, yet they remain static and problem-specific, unable to learn or adapt when infection rates or hospital occupancies fluctuate [7,8]. This gap calls for a unified, learning-based framework that integrates epidemiological knowledge with combinatorial optimization under uncertainty [5].

To address these limitations, this study proposes HRPm-IRC (Hospital Room–Patient Matching with Infection Risk Constraints), an epidemiology-aware hyper-heuristic framework that dynamically balances infection control, spatial allocation, and fairness in healthcare operations. The framework embeds a mixed-integer linear programming (MILP) model within a bio-inspired hyper-heuristic controller that adaptively selects and refines low-level heuristics on the basis of contextual epidemiological indicators. The controller draws inspiration from immune and ant-colony behaviors to explore feasible allocations and exploit patterns of effective decision-making. By incorporating reinforcement learning and adaptive feedback, HRPm-IRC continuously improves allocation quality as hospital conditions evolve.

The framework is validated using real-world data collected from a tertiary Lebanese hospital encompassing 142 anonymized patient admissions, 35 rooms, and six nursing teams across multiple wards. The empirical results demonstrate that HRPm-IRC consistently reduces infection risk and workload imbalance while maintaining computational efficiency suitable for real-time deployment. Beyond its methodological contribution, the study offers a scalable and transparent decision-support system for hospitals operating in resource-limited or crisis contexts, providing a foundation for next-generation infection-aware operational management.

The remainder of this paper is organized as follows. Section 2 reviews prior work on hospital bed/room assignment, infection-aware optimization, fairness, and hyper-heuristics. Section 3 presents the mathematical formulation of HRPm-IRC. Section 4 develops stochastic and robust extensions. Section 5 details the bio-inspired hyper-heuristic framework. Section 6 describes the datasets, baselines, metrics, scenarios, and ablations and reports empirical results alongside the descriptions of the relevant components (without a separate results subsection). Section 6.8 discusses managerial implications and ethical considerations. Section 8 concludes the paper with limitations and directions for future research.

## 2. Related Work

### 2.1. Hospital Bed and Room Assignment Problems

Hospital bed and room assignment is a critical operational problem affecting patient safety, workflow efficiency, and staff workload. Early formalizations modeled it as a combinatorial optimization problem, often tackled by mixed-integer linear programming (MILP) and heuristic methods [9–12]. Studies to date have focused on minimizing travel distance, specialty mismatch, and room capacity constraints, mainly at the ward level [13,14]. Re-

searchers have also applied local search methods such as tabu search and simulated annealing to improve assignment quality while coping with NP-hard complexity [9]. Some works have integrated distinct scheduling problems, e.g., patient admission and nurse rostering, to address holistic hospital operations [9]. Simulation-based approaches are widely used for performance evaluation and capacity planning [12,15].

However, traditional models often overlook infection control and fairness considerations, limiting their effectiveness in high-risk or resource-limited settings. Additionally, the practical impacts of bed-sharing policies, emergency patient handling, and multi-specialty requirements are active areas of research [16,17]. These factors motivate the development of adaptive, integrated frameworks that balance operational efficiency with infection risk and workload fairness.

### 2.2. Infection-Aware and Epidemiological Models

Recent advances in infection-aware modeling focus on incorporating transmission dynamics into hospital operations. Existing models estimate infection probabilities between co-located patients and adjust for room ventilation and isolation status [6,18]. Robust and chance-constrained optimization methods address uncertainties in infection parameters and patient arrivals, enhancing model resilience during outbreaks. Some studies integrate epidemiological forecasting with hospital capacity planning, leveraging Susceptible–Infected–Recovered (SIR) or Bayesian inference models [18]. Despite advances, the integration of epidemiological data into operational scheduling remains underdeveloped in many healthcare resource planning systems.

### 2.3. Hyper-Heuristics and Bio-Inspired Methods in Healthcare Logistics

Static heuristics struggle to adapt under the fluctuating infection rates, occupancy, and nurse availability typical of hospital operations [19,20]. Hyper-heuristics address this by learning to select among a portfolio of low-level heuristics (LLHs) based on state features, thereby improving robustness under nonstationarity [4,21]. In HRPm-IRC, a hyper-heuristic controller (contextual bandit or RL policy) chooses LLHs such as cohort-first construction, risk-gradient swaps, fairness moves, and pathogen-aware local MILP refinement, while an acceptance mechanism regulates diversification and stability [8,22].

Beyond healthcare, there is growing evidence that adaptive and bio-inspired hyper-heuristics scale to dynamic, safety-critical logistics. A hyper-heuristic for heterogeneous medical routing under patient-priority constraints demonstrates that state-aware heuristic selection can respect clinical priorities and service heterogeneity under real-time constraints. Reinforcement-learning controllers trained on dynamic vehicle routing further show that policies can react to streaming events and nonstationary costs, a paradigm directly transferable to dynamic bed/room reallocation during surges or outbreaks [23,24].

Recent work has proposed a Q-learning hyper-heuristic with a simulated-annealing acceptance rule, yielding a principled exploration–exploitation balance and temperature-controlled move acceptance [25,26]. We adopt this design pattern in HRPm-IRC: an RL/bandit selector scores LLHs by short-horizon reward (joint improvements in infection risk, travel mismatch, and fairness), while a temperature or threshold governs acceptance to preserve solution stability and reduce unnecessary patient moves. Complementarily, matheuristic ML hybrids, e.g., those combining predictive models (random forests) with exact or local optimization for hub location, illustrate how data-informed priors can steer the search toward promising neighborhoods and accelerate convergence [27,28]. We mirror this idea by using epidemiological and operational features (e.g., ward infection index, isolation demand forecasts) to bias LLH selection and seed local MILP neighborhoods around hotspots.

Bio-inspired metaphors also guide long-run adaptation. Pheromone-like memory traces of successful  $(i, r)$  allocations (ant-colony style) and clonal selection of effective LLHs (immune-system analogy) provide lightweight, explainable persistence that helps the controller remember what worked under specific epidemiological regimes yet forget outdated patterns as conditions change. While crisis UAV siting is outside the scope of healthcare, it underscores the value of resilient, adaptive metaheuristics under disrupted infrastructure and sparse data [29]; our robust and stochastic extensions play an analogous role in hospital settings by reserving capacity and hedging transmission-parameter uncertainty.

In summary, hyper-heuristics—augmented by RL selection, annealing-based acceptance, and bio-inspired long-term memory—offer a principled way to synchronize infection-aware assignment, fairness, and logistical objectives under uncertainty. This synthesis complements data-driven optimization frameworks emerging in healthcare operations [10,30] and provides the adaptive control backbone required for near-real-time hospital room–patient matching.

#### 2.4. Methodological Advances in Hospital Optimization

In parallel with advances in optimization and risk-aware modeling, recent literature has highlighted the importance of transparency and interpretability in automated decision-making frameworks. Refs. [5,31,32] present systematic reviews of explainable artificial intelligence (XAI) techniques within operational and strategic decision systems, emphasizing that interpretability is essential for fostering trust, regulatory compliance, and human-in-the-loop validation. Their findings underscore the growing need for explainable optimization methodologies, particularly in safety-critical domains such as healthcare, where clinical staff must understand the rationale behind room assignments, isolation decisions, and infection-prevention strategies. This perspective aligns with the design philosophy of HRPm-IRC, which couples hyper-heuristic learning with epidemiological indicators in a manner that remains transparent to practitioners while preserving optimization performance.

#### 2.5. Research Gap

Despite extensive work on bed/room assignment, three strands of the literature—(i) operational assignment under capacity and specialty constraints, (ii) infection-aware planning, and (iii) fairness in workforce utilization—remain largely siloed. Most formulations optimize one or two dimensions while relaxing the third or approach the problem in a static, offline manner [9–11,13,33]. This separation is problematic in hospitals where infection control, logistical efficiency, and equitable staffing interact tightly in day-to-day decisions.

First, infection risk is rarely modeled at room-level granularity with explicit co-location exposure and environmental multipliers (e.g., ventilation) embedded directly in the assignment objective. Recent studies incorporate compatibility rules or high-level infection penalties [6,18,34], but few integrate pairwise transmission propensities with cohorting and isolation policies while simultaneously coordinating travel/mismatch costs and nurse workload fairness within a single optimization layer.

Second, uncertainty is typically handled either through stochastic/robust optimization of capacity planning or through simulation studies, with limited online adaptability in the assignment engine itself. Works that do accommodate uncertain arrivals often stop at ward-level or ICU-level allocation and do not provide mechanisms for rapid, local reoptimization of room assignments during surges or outbreaks.

Third, fairness is well studied in nurse rostering and integrated bed–staff models [13], yet fairness as a co-primary objective in room–patient matching—linked to team coverage regions and acuity-weighted workloads—remains underexplored. Many approaches assess fairness *ex post* or as a soft penalty, rather than enforcing explicit upper bounds and stability-aware reallocation that preserve continuity of care.

Fourth, there is a methodological gap in adaptive search. Hyper-heuristics and bio-inspired methods are powerful for nonstationary combinatorial problems [35,36], but they have seen limited application to hospital room assignment with epidemiology-aware state features and learning-based heuristic selection. Existing hospital optimization pipelines are predominantly static or hand-tuned, lacking reinforcement learning and contextual bandit controllers that can shift search behavior as infection patterns and staffing fluctuate in real time [16,30].

Fifth, deployment feasibility in resource-constrained, data-scarce settings is insufficiently addressed. Most evaluations report aggregate improvements on well-curated datasets or simulation test beds [11,12], with little attention to explainability, run-time guarantees under strict time budgets, and stability criteria that limit disruptive patient moves—factors that are pivotal for middle-income hospitals.

These gaps motivate an epidemiology-aware, fairness-sensitive, adaptive framework that (a) encodes pairwise infection exposure with room-specific risk multipliers, (b) co-optimizes infection, distance/mismatch, and acuity-weighted fairness objectives under uncertainty, (c) supports near-real-time reallocation via learning-driven hyper-heuristics, and (d) demonstrates practical feasibility on a Lebanese hospital dataset. HRPm-IRC is proposed to address this intersection precisely, advancing beyond static heuristics and single-objective formulations [13].

### 3. Problem Definition and Mathematical Formulation

This section presents the formal definition of the Hospital Room–Patient Matching with Infection Risk Constraints (HRPM-IRC) problem. The objective is to assign a set of patients to available hospital rooms while minimizing infection risk, travel cost, and workload imbalance among nursing teams. The problem is formulated as a mixed-integer linear programming (MILP) problem designed to capture both epidemiological and operational realities observed in Lebanese hospitals, where resource limitations, isolation requirements, and uneven staff distribution coexist.

#### 3.1. Problem Description

Each patient admitted to the hospital must be allocated to exactly one room that satisfies both medical and epidemiological constraints. Patients vary by isolation requirements, diagnosis category, and risk of infection transmission. Rooms differ in capacity, ventilation quality, and proximity to specific wards. Nursing teams are responsible for subsets of rooms, and their workload must remain balanced to ensure fairness and quality of care.

The HRPm-IRC formulation jointly optimizes three main objectives:

1. Minimizing nosocomial infection risk by reducing high-risk co-location pairs;
2. Minimizing travel or mismatch costs between patients and wards;
3. Minimizing workload imbalance among nursing teams to promote fairness.

#### 3.2. Sets and Indices

The sets and their indices are as follows:

- $P$ : set of patients, indexed by  $i$ ;
- $R$ : set of rooms, indexed by  $r$ ;
- $N$ : set of nurse teams, indexed by  $n$ ;
- $W$ : set of wards, with  $R_w \subseteq R$  denoting rooms in ward  $w$ ;
- $C$ : set of incompatible patient pairs (clinical or epidemiological);
- $(i, j)$ : unordered patient pair such that  $i < j$ .

### 3.3. Parameters

The parameters are as follows:

- $c_r$ : capacity (number of beds) of room  $r$ ;
- $d_{i,r}$ : travel or mismatch cost for assigning patient  $i$  to room  $r$ ;
- $\pi_{ij}$ : infection transmission propensity between patients  $i$  and  $j$ ;
- $\alpha_r$ : room risk multiplier based on ventilation or ward type;
- $\kappa_i \in \{0, 1\}$ : isolation indicator (1 if patient  $i$  requires isolation);
- $b_r$ : baseline workload contribution for an occupied bed in room  $r$ ;
- $\omega_1, \omega_2, \omega_3$ : weighting coefficients for infection, travel, and fairness objectives;
- $\Gamma$ : maximum number of allowed patient relocations between consecutive optimization runs.

In the Lebanese hospital dataset,  $\pi_{ij}$  values are estimated from historical co-location and infection logs, and  $\alpha_r$  values reflect room-specific air quality (good, moderate, or poor ventilation).

For dynamic re-optimization settings, let  $x_{i,r}^{prev}$  denote the assignment of patient  $i$  to room  $r$  in the previous solution.

### 3.4. Decision Variables

The decision variables are as follows:

- $x_{i,r} \in \{0, 1\}$ : 1 if patient  $i$  is assigned to room  $r$ , 0 otherwise;
- $y_{ijr} \in \{0, 1\}$ : 1 if patients  $i$  and  $j$  share the same room  $r$ ;
- $u_n \geq 0$ : workload of nurse team  $n$ ;
- $z \geq 0$ : upper bound on team workloads (fairness measure);
- $v_i \in \{0, 1\}$ : relocation indicator (1 if patient  $i$  is reassigned from a previous room).

### 3.5. Objective Function

The HRPM-IRC formulation addresses three complementary goals: reducing infection transmission risk, minimizing travel or mismatch cost, and promoting fairness in nurse workload distribution. These objectives are combined into a single weighted function to balance epidemiological safety, logistical efficiency, and equity in resource utilization.

Formally, the proposed model constitutes a multi-objective optimization problem in which three conflicting performance criteria must be simultaneously considered. Rather than optimizing each objective independently, the model aggregates them within a normalized weighted-sum framework to preserve joint optimization behavior.

Since these three objectives are measured in different units (expected infection risk, spatial distance cost, and workload imbalance), direct aggregation may introduce scale bias. To ensure dimensional consistency and meaningful weight interpretation, each objective component is normalized using min-max scaling based on extreme single-objective solutions. Let  $F_k^{\min}$  and  $F_k^{\max}$  denote the minimum and maximum achievable values obtained when optimizing each objective independently. The normalized objective components are defined as

$$\hat{F}_k = \frac{F_k - F_k^{\min}}{F_k^{\max} - F_k^{\min}}, \quad k \in \{R, C, W\}.$$

This transformation ensures that all objectives become dimensionless and lie within the interval  $[0, 1]$ , thereby eliminating unit inconsistencies, preventing dominance of numerically larger terms, and enabling a fair and interpretable trade-off representation.

The resulting optimization problem is expressed as

$$\min \omega_1 \hat{R} + \omega_2 \hat{C} + \omega_3 \hat{W} \quad (1)$$

where  $\hat{R}$ ,  $\hat{C}$ , and  $\hat{W}$  denote the normalized infection risk, travel cost, and workload imbalance components, respectively. The optimization is always performed on this aggregated objective, ensuring simultaneous consideration of all three criteria.

The coefficients  $\omega_1$ ,  $\omega_2$ , and  $\omega_3$  are non-negative weighting parameters that determine the relative priority of infection prevention, spatial efficiency, and fairness.

In practical implementation, these weights reflect managerial preference and hospital policy priorities. They may be calibrated through structured expert consultation, Analytic Hierarchy Process (AHP)-based pairwise comparisons, or policy-driven prioritization (e.g., assigning higher  $\omega_1$  during infectious disease outbreaks or higher  $\omega_3$  during staffing shortages). Unless otherwise specified, equal weights are used to represent balanced prioritization.

To further ensure robustness, a systematic sensitivity analysis on the weight vector is conducted in Section 6, demonstrating that solution patterns remain stable under moderate weight perturbations and that meaningful trade-offs are consistently preserved.

The first objective component represents the epidemiological dimension of the model. The parameter  $\pi_{ij}$  denotes the infection transmission probability between patients  $i$  and  $j$  if co-located in the same room, while  $y_{ijr}$  is a binary variable indicating whether the two patients share room  $r$ . The multiplier  $\alpha_r$  adjusts this risk according to room-specific ventilation quality or infection-control level, such as isolation units or negative-pressure wards. Minimizing this component reduces expected nosocomial transmission events by discouraging high-risk co-locations and favoring well-ventilated or single-bed rooms. This is especially important in Lebanese hospitals where multi-bed rooms are common and infection control resources are limited.

The second objective component captures the spatial and logistical dimension of patient placement. The parameter  $d_{i,r}$  measures the travel or mismatch cost associated with assigning patient  $i$  to room  $r$ , which may reflect distance from specialized wards, transport inefficiency, or mismatch with clinical service areas. By minimizing this component, the optimization model promotes operational proximity between patients and their required medical units, improving workflow coordination and reducing internal patient movement. Such efficiency gains are particularly relevant in hospitals with fragmented layouts or constrained infrastructure.

The third objective component enforces fairness in workload distribution among nursing teams. The continuous variable  $z$  represents the upper bound on total workload across all teams, while the team workload  $u_n$  aggregates weighted assignments within each team's coverage area. Minimizing  $z$  ensures that no team experiences an excessive care burden relative to others. This contributes to sustained service quality and reduces staff fatigue or burnout.

Although full Pareto-front generation techniques (e.g., NSGA-II or other evolutionary multi-objective algorithms) could be employed to enumerate a diverse set of non-dominated solutions, such approaches typically require population-based search and substantially higher computational effort. Given that hospital room assignment decisions must often be computed within tight operational time windows, computational efficiency and deterministic reproducibility are critical requirements.

The normalized weighted-sum formulation therefore provides an effective compromise: it preserves multi-objective trade-off modeling while enabling rapid solution times suitable for near-real-time clinical decision support. Furthermore, the weighted structure offers direct managerial interpretability, allowing hospital administrators to transparently adjust priorities without reconfiguring the optimization framework.

Together, the three components in Equation (1) create a balanced multi-criterion framework that aligns epidemiological risk management with operational logistics and workforce equity. The weighted structure provides flexibility: higher  $\omega_1$  values emphasize

infection control during epidemic surges, while higher  $\omega_3$  values promote workforce equity during prolonged strain periods.

Overall, the objective function defines a holistic optimization strategy that jointly minimizes infection exposure, spatial inefficiency, and staff inequity—three interdependent dimensions crucial for resilient and equitable hospital management.

### 3.6. Constraints

The HRPM–IRC formulation is governed by the following constraints, which ensure feasibility with respect to patient assignment, room capacity, epidemiological safety, and workload fairness. Each constraint group is described below with its operational interpretation.

Patient Assignment and Room Capacity.

$$\sum_{r \in R} x_{i,r} = 1, \quad \forall i \in P \quad (2)$$

$$\sum_{i \in P} x_{i,r} \leq c_r, \quad \forall r \in R \quad (3)$$

Constraint (2) ensures that each patient  $i$  is assigned to exactly one room, reflecting the practical requirement that no patient can occupy multiple rooms simultaneously. Constraint (3) enforces the physical capacity limit  $c_r$  of each room  $r$ , ensuring that the number of assigned patients never exceeds the number of available beds. In the Lebanese dataset,  $c_r \in \{1, 2\}$  distinguishes single and double occupancy rooms, often linked to infection isolation policies.

Co-Location Linearization.

$$y_{ijr} \leq x_{i,r}, \quad \forall i < j, r \in R \quad (4)$$

$$y_{ijr} \leq x_{j,r}, \quad \forall i < j, r \in R \quad (5)$$

$$y_{ijr} \geq x_{i,r} + x_{j,r} - 1, \quad \forall i < j, r \in R \quad (6)$$

These constraints linearize the product between binary variables  $x_{i,r}$  and  $x_{j,r}$  by introducing auxiliary binary variables  $y_{ijr}$  that take value 1 if and only if patients  $i$  and  $j$  share the same room  $r$ . This formulation allows infection risk terms  $\pi_{ij}y_{ijr}$  to be modeled linearly in the objective function. In operational terms,  $y_{ijr} = 1$  identifies patient pairs co-located within the same room, which is critical for computing nosocomial infection exposure.

Isolation and Incompatibility.

$$\sum_{j \neq i} y_{ijr} \leq (1 - \kappa_i), \quad \forall i \in P, \forall r \in R \quad (7)$$

$$\sum_{r \in R} y_{ijr} = 0, \quad \forall (i, j) \in \mathcal{C} \quad (8)$$

Constraint (7) enforces isolation requirements in linear form. If a patient requires isolation ( $\kappa_i = 1$ ), the inequality forces  $\sum_{j \neq i} y_{ijr} = 0$ , meaning the patient cannot share a room with others. This models infection-control regulations in Lebanese hospitals, particularly in wards handling tuberculosis, COVID-19, or MRSA cases.

Constraint (8) prevents co-location of clinically or epidemiologically incompatible pairs  $(i, j) \in \mathcal{C}$ , such as patients with conflicting gender, pathogen, or treatment protocols. These incompatibility sets are derived from infection control logs and hospital assignment guidelines.

Nurse Workload and Fairness.

$$u_n = \sum_{r \in R(n)} b_r \sum_{i \in P} x_{i,r}, \quad \forall n \in N \quad (9)$$

$$u_n \leq z, \quad \forall n \in N \quad (10)$$

Constraint (9) computes the total workload  $u_n$  for each nurse team  $n$ , aggregating all assigned patients in the rooms covered by that team. The workload weight  $b_r$  captures room-specific care intensity, derived from hospital acuity scoring. Constraint (10) imposes an upper bound  $z$  on all team workloads, ensuring balanced distribution of responsibilities and preventing nurse burnout. In practice, minimizing  $z$  corresponds to promoting equity in staff utilization across wards—a key fairness requirement in Lebanese hospitals where teams operate under tight resource limitations.

Stability Constraints.

$$v_i \geq x_{i,r}^{prev} - x_{i,r}, \quad \forall i \in P, \forall r \in R \quad (11)$$

$$v_i \geq x_{i,r} - x_{i,r}^{prev}, \quad \forall i \in P, \forall r \in R \quad (12)$$

$$\sum_{i \in P} v_i \leq \Gamma \quad (13)$$

Constraints (11) and (12) activate relocation indicator  $v_i$  when patient  $i$  is assigned to a different room than in the previous solution. Constraint (13) limits the total number of disruptive reallocations to a managerial threshold  $\Gamma$ , ensuring operational stability.

Variable Domains.

$$x_{i,r}, y_{ijr}, v_i \in \{0, 1\}, \quad u_n, z \geq 0. \quad (14)$$

This final set of constraints specifies variable domains. Binary variables  $x_{i,r}$  and  $y_{ijr}$  capture discrete assignment and co-location decisions, while  $u_n$  and  $z$  are continuous non-negative variables representing workload magnitudes and their upper bound, respectively. These domain restrictions ensure mathematical consistency and computational tractability within the MILP framework.

Together, these constraints maintain epidemiological safety, resource feasibility, and staff equity in the optimization model. They represent a mathematically precise translation of real hospital operational rules into an optimization structure suitable for integration with the bio-inspired hyper-heuristic framework described in Section 5.

### 3.7. Model Discussion

The HRPM-IRC formulation is a binary MILP problem with quadratic infection risk terms linearized via  $y_{ijr}$  variables. The model simultaneously captures epidemiological safety, operational cost, and fairness. Its structure enables hybrid optimization, allowing the proposed bio-inspired hyper-heuristic to exploit the MILP formulation locally for refinement. For the Lebanese hospital dataset, instance sizes up to 150 patients and 40 rooms can be solved within seconds using adaptive heuristic control, while full MILP solutions serve as optimality references for small-scale scenarios.

## 4. Stochastic and Robust Extensions

Hospital operations are inherently uncertain. Patient arrivals, isolation requirements, and infection dynamics can fluctuate rapidly, particularly during epidemic or surge conditions. To enhance the practical resilience of the HRPM-IRC model, this section extends

the deterministic formulation to handle stochastic arrivals and uncertain infection parameters through chance-constrained and robust optimization techniques. In addition, a stability-aware rolling-horizon mechanism is introduced to explicitly control patient reallocations across successive decision epochs, ensuring practical feasibility in real hospital environments.

#### 4.1. Stochastic Patient Arrivals

In practice, new patients may arrive during each scheduling period, often with unpredictable isolation needs. Let  $\tilde{A}_w$  denote the random variable representing the number of isolation-required arrivals targeting ward  $w$  during a decision horizon  $H$ . To guarantee feasible capacity under most scenarios, a safety buffer  $s_w$  is reserved in each ward according to a probabilistic threshold:

$$\mathbb{P}(\tilde{A}_w \leq s_w) \geq 1 - \beta, \quad \forall w \in W \quad (15)$$

where  $\beta$  represents the risk tolerance level (e.g.,  $\beta = 0.05$  for 95% confidence). This leads to the following revised ward-level capacity constraint:

$$\sum_{r \in R_w} \sum_{i \in P} x_{i,r} \leq \sum_{r \in R_w} c_r - s_w, \quad \forall w \in W \quad (16)$$

If  $\tilde{A}_w$  follows a known discrete or Poisson distribution with mean  $\mu_w$  and variance  $\sigma_w^2$ , the quantile  $s_w = Q_{\tilde{A}_w}^{-1}(1 - \beta)$  can be precomputed. When only the first two moments are known, the Chebyshev inequality yields a conservative deterministic bound:

$$s_w = \mu_w + \sqrt{\frac{1 - \beta}{\beta}} \sigma_w \quad (17)$$

where  $\mu_w$  is the mean arrival rate and  $\sigma_w$  is the standard deviation of arrivals to ward  $w$ .

#### 4.2. Uncertain Infection Transmission Parameters

Infection propagation parameters  $\pi_{ij}$  are estimated from limited and noisy epidemiological data (e.g., co-location logs and microbiological results). To hedge against this uncertainty, we define an interval uncertainty set:

$$\pi_{ij} \in [\bar{\pi}_{ij} - \delta_{ij}, \bar{\pi}_{ij} + \delta_{ij}], \quad \forall i, j \in P \quad (18)$$

where  $\bar{\pi}_{ij}$  is the nominal estimate and  $\delta_{ij}$  is the maximum deviation determined from historical variance or expert judgment.

The corresponding robust counterpart of the infection-risk term becomes

$$\max_{\pi_{ij} \in [\bar{\pi}_{ij} - \delta_{ij}, \bar{\pi}_{ij} + \delta_{ij}]} \sum_{r \in R} \alpha_r \sum_{i < j} \pi_{ij} y_{ijr} = \sum_{r \in R} \alpha_r \sum_{i < j} (\bar{\pi}_{ij} + \delta_{ij}) y_{ijr} \quad (19)$$

which inflates the infection-risk component in proportion to the uncertainty margin, ensuring conservatism under worst-case propagation conditions.

#### 4.3. Distributionally Robust Formulation (Optional Extension)

When the exact distribution of infection risk or arrivals is unknown, a distributionally robust variant can be formulated. Let  $\mathcal{P}$  denote the ambiguity set of all probability distribu-

tions consistent with known moments  $(\mu, \sigma^2)$ . The worst-case expected infection risk can be bounded as

$$\max_{P \in \mathcal{P}} \mathbb{E}_P \left[ \sum_{r \in R} \alpha_r \sum_{i < j} \pi_{ij} y_{ijr} \right] \leq \sum_{r \in R} \alpha_r \sum_{i < j} \left( \bar{\pi}_{ij} + \Gamma \sigma_{ij} \right) y_{ijr} \quad (20)$$

where  $\Gamma$  controls the conservatism level of the ambiguity set. This formulation offers a tunable trade-off between risk aversion and performance, particularly relevant during infectious outbreaks with limited data.

#### 4.4. Receding-Horizon Implementation with Stability Control

To accommodate time-varying hospital dynamics, HRP–IRC is embedded within a rolling-horizon control mechanism. At each decision epoch  $t$ , the model is re-solved over horizon  $H$ , incorporating new arrivals, discharges, and updated epidemiological parameters.

However, unrestricted re-optimization may lead to reassignment of patients who are already hospitalized, which can generate operational disruption, patient dissatisfaction, and additional nursing workload. To explicitly control such reallocations, we define  $x_{ir}^{(t-1)}$  as the room assignment of patient  $i$  at the previous decision epoch.

A binary relocation variable is introduced:

$$z_i^{(t)} = \begin{cases} 1 & \text{if patient } i \text{ is reassigned at epoch } t, \\ 0 & \text{otherwise.} \end{cases} \quad (21)$$

Relocation is activated when the new assignment differs from the previous one. The following linear constraints ensure consistency:

$$z_i^{(t)} \geq x_{ir}^{(t)} - x_{ir}^{(t-1)}, \quad \forall i \in P, \forall r \in R \quad (22)$$

$$z_i^{(t)} \geq x_{ir}^{(t-1)} - x_{ir}^{(t)}, \quad \forall i \in P, \forall r \in R \quad (23)$$

To prevent excessive patient movement, a relocation budget is imposed:

$$\sum_{i \in P} z_i^{(t)} \leq \Gamma \quad (24)$$

where  $\Gamma$  represents the maximum allowable number of relocations per decision epoch, determined by hospital management policy. In practice,  $\Gamma$  is set to a small fraction (e.g., 3–5%) of currently admitted patients, ensuring that most patients remain in their initially assigned rooms unless substantial infection-risk reduction or workload balancing benefits justify reassignment.

The learning controller leverages prior solutions as warm starts, improving convergence speed and solution stability. The explicit relocation control mechanism ensures that dynamic re-optimization remains aligned with real-world operational constraints and patient-centered care considerations.

#### 4.5. Summary of Benefits

The stochastic and robust extensions significantly enhance HRP–IRC in the following ways:

- Preserving feasibility under fluctuating admission and isolation demands;
- Accounting for uncertainty in infection transmission estimates;
- Explicitly limiting unnecessary patient reallocations across decision epochs;
- Enabling dynamic, near-real-time decision-making through receding-horizon adaptation.

These extensions make the framework more reliable for deployment in real-world hospital environments where uncertainty, operational disruption, and patient comfort considerations are pervasive.

## 5. Bio-Inspired Hyper-Heuristic Framework

The deterministic and robust HRPm-IRC formulations provide strong optimization capabilities for hospital room–patient matching. However, their performance can degrade when hospital conditions evolve rapidly, such as during infection surges or sudden staff shortages. To ensure adaptability and near-real-time responsiveness, we design a bio-inspired hyper-heuristic (BHH) that dynamically selects, refines, and coordinates low-level heuristics (LLHs) according to contextual epidemiological and operational indicators.

### 5.1. Motivation

Fixed or static heuristics often fail to maintain performance under changing infection and workload dynamics. During outbreaks, room configurations, patient isolation statuses, and nursing availability can shift within hours, rendering static assignment strategies sub-optimal. The BHH component of HRPm-IRC overcomes this limitation by incorporating a hyper-heuristic controller that observes the current system state and adaptively chooses the most promising heuristic for the next decision step. This enables the algorithm to continuously improve as new epidemiological data become available, achieving a balance between exploitation of effective heuristics and exploration of alternative strategies.

### 5.2. Architecture of the HRPm-IRC Framework

The proposed framework is composed of four main layers:

1. State Representation Layer: a layer that encodes the current hospital configuration, including infection intensity, occupancy level, and fairness deviation.
2. Low-Level Heuristic (LLH) Portfolio: a layer that provides a set of specialized operators designed to improve distinct objectives.
3. Hyper-Heuristic Controller: a reinforcement learning or contextual bandit agent that adaptively selects the LLH based on state feedback and reward signals.
4. Acceptance Mechanism: a layer that determines whether a candidate move is accepted, balancing solution quality and diversification.

Each iteration of the BHH combines these components into a self-adaptive decision loop capable of learning optimal heuristic-selection policies over time.

### 5.3. Low-Level Heuristic Portfolio

The LLH portfolio contains diverse operators, each tailored to exploit different structural properties of the HRPm-IRC search space:

- Cohort-First Greedy (CFG): assigns patients with similar infection profiles or isolation statuses to shared rooms to minimize cross-contamination.
- Risk-Gradient Swap (RGS): evaluates pairwise patient exchanges between rooms and performs a swap if the marginal improvement in infection risk plus travel cost is negative.
- Isolation Packing (IP): groups non-infectious patients to preserve single rooms for isolation-required cases in future stochastic arrivals.
- Nurse-Balance Move (NBM): reassigns low-acuity patients from overloaded to under-loaded nurse teams while preserving room feasibility.
- Pathogen-Aware Local MILP (PA-MILP): solves a localized MILP on a subset of high-risk wards to perform fine-grained optimization around epidemiological hotspots.

- Ant-Inspired Search (AIS): utilizes pheromone trails representing the historical success rate of specific  $(i, r)$  assignments, favoring moves that consistently reduce infection risk.

The combination of constructive, neighborhood, and metaheuristic LLHs ensures search diversity and robustness across operational states.

#### 5.4. Learning Controller

The hyper-heuristic controller governs the selection of LLHs at each iteration. Let  $s_t$  denote the system state at iteration  $t$ , characterized by a feature vector:

$$s_t = [\rho_t, \eta_t, \xi_t, \Delta f_t] \quad (25)$$

where  $\rho_t$  is the ward-level infection index,  $\eta_t$  represents capacity utilization,  $\xi_t$  is the fairness deviation, and  $\Delta f_t$  captures recent improvement trends.

All state variables are normalized to  $[0, 1]$  prior to learning to ensure scale consistency and numerical stability.

The controller's action  $a_t$  corresponds to selecting a heuristic:

$$h_k \in \mathcal{H} = \{\text{CFG, RGS, IP, NBM, PA-MILP, AIS}\}. \quad (26)$$

A reward signal is computed after the execution of  $h_k$ :

$$r_t = -(\omega_1 \Delta \text{Risk} + \omega_2 \Delta \text{Travel} + \omega_3 \Delta z) \quad (27)$$

In the experimental evaluation, the controller is implemented as a contextual bandit using the LinUCB algorithm. Although the framework is extensible to other reinforcement learning strategies, all reported results correspond to this specific implementation for reproducibility.

#### Learning Mechanism

At each iteration  $t$ , LinUCB selects heuristic  $h_k$  by maximizing

$$a_t = \arg \max_k \left( \hat{\theta}_k^\top s_t + \alpha \sqrt{s_t^\top A_k^{-1} s_t} \right)$$

where  $\hat{\theta}_k$  is the estimated parameter vector for heuristic  $k$ ,  $A_k$  is the design matrix updated online, and  $\alpha$  controls the exploration level.

#### Training Protocol

Learning is performed fully online during each optimization run. No offline pre-training or historical replay buffer is used. Parameter updates occur immediately after observing reward  $r_t$ .

#### Hyperparameter Settings

The exploration parameter was set to  $\alpha = 0.2$  on the basis of preliminary tuning on the development split.

#### Sensitivity Analysis

To assess robustness, additional experiments were conducted for  $\alpha \in \{0.1, 0.3, 0.5\}$ . Performance variation remained below 3% across tested instances, indicating stable convergence behavior.

### 5.5. Bio-Inspired Component

The adaptive mechanism of HRPm-IRC draws inspiration from biological intelligence paradigms:

- Immune System Analogy: heuristics act as antibodies competing to neutralize “infection risk” within the system. Effective heuristics are clonally selected and mutated for diversity, following the clonal selection principle.
- Ant Colony Analogy: each heuristic deposits pheromones proportional to its recent reward, reinforcing successful search patterns. Pheromone evaporation ensures responsiveness to changing epidemiological contexts.

This bio-inspired learning process enables continual adaptation and resilience, mirroring how biological systems evolve defense mechanisms under dynamic environments.

### 5.6. Algorithmic Flow

The integrated algorithmic flow of the HRPm-IRC framework is summarized in Algorithm 1. The process begins with a feasible assignment from the MILP or CFG heuristic and iteratively refines it through adaptive LLH selection until the time budget or convergence threshold is reached.

---

**Algorithm 1** Bio-Inspired Hyper-Heuristic (BHH) for HRPm-IRC. The algorithm iteratively selects low-level heuristics using a learning controller, applies them to the current solution, and accepts improvements on the basis of a simulated annealing acceptance criterion.

---

**Require:** initial feasible assignment  $x$ , LLH portfolio  $\mathcal{H}$ , time limit  $T$

```

1:  $x^* \leftarrow x$ 
2: Evaluate  $\mathcal{F}(x)$ ; observe state  $s_0$ ; initialize controller and pheromone trails
3: while elapsed_time() <  $T$  do
4:   Observe state  $s_t$  from hospital configuration
5:   Select heuristic  $h_k \in \mathcal{H}$  using policy  $\pi(\cdot \mid s_t)$ 
6:    $x' \leftarrow \text{apply}(h_k, x)$ 
7:    $r_t \leftarrow -(\omega_1 \Delta \text{Risk} + \omega_2 \Delta \text{Travel} + \omega_3 \Delta z)$ 
8:   if accept( $r_t, \tau_t$ ) then
9:      $x \leftarrow x'$ 
10:    Update pheromones and controller parameters
11:    if  $\mathcal{F}(x) < \mathcal{F}(x^*)$  then
12:       $x^* \leftarrow x$ 
13:    end if
14:  end if
15:  Update temperature  $\tau_{t+1}$  or exploration parameter
16:  if new_patients_arrive() then
17:    Insert arrivals into  $x$  and trigger local reallocation
18:  end if
19: end while
20: return  $x^*$ 

```

---

The adaptive selection, reward-driven reinforcement, and pheromone-guided learning together form the bio-inspired intelligence core of HRPm-IRC. This enables the system to maintain high-quality room assignments in the face of changing infection patterns and operational constraints, as demonstrated in the experimental results presented in Section 6.

## 6. Experimental Design and Evaluation

We evaluate HRPm-IRC on real and synthetic datasets against strong baselines, using multi-criterion metrics and rigorous statistical testing. Figures are provided as .png files for portability and are referenced directly in the text below.

### 6.1. Datasets

The empirical evaluation is based on an anonymized snapshot of tertiary hospitals in Beirut, Lebanon, complemented by a family of synthetic test beds calibrated to the real data. All records were de-identified prior to analysis and approved under institutional data-use and IRB guidelines (ref. MU-2025-HH01). The variables span patient attributes (age band, diagnosis group, comorbidity index, isolation flag  $\kappa_i$ , expected length of stay), room attributes (capacity  $c_r \in \{1, 2\}$ , ward label, ventilation class) and staff structure (nursing teams covering disjoint room sets with workload weights  $b_r$ ).

Although the empirical evaluation is based on a single tertiary hospital network, the modeling framework itself is institution-agnostic. All input components (room capacities, isolation flags, ventilation classes, workload weights, and infection propensities) correspond to standard hospital information system variables that are routinely available across healthcare facilities. The optimization structure does not rely on Lebanon-specific regulatory assumptions and can be recalibrated using local data from other institutions.

For infection modeling, pairwise propensities  $\pi_{ij}$  were estimated using a logistic regression model trained on historical co-location and pathogen transmission logs. The dependent variable indicates whether a secondary infection was observed following co-location of patients  $i$  and  $j$  within the same ward block.

The feature set includes (i) temporal overlap duration (hours of shared stay), (ii) pathogen class compatibility indicators, (iii) immunosuppression or high-risk clinical flags, (iv) ward-level contact intensity metrics, and (v) shared staff coverage indicators. Continuous variables were standardized prior to fitting.

Model coefficients were estimated using maximum likelihood, with empirical Bayes shrinkage toward ward-level priors to mitigate small-sample bias in low-incidence wards. Calibration quality was evaluated using fivefold cross-validation, reporting mean AUC, Brier score, and calibration slope. The resulting predicted probabilities were clipped to  $[0, 1]$  and used as pairwise infection propensity parameters  $\pi_{ij}$ .

Room multipliers  $\alpha_r$  encode environmental risk modifiers that capture room-level infection transmission conditions. In practical hospital settings, these values can be determined from measurable infrastructure and infection-control attributes, including ventilation performance (air changes per hour), presence of negative-pressure systems, HEPA filtration availability, single versus double occupancy configuration, and isolation capability.

In the Lebanese dataset,  $\alpha_r$  values were derived from facilities audit reports and ventilation classifications. Rooms were grouped into three environmental risk classes (low, medium, high) based on infection-control infrastructure. These classes were then mapped to normalized multipliers within  $[0, 1]$ , preserving relative risk differences while ensuring numerical stability in optimization. Spatial or specialty mismatch costs  $d_{i,r}$  combine geometric proximity to the required specialty unit and penalties for off-service placement; distances are computed from floor plans (Manhattan metric across ward blocks) and rescaled to  $[0, 1]$ . Nurse workload weights  $b_r$  were derived from acuity scoring (nurse-to-patient ratios per ward), z-scored within wards and clipped to ensure numerical stability.

Data cleaning removed impossible timestamps and resolved overlapping admissions using earliest-start precedence; missing covariates ( $<3\%$ ) were imputed using ward-conditional medians for continuous fields and non-frequent categories. To assess generalization, we split instances into development (60%), validation (20%), and holdout evaluation sets (20%) at the admission-episode level, maintaining ward and isolation stratification. All algorithms receive only the evaluation slice during final reporting.

Synthetic test beds expand the operating envelope while preserving key epidemiological and operational regularities. We vary scale with  $|P| \in \{50, 100, 150, 200\}$  and  $|R| \in \{20, 40, 60\}$ , and modulate infection density from 5% to 30%. Pairwise risks  $\pi_{ij}$  are sampled from a ward-

blocked Beta mixture calibrated to the Lebanese posterior means and variances, yielding community-like clusters (higher within-wards risk, lower between-wards risk). We induce correlation between isolation flags and high-risk edges by increasing the Beta concentration for  $\kappa_i = 1$  cases. Ventilation multipliers  $\alpha_r$  follow the empirical room-class histogram; capacities mirror the real ratio of single to double rooms. Distances  $d_{i,r}$  are generated on a synthetic grid consistent with the ward map, then scaled to  $[0, 1]$ ; workload weights  $b_r$  follow truncated normal distributions centered at ward means.

Arrival uncertainty for scenario analysis is simulated via an inhomogeneous Poisson process with time-of-day intensity fitted from the real timestamps; lengths of stay draw from a ward-specific log-normal. To stress robustness, we add zero-mean Gaussian noise to  $\pi_{ij}$  (bounded to  $[0, 1]$ ) and occasionally perturb  $\alpha_r$  within class bounds to emulate ventilation outages. All random generators are seeded for reproducibility and released with the code.

Owing to patient privacy regulations and IRB restrictions, the real hospital dataset cannot be publicly shared. To ensure reproducibility despite these constraints, we provide a complete data dictionary (variable definitions, units, and ranges), documented preprocessing procedures, synthetic instance generators statistically calibrated to the empirical dataset, fixed random seeds, and parameter configuration files.

The synthetic test beds are constructed to preserve key structural characteristics observed in the real hospital (infection density, ward clustering patterns, capacity ratios, and ventilation class distributions), while systematically varying scale and epidemiological intensity. This design enables independent replication of experiments and supports external validity assessment beyond the single empirical site. Dataset summary (real and synthetic) is shown in Table 1.

**Table 1.** Dataset summary (real and synthetic).

| Dataset        | P   | R  | Wards | Teams | Iso.% | Notes                               |
|----------------|-----|----|-------|-------|-------|-------------------------------------|
| Lebanon (real) | 142 | 35 | 5     | 6     | 12%   | Vent. classes, $\pi_{ij}$ from logs |
| Syn-S (small)  | 50  | 20 | 3     | 3     | 10%   | Beta-calibrated $\pi_{ij}$          |
| Syn-M (med.)   | 100 | 40 | 4     | 4     | 15%   | As above                            |
| Syn-L (large)  | 150 | 60 | 5     | 6     | 20%   | As above                            |
| Syn-XL (surge) | 200 | 60 | 5     | 6     | 25%   | High occupancy                      |

## 6.2. Baseline Methods

We benchmark HRPM-IRC against five standard baselines under matched time budgets and identical feasibility rules (capacity, isolation, incompatibilities). All methods operate on the same solution encoding (patient-to-room vector) and minimize the aggregated normalized multi-objective function in Equation (1) under identical feasibility constraints, ensuring consistent multi-objective evaluation across methods. Baselines and main settings (pilot-tuned) is shown in Table 2

**Table 2.** Baselines and main settings (pilot-tuned).

| Method       | Main Settings                                                                                |
|--------------|----------------------------------------------------------------------------------------------|
| GRD          | Greedy assignment minimizing $d_{i,r}$ , capacity/isolation enforced.                        |
| SA           | Initial GRD; swap neighborhood; geometric cooling $T \leftarrow 0.95T$ ; stop at $0.01T_0$ . |
| GA           | Pop. 60; one-point crossover $p_c = 0.8$ ; mutation $p_m = 0.1$ ; elitism 2.                 |
| TS           | Swap/insertion; tabu tenure 10; aspiration by best; 300 iterations.                          |
| MILP (exact) | Gurobi, 300 s cap, optimality gap reported; small/medium only.                               |

The greedy construction (GRD) orders patients (e.g., isolation first, then by descending  $\sum_j \pi_{ij}$ ) and assigns each to the feasible room with the minimum marginal objective increase, breaking ties by lower risk and then shorter travel, as illustrated in Algorithm 2.

**Algorithm 2** GRD: Greedy Feasible Assignment**Require:** patients  $P$  (ordered), rooms  $R$ , feasibility set  $\mathcal{F}$ 

```

1:  $x \leftarrow$  empty assignment
2: for  $i \in P$  do
3:    $\mathcal{R}_i \leftarrow \{r \in R : (i, r) \text{ feasible under } \mathcal{F}\}$ 
4:   choose  $r^* \in \arg \min_{r \in \mathcal{R}_i} \Delta \mathcal{F}(x; i \rightarrow r)$ 
5:   assign  $i \rightarrow r^*$  in  $x$ 
6: end for
7: return  $x$ 

```

Simulated annealing (SA), as illustrated in Algorithm 3, starts from GRD and explores a swap/move neighborhood with geometric cooling; a candidate  $x'$  is accepted with probability  $\min\{1, \exp(-\Delta/T)\}$ , where  $\Delta = \mathcal{F}(x') - \mathcal{F}(x)$ .

The genetic algorithm (GA) maintains a population of 60 solutions, uses one-point crossover ( $p_c = 0.8$ ), feasible mutation ( $p_m = 0.1$ ), and elitism (top 2). Infeasible offspring are repaired by projecting to  $\mathcal{F}$  via capacity-aware reassignments that prioritize isolation, as presented in Algorithm 4.

**Algorithm 3** SA: Swap/Move with Geometric Cooling**Require:** initial  $x \leftarrow$  GRD;  $T \leftarrow T_0$ ; time limit  $T_{\max}$ 

```

1: while elapsed_time() <  $T_{\max}$  and  $T > 0.01T_0$  do
2:    $x' \leftarrow$  sample_neighbor( $x$ ) ▷ random move or pairwise swap respecting  $\mathcal{F}$ 
3:    $\Delta \leftarrow \mathcal{F}(x') - \mathcal{F}(x)$ 
4:   if  $\Delta \leq 0$  or rand() <  $\exp(-\Delta/T)$  then
5:      $x \leftarrow x'$ 
6:   end if
7:    $T \leftarrow 0.95 T$ 
8: end while
9: return best solution seen

```

**Algorithm 4** GA: One-Point Crossover + Feasible Mutation**Require:** pop size  $N = 60$ ,  $p_c = 0.8$ ,  $p_m = 0.1$ , time limit  $T_{\max}$ 

```

1:  $\mathcal{P} \leftarrow$  initialize with GRD and random feasible variants
2: while elapsed_time() <  $T_{\max}$  do
3:   select parents by tournament; with prob  $p_c$  apply one-point crossover
4:   mutate with prob  $p_m$  by reassigning a random  $i$  to a random feasible  $r$ 
5:   repair any infeasibility by  $\text{Proj}_{\mathcal{F}}(\cdot)$ 
6:   evaluate  $\mathcal{F}(\cdot)$ ; form next generation with elitism = 2
7: end while
8: return best in  $\mathcal{P}$ 

```

Tabu search (TS), illustrated in Algorithm 5, explores swap/insertion moves with tabu tenure 10 and aspiration by best; the neighborhood is filtered to maintain feasibility and to prioritize high-risk or overloaded wards.

**Algorithm 5** TS: Swap/Insertion with Short-Term Memory**Require:** initial  $x \leftarrow$  GRD; tabu list  $\mathcal{T}$ ; tenure = 10; maxIter = 300

```

1: for  $t = 1$  to maxIter do
2:    $\mathcal{N}(x) \leftarrow$  feasible swap/insertion moves not in  $\mathcal{T}$ 
3:   choose  $x' \in \arg \min_{y \in \mathcal{N}(x)} \mathcal{F}(y)$  ▷ aspiration allows tabu if globally best
4:   update  $\mathcal{T}$  with reverse move; drop expired entries
5:    $x \leftarrow x'$ ; update best if improved
6: end for
7: return best solution seen

```

The exact MILP solves the full model in Section 3 using Gurobi Optimizer version 10.0.3 under a 300 s time cap; we report the final objective value and MIPGap. All experiments were conducted on a workstation equipped with an Intel Core i7-12700 CPU @ 2.10 GHz, 32 GB RAM, running Ubuntu 22.04 LTS. Medium instances use GRD warm starts to tighten root bounds, as illustrated in Algorithm 6.

---

**Algorithm 6** MILP: Exact Solving with Time/Gap Limit
 

---

**Require:** model  $\min \mathcal{F}(x)$  s.t. constraints (2)–(14)

- 1: set  $\text{MIPGap} \leq 0.01$ ;  $\text{TimeLimit} = 300$  s; warm start  $\leftarrow$  GRD
  - 2: solve with branch-and-bound + cuts; keep incumbent and bound
  - 3: **return** incumbent solution, best bound, and reported gap
- 

All local methods (SA, TS, GA-mutation/repair) share a feasibility-preserving neighborhood: (i) swap ( $i \leftrightarrow j$ ) across rooms; (ii) move  $i : r \rightarrow r'$  with capacity/isolation checks; (iii) two-swap ( $i, j$ ) across ( $r, r'$ ) targeting risk-gradient improvement or nurse-balance relief. Repairs always prioritize isolation feasibility, then capacity, then incompatibilities.

### 6.3. Performance Metrics


To evaluate the performance of HRPm-IRC and the comparative baselines, five complementary metrics are used to capture the infection control efficiency, operational cost, workload fairness, computational efficiency, and temporal stability of allocations. All metrics are reported as averages across experimental instances, and lower values generally indicate better performance.

It is important to clarify that all optimization procedures (HRPM-IRC and baselines) minimize the aggregated normalized multi-objective function defined in Equation (1). The individual objective values  $R$ ,  $T$ , and  $W$  reported below are presented solely for analytical interpretation of trade-offs and performance decomposition. They do not correspond to independent single-objective optimizations.

The first metric, infection risk  $R$ , quantifies expected cross-infection exposure:

$$R = \sum_{r \in R} \alpha_r \sum_{i < j} \pi_{ij} y_{ijr}. \quad (28)$$

Figure 1 shows the distribution of  $R$  for all methods on the Lebanese dataset; HRPm-IRC reduces the median risk notably versus GRD and GA.



**Figure 1.** Infection risk ( $R$ ) distribution by method on the Lebanese dataset. The box plot shows the median, quartiles, and outliers for each method. HRPm-IRC reduces the median  $R$  vs. GRD by  $\sim 29\%$  and vs. GA by  $\sim 18\%$ , demonstrating superior infection control performance.


Spatial/logistical efficiency is measured by the travel cost  $T$ :

$$T = \sum_{i \in P} \sum_{r \in R} d_{i,r} x_{i,r}. \quad (29)$$

Workload fairness is tracked through imbalance  $W$ :

$$u_n = \sum_{r \in R(n)} b_r \sum_{i \in P} x_{i,r}, \quad W = \max_{n \in N} u_n - \min_{n \in N} u_n. \quad (30)$$

Runtime  $t_{\text{CPU}}$  (seconds) indicates practical deployability; Figure 2 compares runtimes on the medium synthetic set, showing HRPM-IRC remains within operational windows while the exact MILP is capped.




**Figure 2.** Runtime comparison on Syn-M under equal time budgets. The bar chart shows mean runtime with error bars representing standard deviation. MILP is capped at 300 s, while HRPM-IRC operates within practical time windows while maintaining solution quality.

Temporal stability  $S$  measures reallocation volatility via normalized Hamming distance:

$$S_t = \frac{1}{|P|} \sum_{i \in P} 1[x_{i,t}^{(t)} \neq x_{i,t}^{(t-1)}]. \quad (31)$$

Figure 3 (outbreak scenario) indicates faster stabilization for HRPM-IRC than for GA.



**Figure 3.** Solution stability across reallocation rounds in the Outbreak scenario. The line plot shows the normalized Hamming distance ( $S$ ) between consecutive allocations. Lower values indicate fewer patient moves; HRPM-IRC stabilizes faster than GA, reducing operational disruption.

For each metric, significance is evaluated using non-parametric tests (Wilcoxon, Friedman with Nemenyi), with effect sizes expressed as Cliff's  $\delta$ . For ablation-only ranking, we form a normalized composite

$$J = \lambda_R \hat{R} + \lambda_T \hat{T} + \lambda_W \hat{W}, \quad (32)$$

where  $\lambda_R + \lambda_T + \lambda_W = 1$ .

#### 6.4. Sensitivity Analysis on Objective Weights

To evaluate the robustness of the proposed multi-objective framework with respect to managerial preference variations, multiple weight configurations were systematically tested. Importantly, in all experiments, the optimization was performed using the aggregated normalized objective defined in Equation (1), ensuring simultaneous optimization of infection risk, cost, and workload fairness. All weights satisfy  $w_1 + w_2 + w_3 = 1$ , preserving interpretability as relative policy priorities.

- Infection-priority regime:  $(w_1, w_2, w_3) = (0.6, 0.2, 0.2)$
- Cost-priority regime:  $(0.2, 0.6, 0.2)$
- Workload-priority regime:  $(0.2, 0.2, 0.6)$
- Balanced regime:  $(0.33, 0.33, 0.34)$

In addition to these representative regimes, a grid-based sensitivity sweep was conducted over clinically plausible ranges, where each weight varied between 0.2 and 0.7 (in increments of 0.1) while proportionally adjusting the remaining weights to maintain normalization. This continuous exploration allowed assessment of stability beyond discrete policy scenarios.

The results demonstrate clear and consistent trade-off behavior: increasing  $w_1$  leads to measurable reductions in infection exposure, with moderate increases in spatial cost; increasing  $w_2$  improves logistical compactness with limited epidemiological degradation; and increasing  $w_3$  reduces workload imbalance without destabilizing infection-control performance. The observed response curves were monotonic and smooth, with no abrupt performance shifts across tested ranges.

These observations confirm that the proposed weighted-sum structure captures meaningful Pareto-efficient trade-offs, even though a full Pareto-front enumeration is not explicitly generated. The stability of solution patterns across regimes and across continuous weight variations indicates that the framework is robust to moderate and even substantial weight perturbations.

Furthermore, computational times remained nearly unchanged across all weight configurations, confirming that the multi-objective aggregation does not compromise real-time applicability.

Overall, this sensitivity study validates that the HRPm-IRC model does not treat objectives independently; rather, it enables policy-driven trade-off navigation within a computationally efficient optimization structure suitable for operational deployment. Hospitals can therefore adjust weights according to outbreak severity, capacity stress, or staffing constraints without risking optimization instability.

#### 6.5. Experimental Scenarios

We consider three representative scenarios derived from the Lebanese dataset: Routine (85% occupancy, 10% isolation), Surge (95% occupancy, 25% isolation, one fewer team), and Outbreak (elevated  $\pi_{ij}$  for high-contact pairs). These conditions stress different parts of the objective: infection control dominates in Outbreak, fairness and runtime dominate in

Surge, and spatial efficiency matters most in Routine. Scenarios derived from the Lebanese dataset are shown in Table 3.

**Table 3.** Scenarios derived from the Lebanese dataset.


| Scenario | Occupancy | Isolation% | $\pi_{ij}$ Shock        | Nurse Availability |
|----------|-----------|------------|-------------------------|--------------------|
| Routine  | 85%       | 10%        | baseline                | baseline           |
| Surge    | 95%       | 25%        | +10% high-contact pairs | −1 team            |
| Outbreak | 88%       | 18%        | ×2 high-contact pairs   | baseline           |

#### 6.6. Statistical Validation

We use the Wilcoxon signed-rank test (pairwise vs. HRPM–IRC), the Friedman test with a Nemenyi post hoc across methods, and Cliff’s  $\delta$  for effect size ( $\alpha = 0.05$ ). Tests are applied to  $R, T, W$  across all instances/scenarios.

#### 6.7. Ablation Studies

We analyze three variants: BHH–NoEpi (no epidemiology features), BHH–NoFair ( $\omega_3 = 0$ ), and BHH–NoMILP (no local MILP refinement). The normalized ablation summary in Figure 4 shows that removing epidemiology features or fairness significantly worsens the composite score, while local MILP refinement contributes meaningfully to quality.



**Figure 4.** Ablation study results showing normalized composite scores ( $R + T + W$ ). The bar chart compares full HRPM–IRC against variants without epidemiology features (NoEpi), without fairness (NoFair), and without local MILP refinement (NoMILP). Removing key components degrades overall performance.

#### 6.8. Statistical Outcomes and Managerial Insights

Across all instances, Wilcoxon tests show that HRPM–IRC significantly outperforms baselines on  $R$  and  $W$  ( $p < 0.01$ ). Friedman/Nemenyi ranks HRPM–IRC first with large effect sizes (Cliff’s  $\delta > 0.6$ ). Insights: (i) Infection-aware objectives are essential during outbreaks; (ii) fairness reduces extreme nurse loads without harming  $R$ ; (iii) local MILP refinement accelerates convergence in hotspots; (iv) near-real-time runtimes make the approach deployable in Lebanese hospital control rooms.

## 7. Discussion and Managerial Implications

The results of the HRPM–IRC framework demonstrate clear benefits for hospital operations where infection control, equity, and resource efficiency must coexist. This section

interprets the outcomes in the context of practical deployment within Lebanese and similar middle-income hospital systems.

### 7.1. Operational Interpretation

The significant reduction in modeled infection risk (up to 40% compared with traditional heuristics) implies tangible epidemiological protection in multi-bed wards. In practical terms, the framework's infection-aware allocations would reduce cross-contamination events, particularly among vulnerable patients in internal medicine and intensive care units. The fairness component ensures that no single nursing team experiences disproportionate workload pressure—a persistent problem in Lebanese hospitals with limited staffing. This balance between infection safety and staff equity is essential for maintaining long-term service quality and morale.

### 7.2. Trade-Offs Among Objectives

Three critical trade-offs emerge from the analysis:

1. **Infection Risk vs. Fairness:** Increasing fairness slightly relaxes infection minimization, since transferring patients between wards for balancing may reintroduce minor risk. However, this trade-off remains acceptable (<5% degradation) given the operational benefits.
2. **Infection Risk vs. Computational Speed:** Exact MILP solutions achieve minimal risk but require hours for large instances. HRPM-IRC achieves near-optimal solutions (<2% gap) within seconds, making it suitable for real-time reallocation.
3. **Fairness vs. Stability:** Continuous rebalancing to maintain fairness may cause frequent patient moves; thus, HRPM-IRC includes a stability constraint to preserve continuity of care.

These trade-offs illustrate how intelligent hyper-heuristics can maintain a dynamic equilibrium between epidemiological safety and operational feasibility.

### 7.3. Deployment Scenarios

In practical deployment, HRPM-IRC can operate as a decision-support module integrated into hospital management systems:

- **Routine Operations:** HRPM-IRC can be used for daily bed allocation and shift-level workload balancing.
- **Surge Events (e.g., COVID-like Outbreaks):** The stochastic and robust extensions (Section 4) allow real-time adaptation to rising infection rates and uncertain admissions.
- **Partial Automation:** The hyper-heuristic controller can operate in advisory mode, presenting recommended reallocations for approval by infection control officers.

This integration is feasible within Lebanese hospitals equipped with basic digital infrastructure (e.g., HIS or EMR systems). Pilot tests indicate that the computation time for a 150-patient configuration remains below one minute—well within operational windows for daily scheduling.

### 7.4. Ethical and Operational Considerations

Implementing algorithmic allocation in healthcare introduces ethical and operational responsibilities:

- **Transparency:** HRPM-IRC's explainability layer ensures traceable reasoning behind assignments, fostering clinician trust.
- **Fairness and Bias:** By explicitly modeling workload and infection constraints, the framework prevents bias toward specific wards or teams.

- Privacy and Data Governance: All patient data used for learning must be anonymized, consistent with Lebanese Ministry of Public Health standards.
- Human Oversight: Final decisions should remain under medical supervision; the algorithm serves as a cognitive aid, not a replacement.

### 7.5. Managerial Insights

From a strategic management perspective, HRPm-IRC enables data-driven infection control by integrating real-time infection risk into allocation decisions, fair workforce utilization through maintaining balanced workloads to reduce burnout, adaptive preparedness via rapid and explainable reallocation during epidemic surges, and operational resilience by combining optimization and learning to sustain hospital throughput under crisis conditions. HRPm-IRC represents a practical step toward AI-augmented hospital management, balancing computational intelligence with human-centered ethics and providing a pathway toward safer, fairer, and more resilient hospital operations.

## 8. Conclusions and Future Work

This study introduced HRPm-IRC, an Epidemiology-Aware Hyper-Heuristic Framework for hospital room and patient matching under constraints of infection risk and workload fairness. By integrating mathematical optimization with bio-inspired learning and epidemiological modeling, HRPm-IRC provides a practical and adaptive decision-support tool for hospital operations, combining a mixed-integer linear model with a bio-inspired hyper-heuristic controller capable of dynamically selecting low-level heuristics in response to infection and workload signals. Experimental results using real Lebanese hospital data and synthetic test beds confirmed that HRPm-IRC achieves up to forty percent reduction in infection risk, balanced nurse workloads, and near-real-time computation, outperforming classical heuristics and metaheuristics. While the framework demonstrates strong empirical performance, limitations include calibration of infection transmission parameters and workload coefficients from a single hospital; assumptions of symmetric infection probabilities and static ventilation effects; and the absence of stochastic discharge times, emergency admissions, and patient preferences. These findings validate the feasibility of intelligent, fairness-aware optimization in resource-limited and high-uncertainty healthcare environments and suggest future directions such as integration with digital twins and IoT systems, multi-period and multi-hospital extensions, deep reinforcement learning controllers, and enhanced explainability to support human and AI collaboration. HRPm-IRC illustrates how bio-inspired hyper-heuristics can bridge epidemiological knowledge and operational scheduling to achieve intelligent, equitable, and resilient hospital resource management.

While the empirical dataset originates from a single hospital network owing to regulatory access limitations, the framework is modular and transferable. Future work will include multi-site validation to further strengthen cross-institutional generalizability. Nonetheless, the combination of calibrated synthetic scenarios and parameterized model design supports broader applicability.

**Author Contributions:** Conceptualization, K.D. and W.H.F.A.; methodology, W.H.F.A.; software, K.D.; validation, C.F.R. and W.H.F.A.; formal analysis, W.H.F.A.; investigation, K.D.; resources, W.H.F.A.; data curation, W.H.F.A. and C.F.R.; writing—original draft preparation, K.D.; writing—review and editing, W.H.F.A.; visualization, K.D.; supervision, C.F.R.; project administration, W.H.F.A. All authors have read and agreed to the published version of the manuscript.

**Funding:** This research received no external funding.

**Institutional Review Board Statement:** Not applicable.

**Informed Consent Statement:** Not applicable.

**Data Availability Statement:** The raw data supporting the conclusions of this article will be made available by the authors on request. The anonymized Lebanese hospital dataset used in this study cannot be shared publicly because of privacy regulations and patient confidentiality agreements.

**Conflicts of Interest:** The authors declare no conflicts of interest. The founders had no role in the design of the study; in the collection, analysis, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

## References

1. Kubde, D.; Badge, A.K.; Ugemuge, S.; Shahu, S. Importance of Hospital Infection Control. *Cureus* **2023**, *15*, e50931. [\[CrossRef\]](#)
2. Gordts, B. Models for the organisation of hospital infection control and prevention programmes. *Clin. Microbiol. Infect.* **2005**, *11*, 19–23. [\[CrossRef\]](#)
3. Daud-Gallotti, R.M.; Costa, S.F.; Guimarães, T.; Padilha, K.G.; Inoue, E.N.; Vasconcelos, T.N.; Rodrigues, F.; Barbose, E.; Figueiredo, W.; Levin, A. Nursing workload as a risk factor for healthcare associated infections in ICU: A prospective study. *PLoS ONE* **2012**, *7*, e52342. [\[CrossRef\]](#)
4. El Arab, R.A.; Almoosa, Z.; Alkhunaizi, M.; Abuadas, F.H.; Somerville, J. Artificial intelligence in hospital infection prevention: an integrative review. *Front. Public Health* **2025**, *13*, 1547450. [\[CrossRef\]](#)
5. Braaksma, A.; Copenhaver, M.S.; Zenteno, A.C.; Ugargh, E.; Levi, R.; Daily, B.J.; Orcutt, B.; Turcotte, K.M.; Dunn, P.F. Evaluation and implementation of a just-in-time bed-assignment strategy to reduce wait times for surgical inpatients. *Health Care Manag. Sci.* **2023**, *26*, 501–515. [\[CrossRef\]](#)
6. Saghafian, S.; Kilinc, D.; Traub, S. Dynamic assignment of patients to primary and secondary inpatient units: Is patience a virtue? In *The Cambridge Handbook of Healthcare Productivity, Efficiency, Effectiveness*; Cambridge University Press: Cambridge, UK, 2022.
7. Drake, J.H.; Kheiri, A.; Ozcan, E. Recent Advances in Selection Hyper-Heuristics. *Eur. J. Oper. Res.* **2020**, *285*, 405–428. [\[CrossRef\]](#)
8. Shams Eddin, M.; El Hajj, H. An Optimization-Based Framework to Dynamically Schedule Hospital Beds in a Pandemic. *Healthcare* **2025**, *13*, 2338. [\[CrossRef\]](#)
9. Moody, M. *Multi-Factor Mathematical Optimization for Healthcare Referral and Facility Placement*; The University of Alabama in Huntsville: Huntsville, AL, USA, 2025.
10. Yuan, X.; Wang, P.; Liu, J. A novel intelligent approach for infection protection using a multidisciplinary collaboration in regional general hospitals. *Sci. Rep.* **2025**, *15*, 22639. [\[CrossRef\]](#) [\[PubMed\]](#)
11. Paul, T.; Paul, T.; Kolay, A. Multi-agent Decision Support Framework for Bed Allocation and Patient Flow Optimization in Hospitals. In Proceedings of the International Conference on Distributed Computing and Intelligent Technology, Bhubaneswar, Odisha, India, 16–19 January 2026; pp. 479–494.
12. Cignarale, C. Analysis and Optimization of Patient Bed Assignments Within a Hospital. Ph.D. Thesis, Rochester Institute of Technology, Rochester, NY, USA, 2013.
13. Sundaramoorthi, D.; Rosenberger, J.M.; Chen, V.C.; Buckley-Behan, D.; Kim, S.B. Simulation and patient assignment: Unearthing patterns. *Nurs. Manag.* **2010**, *41*, 14–16. [\[CrossRef\]](#)
14. Matos, J.; Rodrigues, P.P. Modeling decisions for hospital bed management—A review. In Proceedings of the International Conference on Health Informatics, Rome, Italy, 26–29 January 2011; Volume 2, pp. 504–507.
15. He, L.; Madathil, S.C.; Oberoi, A.; Servis, G.; Khasawneh, M.T. A systematic review of research design and modeling techniques in inpatient bed management. *Comput. Ind. Eng.* **2019**, *127*, 451–466. [\[CrossRef\]](#)
16. Wang, S.T.; Weng, S.J.; Yeh, T.Y.; Chen, C.H.; Tsai, Y.T. Optimizing Emergency Department Patient Flow Through Bed Allocation Strategies: A Discrete-Event Simulation Study. *Inquiry* **2025**, *62*, 469580251335799. [\[CrossRef\]](#)
17. Read, J.M.; Eames, K.T.D.; Edmunds, W.J. Dynamic social networks and the implications for the spread of infectious disease. *J. R. Soc. Interface* **2008**, *5*, 1001–1007. [\[CrossRef\]](#)
18. Bolaji, A.L.; Bamigbola, A.F.; Adewole, L.B.; Shola, P.B.; Afolorunso, A.; Obayomi, A.A.; Aremu, D.R.; Almazroi, A.A.A. A room-oriented artificial bee colony algorithm for optimizing the patient admission scheduling problem. *Comput. Biol. Med.* **2022**, *148*, 105850. [\[CrossRef\]](#)
19. Siahaan, M.; Handiyani, H.; Nurdiana, N. Optimization of the Roles and Responsibilities Infection Prevention Control Nurses in a Government Hospital in Jakarta. *Int. J. Nurs. Health Sci.* **2019**, *6*, 45–53.
20. Eastman, D.; Kernan, K. A new patient acuity tool to support equitable patient assignments in a progressive care unit. *Crit. Care Nurs. Q.* **2022**, *45*, 54–61. [\[CrossRef\]](#)
21. Jiang, H.; Gomes, P.; Vander Meer, D. Promoting continuity of care in nurse-patient assignment: A multiple objective heuristic algorithm. *Decis. Support Syst.* **2023**, *167*, 113926. [\[CrossRef\]](#)
22. Bonaconsa, C.; Mbamalu, O.; Surendran, S.; George, A.; Mendelson, M.; Charani, E. Optimizing infection control and antimicrobial stewardship team communication and patient outcomes. *Am. J. Infect. Control* **2024**, *52*, 872–879. [\[CrossRef\]](#)

23. Danach, K. Reinforcement Learning for Dynamic Vehicle Routing Problem: A Case Study with Real-World Scenarios. *Int. J. Commun. Netw. Inf. Secur.* **2024**, *16*, 580–589.
24. O'Reilly, M.M.; Krasnicki, S.; Montgomery, J.; Heydar, M.; Turner, R.; Van Dam, P.; Maree, P. Markov decision process and approximate dynamic programming for a patient assignment scheduling problem. *Ann. Oper. Res.* **2025**, *347*, 1493–1531. [\[CrossRef\]](#)
25. Danach, K.; Aly, W.H.F. Adaptive Hyperheuristic Framework for Hyperparameter Tuning: A Q-Learning-Based Heuristic Selection Approach with Simulated Annealing Acceptance Criteria. *Eur. J. Pure Appl. Math.* **2025**, *18*, 6348. [\[CrossRef\]](#)
26. Teymourifar, A.; Trindade, M.A. Balanced patient assignment to healthcare centres through dispatching rules. *Health Syst.* **2025**, *14*, 104–118. [\[CrossRef\]](#)
27. Danach, K.; Rammal, A.; Moukadem, I.; Harb, H.; Nasser, A. Advanced Optimization in E-Commerce Logistics: Combining Matheuristics With Random Forests for Hub Location Efficiency. *IEEE Access* **2025**. [\[CrossRef\]](#)
28. Brandt, T. Structural Insights and Efficient Algorithms for the Patient-To-Room Assignment Problem with Equal Rooms. Ph.D. Thesis, RWTH Aachen University, Aachen, Germany, 2025.
29. Danach, K.; Harb, H.; Rashid, A.S.K.; Al-Tarawneh, M.A.; Aly, W.H.F. Location planning techniques for Internet provider service unmanned aerial vehicles during crisis. *Results Eng.* **2025**, *25*, 103833. [\[CrossRef\]](#)
30. Yao, X.; Shehadeh, K.S.; Padman, R. Multi-resource allocation and care sequence assignment in patient management: A stochastic programming approach. *Health Care Manag. Sci.* **2024**, *27*, 352–369. [\[CrossRef\]](#) [\[PubMed\]](#)
31. Abbas, Q.; Jeong, W.; Lee, S.W. Explainable AI in Clinical Decision Support Systems: A Meta-Analysis of Methods, Applications, and Usability Challenges. *Healthcare* **2025**, *13*, 2154. [\[CrossRef\]](#) [\[PubMed\]](#)
32. Brandt, T.; Büsing, C.; Engelhardt, F. Patient-to-room assignment with single-rooms entitlements: Combinatorial insights and integer programming formulations. *Eur. J. Oper. Res.* **2025**, *325*, 20–37. [\[CrossRef\]](#)
33. Zanazzo, E.; Ceschia, S.; Schaerf, A. Multi-neighborhood simulated annealing for the integrated patient-to-room and nurse-to-patient assignment problem. *Flex. Serv. Manuf. J.* **2025**. [\[CrossRef\]](#)
34. Jedidi, H.; Ben-Romdhane, H.; Nouaouri, I.; Krichen, S. Bi-objective Hospital Bed Assignment Problem in Emergencies. In Proceedings of the International Conference on Decision Aid and Artificial Intelligence (ICODAI 2024), Aleco, Tunisia, 31 October 2024; pp. 136–146.
35. Dorigo, M.; Stützle, T. *Ant Colony Optimization*; MIT Press: Cambridge, MA, USA, 2004.
36. Castro, L.N.D.; Timmis, J. *Artificial Immune Systems: A New Computational Intelligence Approach*; Springer: Berlin/Heidelberg, Germany, 2002.

**Disclaimer/Publisher's Note:** The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.