



Journal of Intelligent Monitoring and Control Systems (JIMCS)

DOI: 10.65098/jimcs.01.2026.56.65



RESEARCH ARTICLE

Bi-Objective Receding-Horizon Scheduling of Nurses and Hospital Service Robots

Khatere Ghorbani-Moghadam¹, Roya Khalili-Amirabadi^{2,*}

¹Mosaheb Institute of Mathematics, Kharazmi University, Tehran 15618, Iran

²Department of Applied Mathematics, Ferdowsi University of Mashhad, Mashhad 91775, Iran

*Corresponding Author Email: roya.khalili.a@gmail.com

This is an open access article distributed under the Creative Commons Attribution License CC BY 4.0, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

ARTICLE DETAILS

Article History:

Received: 31 Aug 2026

Accepted: 09 Oct 2026

Available online: 30 Oct 2026

Online Article Code



ABSTRACT

This paper studies the joint scheduling of nurses and hospital service robots as a bi-objective mixed-integer linear programming problem. The formulation represents heterogeneous nursing skills, alternative skill-coverage modes, nurse availability and temporal rules, workload balance, robot task capability, charging and battery dynamics, and human supervision. Robot workload relief is credited only when an eligible service is actually executed, and explicit supervision-load variables deduct the corresponding nurse capacity from clinical coverage. The receding-horizon formulation also includes a cross-window night-morning constraint and a terminal battery reserve. The computational study is presented as numerical verification on synthetic hospital-like instances rather than clinical validation. Under the proposed formulation, the eight-nurse reference case gives $Z_1 = 0.78$ and $Z_2 = 3.164$. A common-budget comparison on B20 seed 1 yields total shortage 53.70 for the integrated model, compared with 58.03 for sequential nurse-robot scheduling and 59.46 for the nurse-only baseline. Additional experiments examine multiple epsilon budgets on a larger B30 instance, sensitivity to robot-relief/supervision parameters and objective weights, and one-day versus two-day rolling horizons against a static full-horizon reference. The battery and routing representations remain aggregate planning approximations, and the scalability evidence is interpreted accordingly. No patient records or real hospital staff data are used, so the results establish model behavior and comparative numerical evidence rather than practical clinical effectiveness.

KEYWORDS

Nurse Scheduling; Hospital Service Robots; Bi-Objective Optimization; Mixed-Integer Linear Programming; Model Predictive Control

1. INTRODUCTION

Nurse scheduling is not simply a question of meeting headcount. A workable roster must place appropriately qualified nurses in the right wards and shifts while respecting availability, labor rules, workload balance, preferences, and temporal restrictions. These requirements have been studied extensively in the nurse rostering and personnel scheduling literature (Cheang et al., 2003; Burke et al., 2004; Ernst et al., 2004; Van den Bergh et al., 2013; De Bruecker et al., 2015). In hospital settings, skill mix is especially important because two rosters with the same number of nurses can differ substantially in clinical adequacy when critical-care or senior competencies are scarce.

Service robots introduce a second resource class whose usefulness depends on both task capability and internal state. Hospital robots may assist with transport, specimen delivery, supply movement, disinfection, telepresence, and other non-specialist activities. Reviews of healthcare robotics describe a growing range of such applications (Morgan et al., 2022; Huang et al., 2023), and systems such as Lio demonstrate the practical feasibility of supportive robotic assistance in human-centered care environments (Miseikis et al., 2020). Robotics has also been considered for reducing avoidable exposure and supporting logistics during infectious-disease events (Yang et al., 2020). These systems are not substitutes for specialized nursing care. They have limited capabilities, consume energy, require charging, and may depend on human supervision. Work on mixed huma-robot teams also shows that coordination, authority, and workload affect team performance (Gombolay et al., 2015), while autonomous mobile robot planning highlights the close connection between task allocation,

energy, and operational control (Fragapane et al., 2021).

The scheduling problem considered here therefore treats nurses and robots within the same ward-shift environment. The model combines multi-level nursing skills and alternative acceptable coverage patterns with separate objectives for service adequacy and operational quality. Robot presence is distinguished from actual task execution, and battery dynamics connect decisions across time. A receding-horizon control layer then allows the schedule to be updated when demand, availability, or robot state changes. The resulting formulation is intended as a decision-support model: robots can relieve selected support workload, but they cannot replace sensitive clinical competencies.

Model predictive control (MPC) provides a natural structure for this type of repeated decision making because it optimizes over a finite horizon, implements the first decision, observes the updated system, and then solves again (Camacho & Bordons, 2007; Borrelli et al., 2017; Qin & Badgwell, 2003). Related work has used mixed-integer predictive control in fleet systems with charging constraints (Zhang et al., 2016). More recent research has also examined learning-assisted predictive control. In particular, RL-MPC architectures have been surveyed as a way to combine explicit constraint handling with data-driven adaptation (Jalaeian Farimani et al., 2026), while experience-replay methods provide a broad set of mechanisms for reusing past interaction data in reinforcement learning (Khalili Amirabadi et al., 2026). Reinforcement learning is not used in the present study; these directions are relevant as possible extensions when

reliable operational data become available.

The methodological contribution is an integration contribution rather than a claim of a new generic optimization algorithm. Classical nurse-rostering research focuses on staffing, skills, contractual rules, and preferences (Cheang et al., 2003; Burke et al., 2004; Ernst et al., 2004; Van den Bergh et al., 2013; De Bruecker et al., 2015); healthcare-robot studies mainly characterize applications and human-robot interaction (Morgan et al., 2022; Huang et al., 2023; Miseikis et al., 2020; Gombolay et al., 2015); autonomous mobile-robot work emphasizes task allocation, routing, and energy (Fragapane et al., 2021); and predictive-control studies provide the rolling optimization mechanism (Camacho & Bordons, 2007; Borrelli et al., 2017; Zhang et al., 2016). The present model couples nurse assignments, skill coverage, robot task execution, charging, battery state, and supervision in one ward-shift MILP, while keeping clinical/service shortage separate from operational burden. Accordingly, the contribution lies in the joint formulation and state-coupled scheduling architecture rather than in a new generic MILP solution paradigm. Table 1 summarizes the positioning of the proposed formulation relative to representative nurse-rostering, healthcare-robotics, mixed human-robot planning, and predictive-control literature streams.

The mathematical formulation is numerically verified on a small reference case and then examined on replicated synthetic instances and a seven-day rolling-horizon experiment. The term numerical verification is used rather than practical validation because no real hospital dataset, clinical staffing policy, or robot telemetry dataset is used in this study.

2. PROBLEM DESCRIPTION AND NOTATION

Consider a hospital with nurses $\mathcal{N}$, wards $\mathcal{W}$, planning periods $\mathcal{T}$, days $\mathcal{D}$, clinical and support skills $\mathcal{S}$, acceptable skill-coverage modes $\mathcal{G}$, service robots $\mathcal{R}$, and robot service types $\mathcal{L}$. Let $\mathcal{T}_d \subseteq \mathcal{T}$ be the periods belonging to day d , let $\mathcal{T}^N \subseteq \mathcal{T}$ denote night periods, and let $\mathcal{P}$ contain incompatible ordered shift pairs, such as a night shift followed immediately by a morning shift. Battery states are defined on $\mathcal{T}^+ = \{1, \dots, |\mathcal{T}| + 1\}$ so that the terminal state after the final scheduling period is represented explicitly.

A nurse may be assigned to at most one ward in a period and, under the baseline policy used here, to at most one shift per day. Each nurse has an availability profile, a regular-hour threshold, a maximum number of night duties, and a set of skill qualifications. For each ward, period, and skill, the hospital may specify more than one acceptable coverage pattern. For example, a requirement may be met by a smaller number of experts or by a larger number of experienced nurses. One of these modes is selected by the optimization model.

Robots are modeled as supportive resources. A robot can work in at most one ward in a period, charge, or remain idle. When active, it may execute only tasks for which it is qualified, subject to a period task capacity. Its battery is reduced by work and replenished by charging. A robot contributes to nursing workload relief only through an explicitly executed service variable; presence in a ward by itself creates no nursing credit. For specialized ICU, senior, or otherwise sensitive clinical skills, the relief coefficient is set to zero. Human supervision is represented in two ways: a headcount ratio limits simultaneous robot activity, and a continuous supervision-load variable reserves a fraction of assigned nurse capacity so that a supervising nurse is not simultaneously credited as fully available for clinical coverage.

Table 2 summarizes the main sets, and Table 3 lists the principal parameters.

3. MATHEMATICAL FORMULATION

This section translates the joint nurse-robot scheduling problem into a bi-objective mixed-integer linear program. It first defines the decision variables, then develops the staffing and skill-coverage constraints, the robot task, charging, battery, and supervision constraints, and finally the two objective functions used to distinguish service adequacy from operational burden.

3.1 Decision Variables

The principal binary variables are

$$\begin{aligned} x_{nwt} &= 1 && \text{if nurse } n \text{ works in ward } w \text{ at period } t, \\ z_{rwt} &= 1 && \text{if robot } r \text{ is active in ward } w \text{ at period } t, \\ v_{rwtl} &= 1 && \text{if robot } r \text{ executes service } l \text{ in ward } w \text{ at period } t, \\ c_{rt} &= 1 && \text{if robot } r \text{ charges at period } t, \\ m_{wtsg} &= 1 && \text{if coverage mode } g \text{ is selected for } (w, t, s). \end{aligned}$$

Continuous variables $h_{wts}^S \geq 0$ and $u_{wtl}^R \geq 0$ represent human-skill and robot-service shortage, respectively. B_{rt} is the battery state of robot r at the start of state period $t \in \mathcal{T}^+$. R_n and O_n denote nurse work hours and overtime, and $R^{\max}$ and $R^{\min}$ denote the maximum and minimum nurse workloads. The variable $\phi_{nwt} \in [0, 1]$ denotes the fraction of nurse n 's capacity in ward w and period t reserved for robot supervision.

3.2 Nurse Assignment and Skill Coverage

A nurse may work in at most one ward in a period and only when available:

$$\sum_{w \in \mathcal{W}} x_{nwt} \leq A_{nt}, \quad \forall n, t. \quad (1)$$

The baseline daily policy permits at most one shift per nurse:

$$\sum_{t \in \mathcal{T}_d} \sum_{w \in \mathcal{W}} x_{nwt} \leq 1, \quad \forall n, d. \quad (2)$$

Ward headcount is bounded by

$$L_{wt}^{\min} \leq \sum_{n \in \mathcal{N}} x_{nwt} \leq L_{wt}^{\max}, \quad \forall w, t. \quad (3)$$

For every incompatible pair $(t, \tau) \in \mathcal{P}$,

$$\sum_{w \in \mathcal{W}} x_{nwt} + \sum_{w \in \mathcal{W}} x_{nw\tau} \leq 1, \quad \forall n, (t, \tau) \in \mathcal{P}. \quad (4)$$

Night duties satisfy

$$N_n^0 + \sum_{w \in \mathcal{W}} \sum_{t \in \mathcal{T}^N} x_{nwt} \leq B_n^N, \quad \forall n. \quad (5)$$

For each ward-period-skill combination, exactly one acceptable coverage mode is selected:

$$\sum_{g \in \mathcal{G}} m_{wtsg} = 1, \quad \forall w, t, s. \quad (6)$$

Coverage under each candidate mode is enforced by

$$\begin{aligned} \sum_{n \in \mathcal{N}} Q_{nsg}(x_{nwt} - \phi_{nwt}) + \sum_{r \in \mathcal{R}} \sum_{l \in \mathcal{L}} \alpha_{rls} v_{rwtl} + h_{wts}^S \\ \geq H_{wts}^g - M_{wts}(1 - m_{wtsg}), \quad \forall w, t, s, g. \end{aligned} \quad (7)$$

When mode g is selected, Eq. (7) reduces to the corresponding coverage requirement. For nonselected modes, the local big- M term relaxes the constraint. A local value satisfying $M_{wts} \geq \max_{g \in \mathcal{G}} H_{wts}^g$ is sufficient and avoids the numerical weakness of an unnecessarily large global constant. The second term is now based on actual task execution v_{rwtl} , not robot presence z_{rwt} . Consequently, robot presence alone cannot generate nurse-equivalent workload relief, and every credited relief event consumes one unit of the same executed-service capacity that appears in Eq. (16). To prevent a single executed service from being credited independently against several support-skill requirements, the relief coefficients are calibrated to satisfy

$$\sum_{s \in \mathcal{S}} \alpha_{rls} \leq \bar{\alpha}_{rl}, \quad \forall r, l, \quad (8)$$

where $\bar{\alpha}_{rl}$ is the maximum total nurse-equivalent workload that one execution of service l can displace. For sensitive clinical skills, or for robot services that do not displace the corresponding nursing workload, $\alpha_{rls} =$

Table 1 Positioning of the Present Formulation Relative to Selected Literature Streams

Literature stream	Principal elements emphasized	Relation to the present study
Nurse rostering and skill-based workforce planning (Cheang et al., 2003; Burke et al., 2004; Van den Bergh et al., 2013; De Bruecker et al., 2015)	Staffing coverage, skills, contractual rules, preferences, fairness	Provides the nurse-rostering foundation; the cited representative works do not provide the joint nurse-robot state model studied here.
Healthcare/service robotics (Morgan et al., 2022; Huang et al., 2023; Miseikis et al., 2020)	Robot applications, assistance tasks, deployment contexts, human interaction	Establishes feasible hospital robot roles; the cited representative works are not joint nurse-roster optimization models.
Mixed human-robot teams and AMR planning (Gombolay et al., 2015; Fragapane et al., 2021)	Human-robot coordination, task allocation, routing, energy, operational control	Motivates coordination and energy constraints; the present model instead uses an aggregate ward-shift representation coupled to clinical staffing.
Predictive control/fleet MPC (Camacho & Bordons, 2007; Borrelli et al., 2017; Zhang et al., 2016)	Finite-horizon optimization, state propagation, charging-aware fleet decisions	Provides the rolling optimization mechanism, but not heterogeneous nurse-skill coverage.
Present study	Nurse skills, robot service execution, battery/charging, supervision workload, and rolling re-optimization	Integrates these elements in one aggregate ward-shift MILP; the contribution is the joint formulation rather than a new generic MILP algorithm.

Table 2 Sets and Indices

Symbol	Description
$\mathcal{N}$	Nurses, indexed by n
$\mathcal{W}$	Hospital wards, indexed by w
$\mathcal{T}$	Scheduling periods, indexed by t
$\mathcal{T}^+$	Battery-state periods, $\{1, \dots, \mathcal{T} + 1\}$
$\mathcal{D}$	Days; $\mathcal{T}_d \subseteq \mathcal{T}$ contains the periods of day d
$\mathcal{S}$	Nursing skill and support categories, indexed by s
$\mathcal{G}$	Acceptable skill-coverage modes, indexed by g
$\mathcal{R}$	Hospital service robots, indexed by r
$\mathcal{L}$	Robot service types, indexed by l
$\mathcal{T}^N$	Night periods
$\mathcal{P}$	Incompatible ordered shift pairs (t, τ)

0. This calibration restriction is important: the robot-service demand D_{wtl}^R is interpreted as demand for the same completed support activity whose execution may displace nurse workload, rather than as an additional independent activity. The supervision fraction ϕ_{nwt} reduces the effective clinical contribution of an assigned nurse.

If a hard expert-presence rule is required, it is imposed separately:

$$\sum_{n \in \mathcal{N}} Q_{ns}^E x_{nwt} \geq E_{wts}, \quad \forall w, t, s \text{ for which expert presence is required.} \quad (9)$$

Nurse workload is

$$R_n = R_n^0 + \sum_{w \in \mathcal{W}} \sum_{t \in \mathcal{T}} h_t x_{nwt}, \quad \forall n, \quad (10)$$

and overtime satisfies

$$O_n \geq R_n - C_n, \quad O_n \geq 0, \quad \forall n. \quad (11)$$

The workload spread is represented by

$$R^{\max} \geq R_n, \quad R^{\min} \leq R_n, \quad \forall n. \quad (12)$$

3.3 Robot Task Assignment, Charging, and Battery Dynamics

A robot cannot work in more than one ward or work while charging:

$$\sum_{w \in \mathcal{W}} z_{rwt} + c_{rt} \leq 1, \quad \forall r, t. \quad (13)$$

Daily robot utilization is limited by

$$\sum_{t \in \mathcal{T}_d} \sum_{w \in \mathcal{W}} z_{rwt} \leq U_r, \quad \forall r, d. \quad (14)$$

If a minimum charging requirement is imposed,

$$\sum_{t \in \mathcal{T}_d} c_{rt} \geq m_r^C, \quad \forall r, d. \quad (15)$$

Task execution is linked to robot presence and capability:

$$\begin{aligned} v_{rwtl} &\leq z_{rwt}, & \forall r, w, t, l, \\ v_{rwtl} &\leq A_{rl}^R, & \forall r, w, t, l, \\ \sum_{w \in \mathcal{W}} \sum_{l \in \mathcal{L}} v_{rwtl} &\leq K_{rt}, & \forall r, t. \end{aligned} \quad (16)$$

The last inequality prevents a robot from being credited with more services than it can execute during a period. Robot-service demand is covered by

$$\sum_{r \in \mathcal{R}} v_{rwtl} + u_{wtl}^R \geq D_{wtl}^R, \quad \forall w, t, l. \quad (17)$$

Battery dynamics are defined for $t = 1, \dots, |\mathcal{T}|$:

$$B_{r,t+1} = B_{rt} - \sum_{w \in \mathcal{W}} e_{rw} z_{rwt} + c h_r c_{rt}, \quad \forall r, t. \quad (18)$$

Table 3 Main Model Parameters

Parameter	Meaning
Q_{nsg}	1 if nurse n is sufficiently qualified for skill s under coverage mode g , 0 otherwise
H_{wts}^g	Required effective coverage for ward w , period t , skill s , under mode g
M_{wts}	Local deactivation constant for the selected-mode constraints, chosen from the corresponding coverage requirements
A_{nt}	Nurse availability indicator
h_t	Duration of period t in hours
C_n	Regular-hour threshold for nurse n
R_n^0	Historical work hours carried into the current optimization window
N_n^0	Historical night-duty count carried into the current optimization window
B_n^N	Maximum allowed night-duty count for nurse n over the relevant horizon
$L_{wt}^{\min}, L_{wt}^{\max}$	Minimum and maximum nurse headcount in ward w and period t
Q_{ns}^E	1 if nurse n satisfies the expert threshold for sensitive skill s
E_{wts}	Minimum expert presence for a sensitive skill, when required
q_{nwt}	Preference or undesirability penalty for assigning nurse n to ward w in period t
A_{rl}^R	1 if robot r can perform service l
K_{rt}	Maximum number of robot services executable by robot r in period t
D_{wtl}^R	Demand for robot service l in ward w and period t
α_{rls}	Nurse-equivalent workload relief generated when robot r actually executes service l for eligible support skill s ; zero for non-relieving or sensitive clinical skills
$\bar{\alpha}_{rl}$	Maximum total nurse-equivalent workload that one execution of service l by robot r may displace across all support-skill categories
e_{rw}	Battery energy used when robot r works in ward w for one period
ch_r	Energy added to robot r by one charging period
$B_r^{\min}, B_r^{\max}$	Minimum starting level for work and maximum battery capacity
U_r	Maximum active work periods per day for robot r
m_r^C	Minimum charging periods per day, when imposed
ρ_w^R	Maximum number of active robots per assigned nurse in ward w (hard supervision headcount limit)
d_{rw}^{sup}	Fraction of one nurse-period required to supervise active robot r in ward w
ϕ_n	Maximum fraction of nurse n 's assigned period that may be allocated to robot supervision
B_r^{term}	Minimum terminal battery reserve in a receding-horizon window
$\beta_w, \rho_t, \sigma_s, \tau_l$	Priority weights for ward, period, human skill, and robot service
η_R	Relative weight of robot-service shortage in Z_1
$\lambda_B, \lambda_O, \lambda_N, \lambda_P, \lambda_E, \lambda_C$	Weights in the operational objective Z_2

The state is bounded for every $t \in \mathcal{T}^+$:

$$0 \leq B_{rt} \leq B_r^{\max}, \quad \forall r, t \in \mathcal{T}^+. \quad (19)$$

An active robot must have at least its minimum starting battery:

$$B_{rt} \geq B_r^{\min} \sum_{w \in \mathcal{W}} z_{rwt}, \quad \forall r, t \in \mathcal{T}. \quad (20)$$

Together, Eqs. (18)-(20) also prevent a work decision that would drive the next battery state below zero or above capacity.

Human supervision has both a hard headcount condition and an explicit workload condition. The headcount condition remains

$$\sum_{r \in \mathcal{R}} z_{rwt} \leq \rho_w^R \sum_{n \in \mathcal{N}} x_{nwt}, \quad \forall w, t, \quad (21)$$

while supervision workload satisfies

$$\sum_{n \in \mathcal{N}} \phi_{nwt} \geq \sum_{r \in \mathcal{R}} d_{rw}^{sup} z_{rwt}, \quad \forall w, t, \quad (22)$$

$$0 \leq \phi_{nwt} \leq \phi_n x_{nwt}, \quad \forall n, w, t. \quad (23)$$

Thus, robot activity consumes a calibrated fraction of nurse capacity through ϕ_{nwt} , and Eq. (7) no longer counts supervising nurses as fully available for clinical coverage. The coefficient ρ_w^R can be lower in high-acuity wards, while d_{rw}^{sup} can be estimated from observed intervention/supervision time. This remains an aggregate nurse-equivalent approximation; a task-time model would be needed if supervision must be allocated to specific clinical activities. The present formulation also assumes that any assigned nurse is eligible to provide robot supervision. If supervision requires specific training, an additional nurse-robot supervision-eligibility parameter must be introduced before clinical use.

3.4 Objective Functions

The first objective measures clinical and service inadequacy:

$$Z_1 = \sum_{w,t,s} \beta_w \rho_t \sigma_s h_{wts}^S + \eta_R \sum_{w,t,l} \beta_w \rho_t \tau_l u_{wtl}^R. \quad (24)$$

Ward, period, skill, and service priorities are reflected by $\beta_w, \rho_t, \sigma_s$, and τ_l . In practice, critical-care skills may receive much larger shortage weights than general support activities.

The second objective measures operational burden and fairness:

$$Z_2 = \lambda_B (R^{\max} - R^{\min}) + \lambda_O \sum_n O_n + \lambda_N \sum_{n,w,t \in \mathcal{T}^N} x_{nwt} + \lambda_P \sum_{n,w,t} q_{nwt} x_{nwt} + \lambda_E \sum_{r,w,t} e_{rw} z_{rwt} + \lambda_C \sum_{r,t} c_{rt}. \quad (25)$$

The model is therefore bi-objective:

$$\min (Z_1, Z_2). \quad (26)$$

The two objectives are deliberately kept separate because service adequacy and operating burden represent different managerial concerns.

4. SOLUTION METHODOLOGY

The formulation is solved in two complementary ways. The first examines the trade-off between the two objectives through normalized weighted-sum and epsilon-constraint procedures, including lexicographic refinement for the reported efficient solutions. The second embeds the

scheduling model in a receding-horizon framework so that updated workload, availability, demand, and robot battery states can influence subsequent decisions.

4.1 Multi-objective Solution Strategy

Two procedures are used to examine the trade-off between Z_1 and Z_2 . For a normalized weighted sum, let $Z_i^{\min}$ and $Z_i^{\max}$ be finite scaling values obtained from the individual objective anchors. Define

$$\hat{Z}_i = \frac{Z_i - Z_i^{\min}}{\max\{Z_i^{\max} - Z_i^{\min}, \delta\}}, \quad \delta > 0, \quad (27)$$

and solve

$$\min \theta \hat{Z}_1 + (1 - \theta) \hat{Z}_2, \quad 0 \leq \theta \leq 1. \quad (28)$$

The positive constant δ prevents division by zero when an anchor range vanishes and keeps the normalization well defined across instances.

The principal reporting method is the epsilon-constraint approach (Mavrotas, 2009). A shortage-priority lexicographic solution first obtains

$$Z_1^* = \min Z_1,$$

and then solves

$$\min Z_2 \quad \text{s.t.} \quad Z_1 \leq Z_1^* + \varepsilon. \quad (29)$$

For the replicated benchmark, an operational reference is also useful. A pure Z_2 solve first returns an operationally optimal nurse-robot schedule. Because shortage and coverage-mode variables do not appear in Z_2 , their raw values need not be unique. To obtain a reproducible shortage measure for that schedule, the assignment and robot decisions are fixed and Z_1 is minimized over the remaining coverage-mode and shortage variables. The resulting value is denoted Z_1^{op} . The benchmark then minimizes Z_2 subject to $Z_1 \leq f Z_1^{op}$. The imposed budget, rather than the raw shortage slack returned by a Z_2 solve, is used to state the guaranteed shortage reduction.

For Pareto analysis, each epsilon-constrained solution is additionally refined lexicographically to remove weakly efficient points caused by ties in Z_2 . For a fixed budget ε , Stage 1 computes

$$Z_2^*(\varepsilon) = \min\{Z_2 : Z_1 \leq \varepsilon\}, \quad (30)$$

and Stage 2 solves

$$\min Z_1 \quad \text{s.t.} \quad Z_1 \leq \varepsilon, \quad Z_2 \leq Z_2^*(\varepsilon) + \tau, \quad (31)$$

where τ is a small numerical tolerance. The two extreme points are treated analogously: the safety endpoint minimizes Z_1 and then Z_2 , whereas the operational endpoint minimizes Z_2 and then Z_1 . A final pairwise dominance check removes duplicates or dominated points before a frontier is reported. This procedure yields a numerically efficient set up to the stated MILP tolerance rather than relying on arbitrary auxiliary-variable values.

4.2 Receding-Horizon Predictive Control

The static model becomes a deterministic receding-horizon scheduling model when the implemented state is carried from one optimization window to the next. At control iteration p , let $\mathcal{J}_p$ denote the ordered set of periods in the prediction window and let $\mathcal{I}_p \subseteq \mathcal{J}_p$ denote the periods that are actually implemented before the next re-optimization (one day in the reported experiment). Let t_p^0 be the first period in the window, t_p^f the last period in the prediction window, and t_p^l the last implemented period. The state entering iteration p is

$$X_p = (B_{rp}^0, H_{np}, N_{np}, Y_{np}, A_{nt}, D_{wts})_{t \in \mathcal{J}_p}, \quad (32)$$

where B_{rp}^0 is the battery state at the beginning of the window, H_{np} is cumulative implemented nurse work, N_{np} is cumulative implemented night duty, Y_{np} indicates whether nurse n worked the immediately preceding implemented night shift, and A_{nt} and D_{wts} are observed or forecast exogenous inputs over the prediction window. Here “prediction” means that future demand and availability values are supplied over the finite

look-ahead window; no statistical forecasting model is estimated in this paper.

The within-window model uses the same decision variables and constraints as the static MILP, with the initial battery fixed by

$$B_{r,t_p^0} = B_{rp}^0, \quad \forall r. \quad (33)$$

Only decisions in $\mathcal{I}_p$ are committed. Cumulative work and night-duty states are then propagated as

$$H_{n,p+1} = H_{np} + \sum_{t \in \mathcal{J}_p} h_t \sum_{w \in \mathcal{W}} x_{nwt}, \quad (34)$$

$$N_{n,p+1} = N_{np} + \sum_{t \in \mathcal{J}_p} I_t^N \sum_{w \in \mathcal{W}} x_{nwt}, \quad (35)$$

where $I_t^N = 1$ for a night period and 0 otherwise. The battery state passed to the next iteration is the state immediately after the last implemented period,

$$B_{r,p+1}^0 = B_{r,t_p^l+1}, \quad \forall r. \quad (36)$$

The night-history indicator is updated from the last implemented shift: $Y_{n,p+1} = 1$ if nurse n is assigned to a night shift in period t_p^l , and 0 otherwise.

Cross-window rest is imposed explicitly. If t_p^0 is a morning shift that is incompatible with a night shift immediately before the current window, then

$$Y_{np} + \sum_{w \in \mathcal{W}} x_{nwt_p^0} \leq 1, \quad \forall n. \quad (37)$$

Thus, a night assignment implemented in the preceding window cannot be followed by an incompatible morning assignment merely because the optimization horizon has moved. To mitigate artificial battery depletion at the end of the prediction window, the proposed formulation also imposes

$$B_{r,t_p^f+1} \geq B_r^{\text{term}}, \quad \forall r, \quad (38)$$

where B_r^{term} is a prespecified terminal reserve. The reserve should be calibrated from the battery required for plausible near-future service or charging access; alternatively, a soft terminal penalty can be used if a hard reserve is operationally inappropriate.

At each iteration the current state is observed, the finite-horizon MILP is built and solved, only $\mathcal{I}_p$ is implemented, and the state is propagated before the next solve. The prediction horizon therefore controls how far future exogenous inputs and battery consequences can influence the current implemented decisions. A one-day window is the myopic rolling baseline; longer windows provide additional look-ahead at the cost of a larger MILP. This procedure is more accurately described as deterministic rolling-horizon re-optimization with MPC-style state propagation unless an explicit forecasting model is added.

5. COMPUTATIONAL STUDY

The computational study is designed for numerical verification, comparative ablation, and controlled stress testing rather than clinical validation. It first evaluates the reference instance under the proposed formulation, then examines synthetic instances of increasing size, multiple shortage budgets, direct baselines, parameter/weight sensitivity, and prediction-horizon effects. Because no real hospital roster, staffing policy, or robot telemetry dataset is used, the study does not establish deployment performance or clinical benefit.

5.1 Reference Numerical Verification

The reference case contains eight nurses, two wards (ICU and General), three shifts, three nursing skill categories, two robots, and two robot-service types. Both robots can execute transport; only the second robot can execute the disinfection/support service. Workload relief is credited only for an executed eligible service, and supervision consumes nurse capacity through ϕ_{nwt} .

Solving this case under the proposed formulation gives $Z_1 = 0.78$ and $Z_2 = 3.164$. The entire shortage is human-coverage shortage created

by the explicit supervision deduction; robot-service shortage is zero. The returned nurse plan assigns N2 and N8 to morning ICU, N4 and N6 to morning General, N1 to evening ICU, N5 to evening General, N3 to night ICU, and N7 to night General. Robot R1 performs morning ICU transport, remains idle in the evening, and charges at night; R2 charges in the morning, performs evening ICU transport, and performs the General-ward disinfection service at night. The corresponding battery paths are (80, 50, 50, 90) for R1 and (60, 100, 65, 35) for R2. All nurses work eight hours and incur no overtime. These results provide a consistency check for the task-execution, supervision, charging, and battery constraints; they are not a clinical validation exercise.

5.2 Synthetic Benchmark Design

The synthetic benchmark captures heterogeneous wards, skills, and robot capabilities under the proposed formulation. Five independent seeds are reported for B20, B30, and B45; one B60 instance is retained as a larger stress case. The numerical study reports 16 instances: five seeds each for B20, B30, and B45, together with one B60 stress case. Table 4 summarizes the tested dimensions. Each benchmark instance uses one day with three shifts; multi-day behavior is examined separately in Section 5.6.

The generator includes general care, critical care, medication competency, and senior status. Alternative coverage modes represent expert, experienced, and general qualification thresholds. Robot fleets are heterogeneous in capability, task capacity, battery state, charging rate, energy use, and support contribution. These are engineering test parameters rather than clinical recommendations.

5.3 Benchmark and Epsilon-Budget Analysis

The mixed-integer problems were solved with HiGHS (Huangfu & Hall, 2018). For each benchmark instance, an operational reference is computed and then a constrained run imposes $Z_1 \leq 0.60Z_1^{op}$. The reported Z_1 in Table 5 is the canonical achieved shortage of the returned schedule, not the imposed budget. Consequently, the difference $0.60Z_1^{op} - Z_1$ can be read as canonical budget slack. The larger cases are computationally more difficult; therefore, status and MIP-gap information is reported rather than treating every feasible incumbent as a proven optimum.

The 0.60 coefficient is therefore treated as a policy constraint rather than as an independently discovered improvement. In the reported runs, the canonical achieved shortage is below the imposed budget in every reported case. For B20–B45, the mean canonical reduction relative to Z_1^{op} is approximately 41.8–42.8%; the B60 seed-1 incumbent gives 40.5%. Because several B30/B45/B60 runs terminate at the time limit, these values are reported as achieved feasible outcomes rather than universal optima. The largest time-limited MIP gap is 27.99%, so the large-instance results are interpreted as computational stress evidence, not exact frontier points.

To examine sensitivity to multiple shortage-budget levels and to extend the epsilon analysis beyond B20, Table 6 reports a B30 seed-1 sweep. The achieved shortage is canonicalized for each returned schedule. At the 0.85, 0.70, and 0.60 budgets, the canonical slack is positive, so the returned schedules lie strictly inside the corresponding policy bounds. The 0.70 run is time-limited and is therefore treated as a feasible incumbent.

Table 6 and Figure 1 show the expected service-operation trade-off once the budget becomes tighter: reducing canonical shortage below approximately 96 increases Z_2 . Because this is a mixed-integer problem and one intermediate point is time-limited, the table is presented as an epsilon-budget sensitivity study rather than a complete Pareto frontier.

5.4 Direct Baseline Comparison

Three directly comparable formulations were evaluated on B20 seed 1 under the common service standard $Z_1 \leq 60$: (i) a nurse-only model with $z = v = c = 0$; (ii) a sequential model that first fixes the nurse schedule and then allocates robots; and (iii) the integrated joint formulation. Table 7 summarizes the common-budget comparison of the nurse-only, sequential nurse-robot, and integrated joint formulations on B20 seed 1. These controls do not reproduce every specialized algorithm

in the literature, but they isolate the incremental value of robot availability and joint co-optimization under identical staffing, service, and objective definitions.

On this controlled instance, joint optimization reduces total weighted shortage by 9.68% relative to nurse-only scheduling and by 7.47% relative to the sequential approach, while keeping Z_2 comparable. The gain comes primarily from lower robot-service shortage; human shortage alone is not uniformly lower. Accordingly, the result supports the value of joint coordination on this instance but is not presented as a general superiority claim over all nurse-rostering or robot-scheduling algorithms.

5.5 Parameter and Objective-Weight Sensitivity

The workload-relief and supervision coefficients are not directly observable clinical constants. In practice, α_{rls} should be estimated from the nurse time displaced by a completed robot service, while d_{rw}^{sup} should be estimated from time-and-motion observation, robot intervention logs, or structured staff studies. To examine local sensitivity, one fixed synthetic instance was solved under a common shortage budget while scaling the relief or supervision coefficient by 0.5 and 1.5 relative to baseline. Table 8 reports the local sensitivity results for the robot-relief and supervision coefficients under the common shortage budget.

Within this local range, changes in canonical shortage are small and of the same order as the reported solver gaps, so the experiment does not support a strong monotonic sensitivity claim. Its main implication is that these parameters must be calibrated from observed task and supervision times rather than chosen solely for numerical convenience.

A second sensitivity test changes the relative weights in Z_2 . Because changing the weights changes the objective scale itself, Table 9 reports schedule components rather than comparing the numerical Z_2 values across rows.

The schedules change materially when weights are altered, confirming that the coefficients require calibration. The response is not monotone in every component because the terms interact and one setting is time-limited; no universal ranking of weight choices is inferred. In application, a defensible calibration should scale components to comparable reference ranges, elicit relative priorities from staffing and robotics stakeholders, and repeat sensitivity checks around the selected values.

5.6 Prediction-Horizon Comparison and Seven-Day Stress Path

A seven-day deterministic stress path uses 24 nurses, two wards, and three robots, with the same controlled demand, absence, and battery disturbances used to exercise state propagation. The one-day rolling horizon provides a myopic baseline, the two-day horizon provides limited look-ahead, and a static seven-day solve provides a full-horizon reference without feedback re-optimization. Table 10 summarizes cumulative implemented shortage, robot energy, terminal minimum battery, and computation time.

The longer horizon does not dominate the myopic policy on shortage in this single deterministic path: $H = 1$ attains the lowest cumulative Z_1 , whereas $H = 2$ preserves substantially more terminal battery and uses less robot energy. The static full-horizon solution is still more conservative in robot use and has the highest cumulative shortage. This illustrates that horizon length changes the balance between immediate service and resource conservation; it does not establish a universally optimal horizon. Figure 2 shows the day-by-day shortage trajectories for the two rolling policies. The comparison also provides myopic and static full-horizon references under the proposed formulation. The seven-day experiment is not intended as a probabilistic robustness assessment; such an assessment would require repeated random scenarios or an explicit uncertainty set or distribution.

The controlled seven-day path includes a critical-demand increase, short-notice nurse absence, and a robot battery disturbance. It is used to verify state propagation and re-optimization after disturbances, not to estimate event frequencies or clinical robustness.

Table 4 Synthetic Benchmark Dimensions


Class	Nurses	Wards	Robots	Periods	Seeds
B20	20	3	3	3	5
B30	30	4	4	3	5
B45	45	5	6	3	5
B60	60	6	8	3	1

Table 5 Benchmark Results under the $0.60Z_1^{op}$ Shortage Budget

Class	Z_1^{op}	Achieved Z_1	Canonical reduction (%)	Z_2^{op}	Z_2 at budget	ΔZ_2 (%)	Termination
B20	96.72 ± 8.24	55.29 ± 4.77	42.80 ± 2.22	47.51 ± 0.38	49.17 ± 0.66	3.50 ± 1.36	5 optimal
B30	116.23 ± 10.58	67.51 ± 4.79	41.83 ± 1.51	67.27 ± 0.18	70.59 ± 0.65	4.94 ± 0.84	3 optimal, 2 time-limit
B45	117.89 ± 12.88	68.10 ± 6.08	42.13 ± 1.59	82.50 ± 3.19	94.29 ± 11.12	14.19 ± 11.69	5 time-limit
B60 (seed 1)	121.30	72.17	40.51	95.07	104.46	9.88	time-limit

Table 6 Multiple Epsilon Budgets for B30 Seed 1

Factor	Budget ϵ	Achieved Z_1	Slack	Z_2	Solver status
1.00	117.783	117.783	0.000	67.443	optimal
0.85	100.115	95.940	4.175	67.443	optimal
0.70	82.448	77.020	5.428	71.121	time-limit
0.60	70.670	68.345	2.325	71.271	optimal

**Figure 1** Epsilon-Budget Sensitivity for B30 Seed 1. Labels denote the budget factor relative to Z_1^{op} . The 0.70 point (marked with an asterisk) is a time-limited feasible incumbent. The connecting segments are shown only as a visual guide and are not interpreted as a continuous Pareto frontier.

5.7 Computational Complexity and Scalability

The model is a MILP and contains nurse rostering as a special case when robot variables are fixed or removed. Nurse rostering with coverage requirements and forbidden shift sequences is strongly NP-hard (den Hartog et al., 2023); therefore, the integrated problem is also NP-hard. Ignoring lower-order auxiliary variables, the principal binary-variable count is

$$|\mathcal{N}||\mathcal{W}||\mathcal{T}| + |\mathcal{R}||\mathcal{W}||\mathcal{T}| + |\mathcal{R}||\mathcal{W}||\mathcal{T}||\mathcal{L}| + |\mathcal{R}||\mathcal{T}| + |\mathcal{W}||\mathcal{T}||\mathcal{S}||\mathcal{G}|, \quad (39)$$

and the supervision-load variables add $|\mathcal{N}||\mathcal{W}||\mathcal{T}|$ continuous variables. The dominant constraint families scale with the same products, especially robot task assignment and skill-mode coverage.

The one-day benchmark is therefore best viewed as preliminary resource/ward scaling rather than a proof of multi-week scalability. In the constrained runs, B20 terminates with five solver-optimal statuses under the configured tolerances; B30 has three such terminations and two time-limit incumbents; all five B45 runs and the B60 stress case reach the time limit. The largest reported time-limited MIP optimality gap is 27.99%. This is a solver optimality gap, not the LP integrality gap. A formal LP integrality-gap analysis would require the root-relaxation values, which were not recorded in the present experiments. Therefore, a numerical LP integrality gap is not reported. Decomposition, rolling fixing, warm starts,

and Benders-or logic-based separation remain future scalability directions.

6. DISCUSSION

The two objectives answer different managerial questions. Z_1 measures inadequate clinical and support coverage, while Z_2 summarizes workload spread, overtime, night assignments, preferences, robot energy use, and charging. The epsilon-constraint procedure is useful when managers can state a service-shortage tolerance without defining one universal exchange rate between shortage and operational burden.

The task-linked relief and explicit supervision formulation alters the interpretation of robot value. Robot presence alone creates no staffing credit; relief is tied to executed tasks, and supervision is charged against nurse capacity. The direct baseline comparison shows why this distinction matters: on the common-budget B20 test, the integrated model reduces total shortage mainly by reducing robot-service shortage, not by claiming that robots replace specialized nursing work.

The relief and supervision coefficients should be estimated from observed task completion and supervision time. The local sensitivity results should be interpreted cautiously because the observed changes are small relative to the reported solver gaps. Objective-weight sensitivity is stronger: different weight settings produce different schedules and component

Table 7 Common-Budget Baseline Comparison on B20 Seed 1 ($Z_1 \leq 60$)

Method	Human shortage	Robot-service shortage	Total Z_1	Z_2	Reported MIP gap
Nurse-only	39.250	20.205	59.455	50.664	4.99%
Sequential nurse-robot	37.825	20.205	58.030	51.314	0.00%
Integrated joint	47.263	6.435	53.698	50.518	5.00%

Table 8 Local Sensitivity to Robot-Relief and Supervision Coefficients under a Common Shortage Budget

Case	Achieved Z_1	Z_2	Reported MIP gap	Runtime (s)
Relief 0.5×	82.796	29.561	0.69%	5.04
Baseline	82.841	29.690	1.38%	5.76
Relief 1.5×	82.812	30.353	4.15%	4.09
Supervision 0.5×	82.956	29.868	2.37%	3.80
Supervision 1.5×	82.564	29.572	0.82%	6.67

Table 9 Illustrative Objective-Weight Sensitivity on One Fixed Synthetic Instance

Setting	Achieved Z_1	Energy	Charges	Night shifts	Preference sum	Status
Baseline	82.841	63.919	1	3	5	optimal
Fairness-heavy	76.673	84.469	1	5	15	optimal
Energy/charging-heavy	76.080	62.785	0	3	7	optimal
Preference-heavy	83.105	63.919	0	3	12	time-limit

Table 10 Prediction-Horizon and Static Full-Horizon Comparison on the Seven-Day Stress Path

Policy	Cum. human Z_1	Cum. robot Z_1	Cum. total Z_1	Robot energy	Terminal min. battery	Runtime (s)
Rolling $H = 1$	300.962	111.693	412.655	65.047	42.619	7.97
Rolling $H = 2$	301.971	120.992	422.963	16.543	62.640	7.92
Static full horizon	310.459	123.467	433.926	0.000	73.417	33.58

values. Consequently, operational deployment would require component scaling, stakeholder elicitation, and sensitivity checks rather than adopting the synthetic coefficients unchanged.

The seven-day comparison also clarifies the role of “prediction.” The model is deterministic rolling-horizon re-optimization with exogenous look-ahead values, not a learned forecasting system. Increasing the horizon from one to two days preserved more battery on the tested path but did not reduce cumulative shortage. The result is a horizon-sensitivity observation, not evidence that a particular horizon is universally superior. Likewise, the controlled disturbances do not establish stochastic robustness; that would require Monte Carlo testing, scenario-based stochastic programming, or a robust uncertainty set.

The battery and mobility representations remain planning-level approximations. Equation (18) uses constant ward-period energy consumption and a constant charging increment. Robot movement is aggregated into ward-period task capacity and energy, so inter-ward routing, corridor/elevator congestion, collision avoidance, travel-time uncertainty, and charging-station occupancy are omitted. These simplifications can make a ward-level schedule optimistic when travel or charger competition consumes a material fraction of a period. Telemetry-based piecewise-linear or nonlinear battery models and explicit routing would be required for higher-fidelity deployment studies.

Important multi-week nurse-rostering rules are also outside the reported benchmark: maximum consecutive working days, weekly contracted-hour patterns, weekend rules, shift rotations, and longer recovery sequences are not activated. They can be added as standard inter-day constraints, but doing so changes both model size and validation requirements. The complexity analysis and the time-limited B45/B60 runs show why such extensions should be accompanied by larger computational studies and, potentially, decomposition or warm-start strategies.

The evidence remains synthetic. The reference example is a numerical verification case; the larger instances are controlled engineering tests. Real deployment would require de-identified hospital data, clinically approved skill-substitution rules, institution-specific labor rules, supervision eligibility, calibrated objective weights, robot telemetry, and prospective validation with clinical and operational stakeholders.

7. CONCLUSION

This study develops a bi-objective MILP for joint nurse and hospital-service-robot scheduling with heterogeneous nursing skills, alternative coverage modes, task-linked robot assistance, charging and battery dynamics, and explicit human-supervision workload. A deterministic receding-horizon layer carries workload, night-history, and battery states across successive optimization windows.

The task-execution and supervision formulation yields a reference solution of $Z_1 = 0.78$ and $Z_2 = 3.164$. On a common-budget B20 seed-1 comparison, integrated joint optimization obtains total shortage 53.70, compared with 58.03 for sequential nurse-robot scheduling and 59.46 for nurse-only scheduling. Multiple epsilon levels on B30 seed 1 report achieved shortage and budget slack explicitly rather than treating the imposed budget as an observed improvement. Sensitivity experiments show that model outputs depend on parameter and objective-weight choices, and a seven-day comparison shows different service-battery trade-offs for one-day and two-day rolling horizons. The larger instances also expose computational limits: several B45/B60 runs remain time-limited, so the study does not claim exact large-scale optimality or practical hospital validation.

The results support the internal consistency and comparative value of the integrated formulation on controlled synthetic tests, while also defining the evidence still needed for deployment. Future work should use de-identified hospital data, explicit routing and charging-station constraints, telemetry-calibrated battery models, richer multi-week nurse rules, and stochastic or robust uncertainty models, together with scalable decomposition methods for larger horizons.

APPENDIX A. REFERENCE NUMERICAL DATA

Table 11 lists the nurse-skill qualifications used in the reference example. Table 12 reports the corresponding ward- and shift-specific human coverage requirements. Table 13 summarizes the initial battery levels, charging increments, ward-specific energy use, and service capabilities of the reference robots. The reference case is kept separate from the synthetic benchmark generator and is used only to check the consistency of the formulation and computational results. The reference case is kept separate from the synthetic benchmark generator and is used only to check the




Figure 2 Implemented-Day Weighted Shortage Trajectories for the One-Day and Two-Day Rolling Horizons on the Seven-Day Deterministic Stress Path. The static full-horizon policy is summarized in Table 10 because it is solved once rather than as a sequence of daily rolling decisions.

consistency of the formulation and computational results.

Table 11 Nurse Skills in the Reference Example

Nurse	General	ICU	Senior
N1	1	1	1
N2	1	1	1
N3	1	1	0
N4	1	0	1
N5	1	1	0
N6	1	0	0
N7	1	0	0
N8	1	0	0

Table 12 Human Coverage Requirements in the Reference Example

Shift	Ward	General	ICU	Senior
Morning	ICU	2.5	1	1
Morning	General	2.0	0	1
Evening	ICU	1.5	1	1
Evening	General	1.0	0	0
Night	ICU	1.0	1	0
Night	General	1.5	0	0

Table 13 Reference Robot Energy and Capability Data

Robot	B_0	Charge	ICU use	Gen. use	Transport	Disinfection
R1	80	40	30	25	1	0
R2	60	40	35	30	1	1

The robot-service requirements are one transport service in ICU during the morning, one transport service in ICU during the evening, and one disinfection/support service in the General ward at night. Human supervision is set to one active robot per assigned nurse in either ward. Under the proposed formulation, the returned robot plan is R1: ICU transport, idle, charge; R2: charge, ICU transport, General-ward disinfection. The corresponding battery trajectories are (80, 50, 50, 90) for R1 and (60, 100, 65, 35) for R2. The reference objective values are $Z_1 = 0.78$ and $Z_2 = 3.164$.

REFERENCES

Borrelli, F., Bemporad, A., & Morari, M. (2017). *Predictive control for linear and hybrid systems*. Cambridge University Press. <https://doi.org/10.1017/9781139061759>

Burke, E. K., De Causmaecker, P., Vanden Berghe, G., & Van Landeghem, H. (2004). The state of the art of nurse rostering. *Journal of Scheduling*, 7(6), 441–499. <https://doi.org/10.1023/B:JOSH.0000046076.75950.0b>

Camacho, E. F., & Bordons, C. (2007). *Model predictive control* (2nd ed.). Springer. <https://doi.org/10.1007/978-0-85729-398-5>

Cheang, B., Li, H., Lim, A., & Rodrigues, B. (2003). Nurse rostering problems: A bibliographic survey. *European Journal of Operational Research*, 151(3), 447–460. [https://doi.org/10.1016/S0377-2217\(03\)00021-3](https://doi.org/10.1016/S0377-2217(03)00021-3)

De Bruecker, P., Van den Bergh, J., Beliën, J., & Demeulemeester, E. (2015). Workforce planning incorporating skills: State of the art. *European Journal of Operational Research*, 243(1), 1–16. <https://doi.org/10.1016/j.ejor.2014.10.038>

den Hartog, S. J. M., Hoogeveen, H., & van der Zanden, T. C. (2023). On the complexity of nurse rostering problems. *Operations Research Letters*, 51(5), 483–487. <https://doi.org/10.1016/j.orl.2023.07.004>

Ernst, A. T., Jiang, H., Krishnamoorthy, M., & Sier, D. (2004). Staff scheduling and rostering: A review of applications, methods and models. *European Journal of Operational Research*, 153(1), 3–27. [https://doi.org/10.1016/S0377-2217\(03\)00095-X](https://doi.org/10.1016/S0377-2217(03)00095-X)

Fragapane, G., de Koster, R., Sgarbossa, F., & Strandhagen, J. O. (2021). Planning and control of autonomous mobile robots for intralogistics: Literature review and research agenda. *European Journal of Operational Research*, 294(2), 405–426. <https://doi.org/10.1016/j.ejor.2021.01.019>

Gombolay, M. C., Gutierrez, R. A., Clarke, S. G., Sturla, G. F., & Shah, J. A. (2015). Decision-making authority, team efficiency and human worker satisfaction in mixed human-robot teams. *Autonomous Robots*, 39(3), 293–312. <https://doi.org/10.1007/s10514-015-9457-9>

Huang, R., Li, H., Suomi, R., Li, C., & Peltoniemi, T. (2023). Intelligent physical robots in health care: Systematic literature review. *Journal of Medical Internet Research*, 25, Article e39786. <https://doi.org/10.2196/39786>

Huangfu, Q., & Hall, J. A. J. (2018). Parallelizing the dual revised simplex method. *Mathematical Programming Computation*, 10(1), 119–142. <https://doi.org/10.1007/s12532-017-0130-5>

Jalaeian Farimani, M., Khalili Amirabadi, R., Nikkhoy, D., Abdolbaghi, M., Rastegarmoghaddam, M., & Samadzadeh, S. (2026). A systematic review and taxonomy of reinforcement learning-model predictive control integration for linear systems [Preprint]. <https://doi.org/10.48550/arXiv.2604.21030>

Khalili Amirabadi, R., Jalaeian Farimani, M., & Solaymani Fard, O. (2026). Experience replay in reinforcement learning: A systematic review. *Neurocomputing*, 700, Article 134403. <https://doi.org/10.1016/j.neucom.2026.134403>

Mavrotas, G. (2009). Effective implementation of the epsilon-constraint method in multi-objective mathematical programming problems. *Applied Mathematics and Computation*, 213(2), 455–465. <https://doi.org/10.1016/j.amc.2009.03.037>

Miseikis, J., Caroni, P., Duchamp, P., Gasser, A., Marko, R., Miseikiene, N., Zwilling, F., de Castelbajac, C., Eicher, L., Fruh, M., & Fruh, H. (2020). Lio: A personal robot assistant for human-robot interaction and care applications. *IEEE Robotics and Automation Letters*, 5(4), 5339–5346. <https://doi.org/10.1109/LRA.2020.3007462>

Morgan, A. A., Abdi, J., Syed, M. A. Q., El Kohen, G., Barlow, P., & Vizcaychipi, M. P. (2022). Robots in healthcare: A scoping review. *Current Robotics*

- Reports*, 3, 271–280. <https://doi.org/10.1007/s43154-022-00095-4>
- Qin, S. J., & Badgwell, T. A. (2003). A survey of industrial model predictive control technology. *Control Engineering Practice*, 11(7), 733–764. [https://doi.org/10.1016/S0967-0661\(02\)00186-7](https://doi.org/10.1016/S0967-0661(02)00186-7)
- Van den Bergh, J., Beliën, J., De Bruecker, P., Demeulemeester, E., & De Boeck, L. (2013). Personnel scheduling: A literature review. *European Journal of Operational Research*, 226(3), 367–385. <https://doi.org/10.1016/j.ejor.2012.11.029>
- Yang, G.-Z., Nelson, B. J., Murphy, R. R., Choset, H., Christensen, H., Collins, S. H., Dario, P., Goldberg, K., Ikuta, K., Jacobstein, N., Kragic, D., Taylor, R. H., & McNutt, M. (2020). Combating COVID-19: The role of robotics in managing public health and infectious diseases. *Science Robotics*, 5(40), Article eabb5589. <https://doi.org/10.1126/scirobotics.abb5589>
- Zhang, R., Rossi, F., & Pavone, M. (2016). Model predictive control of autonomous mobility-on-demand systems. In *2016 IEEE International Conference on Robotics and Automation* (pp. 1382–1389). IEEE. <https://doi.org/10.1109/ICRA.2016.7487272>