RESEARCH ARTICLE

Energy-Aware Operating Theatres Scheduling Using Metaheuristics

Latifa Dekhici[1,2,*], **Khaled Guerraiche**[2], **Jihen Jlassi**[3] and **Mohammed Adel Djari**[2]

[1] Department of Computer Science, University of Science and Technology of Oran (USTOMB), Oran 31000, Algeria
[2] LDREI Laboratory, Department of Electrical Engineering, Higher School of Electrical and Energy Engineering of Oran (ESGEEO), Oran, Algeria
[3] OLID LAB, Higher Institute of Industrial Management of Sfax (ISGI), University of Sfax, Sfax 3029, Tunisia

Abstract

Operating rooms (ORs) account for a significant fraction of a hospital's energy footprint, yet traditional surgical scheduling models ignore the dynamic energy consumption of heating, ventilation, and air conditioning (HVAC) systems, particularly the thermal inertia and peak restart loads. Furthermore, the interaction with the Post-Anesthesia Care Unit (PACU) creates blocking effects that invalidate isolated OR schedules. This paper introduces a bi-objective mixed-integer linear programming (MILP) model and a metaheuristic framework for the Green Ambulatory Surgery Scheduling Problem with PACU constraints (G-ASSP-P). The objectives minimise both total completion time (makespan) and total energy consumption, explicitly modelling HVAC transition penalties. To solve large, realistic instances, a discrete adaptation of the Interior Search Algorithm (ISA) is adopted and compared against the Non-dominated Sorting Genetic Algorithm II (NSGA-II). Computational experiments on three instance scales (20, 50, 70 surgeries) with 4–12 ORs and 8–25 PACU beds demonstrate that NSGA-II consistently outperforms ISA in terms of hypervolume and solution diversity. The trade-off analysis reveals that substantial energy savings (exceeding 30%) can be achieved with only a marginal increase in makespan, offering hospital managers a practical tool to balance efficiency and sustainability.

Keywords: operating room scheduling, energy efficiency, multi-objective optimisation, HVAC, Healthcare, metaheuristics, NSGA-II, interior search algorithm.

1 Introduction

Healthcare systems are under increasing pressure to reduce their environmental footprint while maintaining high standards of patient care. Operating rooms are among the most resource-intensive areas, consuming 3–6 times more energy per square foot than other hospital spaces due to stringent ventilation requirements. A large portion of this energy is wasted during idle periods when ORs remain in a sterile-ready state despite no surgical activity. By

Submitted: 08 May 2026
Revised: 08 July 2026
Accepted: 06 August 2026
Published: 26 August 2026

Vol. 1, **No.** 1, 2026.
10.62762/ICE.2026.457366

***Corresponding author:**
Latifa Dekhici
latifa.dekhici@univ-usto.dz

Citation
Dekhici, L., Guerraiche, K., Jlassi, J., & Djari, M. A. (2026). Energy-Aware Operating Theatres Scheduling Using Metaheuristics. *Intelligent Computing for Engineering*, 1(1), 17–29.

intelligently scheduling surgeries to create longer idle blocks, HVAC systems can be switched to energy-saving eco modes, leveraging thermal inertia to reduce consumption.

However, scheduling in an ambulatory surgery center (ASC) is not solely an OR problem. The flow of patients into the Post-Anesthesia Care Unit (PACU) introduces a critical blocking constraint: a patient cannot be transferred out of the OR if no PACU bed is available. This coupling transforms the problem into a complex, multi-resource optimisation challenge.

The scheduling of operating theatres has been a central topic in healthcare operational research for decades, as highlighted by Beliën et al. [1] in their comprehensive review of fifty years of OR applied to healthcare. Over the years, models have evolved from simple assignment rules to sophisticated combinatorial optimisation problems involving multiple resources, stages, and specialties. In parallel, the increasing digitalization of hospitals, driven by the Internet of Things (IoT), smart sensors, and cloud computing, has opened new avenues for real-time monitoring and resource allocation. However, the translation of these technologies into energy-aware scheduling tools remains limited, as demonstrated by the scarcity of models that jointly optimise scheduling decisions and physical energy systems such as HVAC [2, 3].

Traditional OR scheduling often aims to minimise makespan or overtime while respecting the availability of surgeons, rooms, and downstream units such as the PACU. El Amraoui et al. [4] proposed a heuristic algorithm for the resource-constrained OR scheduling problem with shared surgical teams, demonstrating an 18–20% overtime reduction compared to manual schedules. Similarly, Vieira et al. [5] developed a random-key optimiser framework for integrated operating room scheduling that decouples the problem encoding from the solving algorithm, allowing the same metaheuristics to be applied across different hospital case studies. Liu et al. [6] addressed a multi-specialty, multi-stage elective surgery scheduling problem and introduced a self-optimised learning algorithm combining reinforcement learning, genetic algorithms, and heuristic rules, achieving substantial improvements in solution accuracy.

Accurate estimation of surgery durations is critical for building reliable schedules. Spence et al. [7] conducted a scoping review of machine learning models for predicting surgical case duration, finding that artificial intelligence consistently outperforms current industry standards in accuracy. Such predictive models can feed into scheduling algorithms to better handle the stochasticity of surgery durations, a critical input for robust OR scheduling. Zhang et al. [8] studied a surgery scheduling problem with setup times and proposed a Q-learning-based artificial bee colony algorithm that dynamically selects local search operators, showing superior performance on makespan minimisation. More recently, machine learning has been integrated with fuzzy optimisation for ambulatory surgery scheduling [9], demonstrating the potential of hybrid approaches in addressing the complexity of surgical scheduling problems.

Several studies have addressed operating room scheduling under various constraints. Dekhici and Khaled [10] model this problem as a hybrid flow shop system without buffers and propose a hybrid approach combining local search and tabu search to satisfy operational constraints while improving performance indicators such as makespan, based on real hospital data. This connection is particularly relevant to recent research on energy-aware hybrid flow-shop scheduling. Dekhici et al. [11] investigated machine-learning-embedded energy cost optimization for hybrid flow-shop scheduling, demonstrating how intelligent scheduling approaches can incorporate energy considerations into production planning. While such approaches have primarily been developed for industrial systems, their underlying scheduling principles are relevant to operating-room environments, where resource allocation, sequencing, timing, and energy consumption must also be jointly considered. In a more recent contribution, Atisattapong and Pongpatcharapun [12] develop a mixed-integer linear programming model combined with a genetic algorithm to generate master surgical schedules that maximise operating room utilisation and reduce idle time, waiting time, and overtime; their results show that exact methods are effective for small instances, while metaheuristics are more suitable for large-scale problems. Similarly, Latorre- Núñez et al.[13] extend the problem by integrating multiple hospital resources such as operating rooms, post-anesthesia beds, medical staff, and equipment, while also considering emergency surgeries; they propose an integer linear programming model and hybrid solution approaches, including constraint programming and genetic algorithms, to handle realistic large-scale instances. Dekhici and Belkadi [14] proposed a bi-objective paediatric operating theatre scheduling problem subject to order and assignment constraints, minimising makespan while maximising constraint satisfaction. They compared firefly algorithm, bat algorithm, particle swarm optimisation and local search, using average deviation from the lower bound to validate effectiveness on a real paediatric hospital in Oran. Belkhamsa et al. [15] tackled a no-wait operating room surgery scheduling problem with resource constraints across preoperative, intraoperative and postoperative stages. Their iterated local search and hybrid genetic algorithm outperformed ant colony optimisation, achieving up to 24% reduction in end time and 70% reduction in total idle time. More recently, Mostofa et al. [16] conducted a narrative review on case sequencing under PACU availability constraints, concluding that sequencing the shortest or least variable cases first is counterproductive when PACU beds are limited. Instead, random sequencing (as often observed in practice) yields more uniform PACU admission rates and minimises peak demand. These findings directly support our model's emphasis on PACU blocking and its impact on both makespan and energy efficiency.

The environmental footprint of healthcare, particularly the energy consumption of operating rooms, has recently attracted attention. Wang et al. [17] provided a systematic review of ventilation and indoor air quality management in ORs, identifying pollutant sources and advanced control strategies that directly affect HVAC energy use. From a scheduling perspective, Chen et al. [18] introduced a prescriptive tree-based model for energy-efficient room scheduling that accounts for uncertainties in both energy consumption and renewable energy generation. Their model generates direct scheduling solutions using contextual information such as temperature, humidity, and sunlight, and was validated on real-world building data. This approach demonstrates the potential of data-driven methods to align surgical scheduling with sustainability goals, yet it does not capture the specific constraints of OR environments such as sterile airflow requirements or PACU blocking.

Recent work has increasingly addressed the joint optimisation of patient flow and downstream resource constraints in surgical scheduling. Zhan et al. [19] addressed a practical surgery scheduling problem with limited PACU capacity and stochastic surgery and recovery durations, proposing a two-stage stochastic programming model that jointly optimises OR usage, surgery-to-OR assignments, and surgery sequencing. Their decomposition algorithm achieves an average operational cost reduction of 29%, providing strong empirical evidence that explicitly modelling PACU capacity constraints leads to more effective utilisation of both ORs and PACU resources. Similarly, Shehadeh and Padman [3] provided a comprehensive review of stochastic optimisation approaches for elective surgery scheduling with downstream capacity constraints, covering stochastic programming, robust optimisation, and distributionally robust optimisation, and underscoring that ignoring PACU and ICU capacity leads to systematically suboptimal and infeasible schedules in practice. Cheng et al. [2] proposed an environmentally friendly elective patient scheduling model that reduces hospital carbon emissions by minimising the number of open operating rooms, embedding machine-learning predictions within a distributionally robust optimisation framework. Their results demonstrate that green scheduling objectives can be pursued without sacrificing service quality, directly motivating the energy-aware formulation developed in the present work.

Despite these advances, the literature still lacks a model that simultaneously: (i) captures the dynamic energy consumption of HVAC systems with explicit thermal inertia and restart penalties, (ii) accounts for PACU blocking effects that force ORs to remain active, and (iii) optimizes both makespan and energy consumption in a single bi-objective framework. Most existing energy-aware OR scheduling studies treat HVAC consumption as a static load or apply post-hoc energy estimation rather than integrating it into the optimization model. Conversely, studies that address PACU-constrained scheduling focus primarily on patient flow and resource utilization without considering energy implications. The integration of these two critical dimensions—PACU-blocking dynamics and HVAC thermal inertia—within a unified optimization framework remains unexplored. The present paper addresses this gap by proposing a bi-objective MILP that incorporates detailed HVAC energy dynamics, PACU blocking, and ambulatory patient flow, and by solving it with a novel discrete Interior Search Algorithm benchmarked against NSGA-II.

This paper makes the following contributions:

1. A novel bi-objective MILP formulation is proposed that captures the trade-off between makespan (including PACU discharge) and total energy consumption, incorporating HVAC peak loads and thermal inertia thresholds.

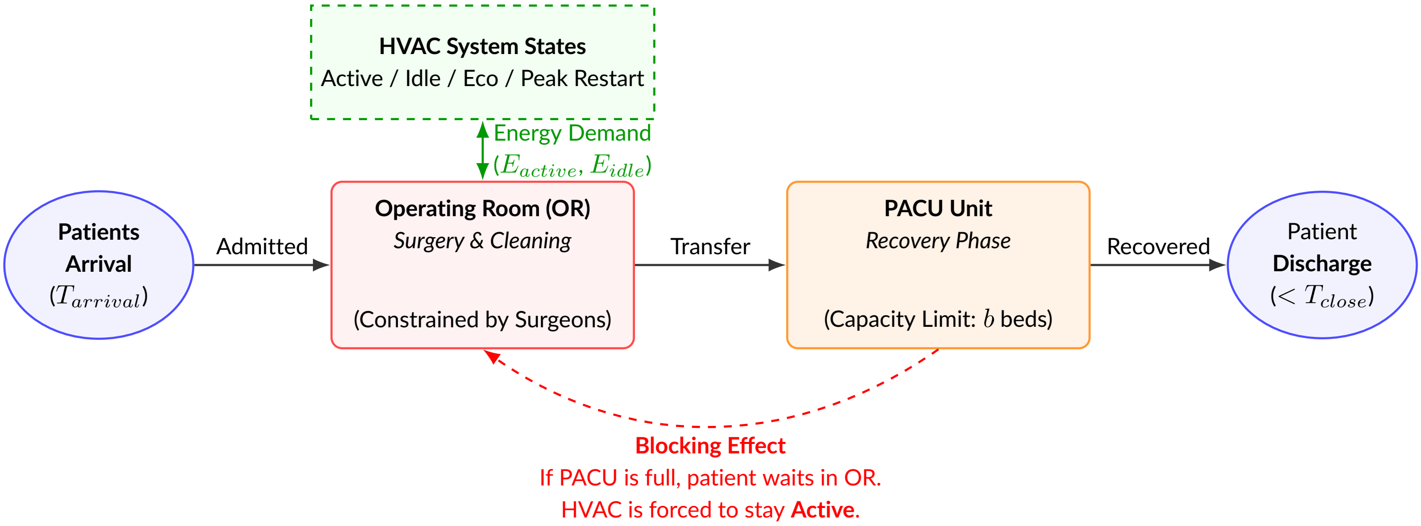

Figure 1. Patient and energy flux in the proposed Green Ambulatory Surgery Scheduling Problem

2. A discrete adaptation of the Interior Search Algorithm (ISA) for solving this scheduling problem is compared with the well-established NSGA-II.

3. Realistic benchmark instances of three sizes (20, 50, 70 surgeries) are generated to provide a comprehensive analysis of the Pareto-optimal trade-off curves, offering guidance on selecting operating points based on managerial priorities.

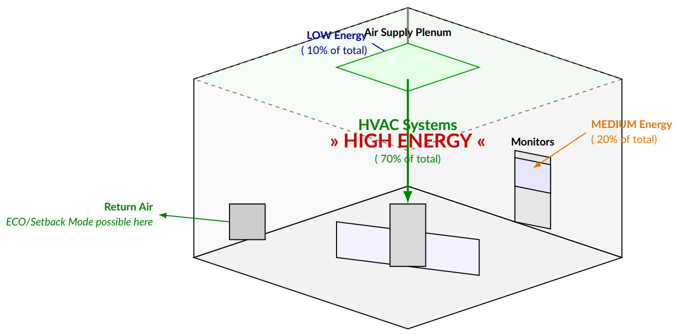

Figure 2. Isometric conceptual mapping of energy consumption within an active operating theatre.

2 Problem Definition

2.1 Energy in Operating Theatre Scheduling

Figure 1 illustrates the integrated patient flow and HVAC energy dynamics in the proposed Green Ambulatory Surgery Scheduling Problem with PACU constraints (G-ASSP-P). The standard patient pathway (blue arrows) begins with patient arrival, proceeds to surgery in an OR, then to recovery in a PACU bed, and finally to discharge. The HVAC system (green) provides conditioned air to each OR and can operate in active, idle, or eco modes. The critical blocking effect (red dashed arrow) occurs when all PACU beds are occupied: a patient cannot leave the OR, forcing the HVAC to remain in active mode and increasing energy consumption. This coupling between OR scheduling and PACU capacity is the core challenge addressed by our bi-objective model.

To understand why HVAC dominates OR energy use, Figure 2 provides an isometric conceptual view of a typical active operating theatre. Surgical lights (bottom-left) and monitoring equipment (right) draw low to medium electrical power, together accounting for roughly 30% of the total energy demand. In contrast, the vast majority ($\approx$70%) is consumed by the HVAC system, which continuously supplies massive volumes of conditioned, sterile air through ceiling plenums and laminar flow diffusers. This "invisible" load is the primary target for energy savings: when ORs are idle, switching to eco mode drastically reduces airflow and power consumption, albeit with a restart penalty due to thermal inertia.

Figure 3 presents the step-function power profile that models HVAC energy consumption over time. During a surgery (Surgery $1, 2, 3$), the system operates at active power P_{active}. After a surgery ends, a short idle gap (duration $\leq \tau^{\text{ramp}}$) keeps the system at P_{idle} to avoid the cost of a restart. If the gap exceeds τ^{ramp}, the system drops to eco mode at P_{eco} (long gap). However, before the next surgery, it must ramp back to active, incurring an extra peak penalty P_{peak} (red area) that accounts for reheating, re-humidification, and re-sterilization. The shaded areas under the curve represent the total energy E_k for an OR k. This piecewise linear model is integrated into our MILP using binary variables to detect long gaps.

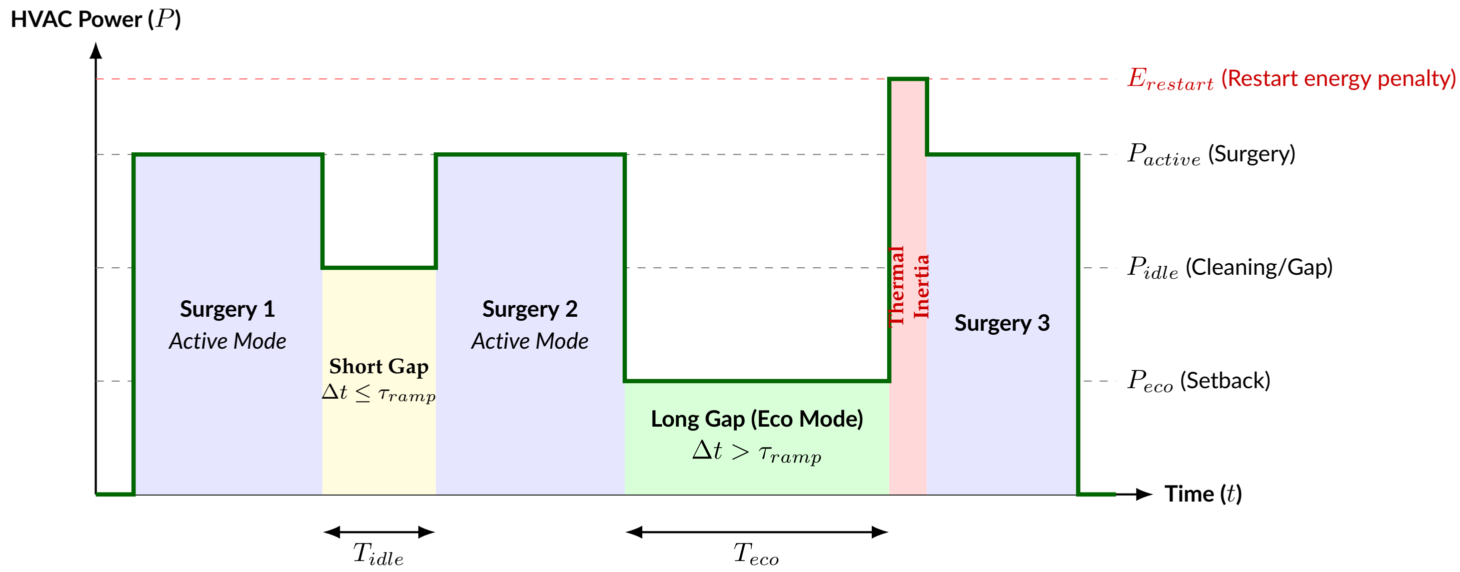

Figure 3. Step-function representation of the dynamic HVAC energy consumption model.

2.2 Mathematical Model

Consider a set of surgeries $\mathcal{P} = \{1, \ldots, n\}$ to be scheduled in a set of ORs $\mathcal{O} = \{1, \ldots, m\}$ and a PACU with beds $\mathcal{B} = \{1, \ldots, b\}$. Each surgery i has a stochastic duration d_i, a PACU length of stay c_i^{PACU}, an arrival window start $T_{\text{arrival},i}$, and is performed by a surgeon from set $\mathcal{S}$. The facility closes at T_{close}.

Each OR k has three power states:

- P_k^{active}: Power during surgery or cleaning (sterile airflow).
- P_k^{idle}: Power when idle but maintained in sterile-ready mode.
- P_k^{eco}: Power in setback mode (reduced airflow).

If an OR is idle for longer than a threshold τ_k^{ramp}, it can transition to eco mode. Restarting from eco to active incurs an extra energy penalty E_k^{restart} to condition the space. The total energy for OR k is:

$$E_k = P_k^{\text{active}} \cdot T_k^{\text{active}} + P_k^{\text{idle}} \cdot T_k^{\text{idle,short}} + P_k^{\text{eco}} \cdot T_k^{\text{idle,long}} + E_k^{\text{restart}} \cdot N_k^{\text{transitions}}$$

2.3 Bi-Objective MILP Formulation

Decision Variables:

- $x_{ik} \in \{0, 1\}$: 1 if surgery i is assigned to OR k.
- $y_{ijk} \in \{0, 1\}$: 1 if surgery i is immediately followed by surgery j in OR k.
- $z_{i\ell} \in \{0, 1\}$: 1 if surgery i is assigned to PACU bed ℓ.
- $t_i^s \in \mathbb{R}_{\geq 0}$: Start time of surgery i.
- $t_i^e \in \mathbb{R}_{\geq 0}$: End time of surgery i (OR release).
- $t_i^{\text{PACU\_in}} \in \mathbb{R}_{\geq 0}$: Admission time to PACU.
- $t_i^{\text{PACU\_out}} \in \mathbb{R}_{\geq 0}$: Discharge time from PACU.

Auxiliary variables δ_{ijk} and ρ_{ijk} for HVAC linearisation are defined in subsection 2.4.

Objectives:

$$\min\ f_1 = \max_{i \in \mathcal{P}} \{t_i^{\text{PACU\_out}}\} \tag{1}$$

$$\min\ f_2 = \sum_{k \in \mathcal{O}} E_k \tag{2}$$

Constraints:

$$\sum_k x_{ik} = 1, \quad \forall i \tag{3}$$

$$t_j^s \geq t_i^e + s_{ij} - M(1 - y_{ij}), \quad \forall i, j \tag{4}$$

$$t_i^{\text{PACU\_in}} \geq t_i^e, \quad \forall i \tag{5}$$

$$\sum_\ell z_{i\ell} = 1, \quad \forall i \tag{6}$$

$$t_j^{\text{PACU\_in}} \geq t_i^{\text{PACU\_out}} - M(2 - z_{i\ell} - z_{j\ell}), \quad \forall i \neq j, \ell \tag{7}$$

$$t_i^s \geq T_{\text{arrival},i}, \quad \forall i \tag{8}$$

$$t_i^{\text{PACU\_out}} \leq T_{\text{close}}, \quad \forall i \tag{9}$$

$$t_j^s \geq t_i^e \quad \text{or} \quad t_i^s \geq t_j^e, \quad \forall i, j \text{ with same surgeon} \tag{10}$$

The energy component E_k is linearized using auxiliary binary variables δ_{ij} indicating whether the gap between consecutive surgeries exceeds τ_k^{ramp}. Due to space constraints, the linearization details are omitted here but follow standard MILP techniques for peak penalties.

2.4 Linearisation of HVAC Energy Consumption

To formulate the energy term E_k in Eq. (2) as a linear expression, we introduce additional parameters, auxiliary binary variables, and Big-M constraints.

Parameters:

- s_{ij}: Setup/cleaning time between surgery i and surgery j when performed consecutively in the same OR.
- M: A sufficiently large constant (e.g., $M = T_{\text{close}} + \max_i d_i$).
- τ_k^{ramp}: Maximum idle duration in OR k before HVAC switches from idle to eco mode (thermal inertia threshold).
- $P_k^{\text{active}}, P_k^{\text{idle}}, P_k^{\text{eco}}$: Power consumption (kW) for active, idle, and eco states, respectively.
- E_k^{restart}: Additional energy penalty (kWh) incurred each time the HVAC system restarts from eco to active mode (includes reheating, re-humidification, and re-sterilization).

Auxiliary binary variables:

- $\delta_{ijk} \in \{0, 1\}$: Equals 1 if surgery i is immediately followed by surgery j in OR k and the idle gap between them exceeds τ_k^{ramp} (triggering an eco-state transition and a restart penalty).
- $\rho_{ijk} \in \mathbb{R}_{\geq 0}$: Continuous variable representing the product $gap_{ijk} \cdot \delta_{ijk}$, used to linearise the energy calculation.

Idle gap definition: For consecutive surgeries i and j assigned to the same OR k (i.e., $y_{ijk} = 1$), the idle gap is:

$$gap_{ijk} = t_j^s - \left(t_i^e + s_{ij}\right).$$

If $y_{ijk} = 0$, the gap is ignored and $\delta_{ijk} = 0$.

Big-M constraints to detect long gaps: To determine whether the gap exceeds the ramp threshold, we impose:

$$gap_{ijk} \leq \tau_k^{\text{ramp}} + M \cdot (1 - \delta_{ijk}), \tag{11}$$

$$gap_{ijk} \geq \tau_k^{\text{ramp}} + \varepsilon - M \cdot \delta_{ijk}, \tag{12}$$

where ε is a small positive constant (e.g., 1 minute) to enforce strict inequality for long gaps. Constraint (11) ensures that if $\delta_{ijk} = 0$, the gap is unconstrained (can be long or short); if $\delta_{ijk} = 1$, the gap must be $\leq \tau$. Constraint (12) ensures that if $\delta_{ijk} = 1$, the gap must be strictly greater than τ; if $\delta_{ijk} = 0$, the constraint is relaxed.

Linearisation of bilinear products: The energy consumed during an idle gap depends on the state (idle or eco):

$$E_{ijk}^{\text{idle}} = \begin{cases} P_k^{\text{idle}} \cdot gap_{ijk}, & \text{if } \delta_{ijk} = 0, \\ P_k^{\text{eco}} \cdot gap_{ijk} + E_k^{\text{restart}}, & \text{if } \delta_{ijk} = 1. \end{cases}$$

To linearise this, we introduce $\rho_{ijk} = gap_{ijk} \cdot \delta_{ijk}$ with the following Big-M constraints:

$$\rho_{ijk} \geq gap_{ijk} - M \cdot (1 - \delta_{ijk}), \tag{13}$$

$$\rho_{ijk} \leq gap_{ijk}, \tag{14}$$

$$\rho_{ijk} \leq M \cdot \delta_{ijk}. \tag{15}$$

If $\delta_{ijk} = 0$, constraints (13)–(15) force $\rho_{ijk} = 0$. If $\delta_{ijk} = 1$, they force $\rho_{ijk} = gap_{ijk}$.

Total energy for OR k: The total energy consumption in OR k is then expressed linearly as:

$$\begin{aligned} E_k = {} & P_k^{\text{active}} \cdot T_k^{\text{active}} \\ & + \sum_{i,j} \Big[P_k^{\text{idle}} \cdot \big(gap_{ijk} - \rho_{ijk}\big) + P_k^{\text{eco}} \cdot \rho_{ijk} \\ & + E_k^{\text{restart}} \cdot \delta_{ijk} \Big], \end{aligned} \tag{16}$$

where $T_k^{\text{active}} = \sum_i d_i \cdot x_{ik}$ is the total surgery duration in OR k (cleaning time is embedded in gap_{ijk} via s_{ij}). This formulation is entirely linear and can be directly embedded into the MILP.

Note: The PACU blocking effect influences t_i^e and t_j^s through Constraint (7); when a patient cannot be transferred, the OR release time is delayed, increasing gap_{ijk} and potentially triggering $\delta_{ijk} = 1$, which in turn increases E_k via the E_k^{restart} penalty.

3 Methodology

The MILP model becomes intractable for $n > 20$ due to the sequencing and PACU blocking constraints. Therefore, two metaheuristics are investigated in this work: NSGA-II (baseline) and a novel discrete Interior Search Algorithm (ISA) [20].

3.1 Solution Encoding and Decoding

A solution is encoded as a tuple (Π, Ω), where Π is a permutation of surgery indices and Ω is a vector of OR assignments. The decoding process is a constructive forward simulation: Energy consumption is then calculated by analysing the idle gaps in each OR's timeline and applying the HVAC state logic.

Algorithm 1 Decoding with PACU Blocking

Initialize OR ready times, surgeon ready times, PACU bed available times.
for each surgery i in order Π **do**
  $k \leftarrow \Omega[i]$
  $earliest \leftarrow \max(\text{OR ready}[k] + S_{\text{clean}}, \text{Surgeon ready}, T_{\text{arrival},i})$
  while true **do**
    $end \leftarrow earliest + d_i$
    Find bed ℓ with earliest availability $\leq end$
    if bed found **then break**
    else
      $earliest \leftarrow$ bed availability $- d_i$
    end if
  end while
  Schedule surgery, update resources and PACU bed ℓ.
end for

3.2 NSGA-II Implementation

Order Crossover (OX) is used for the permutation Π and uniform crossover for Ω. Mutation operators include swap of two positions in Π and random reassignment of an OR in Ω. Non-dominated sorting with crowding distance is applied for selection.

3.3 Discrete Interior Search Algorithm (ISA)

A discrete adaptation of the Interior Search Algorithm (ISA) is used. The population evolves using two operators controlled by a parameter $\alpha \in [0, 1]$:

- **Mirroring** (α): A solution x_i is reflected toward a random member x_a from the non-dominated archive. For the permutation part, a segment from x_a is copied, and the remaining elements are filled in the order of x_i. For OR assignments, each position adopts x_a's value with probability β.
- **Moving** ($1 - \alpha$): A local random walk is performed by swapping two positions in Π and changing one OR assignment.

Algorithm 2 summarises the complete discrete ISA procedure.

Algorithm 2 Discrete Interior Search Algorithm (ISA)

Require: N (population size), $G_{\max}$ (max generations), $\alpha_{\max}, \alpha_{\min}$, β (OR assignment inheritance probability)
Ensure: Pareto archive $\mathcal{A}$
Initialise population $\mathcal{P}op$ of N random solutions (Π, Ω)
Decode each solution using Algorithm 1 to compute (f_1, f_2)
Initialise archive $\mathcal{A}$ with all non-dominated solutions from $\mathcal{P}op$
for $g = 1$ **to** $G_{\max}$ **do**
  $\alpha \leftarrow \alpha_{\max} - (\alpha_{\max} - \alpha_{\min}) \cdot (g/G_{\max})$
  $\mathcal{P}op' \leftarrow \emptyset$
  for $i = 1$ **to** N **do**
    Select x_i from $\mathcal{P}op$ (tournament selection based on crowding distance)
    if rand$(0, 1) \leq \alpha$ **then**
      Select random archive member $x_a \in \mathcal{A}$
      Mirroring to x'
    else
      Moving to x'
    end if
    Decode x'
    if x' is non-dominated w.r.t. x_i **then**
      $\mathcal{P}op' \leftarrow \mathcal{P}op' \cup \{x'\}$
    else if x_i is non-dominated w.r.t. x' **then**
      $\mathcal{P}op' \leftarrow \mathcal{P}op' \cup \{x_i\}$
    else
      $\mathcal{P}op' \leftarrow \mathcal{P}op' \cup \{x'\}$ with probability otherwise x_i
    end if
  end for
  $\mathcal{P}op \leftarrow \mathcal{P}op'$
  Update $\mathcal{A}$ with all non-dominated solutions from $\mathcal{P}op \cup \mathcal{A}$
end for
return $\mathcal{A}$

4 Computational Experiments

4.1 Data

To rigorously evaluate the proposed energy-aware operating theatre scheduling approach, a comprehensive computational benchmark was designed. The experiments aimed to assess the comparative performance of NSGA-II and ISA in optimising the bi-objective problem.

The benchmark dataset consists of three categories of synthetic instances, designed to reflect varying scales of real-world hospital environments:

- **Small (S-20)**: 20 surgeries, 4 ORs, 8 PACU beds, 5 surgeons.
- **Medium (M-50)**: 50 surgeries, 8 ORs, 15 PACU beds, 10 surgeons.
- **Large (L-70)**: 70 surgeries, 12 ORs, 25 PACU beds,

Table 1. Benchmark results: NSGA-II vs. ISA (mean ± std.).

Instance	Algorithm	HV (mean±std)	CPU (s)	|PF| (mean)	Best makespan (min)	Best energy (kWh)
Small	NSGA-II	106136±8009	1.04±0.01	30.0	859	108.1
Small	ISA	71908±8429	0.97±0.02	4.0	928	160.0
Medium	NSGA-II	245088±39792	1.88±0.05	13.0	972	341.8
Medium	ISA	199989±38771	1.82±0.05	3.0	1011	411.8
Large	NSGA-II	242930±45322	3.27±0.12	7.8	1171	785.9
Large	ISA	136825±10983	2.91±0.06	2.8	1201	978.7

15 surgeons.

Surgery durations were sampled from a log-normal distribution ($\mu = \ln(90), \sigma = 0.3$) and bounded between 30 and 240 minutes. PACU stay durations followed a triangular distribution ($a = 30, c = 90, b = 150$ minutes). The energy consumption model utilised the following power parameters: $P_{\text{active}} = 15.0$ kW, $P_{\text{idle}} = 8.0$ kW, $P_{\text{eco}} = 2.0$ kW, and restart energy penalty $E_{\text{restart}} = 5.0$ kWh, with a ramp threshold of 30 minutes. A facility closing time $T_{\text{close}} = 1440$ minutes (24 hours) was enforced.

Both algorithms were executed for 20 independent runs per instance size, with population size 30 and 60 generations/iterations.

4.2 Results

Table 1 summarises the results averaged over 20 independent runs. NSGA-II exhibits superior convergence and diversity across all instances. On the Small instance, NSGA-II attains a mean hypervolume of 106,136 compared to ISA's 71,908, and consistently finds 30 non-dominated solutions while ISA finds only 4 on average. For the Medium instance, NSGA-II's hypervolume reaches 245,088 versus ISA's 199,989, and NSGA-II finds on average 13 non-dominated solutions compared to only 3 for ISA, representing more than four times as many Pareto points. For the Large instance (70 surgeries), both algorithms find feasible schedules within the 24-hour time limit, but NSGA-II again dominates ISA in hypervolume (242,930 vs. 136,825) and solution diversity (7.8 vs. 2.8 points). The CPU times remain very low (under 3.5 seconds for the largest instance), confirming the practicality of both metaheuristics for daily planning.

The convergence behaviour is illustrated in Figure 4, which shows the evolution of the best makespan and energy over generations. NSGA-II rapidly improves both objectives, while ISA stagnates early, especially on the Medium and Large instances.

Figure 5 presents the distribution of hypervolume and CPU time across runs. NSGA-II consistently yields higher and less variable hypervolume on all instances, while both algorithms exhibit similar, negligible CPU times (all under 3.5 seconds).

The core contribution of the bi-objective model is the quantification of energy savings achievable through strategic scheduling. Figure 6 displays the Pareto fronts obtained by both algorithms for each instance (best run per instance). The fronts clearly illustrate the trade-off between makespan and energy consumption.

For the Small instance, the Pareto front obtained by NSGA-II (Figure 6(a)) reveals two extreme solutions: the makespan-optimal solution completes all surgeries in 859 minutes, while the energy-optimal solution on the NSGA-II front achieves a minimum of 108.1 kWh. By contrast, ISA's best makespan is 928 minutes and its best energy is 160.0 kWh—both of which are dominated by NSGA-II's Pareto front. The energy gap between ISA's best-energy solution (160.0 kWh) and NSGA-II's energy-optimal solution (108.1 kWh) corresponds to a reduction of approximately 32%, confirming that NSGA-II consistently discovers superior trade-off solutions. The intermediate points on the NSGA-II front show a clear downward trend in energy consumption as the makespan is allowed to increase beyond its minimum, with energy savings exceeding 30% relative to the ISA solutions. This confirms that the trade-off is genuine, but the quality of the discovered front is highly algorithm-dependent. Similar trade-offs are observed for the Medium and Large instances, where NSGA-II consistently discovers more diverse and higher-quality fronts than ISA.

4.3 Discussion

The experimental results yield several critical insights:

1. **Viability of energy-aware scheduling:** The data unequivocally show that standard operational makespans harbour significant latent inefficiencies. By actively managing idle times and aligning surgery gaps with equipment eco-state thresholds, hospitals can achieve energy

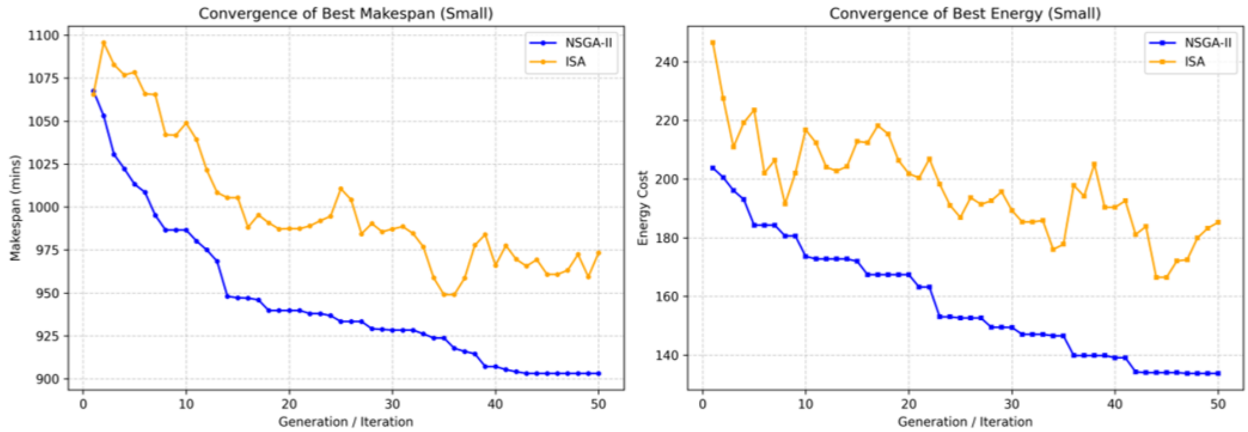

(a) Small instance.

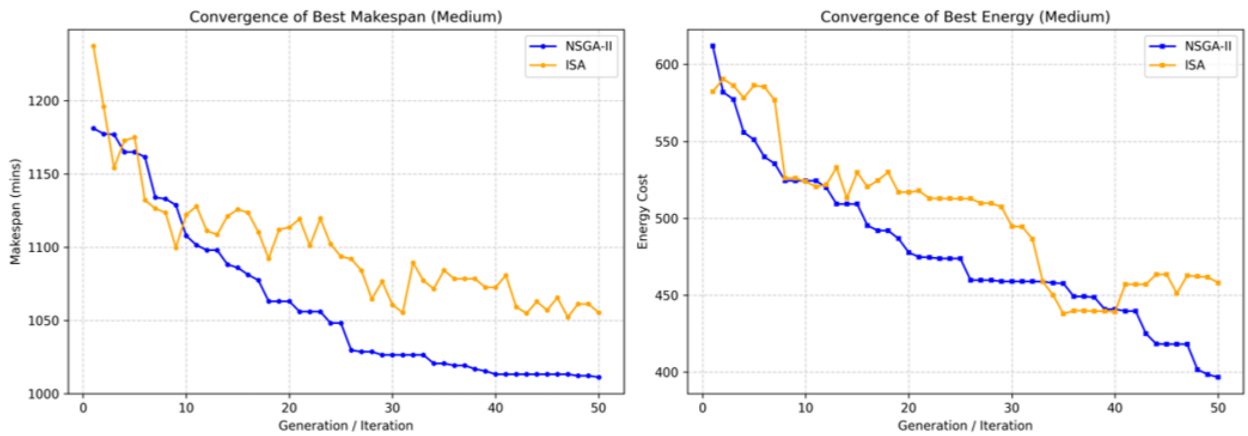

(b) Medium instance.

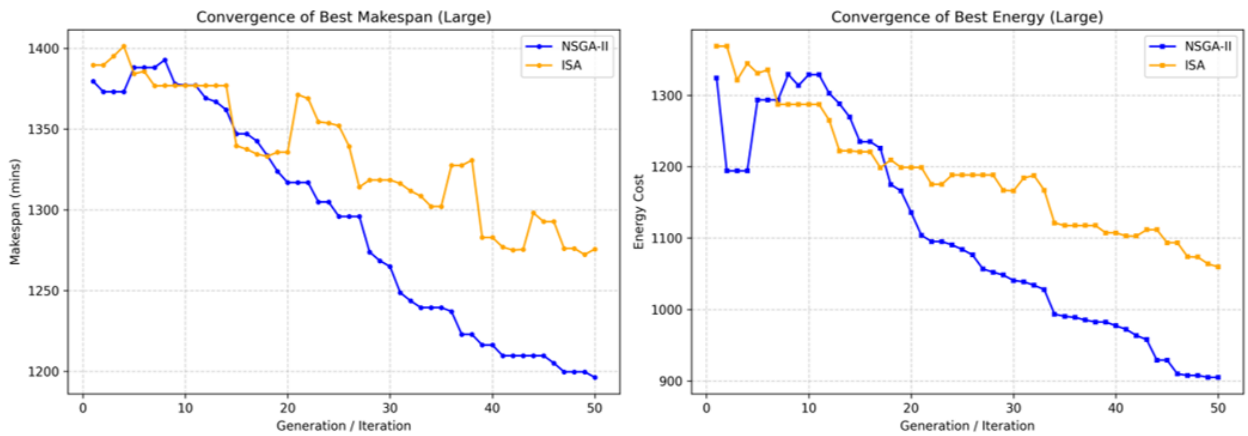

(c) Large instance.

Figure 4. Convergence of best makespan (solid lines) and energy (dashed lines) for NSGA-II and ISA across the three instance sizes.

reductions exceeding 30% with only marginal extensions to the working day. This provides a quantifiable mechanism to balance staff overtime against environmental and energy costs.

2. **Algorithmic supremacy of NSGA-II:** NSGA-II is vastly superior to ISA for this permutation-and-assignment problem. The genetic operators (especially Order Crossover) effectively preserve sequences that respect PACU flow constraints, whereas ISA's perturbations

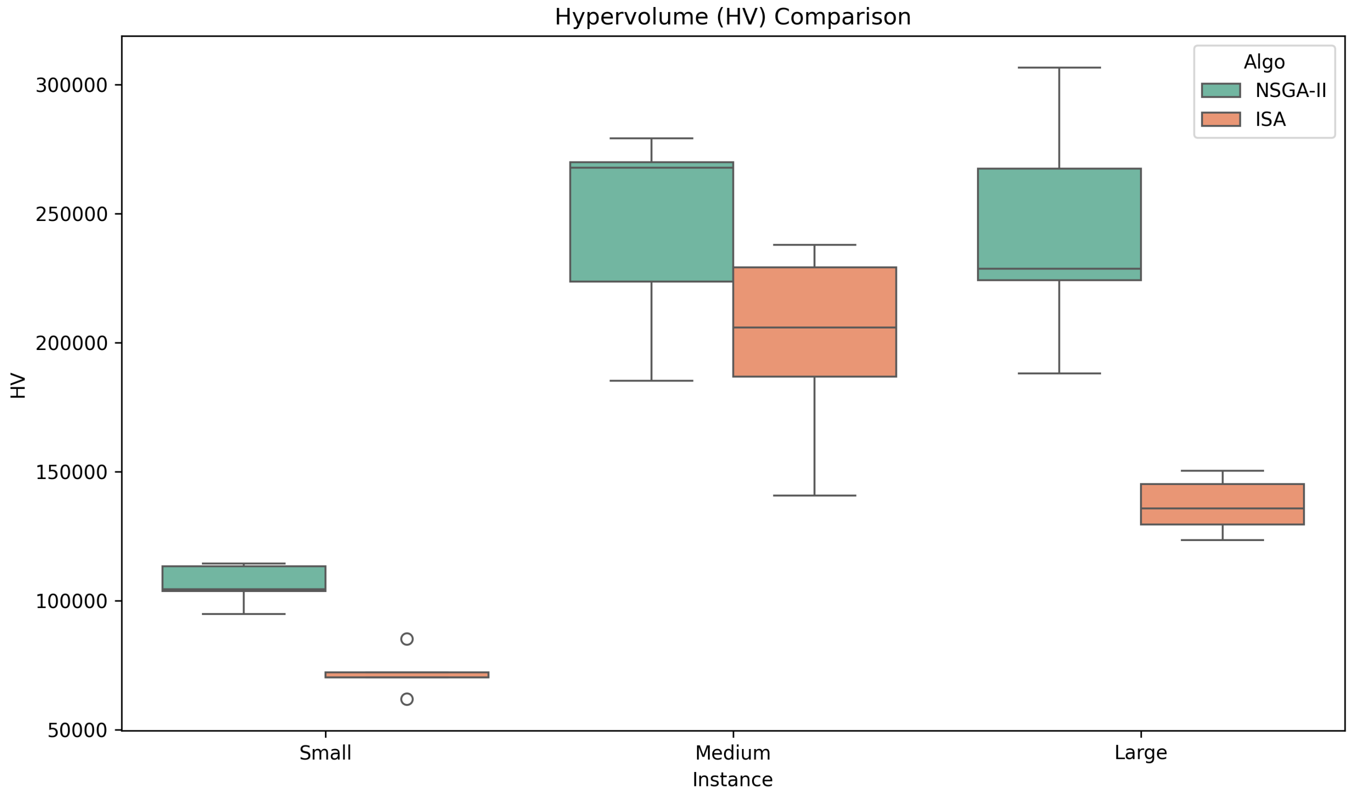

Figure 5. Boxplots of hypervolume (left) and CPU time (right) for NSGA-II and ISA across instance sizes.

appear too disruptive, frequently breaking feasibility and leading to stagnation, especially as problem size increases.

3. **PACU bottlenecks and large-instance feasibility:** Even with 70 surgeries, both algorithms found feasible schedules within the 24-hour limit. However, the significant drop in hypervolume and number of Pareto points for ISA on the Large instance indicates that the problem becomes more challenging as size increases. The PACU capacity remains a critical constraint; further increasing the number of surgeries (e.g., beyond 80) would likely require extending the closing time or adding more PACU beds to maintain feasibility.

5 Conclusion

This paper introduced a bi-objective optimisation framework for green ambulatory surgery scheduling that explicitly accounts for PACU blocking and HVAC thermal inertia—a combination that, to the best of our knowledge, has not been previously addressed in the literature. Extensive benchmarks on three instance sizes (20, 50 and 70 surgeries) demonstrated that NSGA-II consistently outperforms a discrete ISA in both solution quality and convergence speed. The Pareto front analysis reveals that substantial energy savings (exceeding 30%) are achievable with only a modest increase in makespan, offering a quantifiable basis for balancing environmental sustainability against operational efficiency.

The generated Pareto fronts serve as a practical decision-support tool for hospital managers. Rather than forcing a single rigid schedule, the presented approach provides a visual trade-off curve that enables managers to select an operating point based on contextual factors such as time-of-use electricity tariffs, staff overtime costs, or day-specific patient volume targets. This flexibility empowers hospitals to adapt their surgical planning dynamically, aligning operational decisions with both financial and sustainability goals. A natural extension would be the development of an interactive dashboard where schedulers can adjust the weight between objectives and visualise the corresponding schedule.

Despite its promise, several barriers must be addressed before widespread clinical deployment. First, the current model assumes deterministic surgery durations and PACU lengths of stay, whereas real-world environments are characterised by significant stochasticity (e.g., emergency case

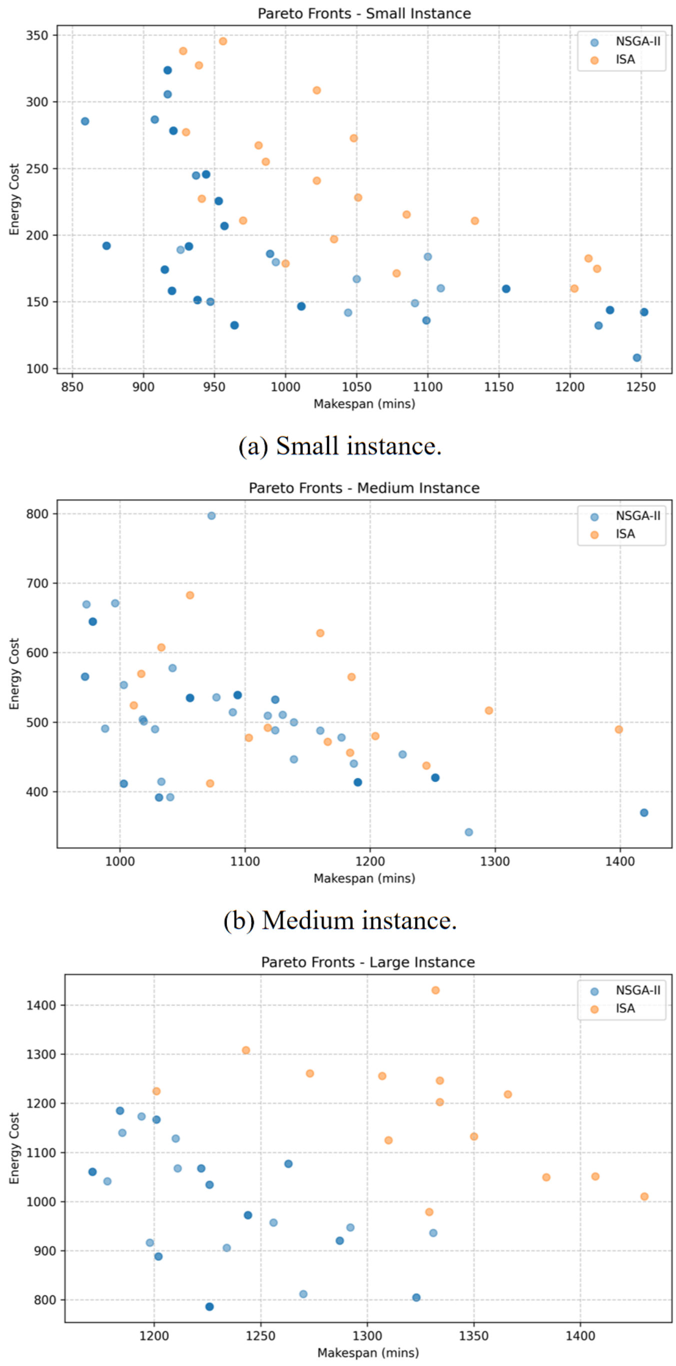

(a) Small instance.

(b) Medium instance.

(c) Large instance.

Figure 6. Pareto fronts obtained by NSGA-II and ISA for the three instance sizes.

insertions, variable recovery times). Second, the HVAC model relies on accurate submetering data for active, idle, eco, and peak power states; many existing hospitals lack the granular energy monitoring infrastructure required to calibrate these parameters. Third, the performance of the metaheuristics depends on careful parameter tuning (e.g., population size, number of generations, crossover/mutation rates), which may vary across hospital configurations and require local expertise. Fourth, the ramp threshold τ_k^{ramp} is treated as a fixed constant, but in reality it may vary with external ambient temperature and HVAC system age. Addressing these limitations will require integrating real-time sensor data and adaptive learning mechanisms.

To bridge the gap between algorithmic research and clinical practice, we propose a practical integration architecture. The optimisation engine can be deployed as a modular decision-support service, interfacing with existing Hospital Information Systems (HIS) and Electronic Medical Records (EMR) platforms (e.g., Epic, Cerner) to automatically ingest patient surgery schedules, estimated durations, and PACU stay predictions. Furthermore, integration with Building Management Systems (BMS) via open protocols would enable closed-loop operation: the BMS can read the optimised schedule and pre-condition ORs to the required active state shortly before surgery start times, while automatically transitioning to eco or idle states during planned idle gaps. This bi-directional communication ensures that the energy savings predicted by the model are physically realised in the building's HVAC infrastructure.

Building on these insights, future research will focus on extending the model to handle stochastic surgery durations through robust optimisation or dynamic rescheduling policies; validating the framework with real-world HVAC data and operating room logs from partner hospitals; and developing constructive heuristics that guarantee PACU-compliant schedules.

Data Availability Statement

Data will be made available on request.

Funding

This work was supported by the Projets de Recherche Formation Universitaire (PRFUs), funded by the General Directorate for Scientific Research and Technological Development (DGRDST), Algeria, under Grant C00L07UN310220230001 (Environmental and Economic Optimization of Production Systems) and Grant A01L07EP310220220001 (Reliability Assessment of the National Electricity Grid Interconnected with the Maghreb Grid). The authors also acknowledge the support of the LDREI Laboratory and the OLIB Laboratory at the Higher Institute of Industrial Management of Sfax (ISGI), Tunisia.

Conflicts of Interest

The authors declare no conflicts of interest.

AI Use Statement

The authors declare that no generative AI was used in the preparation of this manuscript.

Ethical Approval and Consent to Participate

Not applicable.

References

[1] Beliën, J., Brailsford, S., Demeulemeester, E., Demirtas, D., Hans, E. W., & Harper, P. (2025). Fifty years of operational research applied to healthcare. *European journal of operational research, 326*(2), 189-206. [CrossRef]

[2] Cheng, H., Adulyasak, Y., & Rousseau, L.-M. (2025). An environmentally friendly elective patient scheduling under the predict-then-optimize framework. *Annals of Operations Research*. Advance online publication. [CrossRef]

[3] Shehadeh, K. S., & Padman, R. (2022). Stochastic optimization approaches for elective surgery scheduling with downstream capacity constraints: Models, challenges, and opportunities. *Computers & Operations Research, 137*, 105523. [CrossRef]

[4] El Amraoui, A., Belhor, M., & El Amraoui, L. (2026). A new heuristic algorithm for resource-constrained operating room scheduling problem: A case study. *Ain Shams Engineering Journal, 17*(6), 104108. [CrossRef]

[5] Vieira, B. S., Silva, E. M., & Chaves, A. A. (2025). Random-key algorithms for optimizing integrated operating room scheduling. *Applied Soft Computing, 180*, 113368. [CrossRef]

[6] Liu, Y., Huang, Y., Dai, Z., & Gao, Y. (2025). Self-optimized learning algorithm for multi-specialty multi-stage elective surgery scheduling. *Engineering Applications of Artificial Intelligence, 147*, 110346. [CrossRef]

[7] Spence, C., Shah, O. A., Cebula, A., Tucker, K., Sochart, D., Kader, D., & Asopa, V. (2023). Machine learning models to predict surgical case duration compared to current industry standards: scoping review. *BJS Open, 7*(6), zrad113. [CrossRef]

[8] Zhang, R., Yu, H., Gao, K., Fu, Y., & Kim, J. H. (2024). A Q-learning based artificial bee colony algorithm for solving surgery scheduling problems with setup time. *Swarm and Evolutionary Computation, 90*, 101686. [CrossRef]

[9] Dai, Z., & Li, Z.-P. (2025). Intelligent scheduling method for ambulatory surgery based on machine learning and fuzzy optimization. *Annals of Operations Research*. Advance online publication. [CrossRef]

[10] Dekhici, L., & Khaled, B. (2010). Operating theatre scheduling under constraints. *Journal of Applied Sciences, 14*(10), 1380-1388. [CrossRef]

[11] Dekhici, L., Guerraiche, K., Güven, A. F., Bajaj, M., Blazek, V., & Prokop, L. (2026). Machine learning-embedded energy cost optimization for hybrid flow shop scheduling toward low-carbon industrial energy systems. *Unconventional Resources*, 100411. [CrossRef]

[12] Atisattapong, W., & Pongpatcharapun, P. (2024). Master Surgical Scheduling of Multi-Surgeon Surgical Schedules Using a Mixed-Integer Linear Programming Model and Genetic Algorithms. In *Proceedings of the 4th International Conference on Robotics, Automation and Artificial Intelligence* (*RAAI 2024*) (pp. 274-283). IEEE. [CrossRef]

[13] Latorre-Núñez, G., Lüer-Villagra, A., Marianov, V., Obreque, C., Ramis, F., & Neriz, L. (2016). Scheduling operating rooms with consideration of all resources, post anesthesia beds and emergency surgeries. *Computers* & *Industrial Engineering, 97*, 248-257. [CrossRef]

[14] Dekhici, L., & Belkadi, K. (2023). A Bi-Objective Paediatric Operating Theater Scheduling. *International Journal of Healthcare Information Systems and Informatics* (*IJHISI*)*, 18*(1), 1-15. [CrossRef]

[15] Belkhamsa, M., Jarboui, B., & Masmoudi, M. (2018). Two metaheuristics for solving no-wait operating room surgery scheduling problem under various resource constraints. *Computers* & *Industrial Engineering, 126*, 494-506. [CrossRef]

[16] Mostofa, N., Dexter, F., & Titler, S. S. (2026). Influence of sequencing the shortest and least variable cases early in the workday on operating room performance when constrained by post-anesthesia care unit availability: a narrative review. *Perioperative Care and Operating Room Management*, 100654. [CrossRef]

[17] Wang, Y., Wang, J., Yao, W., Ren, J., Chen, L., & Ban, Q. (2026). Ventilation and indoor air quality management in operating rooms: A systematic review of pollutant sources, transmission, and control techniques. *Journal of Building Engineering*, 115427. [CrossRef]

[18] Chen, S., Chiong, R., & Li, D. (2025). A prescriptive tree-based model for energy-efficient room scheduling: Considering uncertainty in energy generation and consumption. *European Journal of Operational Research, 326*(2), 374-388. [CrossRef]

[19] Zhan, Y., Shen, S., Wan, G., & Zhang, Z. (2025). Optimizing surgery scheduling under post-anesthesia care unit capacity constraints and random services durations. *Computers & Operations Research*, 107368. [CrossRef]

[20] Gandomi, A. H. (2014). Interior search algorithm (ISA): a novel approach for global optimization. *ISA Transactions, 53*(4), 1168-1183. [CrossRef]

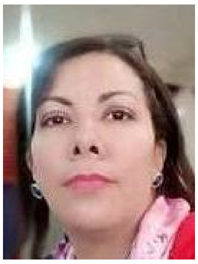

Dekhici Latifa is a full professor of computer science at the Faculty of Mathematics and Computer Science, University of Science and Technology of Oran (USTOMB), Algeria. She is also a researcher at the LDREI Laboratory (Smart Grid Development Laboratory) in the Higher School of Electrical and Energy Engineering of Oran (ESGEEO). She served as Vice Dean for External Relations from 2020 to 2025. Her research interests include healthcare scheduling,power dispatch, power system design and green optimisation using machine learning and metaheuristics. Prof. Dekhici has been a member of the INSIDE Project (Inclusion of Students with Impairments in Distance Education) in 2023, the DIGITAQ Project on Quality Assurance Digitalization in Higher Education, and she leads the Digitalisation Strategy Office at her university. She has participated as a trainer in numerous activities related to sustainability, digitalization, accessibility, cyber responsibility, soft skills, inclusivity, and serious game design. (Email: latifa.dekhici@univ-usto.dz)

Guerraiche Khaled is a full professor in electrical engineering, in the electrical engineering at the High School of Electrical Engineering and Energy. He received his BS degree in electrical engineering from the Electrical Engineering Institute of The University of Sciences and Technology of Oran (USTO) in 2004, the MS degree from the Electrical Engineering Institute of The University of Sciences and Technology of Oran (USTO) in 2008 and the PhD degree from the Electrical Engineering Institute of The University of Sciences and Technology of Oran (USTO) in 2016. His research interests include operations, planning and economics of electric energy systems, as well as optimisation theory and its applications. (Email: guerraiche_khaled@esgee-oran.dz)

Jlassi Jihène is an associate professor of logistics and transport science at the Higher Institute of Industrial Management of Sfax (ISGI Sfax), University of Sfax, Tunisia. She is an active member of the OLID research laboratory (Laboratory of Logistics, Industrial Engineering and Decision). She received her PhD in Industrial Engineering from Paris 8 University, France. Her main research interests lie in logistics and optimisation, with a focus on vehicle routing problems, healthcare logistics, and the application of metaheuristics and operations research techniques. Dr. Jlassi has served on the scientific committees of several international conferences and is a regular reviewer for numerous refereed journals. In 2026, she co-authored the book about modern transportation. (Email: jihene.jlassi@isgis.usf.tn)

Djari Mohamed Adel received the M.Sc. degree in electrical engineering in 2016 and the Ph.D. degree in electrical engineering in 2022. He is currently a researcher and academic affiliated with the Higher School of Electrical Engineering and Energetics of Oran, Algeria, and a member of the Intelligent Electrical Networks Development Laboratory. His research interests include power system stability, renewable energy integration, optimal power flow, FACTS devices, smart grids, and intelligent control of electrical power systems. His current work focuses on the optimisation, monitoring, and reliable operation of modern electrical networks, with particular emphasis on sustainable energy systems and advanced intelligent control techniques for future smart power grids. (Email:djari_adel@esgee-oran.dz)