Expert Systems With Applications 295 (2026) 128803

Contents lists available at ScienceDirect

Expert Systems With Applications

journal homepage: www.elsevier.com/locate/eswa

Application of a sampling and clustering-based heuristic search algorithm to find an efficient staff configuration in an emergency department

Maria Harita [a,*], Alvaro Wong [a], Dolores Rexachs [a], Emilio Luque [a], Eva Bruballa [b], Francisco Epelde [c]

[a] Universitat Autonoma de Barcelona (UAB), Department of Computer Architecture and Operating Systems, Carrer de les Sitges, Cerdanyola del Valles, 08193 Barcelona, Spain
[b] Escoles Universitaries Gimbernat (EUG), Computer Science School, Avinguda de la Generalitat, 202, Sant Cugat del Valles, 08174 Barcelona, Spain
[c] Department of Medicine, Hospital Universitari Parc Taulí, Parc Tauli, 1, Sabadell, 08208 Barcelona, Spain

A R T I C L E I N F O

Keywords:
Agent-based simulation
Classification algorithms
Emergency departments
Heuristic methods
Optimization

A B S T R A C T

Emergency Departments (EDs) are among the most complex areas in healthcare, requiring immediate medical attention for acute and urgent conditions. Optimizing staff configurations to reduce patient Length of Stay (LoS) and improve operational efficiency poses a significant challenge due to the combinatorial and high-dimensional nature of the problem. To identify the most effective staff configuration, we propose a heuristic optimization strategy that is based on the Montecarlo Clustering Search Algorithm (MCSA), which efficiently explores the multidimensional solution space. MCSA leverages an agent-based simulation (ABM) model that evaluates each proposed staff configuration under realistic operational conditions, providing Key Performance Indicator (KPI) feedback values related to each proposed staff configuration. Through this strategy, we explore staff configurations capable of handling patient volumes with varying acuity levels in an ED to optimize the LoS KPI. Results demonstrate that our methodology is capable to find a solution as a staff configuration that reduces LoS compared to a baseline, offering a computationally efficient and practical tool for decision-makers. We identified solutions by exploring less than 1% of the total search space, demonstrating the efficiency of the proposed approach in addressing complex optimization problems. This approach supports informed planning in healthcare environments while maintaining system feasibility and scalability.

1. Introduction

Hospital information systems increasingly provide real-world data, which enhances analysis and insights for ED managers in decision-making. However, leveraging this data effectively adds to the inherent complexity of management, as the optimization of resource allocation and workflows requires exploring large, multidimensional search spaces. This makes the problem not only computationally complex but also organizationally demanding, requiring robust methodologies to identify feasible and efficient solutions.

Simulation and heuristic methods are two complementary strategies widely applied to solve complex optimization problems in fields such as logistics, engineering, and healthcare. Simulation enables the modeling of dynamic systems through computational experimentation, allowing decision-makers to assess the effectiveness of alternative strategies in controlled environments. In contrast, heuristic methods offer practical approaches for exploring high-dimensional solution spaces, particularly when problems are too complex for traditional analytical or exhaustive methods.

Moreover, testing new staff policies in live systems is typically unfeasible, making simulation a vital tool for predicting the impact of staff decisions by analyzing the Key Performance Indicators (KPIs) such as patient Length of Stay (LoS). When paired with heuristic search, simulation provides a robust framework for exploring and evaluating diverse strategies under realistic conditions.

This work presents a search methodology based on the Montecarlo Clustering Search Algorithm (MCSA), a stochastic heuristic that explores discrete, multidimensional staff configurations within a combinatorial search space. Each configuration corresponds to a unique arrangement of senior and junior staff—doctors, nurses, and support personnel—across ED units. As the number of variables increases, the solution space grows exponentially, making exhaustive search methods

* Corresponding author.
E-mail address: MariaDeLosAngeles.Harita@uab.cat (M. Harita).

https://doi.org/10.1016/j.eswa.2025.128803
Received 25 February 2025; Received in revised form 5 June 2025; Accepted 26 June 2025
Available online 1 July 2025

computationally infeasible. To assess their performance, MCSA relies on an independently validated agent-based simulation (ABM) model, which simulates the dynamic behavior of an Emergency Department (ED) under realistic conditions. The ABM—developed by the HPC4EAS research group at Universitat Autònoma de Barcelona (UAB) in collaboration with Hospital Taulí—provides KPIs for each evaluated configuration.

A key feature of MCSA is its ability to progressively reduce the search space. It begins with Montecarlo-based sampling to explore a broad range of solutions and then applies a clustering phase to identify and exploit Feasible Regions, where high-quality configurations tend to concentrate. This iterative refinement improves computational efficiency by focusing the search on the most promising areas. Hospital Taulí, serving over 160,000 patients annually, provides the operational context and real-world data used throughout the simulations (Liu, Cabrera, Rexachs and Luque, 2014).

1.1. Hypotheses

The increasing patient flow in ED presents an ongoing challenge in resource allocation, directly impacting wait times and overall service quality. Despite advancements in modeling and simulation, traditional staff assignment strategies often rely on fixed rules that fail to adapt to fluctuating demand. In this context, there is a growing need to perform stochastic search strategies that leverage agent-based simulation, enabling a coordinated exploration of complex solution spaces in healthcare operations. A promising technique is our MCSA heuristic, which explores staff configurations without resorting to exhaustive search methods. Our methodology is built upon the following hypotheses:

- H1: Effective adjustment of ED staff configurations can improve the patients LoS.
- H2: MCSA Heuristic can identify efficient staff configurations within a high-dimensional space, under real-world constraints.
- H3: MCSA is more computationally efficient than exhaustive search methods in optimizing KPIs related to ED staff allocation.

To test these hypotheses, realistic ED scenarios were simulated, using MCSA to propose configurations and the ED simulator to assess their performance. The methodology aims to demonstrate that effective solutions can be found without exploring the full combinatorial space and that simulation-based feedback enables alignment with real operational goals. Additionally, the simulation provides means to align management objectives, particularly improving operational efficiency.

1.2. Generalization of the optimization problem

This study addresses a combinatorial optimization problem: identifying a configuration of medical staff (e.g., junior/senior doctors, nurses, support personnel) to minimize patient LoS. Each decision variable represents a staff role, bounded by realistic operational constraints. The size of the search space renders exhaustive evaluation infeasible. This problem typifies challenges found in other domains such as production planning, network design, and scheduling Boresta, Giovannelli and Roma (2024), where discrete decisions must be made under constraints. Therefore, the complexity arises from the vast number of possible staff configurations and operational constraints such as personnel availability and resource capacities. Due to the infeasibility of exhaustive evaluation, our approach employs a heuristic based on Montecarlo sampling combined with the ABM ED Simulator to evaluate candidate solutions under realistic ED dynamics. The detailed problem formulation and optimization approach are presented in the methodology section. The remainder of this paper is organized as follows: Section 2 reviews related works in healthcare optimization.

Section 3 details the MCSA algorithm and its application to the ED context. Section 4 presents experimental validation. Section 5 discusses the full implementation and results analysis. Section 6 concludes and outlines future directions.

2. State of the Art

Optimization in the context of EDs has been extensively addressed using both heuristic search (Chen, Huang, Chiang and Yu-Hsin Chen, 2020; Tomar, Bansal and Singh, 2024; Lameesa, Hoque, Alam, Ahmed and Gandomi, 2024) and simulation methodologies, each offering distinct advantages.

Heuristic methods, including evolutionary algorithms and metaheuristics such as Genetic Algorithms (GA), Particle Swarm Optimization (PSO) or Simulated Annealing (SA), among others (Abdalkareem, Amir, Al-Betar, Ekhan and Hammouri, 2021; White, Nano, Nguyen-Ngoc and White, 2007), have been widely used to solve complex staff scheduling problems in healthcare. These approaches, while efficient in identifying near-optimal solutions, often focus on isolated components such as nurse shift assignment or fatigue-aware scheduling under static assumptions. For instance, Ceschia et al. (Ceschia, Di Gaspero, Mazzaracchio, Policante and Schaerf, 2023) proposed a Simulated Annealing model incorporating skills and preferences, applied to real hospitals in Italy. Similarly, Amindoust et al. (Amindoust, Asadpour and Shirmohammadi, 2021) developed a hybrid evolutionary algorithm considering fatigue in pandemic contexts, while Valouxis et al. (Valouxis, Gogos, Goulas, Alefragis and Housos, 2012) introduced a multi- objective evolutionary framework that balances safety, service quality, and staff wellbeing.

Despite these advances, a common limitation of such methods lies in their lack of availability to work with operational data or real-time system dynamics. Most rely on synthetic datasets or simplified constraints, making their adaptation to real-world, high-dimensional ED environments challenging.

On the simulation side, several techniques have been proposed to model hospital workflows. Discrete Event Simulation (DES) has been used to analyze patient flow and queuing performance Forbus and Berleant (2022); Vecillas Martin, Berruezo Fernández and Gento Municio (2025), while System Dynamics (SD) focuses on aggregate trends, offering insights into long-term policy impacts. However, both approaches lack the granularity needed for evaluating the operational consequences of micro-level staff adjustments. Agent-Based Simulation (ABS) (Tisue and Wilensky, 2004), by contrast, models individual interactions between patients, staff, and resources, capturing emergent behaviors in complex systems. The use of ABS in ED contexts, such as in the works of Cabrera et al. (Cabrera, Taboada, Iglesias, Epelde and Luque, 2011) and Bruballa et al. (Bruballa, Wong, Rexachs, Luque and Epelde, 2017), has demonstrated its capability to reflect system variability and support scenario evaluation.

Nonetheless, these simulation-based studies often evaluate predefined scenarios rather than actively optimizing staff configurations. They provide valuable performance insights but do not incorporate search mechanisms for configuration improvement.

To address these limitations, recent studies have begun exploring hybrid models that combine simulation and heuristic optimization. These hybrid approaches leverage the realism of simulation with the search power of heuristics, aiming to find efficient configurations while preserving system fidelity. Applications range from hospital disinfection scheduling Juan, Faulin, Grasman, Rabe and Figueira (2015) to dynamic logistics Wilensky and Rand (2015), Tutsoy and Tanrikulu (2022), demonstrating versatility in complex environments. However, many of these hybrid frameworks either focus on limited dimensions or lack a formal exploration mechanism for large solution spaces.

Unlike previous models that are either search-driven with static assumptions or simulation-based with no optimization, our proposal introduces a coordinated strategy: a heuristic search process (MCSA) guided by feed- back from an independently validated agent-based simulation model of a real ED. This methodology addresses the multi-dimensionality of real EDs, accounting for multiple staff roles, while incorporating constraints such as minimum/maximum staff levels and inter-role dependencies. The proposed method does not merely simulate to obtain KPI outcomes, but actively refines configurations in a feedback loop based on such outcome vallues. The proposed contribution differs from existing literature in key aspects, such as: (1) Stochastic exploration leverages a simulation- based evaluation, providing both solution quality and system fidelity. (2) It addresses real-world dimensionality and constraints, validated on operational data from Hospital Taulí, a high-volume ED in Spain, and (3) It avoids the need for exhaustive search, offering a scalable and traceable optimization method for decision-makers in critical care settings. Furthermore, although our application is healthcare-focused, the underlying methodological framework—namely, hybrid simulation-heuristic coordination—has proven valuable in domains such as manufacturing scheduling, telecommunications, and disaster response planning. By acknowledging these broader contexts, we aim to situate our approach within a wider tradition of solving complex allocation problems using flexible methods. These contributions positions our work at the intersection of advanced heuristic search and high-fidelity simulation, offering a novel framework for data-driven staff optimization.

In contrast to recent general-purpose optimizers such as the RUN algorithm Ahmadianfar, Heidari, Gandomi, Chu and Chen (2021), which leverage mathematical strategies like the Runge-Kutta method to move beyond metaphor- based heuristics Samadi-Koucheksaraee, Shirvani-Hosseini, Ahmadianfar and Gharabaghi (2022), our approach is domain-specific and designed for practical applicability in high-dimensional, simulation-driven environments like Emergency Departments (EDs). While metaphor-free and theoretically grounded algorithms have proven effective across a range of engineering contexts Ahmadianfar, Samadi-Koucheksaraee and Razavi (2023), their use remains limited when complex, stochastic simulations are required for solution evaluation. Our methodology does not focus on outperforming optimization algorithms, but on providing healthcare decision-makers with a scalable framework for real-world operational planning.

Our approach does not assume a predefined parametric model; instead, it employs a heuristic (MCSA) guided by simulation-based evaluations to identify viable staff configurations under complex, dynamic, and uncertain conditions. Unlike models commonly used in fields such as population genetics and epidemiology, our methodology focuses on local resource allocation (e.g., doctors, nurses) and optimizing performance metrics. This distinction in scale and scope underscores the nature of our work: providing tactical optimization tools for hospital decision-makers aiming to enhance efficiency in highly variable environments.

To illustrate the structure and dynamics of the simulation model employed in this study, Fig. 1 provides a schematic overview of the ED simulator as used in conjunction with our methodology. The simulator is treated as a calibrated system, previously validated using empirical data from Hospital Taulí. It operates with two types of input:

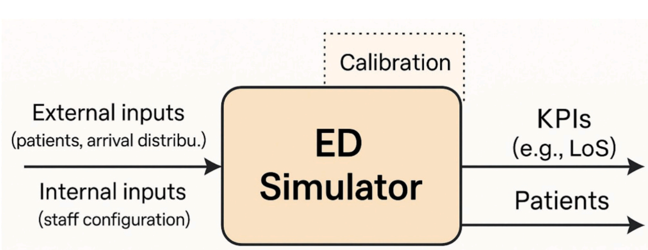

Fig. 1. This schematic highlights how the simulation plays an important role in bridging the gap between abstract search (MCSA) and real-world performance evaluation.

- External inputs: Representing real-world elements outside the control of decision-makers, such as patient inflow patterns and arrival distributions. These values are derived from historical data and reflect observed fluctuations across time (e.g., hourly, daily, or seasonal variations).
- Internal inputs: Defined by the user and representing the configuration of medical personnel (e.g., numbers of senior and junior doctors, nurses, and technicians). These variables form the search space over which our heuristic algorithm (MCSA) operates.

On the output side, the simulator provides:

- A list of KPIs, including the LoS, staff occupancy, and patients queue in admission and triage phases. These metrics are used to evaluate the quality of each configuration.
- Other outputs, such as the simulated patient dispatch behavior, are generated by the simulator to reflect system dynamics, although they are not directly utilized in our methodology.

3. Montecarlo clustering search algorithm (MCSA)

MCSA is a heuristic designed to explore high-dimensional search spaces in combinatorial optimization tasks. In this context, the problem of identifying optimal ED staff configurations is reformulated as a discrete optimization problem. The search begins by sampling diverse combinations of staff allocations, treating each as a candidate solution. However, evaluating feasibility and performance is essential. To address this, each configuration is assessed within a realistic and dynamic environment using simulation, which provides quantitative feedback (e.g., LoS) for guiding the search process.

The MCSA methodology consists of two main phases. During the *Exploration Phase*, random sampling is performed within an *n*-dimensional space, guided by a fitness function to identify promising configurations. In the subsequent *Exploitation Phase*, clustering techniques refine the search. A density-based analysis first locates high- potential regions, followed by centroid-based clustering to focus the optimization effort. The method has been enhanced for scalability to high-dimensional, combinatorially complex problems in prior studies Harita, Wong, Rexachs del Rosario and Luque Fadón (2021), enabling efficient identification of high-quality solutions.

The algorithm was validated using standard mathematical benchmark functions (Mohamed, Sallam, Agrawal, Hadi and Mohamed, 2023), assessing both performance and solution quality. Robustness was evaluated through over 200 independent runs with varying random seeds, consistently converging to the same solutions. This repeatability, as reported in Harita et al. (2024), confirms the algorithm's stability. Scalability was examined by incrementally expanding variable ranges, with MCSA maintaining consistent convergence behavior. Fig. 2 illustrates the application of the algorithm to the Michalewicz and Griewank functions as two representative benchmark cases.

The methodology was subsequently applied to combinatorial optimization problems, such as the Knapsack Problem (KP01) (Karp, 1972; Penn, Hasson and Avriel, 1994). This required redefining the search space in MCSA terms to capture all unique solution combinations. Each variable contributes to the dimensionality of the space, and the total number of configurations grows exponentially, e.g., 2^n in binary cases or $\Pi^{n} R_i$ in more general scenarios where.

R_i represents the range of the i*th* variable. In decision problems like the Knapsack Problem, each subset of items represents a potential solution, with binary inclusion decisions. The exponential growth of the search space quickly renders exhaustive methods computationally infeasible for large instances. As a result, heuristic approaches become essential to efficiently navigate the space by focusing on high-potential regions, as discussed in Harita, Wong, Rexachs and Luque (2022).

Other recent metaheuristic methods such as the Gradient-Based Optimizer (GBO) Ahmadianfar, Bozorg-Haddad and Chu (2020) and

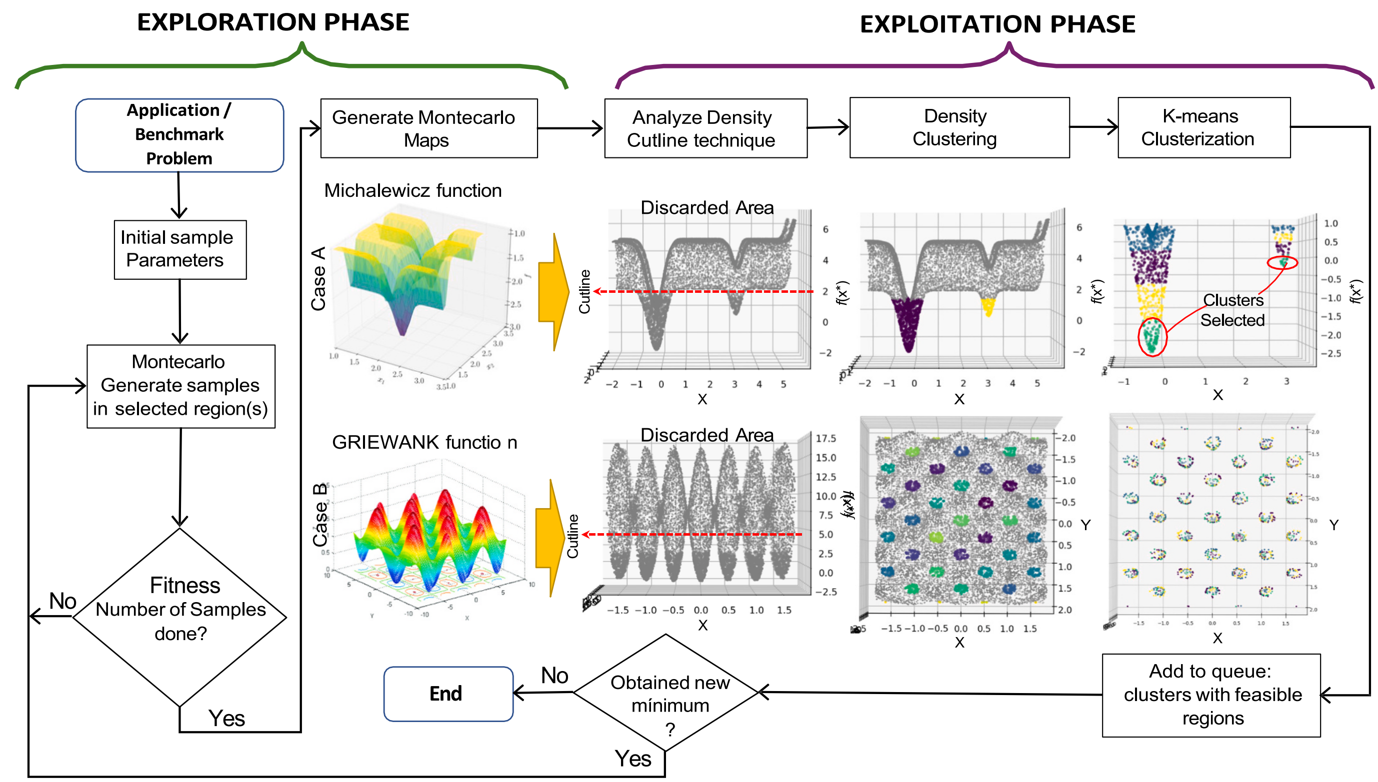

Fig. 2. Flowchart of the MCSA Heuristic, illustrated in the context of its evaluation using a benchmark function. The methodology and algorithmic details are extensively described in Harita et al. (2024), and this diagram provides a visual overview of the main components and flow of the approach during the optimization process.

the Weighted Mean of Vectors algorithm (INFO) Ahmadianfar, Heidari, Noshadian, Chen and Gandomi (2022) have shown strong performance in abstract mathematical benchmarks and constrained engineering problems. These algorithms emphasize enhanced exploration and convergence strategies through hybrid mechanisms. However, their evaluation is typically performed on artificial test functions or synthetic design tasks, focusing on generic performance metrics like convergence rate or success ratio. In contrast, our methodology is not aimed at generalized benchmarking but at solving a domain-specific, real-world problem: staff configuration in a hospital ED. Here, optimization is achieved through iterative interaction with a validated simulation model, using operational KPIs as evaluation criteria under realistic constraints and uncertainty. This differentiates our contribution by focusing on practical applicability and decision support, rather than algorithmic novelty alone.

3.1. MCSA and the ED simulator

This work aims to demonstrate the feasibility of generating practical and implementable ED staff configurations by means of the MCSA heuristic leveraging a validated ABM ED simulator. This hybrid approach leverages stochastic exploration to efficiently navigate the high-dimensional search space, while simulation captures the operational complexity of the system. The optimization focuses on minimizing patient LoS, an essential KPI in emergency care.

The ED simulator, validated in prior studies such as Taboada González (2013), reproduces realistic system behavior under diverse resource configurations. MCSA iteratively uses simulation outputs to

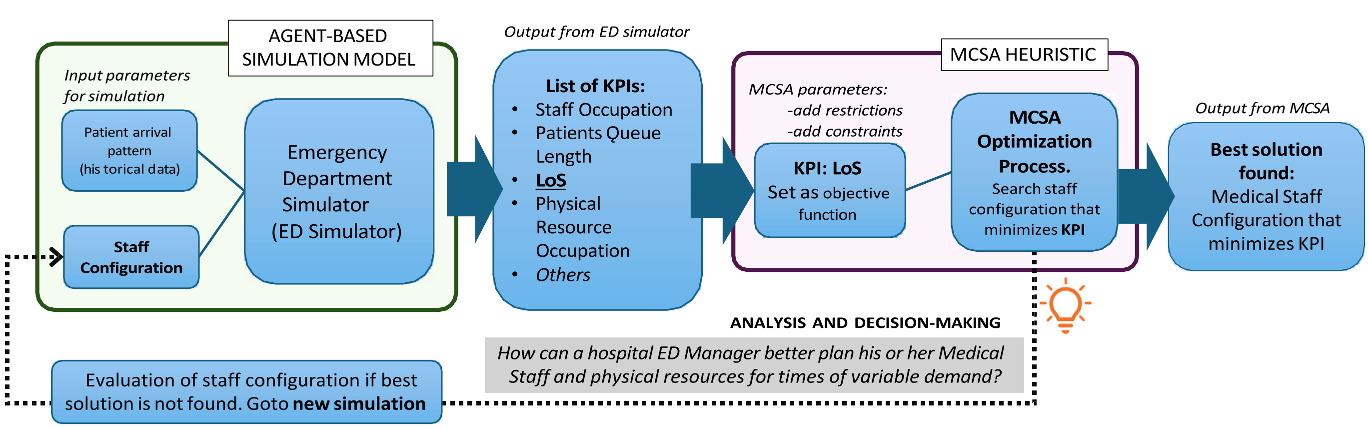

Fig. 3. Methodology flowchart showing the interaction between ABM and MCSA. The diagram outlines the feedback loop in which MCSA proposes staff configurations as input to the simulator. The simulator evaluates each configuration and returns KPI outputs, which guide the next iteration of the exploration process. This structure reflects the core logic of the optimization framework.

guide its search, refining candidate solutions through a feedback loop. As shown in Fig. 3, the simulator evaluates each configuration and returns KPI metrics, which the heuristic uses to generate improved staff scenarios. This closed-loop interaction ensures that the optimization process remains data-driven and context-aware.

3.2. Framing the ED staff configuration problem for MCSA: Exploration phase

This study formulates the ED staff problem as a combinatorial optimization task, addressing both human and physical resource allocation. Each candidate solution is represented as a vector of discrete decision variables, corresponding to personnel types (e.g., doctors, nurses, technicians) and critical resources (e.g., careboxes). The dimensionality of the problem, combined with interdependent constraints, generates a vast and complex search space. To efficiently explore this space, MCSA employs Montecarlo sampling to generate candidate configurations. Each sampled point represents a potential resource allocation, constrained within predefined bounds that reflect real-world operational limits. These configurations are then evaluated through the ED simulator, allowing the heuristic to assess the impact of each on KPIs. The primary optimization objective is to minimize the average Length of Stay (LoS), a recognized indicator of patient flow efficiency and service quality Mehrolhassani, Behzadi and Asadipour (2025). Constraints such as staff availability, space limits, and operational rules are embedded directly into the sampling ranges to ensure feasibility. For the proof-of-concept, personnel availability constraints were embedded directly into the sampling ranges to exclude non-viable regions of the search space.

3.3. Optimization problem formulation

This study addresses a combinatorial optimization problem aimed at identifying ED staff configurations that minimize operational delays, specifically measured by LoS. The decision variables depict a high-dimensional, discrete search space encompassing the number of staff per role and physical resources. Due to the inherent complexity and stochastic interactions within patient flows and resource dynamics, deriving analytical or closed-form models is not tractable. Instead, an externally validated agent-based ED simulator is used to evaluate each configuration's performance by producing relevant KPIs.

To formalize the proposed approach, we introduce here the mathematical structure of the problem, including the definition of decision variables, the stochastic evaluation process, and the formulation of the objective function used in the optimization phase. Unlike classical optimization with closed-form objective functions, here the objective function is implicitly defined by the simulation outputs. To clarify this relationship, we define two interconnected mappings:

$$y = f(internalparameters, externalparameters)$$

here, the function *f* represents the simulation model where the **internal parameters** are fixed calibration factors of the simulation model (e.g., agents interaction times) and other internal aspects. In contrast, **the external parameters** represent operational decision variables including the number of doctors, nurses, technicians, and the physical resources. These variables define the search space and will be adjusted by MCSA, thereby narrowing the range of exploration to viable scenarios. Once a configuration is generated, the simulator computes the relevant KPIs, which serve as performance measures for that configuration. MCSA then applies an objective function, denoted here as:

$$y = g(KPI)$$

The function *g* represents the optimization objective used by MCSA to evaluate and compare configurations, usually with the goal of minimizing a target KPI. In this setup, the simulator encapsulates the operational constraints and dynamics of the ED, while MCSA guides the search toward configurations that enhance the performance metric returned by the simulation. This iterative framework allows the search to focus on high-potential regions of the solution space, enabling efficient refinement toward configurations that yield improved operational outcomes.

For each configuration i, the simulator is executed n = 15 times to account for the stochastic nature of the system. Each execution returns a single LoS value, which internally represents the average across five patient types. These values form the set:

$$\frac{LoS^i = \{LoS^i, LoS^i, \ldots LoS^i\}}{MCn\ 1\ 2\ n}$$

The arithmetic mean of this set represents the expected performance of configuration i, and is defined as:

$$\mu LoS(i) = \frac{1 \sum LoS^i_j}{n_{j=1}}$$

This average is used as the objective function in the optimization process:

$$f(i) = \mu_{LoS}(i)$$

Formally, the objective function used by the heuristic can be expressed as:

$$f(i) = \frac{1 \sum LoS^i_j}{n_{j=1}} \tag{1}$$

where:

- i: proposed staff configuration.
- n = 15: number of independent Montecarlo simulation runs per configuration.
- LoS: LoS value obtained from the j-th simulation run for configuration i, internally averaged over patient types.
- *f*(i): objective function value, defined as the average of LoS across the n simulations.

The complete process ensures that each evaluation yields a statistically stable performance estimate, enabling meaningful comparisons between configurations despite the stochastic nature of the system. Fig. 4 illustrates the MCSA sampling procedure within a multidimensional decision space, simplified here to two dimensions for clarity.

In the diagram, each axis corresponds to a decision variable (such as the number of doctors or nurses) allocated. Due to the combinatorial structure of the problem, the search space grows exponentially with the number and range of variables, underscoring the complexity of the task. The figure highlights how each candidate configuration is sampled and evaluated using the ED simulator, which returns KPIs that inform the heuristic's next steps. This process of iterative sampling and evaluation continues until the algorithm converges to a stable Montecarlo map of candidate solutions.

Algorithm 1 get_func_val(samples)

0: **if** number_of_generated_samples < SamplesPerItMC **then**
0: **if** list is valid according to verify_list (restrictions) **then**
0: Append list to ListCoord
0: **end if**
0: **end if**
0: **if** number_of_generated_samples = SamplesPerItMC **then**
0: Write ListCoord to disk using WriteEDSamples
0: Initialize empty list of processes
0: Build command to execute all samples with EdLauncher.sh to run N ED simulators
0: Launch the command
0: Initialize empty KPIList
0:**for** each process launched **do**
0: Wait until it finishes
0: **for** each sample processed **do**
0: Process sensor file to obtain queues, occupancy, and LoS

(*continued on next page*)

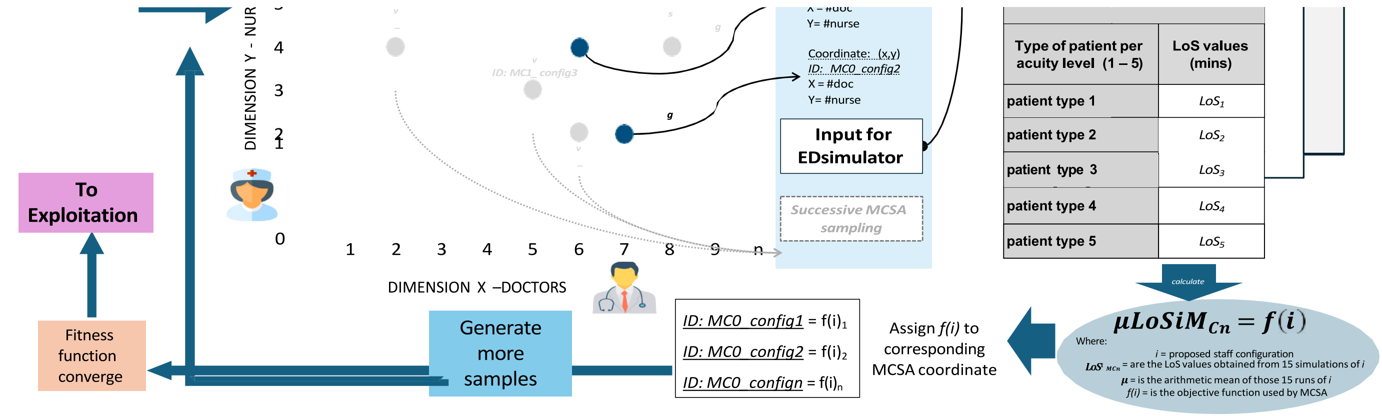

Fig. 4. Diagram illustrating MCSA sampling process of staff configurations and the subsequent evaluation by the ABM which outputs corresponding KPI values. MCSA iteratively generates and refines samples, until convergence.

(*continued*)

Algorithm 1 get_func_val(samples)

```
    Compute KPI using SelectKPI
    Append KPI to KPIList
  end for
end for
Increment Montecarlo iteration number (HIteration)
return ListCoord, KPIList
else
  return None,
None end if = 0
```

From the ED simulator, a list of KPI values (e.g., LoS) is returned after multiple executions. These values are averaged to compute the objective function *f*(**i**), which is then fed back to MCSA, allowing it to assign a performance score to that specific sample. This iterative feedback loop is implemented as shown in the Algorithm pseudocode 1, which outlines the function responsible for orchestrating the simulation calls, waiting for the results, and retrieving the averaged KPI value used in the optimization process, guiding the search toward feasible regions over time.

The simulation time of each sample plays an important role in linking the heuristic search with the simulation process. Running a full simulation for every possible configuration would be computationally expensive; MCSA strategically selects a limited number of samples, optimizing resource usage. By managing the number of evaluations and executing multiple configurations simultaneously in batch mode, MCSA ensures an efficient search process without requiring exhaustive exploration. As the simulation progresses, staff and resource agents interact with patient flow and triggers events that directly influence KPIs. This iterative process illustrates how resource variability dynamically affects system performance over time, emphasizing the necessity of optimizing staff configurations to maintain operational efficiency.

Upon completion of the simulation, the KPI output from each run is assigned as the function value $(i) = KPI$ to its corresponding coordinate within the MCSA search space. As our primary performance index, we define the global LoS as μLoS_{KPI}, representing the aggregated mean LoS across all patient types. For each candidate configuration, the simulator is executed 15 times to account for stochastic variability. In each of the 15 simulation runs, the LoS metric is computed individually for each of the five patient types defined in the model. These five values are then weighted and averaged within each run to produce a single representative LoS value per simulation. The global μLoS_{KPI} for a given configuration is then calculated as the mean of these representative values across all 15 repetitions. This global score is used as the performance indicator assigned to the evaluated configuration during the optimization process.

The distribution of the resulting LoS values across simulations approximates normality, thereby supporting the use of μLoS_{KPI} as a stable and representative performance metric.

Successive sets of samples are produced and their adequacy assessed (comparing the mean values of each set)

through the fitness function *F*, which evaluates the similarity between the average values of different sample sets. Specifically, the sampling process stops when the difference between the mean values of two consecutive sets is negligible, this is: $\mu_i - \mu_i^{-1} < \epsilon$, where μ_i and μ_i^{-1} are the mean values of the current and previous sample sets, respectively, and ϵ represents a predefined threshold (e.g., convergence within one or two decimal places). Once this convergence criterion is met, indicating that additional sampling would not significantly improve the results, the final Montecarlo map is established and MCSA transitions to the Exploitation phase.

3.4. Constraints

The optimization model incorporates a set of constraints to ensure that all evaluated configurations are operationally feasible within an ED environment. These constraints reflect practical limits on staff availability and resource capacity, as defined by institutional policies or operational guidelines.

1. **Minimum and maximum staff levels:** Each medical staff type (e.g., doctors, nurses, technicians) must be assigned in quantities that fall within a predefined valid range. Both junior and senior personnel for each role must adhere to these bounds:

$$0 \leqslant P_{type,junior} \leqslant Max_{type}\ \forall type \in MedicalStaff \tag{2}$$

$$0 \leqslant Ptype, senior \leqslant Maxtype\quad \forall type \in MedicalStaff \tag{3}$$

2. **Total staff capacity per role:** For each medical staff type, the combined number of junior and senior personnel must not exceed the maximum capacity defined for that role:

$$P_{type,junior,i} + P_{type,senior,i} \leqslant Max_{type} \forall type \in MedicalStaff,\ \forall i \in StaffConfigurations \tag{4}$$

where:

- $P_{type,junior}$ denotes the number of junior personnel of a given type (e.g., junior nurse).
- $P_{type,senior}$ denotes the number of senior personnel of that same type (e.g., senior nurse).
- Max_{type} represents the maximum allowable total number of staff for that role, regardless of seniority.

It is important to clarify that, while junior and senior personnel within each staff category may individually be assigned a value of zero, their combined total must always remain within operationally acceptable bounds. This constraint maintains flexibility in workforce allocation while ensuring that the ED staff remains functional.

The inclusion of zero in the parameter ranges reflects the structure of the search space axes rather than implying that all-zero configurations are feasible or desirable. Even if certain roles are temporarily assigned zero personnel, overall operational coverage is preserved through redundancy in other staff categories. MCSA internally enforces these constraints and automatically filters out infeasible configurations to ensure that the optimization process focuses exclusively on solutions that are both realistic and practically implementable within an ED environment. This modeling strategy enables a robust and flexible search while preserving operational integrity.

3.5. Exploitation phase

The simulation process offers a deeper insight into how various strategies impact system performance. KPIs like occupancy levels help evaluate the effectiveness of the solution. Post-simulation analysis ensures the selected configuration works in practice, without the need for a traditional objective function, as KPIs are directly calculated by the simulation. Once the sample map is generated, MCSA moves to the second phase, Exploitation, where the search is refined by focusing on promising areas.

A key challenge in this phase is identifying promising areas in the search space, which is where clustering comes in. Montecarlo simulations generate a dense map of sample points, and DBSCAN Ester, Kriegel, Sander and Xu (1996) clustering is used to group them based on spatial *density*. High-density regions suggest potential solutions, while sparse areas are discarded. This approach focuses the search on viable regions, improving efficiency by filtering out unlikely configurations. Samples are grouped into clusters and form the Feasible Regions, allowing for the identification of high-quality solutions.

The next step involves partitioning the Montecarlo map into intervals based on the distribution of sampled points to facilitate further refinement. The Montecarlo map is segmented into intervals along the (**i**) axis, ranging from the lowest to the highest observed objective function values. Fig. 5 illustrates this segmentation process. The left panel displays the complete Montecarlo map, while the right panel presents a vertical plane view, where function values $f(\mathbf{i}) = \mu(LoS_{MCn})$ are plotted against the corresponding configuration points. These function values, denoted as μLoS, represent the average LoS (in minutes) obtained from multiple simulation runs. The figure also highlights the lower and upper interval boundaries along the f (**i**) axis. Intervals are analyzed sequentially from the lowest values upward, guided by two predefined stopping criteria:

- Criterion 1: If the upper limit of the interval is reached, the process stops, and the cut is made.
- Criterion 2: If more than one high-density region (cluster) is found in two consecutive intervals, the process steps back two intervals to the point where the highest number of clusters was observed. At this point, a counter is reset to 1, meaning the process restarts from this previous interval to reevaluate the area with significant cluster density. This mechanism avoids continuing into regions of lower density and ensures the analysis focuses on meaningful areas without requiring a complete examination of the Montecarlo map.

Once high-density regions are identified, indicating potential solutions with desirable properties, these regions are prioritized for further exploration. Less promising areas above the upper cut limit are discarded, while those below the lower cut limit, with higher sample density, are retained.

To further refine the groups, we apply the k-means algorithm MacQueen (1967); Ikotun, Almutari and Ezugwu (2021); Yastrebov, Kubuś and Poczeta (2023) to dense areas, dividing the data into more precise clusters based on Euclidean distance. This two-step strategy first captures dense regions and then optimizes the positions and boundaries of the clusters, as shown in Fig. 6. Empirical testing set $k = 4$, balancing computational efficiency with effective differentiation. The algorithm shifts centroids towards areas with higher sample density, subdividing clusters into meaningful subgroups. The k-clusters containing values that represent potential solutions, with multiple clusters indicating multiple global minima at different locations.

At this point, the complete execution of MCSA has ended. However, the iterative nature of the algorithm comes into play. We shift our focus towards Feasible Regions to initiate a new exploration process, with these regions becoming the new search space. The goal is to continue finding and exploiting new solutions within the reduced space, aiming to improve results and avoid getting trapped in local minima. Additionally, each time we obtain a new solution, it will be compared with the previous one. This strategy of continuous improvement refines the search space and improves the quality of the solutions until it can no longer be enhanced.

4. Experimental validation

4.1. Preamble

To assess the effectiveness of MCSA in identifying staff

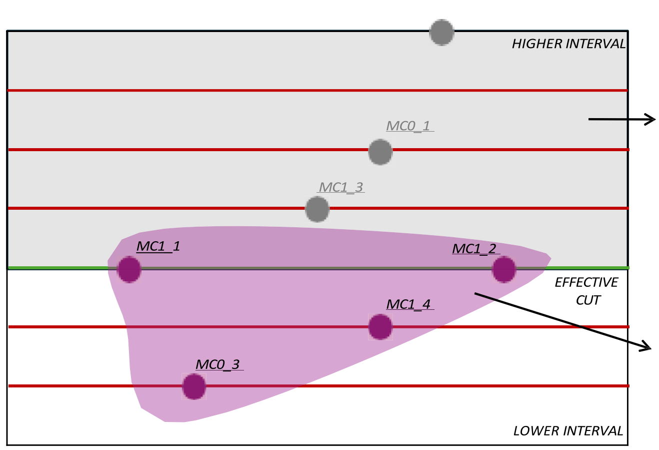

Fig. 5. Visualization of the search space before and after detecting dense regions, or regions with a high concentration of samples. To perform the cutting process, we begin by dividing the map along the f i axis to locate what we call high-density regions.

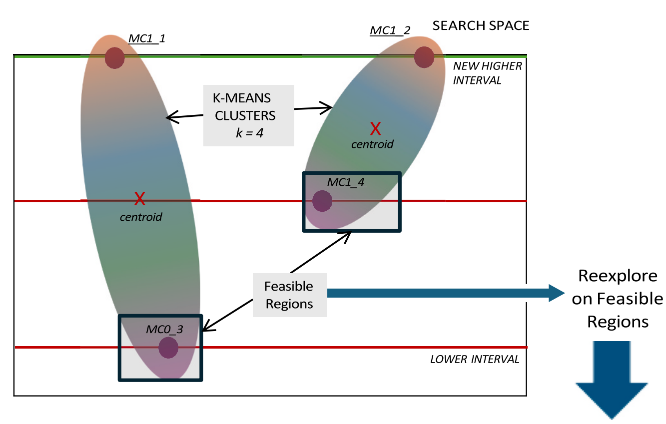

Fig. 6. After identifying dense regions, k-means clustering is applied, dividing the data into 4 k-clusters.

configurations, we implemented a structured experimental framework. The optimization process includes both phases—exploration and exploitation—and is conducted without prior knowledge or predefined assumptions.

Although local sensitivity analysis (e.g., small variations in nurse or capacity levels) is feasible, it is computationally expensive due to the cost of simulation per change. Rather than exhaustively evaluating the entire search space, MCSA focuses on progressively refining the search toward high-potential regions. Fig. 7 illustrates the patients flow through the ED modeled in the simulator, a configurable input that allows adaptation to various real-world conditions or hypothetical scenarios. For the primary optimization, the patient arrival pattern is assumed constant, based on typical hospital records across different times of the day and week. This simplifies the optimization of staff configurations without the complexity of fluctuating arrivals. The simulation runs over 90 days, offering a comprehensive analysis within a realistic operational time-frame.

4.2. Experimental setup

The experiment was conducted within a combinatorial search space defined by 12 discrete variables, each ranging over six integer values from 0 to 5 (see Table 1). This corresponds to a theoretical space of $6^{12} = 2,176,782,336$ possible configurations; however, operational constraints reduce the feasible space considerably. A key constraint imposed was the mandatory presence of at least one staff member in each area, effectively limiting the search to $5^{12} = 244,140,625$ configurations. Although the original search space is vast, MCSA applies sample restrictions and heuristic guidance to focus exploration on valid and promising configurations, avoiding exhaustive enumeration which is computationally prohibitive.

The ED simulator employed is a preconfigured, standalone model capable of processing multiple input configu- rations independently. This feature enables concurrent evaluation of samples by launching multiple instances of the simulator, facilitating efficient assessment of candidate staff allocations within the heuristic search. To manage the computational resources required for this study, we used the Department's 8-node cluster with 512 AMD Opteron™ Processor 6262 HE cores, and 2 TB RAM.

Table 1
Minimum and maximum Axis values for each type of staff (Sr and Jr are separate dimensions). The restriction $Sr+Jr \leq 10$ applies per category.

Staff/Resource	Axis Min	Axis Max	Adjusted Range [min, max]
Sr Admission	0	5	[1, 5]
Jr Admission	0	5	[1, 5]
Sr Triage	0	5	[1, 5]
Jr Triage	0	5	[1, 5]
Sr Doctor (Area A)	0	5	[1, 5]
Jr Doctor (Area A)	0	5	[1, 5]
Sr Doctor (Area B)	0	5	[1, 5]
Jr Doctor (Area B)	0	5	[1, 5]
Sr Nurse (Area A)	0	5	[1, 5]
Jr Nurse (Area A)	0	5	[1, 5]
Sr Nurse (Area B)	0	5	[1, 5]
Jr Nurse (Area B)	0	5	[1, 5]

4.3. Experimental results

During the exploration phase, the MCSA heuristic generated and simulated a variety of staff configurations, collecting the corresponding LoS indicators. This was followed by a refinement step employing two-phase clustering to focus the search on the most promising regions of the solution space. The iterative interaction between MCSA and the ED simulator enabled efficient identification of configurations that progressively improved LoS. Ultimately, MCSA outputs a staff configuration represented as a coordinate in the search space that optimizes the LoS metric. It is then the responsibility of decision-makers to assess the proposed configuration's feasibility and alignment with system-specific objectives and constraints.

The results are summarized in Table 2. The objective function aims to

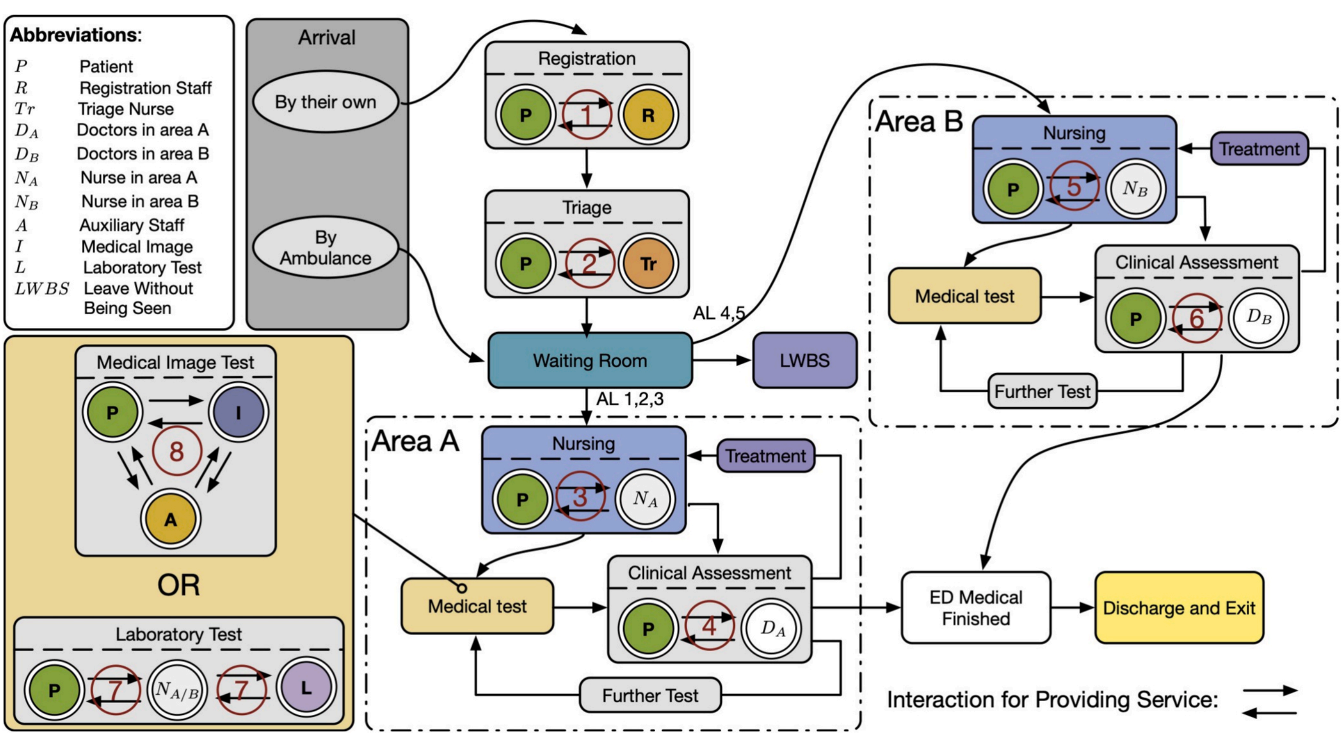

Fig. 7. Patient flow through the ED. Eight service processes, numbered in the circles, govern the patient flow. These services are interdependent, with varying service durations. Note: Areas A and B are designated for urgent and non-urgent patients, respectively, with separate staff and independent operations.).
Source: (Liu et al., 2017

Table 2
Summary of key problem parameters and resulting metrics from the MCSA heuristic. It includes the search space definition, simulation times, sampling effort, and efficient configuration identified (associated LoS).

Problem parameter	Value
f (i) = μLoSMCn	Time in minutes
Range	[0, 5]
# Dim	12
Problem Size (Psize)	6^{12}
Possible Combinations	2,176,782,336
SimTime EDSim (in seconds)	600 approximate
# samples MCSA Heuristic	6,350
% vs Psize	0.00029
Result	**Value**
SimTime MCSA samples (min.)	1,557
Improved LoS found by MCSA (min)	188
SimTime Exhaustive (min.)	48,828,125
Final configuration: Coord. form	[4, 4, 4, 2, 2, 5, 4, 3, 3, 3, 5, 5]

minimize LoS, a key metric. Although each patient type has its own LoS, a single representative value is used to reflect overall system performance, consistent with the global perspective discussed in Section 3. Notably, MCSA explores the solution space without prior knowledge of which configurations will improve LoS; instead, it iteratively optimizes the metric based on data generated through sampling. Each proposed configuration results from 15 independent executions of the ED simulator, each with different random seeds and the staff configuration suggested by MCSA. While the global optimum cannot be guaranteed, prior evaluations of MCSA on standard benchmark functions demonstrate its consistent near-optimal performance, supporting its robustness in this context.

The ED simulator employed in this study was previously calibrated and validated using real data from Hospital Parc Taulí (Sabadell), enhancing the reliability of the results. The average simulation time per configuration (SimTime EDSim), expressed in seconds, provides an estimate of the computational cost associated with each scenario, which varies according to configuration complexity. To improve efficiency, simulations were executed in parallel batches, enabling evaluation of 6,350 configurations in 25 h, compared to over 1,058 h if processed sequentially. This underscores the scalability and practical feasibility of the approach. The final configuration identified by MCSA, detailed in Table 3, achieved a significant reduction in LoS without requiring exhaustive evaluation of the entire solution space. These results demonstrate MCSA capability to derive operationally effective staff configurations through guided simulation, validating the proposed heuristic framework.

Table 3
Final staff configuration identified by MCSA. Each coordinate indicates the number of assigned staff per role and area, corresponding to the dimensions explored during the optimization.

Staff Type or Resource	MCSA coord.
Senior Admission	4
Junior Admission	4
Senior Triage	4
Junior Triage	2
Senior Doctor (Area A)	2
Junior Doctor (Area A)	5
Senior Doctor (Area B)	4
Junior Doctor (Area B)	3
Senior Nurse (Area A)	3
Junior Nurse (Area A)	3
Senior Nurse (Area B)	5
Junior Nurse (Area B)	5

5. Experimental evaluation including human and physical resource variables

5.1. Preamble

While the core experimental structure remained unchanged, the model was expanded to include additional parameters, increasing its complexity. This extended phase considers 17 variables instead of the original 12, enabling a more accurate representation of the operational intricacies within an ED. Through simulation, a wide range of staff and resource configurations are evaluated using the LoS metric, which serves as the primary indicator for assessing system performance and guiding the optimization process. The complete list of variables and their respective ranges is provided in Table 4.

The staff categories, Senior and Junior, are listed separately in the table to explicitly reflect the model's constraints. Although treated as distinct dimensions in the search space, their combined levels are subject to two key rules: (1) each category must include at least one staff member (i.e., [0, 0] is invalid), and (2) the sum of Senior and Junior staff cannot exceed the predefined upper limit (e.g., [5, 4] is invalid if it exceeds this threshold). This representation facilitates constraint visualization while ensuring that MCSA explores only feasible configurations, maintaining the integrity of the multidimensional model.

5.2. Set-up and execution

The experiment begins by defining a multidimensional search space, where each staff type is treated as an independent decision variable (see Table 5). To reflect regulatory and hospital-specific constraints, the range of values for each variable has been partially restricted. For instance, the capacity of Area B (*AreaBcap*) is bounded between 5 and 52 units, with the lower bound imposed to ensure operational safety and continuity of care. These constraints delineate a feasible region within the search space that aligns with realistic staff and infrastructure limits. Such bounds do not eliminate variability but rather guide the optimization process toward practical and implementable solutions. By constraining the search space, the algorithm avoids exploring configurations that are either operationally infeasible or clinically unsafe, while maintaining the diversity necessary for robust optimization.

To simulate increased operational pressure, a 10 % surge in patient inflow was introduced across all time intervals. This adjustment aims to replicate heightened demand scenarios and evaluate the algorithm's

Table 4
Definition of the multidimensional search space used in the MCSA heuristic. The table lists all decision variables (medical staff and physical resources), along with their full axis limits and the adjusted operational bounds applied during the optimization process.

Staff or Type of Resource	Axis Min.	Axis Max.	Adjusted Staff + Resource	
			Min	Max
Senior Admission	0	8	3	8
Junior Admission	0	8	3	8
Senior Triage	0	8	3	8
Junior Triage	0	8	3	8
Senior Doctor (Area A)	0	8	3	8
Junior Doctor (Area A)	0	8	3	8
Senior Doctor (Area B)	0	8	3	8
Junior Doctor (Area B)	0	8	3	8
Senior Nurse (Area A)	0	8	3	8
Junior Nurse (Area A)	0	8	3	8
Senior Nurse (Area B)	0	8	3	8
Junior Nurse (Area B)	0	8	3	8
Area B Capacity	0	52	5	52
Carebox Capacity	0	49	4	49
Internal Test Capacity	0	6	3	6
Lab Test Capacity	0	6	3	6
Assistants	0	3	1	3

Table 5
Summary of key parameters and results from the MCSA optimization process. The table outlines the dimensionality, effective search space, and the number of evaluations required to identify a configuration. The final result is expressed in coordinate form, representing a feasible solution that achieved the lowest observed LoS.

Problem parameter	Value
f (i) = μLoSMCn	Time in minutes
Range	[3 to 8(Staff),5 to 52(BCap),4 to 49(Carebox),3 to 6 (Int, Lab),1 to 3(Assistants)]
Effective Range	$6^{12} \times 48 \times 46 \times 4 \times 4 \times 3$
# Dim	17
Theoretical Problem Size (Psize)	$9^9 \times 53 \times 50 \times 7 \times 7 \times 4$
Possible Combinations	249,609,796,080,000
SimTime (in seconds)	358 approximate
# samples Heuristic	2586
% vs Psize	1.12 E-09
Result	**Value**
Improved LoS found by MCSA (min)	145.32
Final configuration: Coord. form	[1, 3, 2, 3, 3, 2, 4, 3, 0, 7, 5, 3, 51, 27, 1, 3, 2]

performance under stress conditions. All other simulation parameters remained consistent with prior experiments to ensure comparability.

A key difference in this scenario lies in the increased problem dimensionality, which significantly increases the search space. To manage the associated computational load, the Amazon Web Services (AWS) Amazon Web Services (2024) were employed, leveraging its cloud infrastructure for distributed simulation execution. Specifically, simulations were run in parallel using c5.18xlarge instances, each offering 72 vCPUs and 144 GB of RAM.

This infrastructure enabled concurrent execution of up to 50 simulations, reducing the total runtime for 2,586 configuration evaluations from over 258 h (if run sequentially) to approximately 5 h. Each batch of 50 simulations completed in an average of 6 min. This highlights the critical role of parallel computing in enabling high-throughput simulation-based optimization and underscores the scalability of our methodological framework.

The inclusion of additional decision variables significantly expands the combinatorial search space. Each variable contributes a different range of values—some relatively narrow (e.g., 1 to 8 staff members), others broader (e.g., 5 to 52 resources)—leading to an uneven and high-dimensional landscape. This diversity amplifies the complexity of the optimization task, particularly as variable interactions may produce conflicting effects on system performance.

In such scenarios, random or uniform exploration becomes computationally inefficient. To address this, our methodology combines guided exploration with adaptive clustering to focus search efforts on high-potential regions of the solution space. While exploration ensures broad coverage, clustering helps refine the search by identifying regions where configurations demonstrate promising performance with respect to the optimization objective. This dual strategy enables effective pruning of unpromising configurations, reducing unnecessary evaluations and accelerating convergence.

Despite the existence of a known baseline configuration currently used by hospital management, the MCSA algorithm initiates the search process independently without leveraging any prior knowledge of existing solutions. The theoretical size of the search space, considering all possible combinations within the defined parameter ranges, exceeds 249 trillion configurations (2.49×10^{14}). In this context, brute-force methods become computationally infeasible, further underscoring the necessity of a heuristic-based approach.

To manage this complexity, an adjusted range criterion was applied, restricting the search to configurations that are operationally valid and clinically safe. This a priori pruning technique significantly reduces the effective search space and helps steer the algorithm toward viable regions. While not solely responsible for optimization, this constraint-driven refinement enables faster convergence toward high-quality solutions.

The final configuration identified by MCSA was benchmarked against the hospital's baseline setup. Results demonstrate that the proposed solution consistently outperformed the baseline across multiple KPIs, particularly in terms of average LoS. These findings validate the effectiveness of our hybrid simulation–optimization methodology. More broadly, this approach offers a practical, data-driven decision support tool for medium- and long-term resource planning—especially under seasonal variations, anticipated demand surges, or temporary staff shortages due to leaves or sick absences.

5.3. Discussion of the results

The comparison between the optimized staff configuration obtained by the MCSA heuristic and the hospital's baseline configuration is shown in Tables 6(a) and 6(b). The results demonstrate a significant improvement in system performance, particularly a 23.94 % reduction in the average LoS, from 191.05 to 145.32 min. This supports our first hypothesis (H1): that targeted adjustments in staff allocation can lead to meaningful gains in operational efficiency. These improvements were

Table 6
Comparison of hospital base configuration vs MCSA solution (left: resource allocation, right: performance KPIs).

(a) Staff and Resource Allocation

Staff/Resource	Base	MCSA
Senior Admission	0	1
Junior Admission Senior Triage Junior Triage Senior Doctor (A) Junior Doctor (A) Senior Doctor (B) Junior Doctor (B) Senior Nurse (A) Junior Nurse (A) Senior Nurse (B) Junior Nurse (B) AreaBcap Carebox	3	3
Internal Test	1	2
Lab Test Assistants	2	3
	2	3
	2	2
	2	4
	3	3
	2	0
	2	7
	5	5
	7	3
	52	51
	49	27
	6	1
	6	3
	3	2

(b) Key Performance Indicators

KPI	Base Config.	MCSA Config.
Queue Lengths (number of patients)		
Admission wait Queue	19	6
Triage wait Queue	33	9
Carebox wait Queue	0	0
AreaB wait Queue	0	19
Max. Occupancy (%)		
Admission Staff	100.00	98.95
Triage Nurse	100.00	100.00
Doctor AreaA	100.00	100.00
Doctor AreaB	100.00	100.00
Nurse AreaA	100.00	87.26
Nurse AreaB	82.36	97.70
Auxiliary Staff	78.88	100.00
Laboratory	100.00	100.00
Imaging	100.00	100.00
Lenght of Stay (LoS) per patient Type (minutes)		
LoS (patient Type 1)	211.02	112.67
LoS (patient Type 2)	238.72	133.49
LoS (patient Type 3)	232.29	129.96
LoS (patient Type 4)	143.83	183.28
LoS (patient Type 5)	129.42	167.19
Average LoS	191.06	145.32

achieved while exploring a high-dimensional solution space under realistic constraints, validating the second hypothesis (H2). Despite the theoretical problem size exceeding 2.49×10^{14} combinations, only 2,586 samples were required (less than 1.12E-09 of the total space). This efficient convergence illustrates MCSA's ability to focus the search on high-potential regions and confirms our third hypothesis (H3): that MCSA is computationally more efficient than exhaustive search strategies in identifying viable solutions.

In addition to LoS, further analysis reveals that MCSA preserved or improved staff occupancy rates in critical service areas such as triage and admission. This indicates the method's ability to balance workload while reducing patient queues without necessarily increasing overall staff levels. The configuration identified by MCSA thus not only improves throughput but maintains resource utilization at operationally sustainable levels.

Moreover, the results emphasize the relevance of strategic variation in staff composition. The presence of senior staff in early care stages (admission and triage) was also associated with smoother patient flow and reduced bottlenecks. While the current study focuses on minimizing average LoS as a primary performance metric, the strong interdependence among KPIs suggests that future efforts should adopt a multi-objective perspective. For example, disaggregated LoS by patient severity levels (types 1–5) would offer more granular insights into service-specific performance and support finer-grained optimization goals.

In summary, the findings validate the feasibility and robustness of our simulation-guided heuristic approach. Compared to traditional methods that are often rigid or computationally intensive, MCSA offers a flexible, data-driven alternative capable of providing actionable insights with minimal evaluation cost. This makes it a valuable tool for decision-makers seeking to enhance resource allocation in dynamic and complex healthcare settings.

6. Conclusions

This study demonstrates the potential of combining heuristic optimization with simulation to support strategic staff allocation in hospital EDs. By performing a stochastic search aligning our MCSA heuristic with ABM simulation, an effective staff configuration was identified that significantly improved the overall performance of the system.

The results support the hypothesis that informed adjustments to staff configurations guided by simulation and heuristic analysis can lead to measurable performance improvements. MCSA can efficiently explore high-dimensional search spaces under real-world constraints, and it does so without incurring the computational cost of exhaustive methods. Notably, our approach achieved a 23.94 % reduction in average LoS compared to the baseline configuration, while evaluating only a negligible fraction of the total solution space. Rather than providing a fixed staff configuration, the proposed methodology provides decision-makers with data-driven insights into how different staff configurations may impact operational performance, such as LoS, enhancing preparedness for dynamic scenarios such as seasonal surges or staff shortages.

Future work will extend this approach toward a multi-objective framework, balancing LoS reduction with other critical KPIs such as staff occupation. Addressing this trade-off is key to supporting robust, data-informed decisions in complex healthcare environments.

7. Declaration statement

During the preparation of this work the author(s) used the Grammarly application to improve readability and language in the writing process. After using this tool/service, the author(s) reviewed and edited the content as needed and take(s) full responsibility for the content of the publication.

CRediT authorship contribution statement

Maria Harita: Investigation, Conceptualization, Methodology. **Alvaro Wong:** Data curation, Methodology. **Dolores Rexachs:** Project administration, Formal analysis. **Emilio Luque:** Supervision. **Eva Bruballa:** Supervision. **Francisco Epelde:** Supervision, Conceptualization.

Funding

This publication is supported under contract PID2020-112496GB-I00, funded by the Agencia Estatal de Investigacion (AEI), Spain and the Fondo Europeo de Desarrollo Regional (FEDER) UE, and under contract PID2023-147955NB-I00, funded by the Science Innovation and Universities Ministry.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

References

Abdalkareem, Z., Amir, A., Al-Betar, M., Ekhan, P., & Hammouri, A. (2021). Healthcare scheduling in optimization context: A review. *Health and Technology, 11*, 445–469.

Ahmadianfar, I., Bozorg-Haddad, O., & Chu, X. (2020). Gradient-based optimizer: A new metaheuristic optimization algorithm. *Information Sciences, 540*, 131–159. https://doi.org/10.1016/j.ins.2020.06.037. URL: https://www.sciencedirect.com/science/article/pii/S0020025520306241.

Ahmadianfar, I., Heidari, A. A., Gandomi, A. H., Chu, X., & Chen, H. (2021). Run beyond the metaphor: An efficient optimization algorithm based on runge kutta method. *Expert Systems with Applications, 181*. https://doi.org/10.1016/j.eswa.2021.115079. URL: https://www.sciencedirect.com/science/article/pii/S0957417421005200.

Ahmadianfar, I., Heidari, A. A., Noshadian, S., Chen, H., & Gandomi, A. H. (2022). Info: An efficient optimization algorithm based on weighted mean of vectors. *Expert Systems with Applications, 195*. https://doi.org/10.1016/j.eswa.2022.116516. URL: https://www.sciencedirect.com/science/article/pii/S0957417422000173.

Ahmadianfar, I., Samadi-Koucheksaraee, A., & Razavi, S. (2023). Design of optimal operating rule curves for hydropower multi-reservoir systems by an influential optimization method. *Renewable Energy, 211*, 508–521. https://doi.org/10.1016/j.renene.2023.04.113. URL: https://www.sciencedirect.com/science/article/pii/S0960148123005797.

Amazon Web Services, I., 2024. Amazon web services (aws). URL: https://aws.amazon.com. accessed: 2024-11-13.

Amindoust, A., Asadpour, M., & Shirmohammadi, S. (2021). A hybrid genetic algorithm for nurse scheduling problem considering the fatigue factor. *Journal of Healthcare Engineering, 2021*, Article 5563651. https://doi.org/10.1155/2021/5563651

Boresta, M., Giovannelli, T., & Roma, M. (2024). Managing low–acuity patients in an emergency department through simulation–based multiobjective optimization using a neural network metamodel. *Health Care Management Science, 27*, 415–435. URL: https://api.semanticscholar. org/CorpusID:270369303.

Bruballa, E., Wong, A., Rexachs, D., Luque, E., & Epelde, F. (2017). Evaluation of response capacity to patient attention demand in an emergency department. *International Journal on Advancesin Systems and Measurements, 10*, 11–22.

Cabrera, E., Taboada, M., Iglesias, M., Epelde, F., & Luque, E. (2011). Optimization of healthcare emergency departments by agent-based simulation. *Procedia Computer Science, 4*, 1880–1889.

Ceschia, S., Di Gaspero, L., Mazzaracchio, V., Policante, G., & Schaerf, A. (2023). Solving a real-world nurse rostering problem by simulated annealing. *Operations Research for Health Care, 36*. https://doi.org/10.1016/j.orhc.2023.100379

Chen, P., Huang, W., Chiang, T., & Yu-Hsin Chen, G. (2020). Applying heuristic algorithms to solve inter-hospital hierarchical allocation and scheduling problems of medical staff. *International Journal of Computational Intelligence Systems, 13*, 318–331.

Ester, M., Kriegel, H., Sander, J., & Xu, X. (1996). A density-based algorithm for discovering clusters in large spatial databases with noise. In *Proceedings of 2nd International Conference on Knowledge, Discovery and Data Mining* (pp. 226–231).

Forbus, J. J., & Berleant, D. (2022). Discrete-event simulation in healthcare settings: A review. *Modelling, 3*, 417–433.

Harita, M., Wong, A., Rexachs, D., & Luque, E. (2022). Kp01 solved by an n-dimensional sampling and clustering heuristic. *International Conference on Computational Science, Springer International Publishing.*, 229–236.

Harita, M., Wong, A., Rexachs del Rosario, D., & Luque Fadón, E. (2021). Evaluation of a heuristic search algorithm based on sampling and clustering. *IX Jornadas de Cloud Computing, Big Data & Emerging Topics (Modalidad virtual, 22 al 25 de junio de 2021)*.

Harita, M., Wong, A., Suppi, R., Rexachs, D., & Luque, E. (2024). A metaheuristic search algorithm based on sampling and clustering. *IEEE Access, 12*, 15493–15508. https://doi.org/10.1109/ACCESS.2024.3354714

Ikotun, A. M., Almutari, M. S., & Ezugwu, A. E. (2021). K-means-based nature-inspired metaheuristic algorithms for automatic data clustering problems: Recent advances and future directions. *Applied Sciences, 11*. https://doi.org/10.3390/app112311246

Juan, A. A., Faulin, J., Grasman, S. E., Rabe, M., & Figueira, G. (2015). A review of simheuristics: Extending metaheuristics to deal with stochastic combinatorial optimization problems. *Operations Research Perspectives, 2*, 62–72. https://doi.org/10.1016/j.orp.2015.03.001. URL: https://www.sciencedirect.com/science/ article/pii/S221471601500007X.

Karp, R. M. (1972). Reducibility among combinatorial problems. *Complexity of Computer Computations, Springer.*, 85–103.

Lameesa, A., Hoque, M., Alam, M., Ahmed, S., & Gandomi, A. (2024). Role of metaheuristic algorithms in healthcare: A comprehensive investigation across clinical diagnosis, medical imaging, operations management, and public health. *Journal of Computational Design and Engineering, 11*, 223–247.

Liu, Z., Cabrera, E., Rexachs, D., & Luque, E. (2014). A generalized agent-based model to simulate emergency departments. In *Sixth Int Conf Adv Syst Simul* (pp. 65–70).

Liu, Z., Rexachs, D., Epelde, F., & Luque, E. (2017). A simulation and optimization based method for calibrating agent-based emergency department models under data scarcity. *Computers and Industrial Engineering, 103*, 300–309.

MacQueen, J. (1967). Some methods for classification and analysis of multivariate observations. *Proceedings of 5-th Berkeley symposium on mathematical statistics and probability*. University of California Press.

Mehrolhassani, M. H., Behzadi, A., & Asadipour, E. (2025). Key performance indicators in emergency department simulation: A scoping review. *Scandinavian Journal of Trauma, Resuscitation and Emergency Medicine, 33*, 15.

Mohamed, A., Sallam, K., Agrawal, P., Hadi, A., & Mohamed, A. (2023). Evaluating the performance of meta-heuristic algorithms on cec 2021 benchmark problems. *Neural Computing and Applications, 35*, 1493–1517. https://doi.org/10.1007/s00521-022-07788-z

Penn, M., Hasson, D., & Avriel, M. (1994). Solving the 0–1 proportional knapsack problem by sampling. *Journal of Optimization Theory and Applications, 80*, 289–297. https://doi.org/10.1007/BF02192937

Samadi-Koucheksaraee, A., Shirvani-Hosseini, S., Ahmadianfar, I., & Gharabaghi, B. (2022). Optimization algorithms surpassing metaphor. In *Computational intelligence for water and environmental sciences* (pp. 3–33). Springer.

Taboada González, M. (2013). *Simulación del Servicio de Urgencias Hospitalarias: una aproximación computacional desarrollada mediante técnicas de Modelado Orientadas al Individuo (MoI)*. Universitat Autònoma de Barcelona. Barcelona. Ph.D. thesis. URL: https://ddd.uab.cat/ record/126902. tesi doctoral - Universitat Autònoma de Barcelona. Departament d'Arquitectura de Computadors i Sistemes Operatius.

Tisue, S., & Wilensky, U. (2004). Netlogo: A simple environment for modeling complexity. *International conference on complex systems*, 16–21.

Tomar, V., Bansal, M., & Singh, P. (2024). Metaheuristic algorithms for optimization: A brief review. *Engineering Proceedings, 59*, 238.

Tutsoy, O., & Tanrikulu, M. Y. (2022). Priority and age specific vaccination algorithm for the pandemic diseases: A comprehensive parametric prediction model. *BMC Medical Informatics and Decision Making, 22*, 4. https://doi.org/10.1186/s12911-021-01720-6

Valouxis, C., Gogos, C., Goulas, G., Alefragis, P., & Housos, E. (2012). A systematic two phase approach for the nurse rostering problem. *European Journal of Operational Research, 219*, 425–433.

Vecillas Martin, D., Berruezo Fernández, C., & Gento Municio, A. M. (2025). Systematic review of discrete event simulation in healthcare and statistics distributions. *Applied Sciences, 15*, 1861.

White, C., Nano, E., Nguyen-Ngoc, D., & White, G. (2007). An evaluation of certain heuristic optimization algorithms in scheduling medical doctors and medical students. In *Practice and Theory of Automated Timetabling VI: 6th International Conference, PATAT 2006 Brno, Czech Republic, August 30–September 1, 2006 Revised Selected Papers* (pp. 105–115). Berlin Heidelberg: Springer.

Wilensky, U., & Rand, W. (2015). *An introduction to agent-based modeling: modeling natural, social, and engineered complex systems with NetLogo*. MIT Press.

Yastrebov, A., Kubuś, Ł., & Poczeta, K. (2023). Multiobjective evolutionary algorithm idea and k-means clustering for modeling multidimenional medical data based on fuzzy cognitive maps. *Natural Computing, 22*, 601–611.

Maria Harita is a Ph.D. student in High-Performance Computing (HPC) at the Autonomous University of Barcelona. Her research focuses on the design and development of heuristic algorithms, with a particular interest in optimization techniques applied to complex problems. Among the most notable contributions is the development of an optimization algorithm designed to tackle large-scale problems. She was also a collaborating teacher in subjects such as Distributed Systems. Her research is supported by the Spanish government.

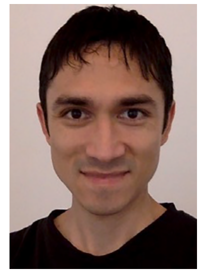

Alvaro Wong is an Associate Researcher at the Computer Architecture and Operating System Department at University Autonoma of Barcelona, Spain. He has worked in performance prediction of HPC applications in the ITEA 2 European Project No 09011, research centers & industries. He has co-authored a total of 20 full-reviewed technical papers in journals and conference proceedings.

Dolores Rexachs is an Associate Professor at the Computer Architecture and Operating System Department at University Autonoma of Barcelona (UAB), Spain. She has been the supervisor of 13 PhD theses and has been an invited lecturer in Universities in Argentina, Brazil, Chile and Paraguay. Her research interests include parallel computer architecture, parallel I/O subsystem, fault tolerance in parallel computers, tools to evaluate, predict, and improve the performance in parallel computers. She has co-authored more than 70 fully reviewed technical papers in journals and conference proceedings.

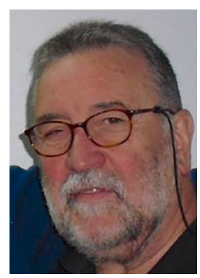

Emilio Luque is an Emeritus Professor at the Computer Architecture and Operating System Department at University Autonoma of Barcelona, Spain. Invited lecturer at universities in the USA, South America, Europe and Asia, key note speaker in several conferences and leader in several research projects funded by the European Union (EU), Spanish government and different industries. His major research areas are: parallel and distributed simulation, performance prediction and efficient management of multicluster-multicore systems and fault tolerance in parallel computers. He has supervised 20 PhD theses and co-authored more than 230 technical papers in journals and conference proceedings. conference proceedings.

Eva Bruballa is currently an Associate Professor with the Computer Science School, Escoles Universitaries Gimbernat i Tomàs Cerdà, University Autonoma of Barcelona (UAB), and an Associate Researcher with the Computer Architecture and Operating System Department, UAB, Spain. She has coauthored eight fully reviewed articles. Her research is on the use of simulation to improve ED operation.

Francisco Epelde is a specialist in internal medicine, a Master in ischemic heart disease, hospital management, cardiovascular diseases, and the patient of advanced and pluripatological age with the Central University of Barcelona (UB), and a Master in clinical management with the Catholic University of Murcia. He is currently the Medical Coordinator of the Short Stay Unit of the Hospital Universitari Parc Taulí and the Emergency Working Group of the Spanish Society of Internal Medicine. He is an Associate Professor with the Department of Medicine, Universitat Autonoma de Barcelona (UAB), where he is responsible for emergency medicine. He has been an external consultant in emergency medicine with the WHO. He has authored over 250 publications and seven book chapters.