

WILEY

INTERNATIONAL
TRANSACTIONS
IN OPERATIONAL
RESEARCHIntl. Trans. in Op. Res. 0 (2025) 1–47
DOI: 10.1111/itor.70140

A unified formulation for home healthcare routing and scheduling problems

Sara Ceschia^a , Francesca Da Ros^a , Luca Di Gaspero^{a,*} ,
Simona Mancini^b , Vittorio Maniezzo^c , Roberto Montemanni^d ,
Roberto Maria Rosati^{a,e,f}  and Andrea Schaerf^a 

^aIntelligent Optimization Laboratory, Polytechnic Department of Engineering and Architecture, University of Udine, Via delle Scienze 206, Udine I-33100, Italy

^bDepartment of Engineering, University of Palermo, Viale delle Scienze, Palermo I-90128, Italy

^cDepartment of Computer Science, University of Bologna, Via dell'Università 50, Cesena I-47521, Italy

^dDepartment of Sciences and Methods for Engineering, University of Modena and Reggio Emilia, Via Amendola, 2, Reggio Emilia I-42122, Italy

^eResearch Institute for Supply Chain Management, WU Vienna University of Economics and Business, Vienna 1020, Austria

^fDepartment of Operations, Energy, and Environmental Management, University of Klagenfurt, Klagenfurt 9020, Austria

E-mail: sara.ceschia@uniud.it [Ceschia]; francesca.daros@uniud.it [Da Ros]; luca.digaspero@uniud.it [Gaspero];

simona.mancini@unipa.it [Mancini]; vittorio.maniezzo@unibo.it [Maniezzo];

roberto.montemanni@unimore.it [Montemanni]; robertomaria.rosati@wu.ac.at [Rosati];

andrea.schaerf@uniud.it [Schaerf]

Received 8 April 2025; received in revised form 27 November 2025; accepted 1 December 2025

Abstract

Home Healthcare is an essential component of healthcare systems, where caregivers visit patients' homes to deliver services. While presenting advantages with respect to institutional care, such as being cost-effective and alleviating family burdens, it presents challenges in scheduling and routing caregivers efficiently. While various formulations of such a problem—the home healthcare routing and scheduling problem—have been proposed, many fail to include key real-world features, limiting their applicability in practice. This paper consolidates relevant features from existing models and significantly extends the formulation by Mankowska et al. (2014) to develop a unified framework incorporating as many real-world aspects as possible. We introduce a new instance generator, a set of real-world features to describe these instances, and a comprehensive infrastructure for validating and comparing solutions. Furthermore, we extend two state-of-the-art solution methods—simulated annealing and one based on a compact mixed integer linear programming model solved by a state-of-the-art software – to integrate all these features and provide a comparative analysis of their performance. Results show that our general solution methods are able to outperform methods specialized for a specific formulation on many instances.

*Corresponding author.

© 2025 The Author(s).

International Transactions in Operational Research published by John Wiley & Sons Ltd on behalf of International Federation of Operational Research Societies.

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

Keywords: home healthcare; routing and scheduling; benchmarking; simulated annealing; mixed integer linear programming

1. Introduction

Home Healthcare (HHC) constitutes an essential component of modern healthcare systems. It provides patients with personalized, accessible, and continuous care in the comfort of their homes, extending services beyond traditional medical facilities such as hospitals and nursing homes. HHC is particularly beneficial for seniors, individuals with disabilities, and those recovering from illness or surgery, as it includes services such as medical care (e.g., nursing, physical therapy), assistance with activities of daily living (e.g., bathing, meal preparation), and emotional support through companionship. Compared to traditional institutional care, HHC offers several advantages: it is typically more cost-effective (Genet et al., 2012), reduces the burden on family caregivers, and can be adapted to address evolving patient needs while enabling individuals to receive care in familiar environments.

The operational model of HHC involves specialized caregivers visiting patients' homes to deliver required services before proceeding to subsequent appointments. Therefore, unlike traditional healthcare scheduling (Vieira et al., 2021; Ceschia et al., 2023), these systems must account for both the scheduling and routing of caregivers. This optimization problem, known as the Home Healthcare Routing and Scheduling Problem (HHCRSP), has emerged as a significant and complex topic within the Operations Research (OR) community.

The literature presents numerous formulations of HHCRSP, each with distinct constraints and objectives (see, e.g., Euchi et al., 2022; Fikar and Hirsch, 2017). However, relatively few formulations incorporate multiple real-world features, limiting their broader practical applicability. Furthermore, problem data and benchmark instances are rarely available to the research community, primarily due to significant privacy concerns associated with real-world healthcare data. This scarcity of accessible test instances creates additional challenges for comparative analysis and validation of solution approaches across studies.

Among the HHCRSP formulations, the specification proposed by Mankowska et al. (2014) has attracted the attention of several researchers (e.g., Lasfargeas et al., 2019; Kummer et al., 2020, 2024)—possibly because it represents a significant attempt to address the availability of test instances and therefore allows a sound comparison of solution methods. This formulation involves scheduling caregiver visits to patients, where each patient requires specialized services and specifies preferred service times. The model optimizes daily scheduling decisions by minimizing caregiver travel times between patient locations and reducing deviations from patients' preferred time windows. Researchers have addressed this problem using various solution methods, with state-of-the-art approaches including a local search method based on simulated annealing (SA) by Ceschia et al. (2026) and a mixed integer linear programming (MILP) model by Montemanni et al. (2025).

The public dataset introduced by Mankowska et al. (2014) has been later supplemented by contributions from Kummer (2021) and Ceschia et al. (2026). All three datasets are artificial, though with varying characteristics. The original dataset exhibits relatively limited variability among instances. In contrast, the subsequent datasets offer a wider range of features. Specifically, Kummer's

dataset provides increased diversity by incorporating realistic travel times, although the instances remain relatively small in scale. The dataset presented by Ceschia et al. enhances realism further by incorporating actual population distributions.

While the HHCRSP formulation by Mankowska et al. (2014) has served as a foundation for benchmarking optimization methods, it still lacks several important real-world features, making the proposed model not completely realistic from a practical perspective. For instance, it does not account for multiple departure points for caregivers, caregiver working shifts (including overtime and general working time conditions), optional patient visits, or incompatibilities between patients and caregivers. Although other problem formulations have addressed these aspects individually, they have not been considered simultaneously in published research, and related instances are either unpublished, no longer available, or not based on realistic data.

This work aims to extend the specifications from Mankowska et al. (2014), for which public benchmarking instances are available, by incorporating the most significant and practically relevant features proposed across various HHCRSP formulations in the literature. We present a unified framework that positions the existing specifications as subclasses, while enriching them with additional entities and concepts that address a broader spectrum of real-world constraints. To support this enhanced formulation, we provide a comprehensive collection of artificial yet realistic instances that facilitate rigorous benchmarking within this unifying framework.

The main contributions of this work are as follows:

- We present a unified formulation for the HHCRSP that comprehensively integrates diverse constraints and features from existing literature, creating a standardized framework for addressing real-world HHC routing and scheduling problems.
- We establish a robust testing infrastructure to facilitate fair comparison of solution methods, including an instance generator, feature extractor, and solution validator. This infrastructure uses a flexible json file format we developed to capture the full range of problem characteristics in a structured, yet readable way.
- We convert existing benchmark datasets to our proposed file format and introduce a new dataset demonstrating our instance generator's capability to create diverse problem instances. These enable direct comparisons with previous research and provide a foundation for future comparative studies.
- We define a comprehensive set of features and use them for analyzing HHCRSP instances. This analysis confirms that our generated instances span the entire feature space of existing literature formulations and extend into novel combinations of features, providing more comprehensive coverage of potential real-world scenarios and establishing a richer benchmark for evaluating solution methods.
- We extend existing solution methods, a Simulated Annealing (SA) (Ceschia et al., 2026) and a Mixed Integer Linear Programming (MILP) (Montemanni et al., 2025), to incorporate additional real-world features that were absent in their original formulations and are included in our new unified formulation. We evaluate these enhanced methods on the translated benchmarks and our newly generated instances, presenting a comprehensive comparative analysis that assesses their performance across the expanded feature space.

All datasets, solution files, and supporting tools are publicly accessible on the Zenodo repository (Ceschia et al., 2025) and the accompanying GitHub project (<https://github.com/iolab-uniud/uohc>) to promote reproducibility and comparison with future studies.

While our work tackles a common challenge in optimization problems—namely, the diversity of both formulations and input formats across related problem variants—it also presents some limitations. As discussed in Section 3, we consider features that are either directly associated with the original benchmark formulation from Mankowska et al. (2014) or derived from publicly available datasets (or those available upon request). Our efforts have included more than 10 real-world aspects; however, some features, such as considering a multiday time horizon, have been excluded. Introducing such changes would significantly move away from the original datasets, affecting their readability and usability. Nevertheless, we acknowledge the importance of these features and will address them in future work (Section 8).

The remainder of this work is structured as follows. In Section 2, we review related research in the field. Then, in Section 3, we introduce the basic notions of the HHCRSP and all the real-world features we consider. Section 4 presents a mathematical model capturing these features. In Section 5, we describe the process for representing the data in the new formulation, analyze instance features and their distribution, and present our software toolbox, for instance, generation, validation, and analysis. We present the extension of the two solution methods in Section 6. In Section 7, we report our experimental results and provide some managerial insights. Finally, in Section 8, we draw some conclusions and discuss future work, including managing missing features. The appendices contain a detailed description of the data format (Appendix A) and of the tuning process of the SA solver (Appendix B).

2. Related work

The problem of scheduling HHC services has been extensively studied over the past three decades. Begur et al. (1997) were the first to formally introduce the problem, addressing a simplified version which they solved using a Clarke & Wright routing heuristic. One year later, Cheng and Rich (1998) proposed the first MILP formulation for the problem. Later approaches modeled this problem as a *set covering problem*. Notably, Rasmussen et al. (2012) made a significant contribution by being the first to incorporate temporal dependencies between activities into their formulation. Bertels and Fahle (2006), Rendl et al. (2012), and Hiermann et al. (2015) investigated scenarios in metropolitan settings where caregivers use public transportation networks for traveling between patients' locations, rather than using dedicated vehicles. Additionally, Rendl et al. (2012) and Hiermann et al. (2015) expanded this concept by incorporating transportation mode changes during trips, enabling multi-modal journeys.

Di Gaspero and Urli (2014) addressed a similar setting, albeit from a different perspective. Their study examined a multiday horizon, aiming to minimize caregivers' overtime while ensuring balanced workloads. Furthermore, they introduced the option of leaving certain patients unscheduled, which more accurately reflects overloaded systems where covering all demand with existing staff is unfeasible, suggesting the potential need for hiring temporary caregivers. To tackle this problem, the authors developed a Constraint Programming (CP) framework featuring specialized branching heuristics and solved the problem with a Large Neighborhood Search (LNS) algorithm.

Mankowska et al. (2014) proposed a standardization of the problem and provided a set of benchmark instances. Furthermore, they are the first to consider synchronization constraints at patients' homes, allowing for a maximum of two services to be simultaneously performed. The goal is to minimize the total travel time, the overall tardiness, and the highest individual tardiness. By incorporating the highest tardiness component, they aim to enhance fairness among patients, ensuring that no individual experiences excessive delays compared to others. To efficiently handle the problem, they designed an Adaptive Variable Neighborhood Search (AVNS) algorithm exploiting eight neighborhood operators.

The same formulation has attracted interest also from Kummer et al. (2020, 2024) and Kummer et al. (2021), who designed a Biased Random Key Genetic Algorithms (BRKGA), which outperformed the results from Mankowska et al. (2014). They also extend the original dataset (Mankowska et al., 2014) with new instances that consider realistic travel times (Kummer, 2021). The results provided by Kummer (2021) have recently been outperformed by Ceschia et al. (2026), in which the authors implemented a Multi-Neighborhood local search guided by SA. Furthermore, Ceschia et al. (2026) also introduced a new benchmark of realistic instances, developed by considering demographic data and for which Open Source Routing Machine (OSRM) has been used to compute actual travel times on the road network. The dataset by Mankowska et al. (2014) has also been used by Lasfargeas et al. (2019), who proposed a Variable Neighborhood Search (VNS) algorithm to address a multi-period HHCRSP on derived instances, and by Montemanni et al. (2025), who solved the problem with a compact MILP model. Further details about the datasets by Kummer (2021) and Ceschia et al. (2026) are presented in Section 5.3.

The research by Ait Haddadene et al. (2016) also considered patient availability time windows but differed from Mankowska et al. (2014) by treating them as strict constraints. Their objective focused on minimizing total travel times and penalties for unsuitable caregiver-patient assignments. This work was the first to consider integrating patient preferences into caregiver assignments—a crucial practical consideration, as effective HHC depends significantly on strong patient–caregiver relationships. A Mixed Integer Programming (MIP) model, a Greedy Randomized Adaptive Search Procedure (GRASP), and an Iterated Local Search (ILS) are proposed to handle this problem. The performances of the approaches have been tested on the benchmark instances by Bredström and Rönnqvist (2008) for the vehicle routing and scheduling with temporal precedence and synchronization constraints, suitably adapted to fit with the HHC application, in which multiple types of services are considered. The same setting has been addressed in Masmoudi et al. (2023), where several metaheuristic approaches are presented.

Several population-based approaches have been presented in the literature. Decerle et al. (2018) and Grenouilleau et al. (2019) developed memetic algorithms, combining genetic algorithm (GA) with local search procedures, whereas a model-based evolutionary algorithm has been proposed by Clapper et al. (2023). Similar to Di Gaspero and Urli (2014), Grenouilleau et al. (2019) considered a multiday horizon, aiming at minimizing total travel time plus penalties for overtime, while ensuring workload balance among caregivers. Their setting is more complex than the one addressed by Di Gaspero and Urli (2014), as they incorporate time-dependent travel times between patients' locations.

More recently, Xiang et al. (2021) and Oladzad-Abbasabady et al. (2023) considered a multi-objective version of the problem, looking for a balance between operational costs and patients' and caregivers' satisfaction, exploiting a Non-dominated Sorting Genetic Algorithm II (NSGA-II)

and ILS algorithms, respectively. Kordi et al. (2023) also addressed the problem with a Multi-Objective VNS approach employing four objectives, that is, total cost, emissions, workload balance, and service quality.

Rich formulations dealing with real-world constraints and features include the following works. Liu et al. (2017) were the first to consider lunch breaks for caregivers. In addition, Liu et al. (2021) considered flexible departure locations (hospital or caregiver's home) and task synchronization. Conversely, Bazirha et al. (2023a) consider multiple time windows for the patients in addition to synchronization and desynchronization requirements. de Aguiar et al. (2023) analyzed a context with a limited number of vehicles relative to the number of caregivers. This constraint means that a single car should be shared by two caregivers, requiring routes that accommodate patient visits for both staff members. Yadav and Tanksale (2022) considered a very rich and complex pandemic setting in which patients select the level of safety measures (in terms of contact limitations) required (based on their pathology or their health conditions) and may impose requirements on the gender and the language spoken by the operator. The goal is to maximize the number of patients served and the revenue obtained. Unfortunately, their datasets are not publicly available. Varas et al. (2024) incorporated several additional real-world features. Specifically, caregivers are required to return to the hospital for lunch breaks, some services involve the collection of perishable biological samples that must be delivered to the hospital within a maximum time frame, and patients express multiple availabilities.

Stochastic travel and service times have been introduced by Bazirha et al. (2023b). The problem is formulated as a two-stage stochastic programming model, in which routing and scheduling plans are provided first. At the same time, penalty costs incurred for delayed services to patients and caregivers' overtime are computed at the second stage. The authors provided a Genetic Algorithm (GA) to solve a deterministic version of the problem. The algorithm is embedded with a Monte Carlo simulation to compute the expected second-stage costs of the first-stage solution corresponding to each chromosome in the GA. This approach has been tested on a new dataset inspired by the benchmark of Mankowska et al. (2014), properly extended to consider stochastic travel and service times. Liu et al. (2019) proposed a chance-constrained model that handles stochastic travel and service times. Specifically, they employ a route-based deterministic formulation in which a route is feasible if it respects patients' time windows in at least a minimum percentage of a set of stochastic scenarios. A similar approach was adopted by Ma et al. (2022) and Zhang et al. (2023). Shahnejat-Bushehri et al. (2021), instead, proposed a robust approach.

Table 1 summarizes the objectives and features of the single-period HHCRSP formulations introduced in recent literature. Table 2 reports the abbreviations employed in Table 1. The interested reader is forwarded to the following surveys on HHCRSP (Fikar and Hirsch, 2017; Cissé et al., 2017; Grieco et al., 2021; Di Mascolo et al., 2021; Euchi et al., 2022).

3. A comprehensive formulation for HHCRSP

We propose a unified and comprehensive formulation of the HHCRSP that integrates the diverse variants found in the literature. We first define the essential entities of any HHCRSP model (Section 3.1). These constitute the core elements upon which all problem variants are built, regardless of their specific operational contexts or additional constraints. Then, we describe the additional

Table 1
Objectives and features of recent single-period HHCRRSP formulations

Formulation	Objectives													Features												
	FA	OT	TA	TC	TT	UP	WT	Other	CB	CC	MD	MS	PR	PC	SY	SK	TW	UN	WR							
Mankowska et al. (2014)			✓		✓						✓				P-S-D	H	SI-H		H							
Hiermann et al. (2015)			✓		✓		✓	skills violation		✓			S		S	S	SI-S		H							
Ait Haddadene et al. (2016)					✓						✓		S		P-S	H	SI-H		H							
Liu et al. (2017)				✓		✓			P								SI-H		H							
de Aguiar et al. (2023)				✓				working time, number of caregivers	P		✓			S		H	SI-H		H							
Decerle et al. (2018)				✓						✓				S		H	SI-H	TS	H							
Liu et al. (2019)				✓						✓						H	SI-S	TS	H							
Liu et al. (2021)				✓					P		✓					H	SI-H		H							
Ma et al. (2022)				✓						✓						H	SI-S	TS	S							
Shahnejat-Bushehri et al. (2021)	OR		✓				✓	skills violation, number of caregivers			✓				S	S	SI-H	TS								
Xiang et al. (2021)				✓												H	SI-H		H							
Yadav and Tanksale (2022)				✓				max number served patients, revenue	P		✓		H			H	SI-H		H							
Bazirha et al. (2023b)			✓												P-S	H	SI-H	TS								
Bazirha et al. (2023a)	OR			✓			✓							S		H	SI-H		H							
Clapper et al. (2023)			✓				✓									H	SI-H		H							
Kordi et al. (2023)	OR		✓				✓	patients satisfaction								H	SI-H		H							
Oladzad-Abbasabady et al. (2023)			✓				✓			✓			S	P-S-D	H	SI-S		H								
Varas et al. (2024)	OR			✓			✓	patients satisfaction	H			✓	✓	H	S	M-H			H							
Zhang et al. (2023)				✓			✓			✓						H	SI-H	S	H							

elements that capture the operational complexities of real-world HHC settings (Section 3.2) and bridge the gap between the academic formulations and the practice. Finally, we identify and categorize the cost components and constraint violations that shape the objective function (Section 3.3).

3.1. Essential entities

3.1.1. Patients

Patients are the recipients of HHC services, each characterized by a set of attributes that may influence service delivery planning. Each patient is associated with a specific geographic position (i.e., their home), which is the destination for caregivers' visits. The spatial distribution of patients, combined with the underlying road and transportation networks, determines travel times between locations.

The patients' temporal availability is represented by at least one time window (Section 3.2.1), depending on personal schedules, preferences, or medical requirements. Time windows establish constraints within which services must be delivered.

3.1.2. Services

Services represent the specific care activities that must be delivered to patients, spanning a broad spectrum of medical, therapeutic, and supportive interventions. Each service is characterized by its nature and purpose, ranging from clinical procedures such as medication administration or wound care to supportive functions like personal hygiene assistance or rehabilitation exercises. The diversity of services reflects the multidimensional nature of HHC delivery and requires the appropriate matching of patient needs with caregiver capabilities (Section 3.1.3).

Several parameters define the temporal dimension of services. Service duration specifies the time required to complete the activity, which may vary based on patient condition, service complexity, or caregiver experience. For time-sensitive services, such as insulin administration or collecting biological samples, precise timing is critical to ensure clinical efficacy or specimen viability. Patients often require multiple services, creating interdependencies that must be carefully coordinated. Services may be subject to precedence relationships, where certain activities must precede others to maintain care continuity and effectiveness. The synchronization of services (Section 3.2.2), mainly when multiple caregivers must work in tandem, introduces additional complexity to the scheduling problem.

3.1.3. Caregivers

Caregivers are the workforce responsible for delivering HHC services to patients. Each caregiver has a unique professional profile defined by their qualifications (also called specializations, competencies, or skills). Qualifications are essential in HHC planning, as not all caregivers can perform all services. For instance, a nurse caregiver specialized in wound care may not be qualified to conduct physical therapy, while a physiotherapist caregiver may lack the credentials to administer certain medications.

Each caregiver typically operates within defined working time shifts (Section 3.2.3), establishing their availability throughout the planning horizon. These shifts may incorporate mandatory lunch breaks (Section 3.2.4) and allow for overtime under the limits of labor regulations.

Mobility aspects further characterize caregivers within the HHC system. Each caregiver begins and ends their workday at designated locations, which may be their home, a healthcare facility, or another fixed point (Section 3.2.5). Their movement between patient locations is governed by available transportation modes, which influence travel times and operational costs. In this formulation, we consider only one transportation mode, that is, individual cars.

3.1.4. Scheduling horizon

The scheduling horizon represents the temporal framework within which HHC operations are planned and executed. In this formulation, we restrict it to a single day and time granularity based on minutes.

3.1.5. Core feasibility criteria

The essential feasibility criteria for a solution to the HHCRSP problem are as follows. A service cannot be provided by a caregiver who is not qualified for it. This ensures that patients receive care only from caregivers with the appropriate skills and certifications required for their medical needs. Moreover, a service cannot start before the patient's time window begins. In case of early arrival, the caregiver has to wait until the time window starts. This temporal constraint respects patient availability and preferences while ensuring service delivery occurs when patients are prepared to receive care.

3.2. Additional elements

3.2.1. Time windows

Patient availability is defined through either *single* or *multiple* time windows. When patients have multiple time windows and require multiple services, as in Bazirha et al. (2023a, 2023b), all services must be delivered within a single selected time window, rather than distributed across different periods.

Time window enforcement can be approached in two primary ways: the standard approach requires only that the service begins within the designated time window, allowing the service to potentially extend beyond the window's end (Mankowska et al., 2014); alternatively, both the beginning and end of the service must occur within the specified time frame (*strict time windows*, for example, Bazirha et al., 2023a &b).

3.2.2. Synchronization

For patients necessitating multiple services, the delivery pattern may take various forms to accommodate clinical requirements and operational constraints. Services may be provided *simultaneously*, with multiple caregivers attending to the patient at the same time; *sequentially*, where one service must follow another; or *independently*, allowing services to be scheduled without interdependence.

Sequential double-service patients require a specific minimum and maximum time gap between services. This constraint expresses a relationship between services; for example, a medication must be administered at least a certain time after a meal but before a specific elapsed time to ensure therapeutic efficacy.

3.2.3. Work shift

The caregivers' schedule is structured around defined shifts, which may either consist of a complete workday (full-time) or segmented periods (part-time). Each caregiver begins their service route from a designated location at or after their shift's start and is expected to return by the shift's end. When a caregiver's return time extends beyond their scheduled shift end, this constitutes *overtime* (Section 3.3).

3.2.4. Caregiver break

Caregiver schedules must accommodate rest periods, particularly during extended shifts. Caregivers working through midday require a *lunch break* that must begin within a specified time window and last for a predetermined duration. These breaks may occur near client locations or at designated facilities, with any necessary movement incorporated into the overall travel time calculations.

3.2.5. Departure and arrival points

In practical HHC operations, it is unrealistic to assume that all caregivers begin their shifts at a central office location. Consequently, we allow caregivers to start their routes from either the central healthcare facility or their residences and subsequently reach their arrival point upon shift completion.

3.2.6. Patients / caregivers relationships

The relationship between caregivers and patients extends beyond simple service provision, including interpersonal dynamics that can impact care quality and patient satisfaction. These relationships are expressed in the form of incompatibilities and preferences.

Incompatibility: Specific patient–caregiver pairings may be unsuitable due to various personal or medical factors. For instance, a nurse with a cat allergy would be unable to provide care in a home with feline pets, or a patient might request caregivers of a specific gender for personal care services as a firm requirement. These restrictions are included through specific patient–caregiver combinations declared as *incompatible*, thereby preventing these assignments in the resulting schedule.

Preferences: Beyond strict incompatibilities, the quality of HHC is significantly influenced by patient preferences for particular caregivers. These preferences may stem from various factors, for example, communication style, and are defined as sets of preferred caregivers for a given patient.

3.2.7. Optional patients

Often, real-world cases involve situations where the available caregiver capacity cannot satisfy the total patient demand. Within this context, we distinguish between two patient categories based on service urgency. *Mandatory* patients must receive all scheduled services within the current planning horizon without exception. Conversely, *optional* patients may have their care postponed to subsequent planning periods when service capacity is insufficient. Partial service delivery is not possible: a patient is adequately deemed served when all their required services have been successfully scheduled.

3.3. Cost components and violations

In a general formulation of the HHCRSP, various operational aspects can be incorporated either as *hard* constraints that must be satisfied for a solution to be feasible or as *soft* constraints that may be violated with associated penalty costs in the objective function. The following cost components and potential constraint violations capture operational efficiency metrics and quality-of-service factors. The specific implementation of these elements as either hard constraints or weighted penalty terms depends on organizational priorities and operational requirements. Section 5 illustrates specific modeling choices and provides a flexible framework that accommodates various approaches to handling these constraints.

Travel time: It represents the duration caregivers need to move between patient locations. It includes all movement within a caregiver's route, including travel from the starting location to the first patient, between consecutive patients, and from the final patient back to the ending location. This component directly influences operational costs through fuel consumption, vehicle wear, and productive time utilization.

Waiting time: It occurs when a caregiver arrives at a patient's location before their time window begins. During this period, the caregiver must wait until the earliest permissible service start time, resulting in unproductive time that nonetheless counts toward their work hours. This component represents inefficiency in schedule coordination between caregiver availability and patient time windows.

Tardiness: It measures the delay in service delivery beyond the scheduled or promised time. This can occur when a caregiver arrives (in relaxed time window scenarios) or completes (in strict time window scenarios) a service after the end of the time window. Tardiness negatively impacts patient satisfaction and may compromise care quality for time-sensitive services.

Overtime: It represents work performed beyond a caregiver's shift duration. It occurs when a caregiver's return time to their ending location exceeds their shift end time. This component typically incurs premium labor costs and may contribute to caregiver fatigue and decreased job satisfaction.

Idle time: It represents the sum of all time periods within a caregiver's shift that are not used for patient service, travel, or required lunch break. It is calculated as the total shift duration minus the combined service, travel, and lunch time. Specifically, idle time includes three components: waiting time that occurs when a caregiver arrives early at a patient location; the time interval between shift start and departure to the first patient; and the period between returning to the ending location and the conclusion of the shift (only when this return occurs before the shift ends). This measure captures all nonproductive periods within a caregiver's scheduled shift and represents potential inefficiency in resource utilization. It also serves as an indicator of fairness when considering its distribution across caregivers. Significant disparities in idle time may lead to unequal work experiences, where some caregivers face fragmented or underutilized shifts, impacting job satisfaction and perceived equity.

Unscheduled optional patients: It quantifies the penalty for not providing service to optional patients within the current planning horizon. While these patients can have their care postponed, each postponement potentially impacts their satisfaction.

Table 3

Attribution of cost components: HHC stakeholders and classification by operational goals. The column “MILP” reports a reference to the cost components of the MILP model described in Section 4

Cost component	MILP	Stakeholder			Operational goals		
		Caregiver	Patient	Organization	Efficiency	Fairness	Well-being
Travel time	(CC1)			✓	✓		
Overtime	(CC4)	✓			✓		✓
Tardiness	(CC2), (CC3)		✓			✓	✓
Waiting time	(CC5)			✓	✓		
Idle time	(CC6)			✓	✓	✓	
Unscheduled optional patients	(CC10)		✓		✓		✓
Unmet preferences	(CC8)		✓				✓
Ignored incompatibilities		✓	✓				✓
Qualifications requirements			✓				✓
Missed lunch break	(CC9)	✓					✓
Workload balance	(CC7)	✓				✓	

Unmet preferences: It measures violations of patient preferences regarding their caregivers. When a patient is assigned a non-preferred caregiver despite having expressed specific preferences, this component captures the resulting patient dissatisfaction or potential negative impact on care effectiveness.

Ignored incompatibilities: It measures violations of defined incompatibility relationships between patients and caregivers. Similar to unmet preferences, these violations result in patient dissatisfaction, but typically with more severe consequences for the care relationship.

Qualifications requirements: It measures violations for assigning caregivers to services for which they lack the necessary qualifications. Such assignments compromise care quality, patient safety, and regulatory compliance.

Missed lunch break: It represents the penalty incurred when a caregiver’s schedule fails to accommodate the required lunch break within the designated time window. As this affects caregiver well-being and may violate labor regulations, the associated cost reflects regulatory compliance risks and potential negative impacts on caregiver performance and retention.

Workload balance: It is a fairness-oriented cost component that considers the deviation of individual caregiver’s working time with respect to the average working time among all caregivers. This component promotes a more equitable distribution of work.

Table 3 categorizes the various cost components according to the stakeholders primarily affected by them. Each component is attributed to one or more of three key stakeholder groups: that is, caregivers, patients, and the healthcare organization itself. In addition, each component is linked to one or more operational objectives, classified into three categories: efficiency, which relates to overall costs and resource usage; fairness, which concerns the equitable distribution of workload and service among individuals (both, caregivers and patients); and individual well-being, which reflects the quality of experience or outcomes for each person involved, particularly focusing on patient satisfaction and caregiver comfort. Specifically, patients’ and caregivers’ preferences are handled through multiple modeling choices, which are briefly discussed in Section 4.3.

In Section 7.3, we provide some managerial insights gained by ad hoc experiments with alternative weights for the different cost components, specifically tailored to shed light on the various trade-offs among different stakeholders and operational objectives.

4. A formal model for HHCRSP

Following the conceptual framework established in the previous section, we develop a rigorous mathematical formulation of the HHCRSP. A compact MILP model describing the problem is introduced in Section 4.2, after the systematic introduction of the necessary notation and variable definitions in Section 4.1. A final note on how patients' and caregivers' preferences are handled in our model is given in Section 4.3.

4.1. Notation

As described in Section 3, the problem is defined by three main components: a set of patients C' , a set of service types S , and a set of caregivers V . Each caregiver $v \in V$ can perform only a subset of the services in S . This is specified by the binary parameter γ_{vs} , which equals 1 if caregiver v is qualified to provide service $s \in S$, and 0 otherwise. In addition, each caregiver $v \in V$ is associated with a location σ_v , representing the point where the caregiver both starts and ends their working shift. The working shift of caregiver v is defined by the time interval $[\alpha_v, \beta_v]$. A set $V_{lb} \subseteq V$ of caregivers is assigned a time window for a lunch break.

We define the node set $C'_0 = C' \cup \{0\}$, which includes the locations of all patients and an additional node 0, used as a placeholder to represent the start and end location of the caregiver under consideration. For any two nodes $i, j \in C'_0$, the travel time between them is denoted by d_{ij} when both i and j correspond to patient locations (i.e., $i, j \neq 0$). If one of the indices is 0, the placeholder is replaced by the actual starting/ending location σ_v of the relevant caregiver.

The set $C'_{opt} \subseteq C'$ represents the optional patients. To account for lunch breaks, an additional set $C^{lb} = \{c_v^{lb} : v \in V_{lb}\}$ is included in C'_{opt} , where each element corresponds to an artificial patient associated with a caregiver. These artificial patients are defined so that their service time matches the duration of the lunch break, their time windows coincide with the allocated lunch break interval, and they can only be served by their respective caregiver. Furthermore, they are assumed to be located at distance zero from all patients and depot locations, and their service must strictly begin within the assigned time window, without any delay.

The set of patients C' is divided into those requiring a single service C_1 and those requiring two distinct services C_2 . For C_2 patients, two different caregivers must provide the services. For each patient $i \in C'$, let $S_i \subseteq S$ denote the set of services required by that patient. By assumption, each patient requires at least one and at most two services, that is, $1 \leq |S_i| \leq 2$. For technical convenience, we define $S_0 = S$. Furthermore, let $V_i \subseteq V$ be the set of caregivers qualified to provide at least one of the services in S_i , taking into account the patient–caregiver incompatibilities specified in the input.

For each patient $i \in C_2$ who requires two services, we impose both a minimum and a maximum separation time (δ_i^{min} and $\delta_i^{max} \geq \delta_i^{min}$) between the two services. Specifically, if the first service begins at time t , then the second must begin within the interval $[t + \delta_i^{min}, t + \delta_i^{max}]$. A precedence relation

is therefore established between the two services. In the special case where $\delta_i^{\min} = \delta_i^{\max} = 0$, the two services must start simultaneously. Moreover, it is assumed that the two services required by a double-service patient must be carried out by different caregivers.

Given a patient $i \in C'$ and a required service $s \in S_i$, the corresponding service duration is denoted by p'_{is} . Patients may also specify a set of preferred caregivers P_i ; if no preference is given, P_i is assumed to include all caregivers qualified to provide service to i . Each patient i is further associated with a nonempty set of alternative time windows $T_i = \{[e_i^1, l_i^1], \dots, [e_i^k, l_i^k]\}$, which indicate the intervals during which service must begin. If a caregiver arrives before the opening time e_i^k , they are required to wait until the time window starts. All services assigned to the same patient must take place within the same time window.

The problem addresses the spatial and temporal planning of caregivers' tours to ensure that patients receive the required treatments, while minimizing a multicomponent objective function. The specific cost components and their weights can be flexibly adjusted to different operational contexts, allowing the formulation to capture most variants found in the literature through suitable parameter choices.

The cost components we consider are the following:

- (CC1) Total time distance travelled by the caregivers.
- (CC2) Total tardiness (delay) in treatments for all patients.
- (CC3) Maximum tardiness in treatments for all patients.
- (CC4) Total overtime carried out by the caregivers.
- (CC5) Total waiting time at patients' locations spent by the caregivers.
- (CC6) Maximum idle time over all caregivers.
- (CC7) Total absolute deviation of working time of caregivers with respect to the average working time.
- (CC8) Total penalty paid due to patients visited by caregivers different from the preferred ones.
- (CC9) Total penalty accumulated by caregivers who are skipping their lunch breaks.
- (CC10) Total penalty accumulated by not visiting optional patients.

While our formulation provides flexibility in defining the objective function, certain aspects described in Section 3.3 are treated as hard constraints rather than penalty terms. Specifically, patients/caregiver incompatibilities and qualification requirements are mandatory restrictions that must be satisfied in any feasible solution. On the other hand, other aspects, such as caregiver late arrivals with respect to patient's time windows, overtime, and missed lunch breaks, are accounted for as penalties.

4.2. Model

The model presented is an extension of that originally proposed in Montemanni et al. (2025) and Montemanni (2025). Tables 4 and 5 summarize the updated notation.

By observing that in the model, we have no benefit associated with the concept of patients requiring two visits, a simplified model can be derived in duplicating the nodes of the set C_2 , those associated with patients requiring two services. Formally, we define $C = C' \cup C_3$, where C_3 has the same cardinality of C_2 and for each node $i \in C_2$ there is exactly one node $j \in C_3$, representing both services required by the same patient. Notice that, when $i \in C'_{opt}$, then the new node j is inserted in

Table 4
Sets of the MILP model.

Symbol	Reference set(s)	Description
C'		Original patients
V		Caregivers
$S = S_0$		Services
C'_0	$C' \cup \{0\}$	Locations of the original patients plus a placeholder location for the caregivers
V_{lb}	$V_{lb} \subseteq V$	Caregivers assigned with lunch break
C'^{lb}	V_{lb}	Artificial patients modeling the lunch breaks
C'_{opt}	$C'_{opt} \subseteq C'$	Optional patients considering the initial set of patients
C_1	$C_1 \subseteq C'$	Patients requiring one service
C_2	$C_2 \subseteq C'$	Patients requiring two services
C_3	$C_3 \subseteq C$	Set with the same cardinality as C_2 , used to divide the double-services of C_2
C	$C' \cup C_3$	All patients real and fictitious, including those with divided double-service and the artificial ones used to map the lunch breaks
C_0	$C \cup \{0\}$	Locations of the patients plus a placeholder location for the caregivers
C_{opt}	$C_{opt} \subseteq C$	Optional patients
V_i	$V_i \subseteq V, i \in C$	Caregivers that can provide services to patient i
T_i	$i \in C$	Time windows in which patient i is available
S_i	$S_i \subseteq S, i \in C$	Services required by patient i
P_i	$P_i \subseteq V, i \in C$	Caregivers preferred by patient i
C^v	$\{i \in C : v \in S_i\},$ $v \in V$	Patients that might be served by caregiver v
C^v_0	$C^v \cup \{0\}, v \in V$	Locations of the patients that might be served by caregiver v plus a placeholder location for the caregiver

Table 5
Parameters of the MILP model

Symbol	Reference set(s)	Description
p'_{is}	$i \in C', s \in S$	Duration of service s for patient i
p_i	$i \in C$	Duration of service for patient i
e^k_i	$i \in C, k \in \{1, \dots, T_i \}$	Start of time window k of patient i
l^k_i	$i \in C, k \in \{1, \dots, T_i \}$	End of time window k of patient i
δ_i^{min}	$i \in C$	Minimum time separation between the services of patient i (if double service)
δ_i^{max}	$i \in C$	Maximum time separation between the services of patient i (if double service)
σ_v	$v \in V$	Location of caregiver v
α_v	$v \in V$	Start time of the working shift of caregiver v
β_v	$v \in V$	End time of the working shift of caregiver v
γ_{vs}	$v \in V, s \in S$	1, if caregiver v can perform service s , 0 otherwise
d_{ij}	$i, j \in C_0$	Travel time between generic nodes i and j
λ_i		Weight of cost component (CC i)

$C_{opt} \subseteq C$ (C_{opt} is the update of C'_{opt} with the fictitious second services). A precedence $c_i \prec c_j$ is also imposed between the two nodes to connect them. Notice that if the two services associated with c_i and c_j have to be executed simultaneously, the precedence is fictitious and only indicates the relation between the nodes. As a consequence of this transformation, the processing time p_i can now be directly associated with each node $i \in C$. The sets defining the problem are redefined consequently

Table 6
Variables of the MILP model

Symbol	Domain	Description
x_{ij}^v	$\{0, 1\}$	1, if caregiver v visits patient j after patient i ; 0, otherwise
t_i	$[0, \infty]$	Start time of the visit at patient i
r_{ik}	$\{0, 1\}$	1, if patient i is visited during their time window k
z_i	$[0, \infty]$	Tardiness in the execution of the service for patient i
D_{max}	$[0, \infty]$	Maximum tardiness
v_i	$\{0, 1\}$	1, if optional patient i is skipped; 0, otherwise
w_i	$[0, \infty]$	Waiting time spent at patient i
a_v	$[0, \infty]$	Active working time of caregiver v
s_v	$[0, \infty]$	Deviation of the working time of caregiver v w.r.t. the average
y_v	$[0, \infty]$	Extra time worked by caregiver v
b_v	$[0, \infty]$	Transportation time of caregiver v from the patients before and after the lunch break
H_{max}	$[0, \infty]$	Maximum idle time

as $C_0 = C \cup \{0\}$, $C^v = \{i \in C : v \in S_i\}$ and $C_0^v = C^v \cup \{0\}$, where the definition of each S_i and V_i is updated to the new set of nodes C .

The problem can be formulated as a vehicle routing problem with soft time windows (Koskosidis et al., 1992) featuring a complex objective function and additional constraints regarding

- (i) the assignment of service providers to each node,
- (ii) strict temporal constraints about visit times for node pairs linked to the same customer (i.e., patient),
- (iii) management of idle and waiting periods, overtime, and tardiness considerations.

The weighting coefficients λ_i that determine the relative importance of different cost components are not fixed parameters but rather depend on the specific problem instances, providing the necessary flexibility to adapt the cost function to different scenarios and requirements.

The variables of the new model are defined as follows (see also Table 6):

- x_{ij}^v is a binary variable equal to 1 if caregiver $v \in V$ visits patient $j \in C^v$ immediately after patient $i \in C^v$, and 0 otherwise; in particular, $x_{ii}^v = 1$ denotes that patient i is not visited by caregiver v ;
- t_i denotes the starting time of the visit at patient $i \in C$ (possibly greater than the arrival time at the location of patient i);
- v_i is a binary variable taking value 1 if the optional patient $i \in C_{opt}$ is skipped, 0 otherwise;
- r_{ik} is a binary variable that takes value 1 if patient $i \in C$ is visited during time window $k \in \{1, \dots, |T_i|\}$;
- a_v denotes the active working time of caregiver $v \in V$, calculated as service time plus transportation time and lunch time, if applicable;
- s_v denotes the absolute deviation of the working time of caregiver $v \in V$ with respect to the average working time of all caregivers;
- z_i denotes the tardiness (delay) in the execution of treatment at patient $i \in C$ with respect to the end of the time window l_i of the patient;
- y_v denotes extra time worked by caregiver $v \in V$ beyond β_v , the theoretical end of their shift;

- b_v denotes the transportation time spent by a caregiver $v \in V_{lb}$ to move from the patient immediately before the lunch break to that immediately after. The variable takes the value 0 if the lunch break is skipped;
- w_i denotes the waiting time at patient i , with i not being the first patient visited by a caregiver;
- D_{max} quantifies the maximum tardiness experienced when visiting patients;
- H_{max} quantifies the maximum idle time experienced by caregivers.

The mathematical model is presented below. For clarity, we have organized the constraints into separate paragraphs, each containing the formal presentation and explanatory comments.

Objective function.

$$\begin{aligned} \min \lambda_1 & \left(\sum_{i \in C_0^v} \sum_{j \in C_0, j \neq i} d_{ij} \sum_{v \in V_i \cap V_j} x_{ij}^v + \sum_{v \in V_{lb}} b_v \right) + \lambda_2 \sum_{i \in C} z_i \\ & + \lambda_3 D_{max} + \lambda_4 \sum_{v \in V} y_v + \lambda_5 \sum_{i \in C} w_i + \lambda_6 H_{max} + \lambda_7 \sum_{v \in V} s_v \\ & + \lambda_8 \sum_{i \in C} \sum_{v \in V, v \notin P_i} (1 - x_{ii}^v) + \lambda_9 \sum_{v \in V_{lb}} x_{c^v l b c^v}^v + \lambda_{10} \sum_{i \in C_{opt}} v_i. \end{aligned} \quad (1)$$

The objective function (1) optimizes the problem by combining the 10 targets defined in Section 4.1, following the same order as in the presentation. As explained in Section 5, the weights λ_i (with $i = 1, \dots, 10$) are instance-dependent and written in the input file. The optimization model is subject to the constraints described in the following paragraphs.

Routing constraints.

$$\sum_{j \in C_0^v} x_{ji}^v = 1 \quad v \in V, i \in C_0^v, \quad (2)$$

$$\sum_{v \in V: i \in C^v} (1 - x_{ii}^v) \underbrace{+ v_i}_{\text{if } i \in C_{opt}} = 1 \quad i \in C, \quad (3)$$

$$\sum_{j \in C_0^v, j \neq i} x_{ji}^v = \sum_{j \in C_0^v, j \neq i} x_{ij}^v \quad v \in V, i \in C_0^v, \quad (4)$$

$$x_{00}^v = 1 \Rightarrow x_{ii}^v = 1 \quad v \in V, i \in C_0^v. \quad (5)$$

Equations (2) state that for each caregiver v and each location i , exactly one x variable must be active. Constraints (3) state that exactly one caregiver must visit each compulsory patient, while optional patients can be skipped. Constraints (4) are flow-conservation constraints to impose, together with the timing constraints defined later, feasibility for the tours of the caregivers. Constraints (5) state that if a caregiver does not operate (x_{00}^v), then no patient can be assigned to their route.

Time-related constraints.

$$r_{ik} = 1 \Rightarrow z_i \geq t_i - l_i^k \quad i \in C, k \in \{1, \dots, |T_i|\}, \quad (6)$$

$$x_{i0}^v = 1 \Rightarrow y_v \geq t_i + p_i + d_{i\sigma_v} - \beta_v \quad v \in V, i \in C_0^v, \quad (7)$$

$$x_{0j}^v = 1 \Rightarrow t_j \geq \alpha_v + d_{\sigma_v j} + w_j \underbrace{+b_v}_{\text{if } v \in V_{lb} \wedge j = c_v^{lb}} \quad v \in V, j \in C^v, \quad (8)$$

$$x_{ij}^v = 1 \Rightarrow t_j = t_i + p_i + d_{ij} + w_j \underbrace{+b_v}_{\text{if } v \in V_{lb} \wedge j = c_v^{lb}} \quad v \in V, i, j \in C^v, i \neq j, \quad (9)$$

$$\delta_i^{min} \leq t_j - t_i \leq \delta_i^{max} \quad i, j \in C : i < j. \quad (10)$$

Inequalities (6) define the tardiness at each patient i as the eventual excess over l_i^k of the starting time of service at patient i during time window k . Constraints (7) define the eventual overtime of each caregiver v as the arrival time back at the ending location—described as the starting time at the last patient i plus the operating time at i plus the travel time from i to σ_v —minus the theoretical end of the shift β_v . Notice that the constraint is activated only if $x_{i0}^v = 1$. Such a conditional expression can be easily implemented in modern MILP solvers through indicator constraints (IBM, 2025; Perron and Didier, 2025). This is also true for similar constraints in the remainder of the model. Equations (8) take care of timing for the first patient of each caregiver. In particular, the starting time at the first patient j of caregiver v is given by the start time α_v of the shift of v plus the eventual waiting time at the base w_v plus the travel time from the base to the first patient j . We assume that a caregiver waits at their starting location σ_v rather than traveling to the first patient and waiting there. As a result, any waiting time for the first patient is disregarded. The corresponding constraint is enforced only when $x_{0j}^v = 1$. Constraints (9) take care of timing in case of the move of a caregiver from one patient to another, or back to the base. The starting time at the patient j of caregiver v is given by the start time t_i at patient i plus the service time at patient i plus the travel time from patient i to patient j plus the eventual waiting time w_j at patient j . The constraint is activated only if $x_{ij}^v = 1$. Constraints (10) impose the hard time-windows associated with double visits to the same patient (here split into virtual patients i and j).

Double-service and lunch-break constraints.

$$x_{ii}^v + x_{jj}^v \geq 1 \quad v \in V, i, j \in C_0^v : i < j, \quad (11)$$

$$\left((x_{ic_v^{lb}}^v = 1) \wedge (x_{c_v^{lb}j}^v = 1) \right) \Rightarrow b_v = d_{ij} \quad v \in V_{lb}, i, j \in C_0^v. \quad (12)$$

Inequalities (11) force each caregiver to do at most one service¹ at the same patient in need of a double visit (here split into virtual patients i and j). Constraints (12) set the value of variable b_v , expressly inserted to collect the travel time around the possible lunch break according to the patients visited right before and right it. Namely, if the artificial patient c_v^{lb} is visited in between patients i and j by the caregiver v , then the variable b_v , expressly inserted to collect the travel time

¹We recall that x_{ii}^v is 1 when patient i is not visited by caregiver v .

between the lunch-time customers, is assigned to d_{ij} . The value of b_v will be used when calculating the active working time of caregiver v in constraints (13).

Working time constraints.

$$a_v = \sum_{i \in C^v} \left(p_i(1 - x_{ii}^v) + d_{\sigma_v, i} x_{0i}^v + d_{i\sigma_v} x_{i0}^v + \sum_{j \in C^v, j \neq i} d_{ij} x_{ij}^v \underbrace{+ b_v}_{\text{if } v \in V_{lb}} \right) \quad v \in V, \quad (13)$$

$$s_v \geq a_v - \frac{\sum_{k \in V} a_k}{|V|} \quad v \in V, \quad (14)$$

$$s_v \geq \frac{\sum_{k \in V} a_k}{|V|} - a_v \quad v \in V. \quad (15)$$

Constraints (13) define the total active working time of caregiver $v \in V$, calculated as the sum of service time at patients, the traveling times from the starting point and to the ending point, the travel times among customers, and the lunch break in case it is taken. Inequalities (14) and (15) calculate, for each caregiver, the absolute deviation of the associated working time with respect to the average one. Specifically, given a caregiver v , inequality (14) captures the excess of working time with respect to the average value among all the caregivers, while inequality (15) captures its deficiency.

Time window constraints.

$$\sum_{k \in \{1, \dots, |T_i|\}} r_{ik} \underbrace{+ v_i}_{\text{if } i \in C_{opt}} = 1 \quad i \in C, \quad (16)$$

$$r_{ik} = r_{jk} \quad i, j \in C : i < j, k \in \{1, \dots, |T_i|\}, \quad (17)$$

$$r_{ik} = 1 \Rightarrow t_i \geq e_i^k \quad i \in C, k \in \{1, \dots, |T_i|\}, \quad (18)$$

$$r_{ik} = 1 \Rightarrow t_i \leq l_i^k \quad i \in C, k \in \{1, \dots, |T_i|\}. \quad (19)$$

Equalities (16) impose that for each mandatory patient, exactly one time window has to be selected, or none for optional patients. Equalities (17) force the same time window for all the services requested by the same patient. Inequalities (18) and (19) constrain the visit time at each patient based on the time window selected to service them.

Maximum value definition constraints.

$$D_{max} \geq z_i \quad i \in C, \quad (20)$$

$$H_{max} \geq \beta_v + y_v - \alpha_v - \sum_{i \in C^v} \left(d_{\sigma_v, i} x_{0i}^v + (p_i + d_{i\sigma_v}) x_{i0}^v + \sum_{j \in C^v, j \neq i} (p_i + d_{ij}) x_{ij}^v \right) - b_v \quad v \in V. \quad (21)$$

Constraints (20) assign the maximum tardiness to patients, D_{max} . Analogously, constraints (21) are used to assign the maximum idle time among caregivers, H_{max} . The idle time for each caregiver $v \in V$ is defined as the length of the working shift of v , eventually extended by overtime y_v , minus the time spent traveling or operating by the caregiver.

4.3. A note on preferences

Given the importance of patient-centered care and staff well-being, our model incorporates features that bring the formulation closer to real-world HHC operations. Unlike Mankowska et al. (2014), which neglects patient–caregiver preferences, our model considers them as follows:

Incompatibilities : for each patient i , the set of compatible caregivers V_i (Table 4) allows us to define hard constraints that exclude infeasible assignments.

Preferences : for each patient i , the set of preferred caregivers P_i Table 4 allows us to define soft constraints whose violations are penalized in the objective function via cost component (CC8).

Consequently, preferences are represented on a three-level scale:

- (i) incompatibilities (hard exclusions),
- (ii) preferred assignments (soft constraints; violations penalized), and
- (iii) neutral assignments (neither preferred nor incompatible, incurring no reward or penalty).

This approach is more practical than fine-grained scoring schemes (e.g., Ait Haddadene et al., 2016).

The model also incorporates staff preferences, which were absent in prior works:

Home-based routes : Caregivers may start and end their routes at home rather than at the central office (Section 3.2.5), represented by σ_v in Table 5 and in the extended set of locations by the signpost $\{0\}$.

Working hours : Shift types specify preferred working hours, represented by α_v and β_v in Table 5 and accounted for in Constraints (8) and in cost component (CC4).

5. Representation of problem characteristics

In this section, we describe the motivations for a unified representation for both HHCRSP instances and solutions (Section 5.1). Then, we introduce a set of instance features to address their differences and examine their distribution in the instance space (Section 5.2). Next, we describe the process of adapting benchmark instances from the literature to ensure compatibility with our representation (Section 5.3). In addition to these existing benchmarks, we present a newly generated dataset tailored to explore specific aspects of the problem and the instance generator used to create it (Section 5.4). Finally, we describe the system that orchestrates the instances and the solutions (Section 5.5).

5.1. Instance and solution representation

Problem instances are essential for combinatorial optimization, modeling real-world scenarios, and evaluating solution methods (Queiroz et al., 2024). For the HHCRSP, the accessibility to well-defined and diverse instances faces several challenges: varying problem formulations with different objectives and constraints (Section 2), a lack of a standardized format for instances and solutions, and limited availability to the research community due to privacy concerns.

Despite growing research interest, existing HHCRSP datasets remain fragmented with heterogeneous formats, creating two significant limitations: difficulty accommodating evolving healthcare requirements and regulatory changes, and complications in implementing and comparing solution approaches across studies. Text-based formats also introduce fragility issues due to a lack of error-checking mechanisms.

We propose a unified json-based representation for HHCRSP to address these challenges. This format is widely used, human-readable, and accommodates new problem-specific features as formulations evolve. json files are easily parsed by most programming languages, ensuring compatibility with different solver implementations. While we provide a dedicated semantic validator (Section 5.5), users can independently verify syntactic compliance against a json schema (Pezoa et al., 2016) specification.

The format's full specification and the entities' description are reported in Appendix A. The json schema is provided along with the data and code accompanying this paper. We note here that the file format specification accommodates different formulations by allowing cost components to be designated as either hard or soft constraints, with their respective weighting coefficients λ_i adjustable according to the specific scenario. This information is explicitly encoded within the problem instance specification to ensure accurate calculation of both objective values and constraint validation when evaluating solutions.

5.2. Instance features

Algorithms should be tested on diverse sets of instances to reveal their strengths and weaknesses (Smith-Miles and Muñoz, 2023). The key rationale is that similar instances are unlikely to provide new insights or meaningful additional knowledge.

To analyze instance characteristics, we define a set of 272 direct features for HHCRSP instances, meaning they are either extracted directly from the instance or require minimal computation. Examples include the number of patients and caregivers, the proportion of patients requiring multiple services, caregiver shift durations, and the overlap between caregiver shifts and patient availability.

Some features are represented by a single value, while others require aggregation to capture their statistical distribution. For the latter, we compute the minimum (min), maximum (max), range (range = max – min), mean, quartiles (q_1 , q_2 , q_3), and standard deviation (std). For example, patient time window lengths require aggregation at the patient level. Note that the counter on the number of features (272) also includes these aggregators.

In this work, these features are used to evaluate the capabilities of our generator (Section 5.4). In future research, they could support instance space analysis (Smith-Miles and Muñoz, 2023), algorithm selection, and other analytical tasks.

5.3. Existing instances

Various studies of the HHCRSP have explored different constraints and objectives (Section 2), but confidentiality agreements frequently prevent public release of instances. Moreover, when instances are available, they often employ diverse formats typically tailored to particular problem formulations, restricting their extensibility to accommodate new elements (Section A.1).

Table 7 presents an overview of the publicly available datasets for the single-period HHCRSP. Each column corresponds to a dataset, while each row represents a specific characteristic. The listed features are examples of those introduced in Section 5.2. Additionally, we report the number of instances in each dataset, whether they incorporate real-time distances and population density (indicated with a ✓), the file format, and their availability. For Kummer (2021) and our results (This work), we report both the number of validation instances (160 and 24, respectively) and their total number (15,040 and 465, respectively).

The size of the instances ranges from as few as ten patients to several hundred, with varying numbers of caregivers and services depending on the formulation. In most cases, time windows are fixed at 120 minutes, with some exceptions (Hiermann et al., 2015; Liu et al., 2019; Varas et al., 2024; Ceschia et al., 2026). Not all datasets incorporate real-world characteristics, and only Varas et al. (2024) and Ceschia et al. (2026) include real-time distances and population density. The file format differs across datasets, except for those by Bazirha et al. (2023a, 2023b), which share a subset of authors and therefore adopt the same format. While Hiermann et al. (2015)'s dataset includes only four randomly generated instances (with realistic time distances based on Vienna, Austria), their original study also examined real-world instances. However, these real-world instances remain unavailable to the public due to confidentiality agreements.

Among the publicly available datasets, we selected and translated into our unified format those provided by Mankowska et al. (2014), Hiermann et al. (2015), Kummer (2021), Bazirha et al. (2023a, 2023b), and Ceschia et al. (2026). These datasets were chosen because they either align closely with the formulation by Mankowska et al. (2014), introduce additional real-world extensions to previous formulations, or include many instances. We exclude from our conversion the single instance provided by Varas et al. (2024), due to its limited scope, and the dataset by Liu et al. (2019), as it addresses a stochastic formulation of the problem, which falls outside the scope of our analysis. Note that the datasets by Mankowska et al. (2014), Kummer (2021), Bazirha et al. (2023b), and Ceschia et al. (2026) refer to the same formulation. Whereas Bazirha et al. (2023a) also report the possibility for patients to express more than one availability time window, and Hiermann et al. (2015) include the possibility of patients being optional, the incompatibilities between patients and caregivers, and multimodal means of transportation (since our search methods do not handle this, we omit this dataset in the experimental evaluation of Section 7). For specific details on their conversion to our json format, we refer the interested reader to Section A.4.

5.4. Generator and newly generated instances

While in the literature, many formulations exist that include a limited set of real-world aspects each, this is not the case for the instances: as reported in Section 5.3, there exist only a limited number of datasets that are available, which are based on real-world features (like time distances), and that

Table 7
Publicly available datasets for single-period HHCRSP

	Mankowska et al. (2014)	Hiermann et al. (2015)	Liu et al. (2019)	Kummer (2021)	Bazirha et al. (2023b)	Bazirha et al. (2023a)	Varas et al. (2024)	Ceschia et al. (2026)	This work
Instances [#]	70	4	30	160/15040	42	40	1	30	24/465
Patients [#]	10–300	419–424	30, 40, 50	10–400	10–52	10–200	141	44–378	10–700
Caregivers [#]	3–40	493–518	5, 7, 9	3–40	3–10	3–40	17	7–51	2–124
Services [#]	6	5	30, 40, 50	6	6	6	155	2–4	4–18
Double-service patients [%]	30	55–77 *		30	0–32	0–60**	11	3–48	0–100
Time windows [min]	120	25–385	60–120	120	120	120	120–300	60–180	20–480
Variable service durations		✓	✓	✓	✓	✓	✓	✓	✓
Real time distances		✓		✓			✓	✓	✓
Real population							✓	✓	✓
File format	txt	dm, conf	c++	txt	xlsx	xlsx	csv	json	json
Availability	<i>i</i>	<i>ii</i>	<i>iii</i>	<i>iv</i>	<i>v</i>	<i>vi</i>	<i>vii</i>	<i>viii</i>	<i>ix</i>

* Patients with 3 services are present (0–2%).

** Patients with 3 services are present (0–9%).

i https://prodlog.wiwi.uni-halle.de/forschung/research_data/hhcrsp/.

ii <https://www.ac.tuwien.ac.at/research/problem-instances/>.

iii <https://drive.google.com/file/d/1ppfHc0AZfb2HpJmVCEsxRpwEiLNE4f5b/view>.

iv <https://github.com/afkummer/hhcrsp-dataset-2021>.

v On request. Translated at vii.

vi On request. Translated at vii.

vii On appendix of the article and on request.

viii <https://github.com/iolab-uniud/hhcrsp>.

ix <https://github.com/iolab-uniud/uhhc>.

employ some of the realistic aspects reported in Section 3. For this reason, we need to have an instance generator that

- (i) is modeled with real-world features;
- (ii) can account for all the realistic aspects;
- (iii) is highly flexible and customizable.

The generator takes several inputs to define the characteristics of an instance. These inputs include *geographic information*, such as the target city and the radius within which patients are located; *service-related parameters*, including the type and the total number of services per each type; *patient characteristics*, such as the total number of patients, the proportion of single- and double-service patients (distinguishing between independent, synchronous, and sequential services), the share of patients with incompatible or preferred caregivers, optional patients, and those requiring multiple time windows; *caregiver parameters*, including shift durations, lunch break specifications, and the configuration of departure and arrival points (e.g., whether they are fixed, flexible, or different for start and end locations); *temporal settings*, such as the temporal horizon, the interpretation of time windows (whether they should be considered as strict or not), time window durations, service durations (with bias options for selection), the temporal gap between consecutive services, and the granularity of time windows. Additionally, the generator allows customization of *cost components*, enabling users to define their mix and weights. Only two assumptions are made:

- (i) A patient can require at most two services.
- (ii) A patient can express at most two preferred time-windows.

Nevertheless, the generator allows for a high degree of customization, enabling users to model different scenarios. In fact, certain aspects can be omitted if not required, for example, lunch break specifications or setting the proportion of single-service patients to zero. This flexibility ensures that instances can be tailored to different problem formulations, regulatory constraints, and practical considerations.

The generator operates in two sequential phases. First, it generates geographic information (patient and departing/arrival locations along with travel times), followed by other entities such as temporal distribution and patient characteristics. The generator leverages real-world city networks and population patterns for spatial distribution modeling. It employs GEOSTAT Census 2021 datasets (Eurostat, 2024) to represent population data in specific geographical areas (at 1 km² resolution), while OSRM (Luxen and Vetter, 2011; OSRM, 2025) calculates travel times using actual road networks and a car as a means of transportation. Geographic data can be downloaded from these reference sites when required, optimizing resource usage. All other aspects are randomly generated according to user-specified parameters provided through command line inputs and configuration files.

The enhanced generator builds upon the one proposed by Ceschia et al. (2026), with both versions integrating actual road travel times and population distribution data to create more realistic models. However, our generator enhances flexibility by exposing all adjustable parameters to the user, allowing for greater customization. Additionally, it provides a comprehensive representation of real-world aspects that were neglected before.

We generate 465 new instances, whose characteristics are reported in Table 7, alongside those of the existing datasets. For the experiments reported in Section 7, we randomly select 24 instances.

The new dataset is designed to include all the realistic characteristics described in Section 3. The weights of the cost components are assigned based on a predefined range (note that a user can generate instances with their weight setting and mix). We utilize a Hammersley point set (Hammersley and Handscomb, 1964) to systematically sample weight values from the following ranges. Travel time, total tardiness, total waiting time, and total extra time weights range from 1 to 10, with a step size of 1. The highest tardiness weight varies between 0 and 5, with a step of 1. The highest idle time weight is assigned within the broader range of 0–100, while caregiver preferences range from 0 to 20, with a step of 1. The weight for optional patients is selected from 100 to 200 in increments of 10, whereas the penalty for missed lunch breaks follows a range of 50–100, also increasing in steps of 10. This approach ensures a diverse set of weight configurations, balancing different aspects of the problem and stressing the algorithms differently.

With these new instances, we aim to establish the following key points:

- (i) Our generator can produce instances comparable to those in existing literature, particularly those exhibiting realistic characteristics (e.g., realistic travel times).
- (ii) It can produce instances that address gaps in current datasets.
- (iii) It supports additional features described in problem formulations for which no publicly accessible instances are currently available.

While point 5.4 is inherently achieved through the generator's configurability (e.g., enabling lunch breaks, optional patients, etc.), the first two require a more detailed investigation. To address these, we adopt the following methodology: We group the existing datasets based on the problem formulations they reflect and perform a principal component analysis (PCA) on each group. We then project our newly generated instances onto the resulting Principle Component Analysis (PCA) space to visually and quantitatively evaluate how they align with and differ from the existing datasets.

A PCA (Gewers et al., 2021) is a dimensionality reduction approach that converts high-dimensional datasets into lower dimensional representations while maximizing the retention of variance in the original data (Ivan Chang, 2025). It identifies the principal components—orthogonal directions along which the data varies the most—allowing for a more interpretable visualization and analysis of datasets. To perform the analysis, we use the `pca` v.2.0.9 Python package (Taskesen, 2024).

The examination of existing datasets concerning their underlying formulations is essential, as certain features are exclusive to particular formulations (e.g., the instances by Ceschia et al., 2026, lack optional patients. In contrast, those by Hiermann et al., 2015, include them). As a result, these features are assigned a default value of 0 in datasets that do not incorporate them.

Figure 1 displays two PCA analyses: the left plot shows instances from Kummer (2021) and Ceschia et al. (2026), whereas the right plot presents those from Hiermann et al. (2015). Literature instances are shown as dark scatter points, while our generated instances are visualized through density distributions across the space.

In the datasets by Kummer (2021) and Ceschia et al. (2026), four principal components account for 80% of the variance. These components are mainly linked to the size of the patients' time win-




Fig. 1. PCA of existing datasets, distinguishing by problem formulation.

dows and include the total length of all patient time windows (i.e., the sum of the duration of all the time windows), the minimum time window duration, the average number of services required per patient, and the median patient availability (calculated as the total duration of all time windows assigned to a patient). In contrast, the variability in the instances from Hiermann et al. (2015) is primarily explained by two principal components, mainly composed of the following features: the total number of time windows and the median duration of caregiver shifts.

In all cases, our instances span the feature space of the existing datasets. It is important to note that this approach projects our instances onto a simpler feature space not only because of the dimensionality reduction but also because the literature instances do not account for some aspects. Consequently, we exclude from the PCA some characteristics present in our generated instances that are absent in the existing datasets.

5.5. HHCRSP toolbox

We developed a Python package for handling HHCRSP instances and solutions, providing a comprehensive tool for instance generation, validation, and assessment. All components of the toolbox, along with the datasets and additional material, are available at <https://github.com/iolab-uniud/uhhc>, with the full archive preserved on Zenodo (Ceschia et al., 2025).

The package includes the generator for creating problem instances described in Section 5.4, a validator for verifying both instances and solutions in terms of format (Appendices A.1 and A.2) and validity (e.g., assessment of solutions in relation to the problem formulation at hand), and additional utilities to facilitate analysis and benchmarking, such as the calculation of the cost components, the extraction of the features (Section 5.2), the visualization and inspection of solution (Fig. A1), and the conversion of existing literature instances (Section 5.3). Each element in the instance and solution is defined as a separate class: for example, a class defines the patient, a class defines the caregiver, etc. All the components and methods can also be accessed through the command line interface, ensuring ease of use for different workflows. Figure 2 reports the general scheme of the toolbox and the different tasks a user can perform.

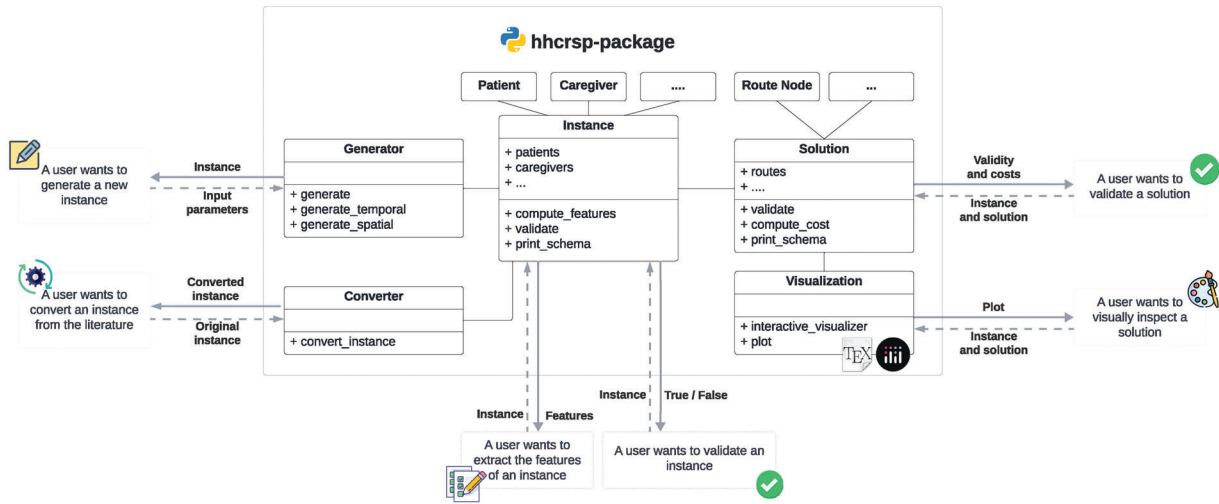


Fig. 2. General framework of the HHCRSP toolbox.

6. Solution methods

To accommodate the aspects introduced in Section 3.2, we have extended two existing solution methods, a MILP model (Montemanni et al., 2025) and a local search algorithm guided by a SA. In this section, we report the updates we have introduced and refer the interested reader to the original publications for further details.

6.1. Mixed integer linear programming

To solve the problem using the model introduced in Section 4.2, we formulated the model in the syntax required by a constraint satisfaction solver and processed it using a black-box solver engine. Building on prior research on similar problems (e.g., Montemanni and Dell’Amico, 2023; Montemanni et al., 2024) and a specialized variant of our problem (Montemanni et al., 2025), we chose the Google OR-Tools CP-SAT solver (Perron and Didier, 2025), specifically version 9.12, for our experiments.

6.2. Simulated annealing

The starting point of our SA method is the one proposed by Ceschia et al. (2026) for the solution of the basic formulation by Mankowska et al. (2014). The search space is composed of a global permutation of the patients (i.e., a sequence without repetitions) and the assignment of the caregivers in the appropriate number (one or two) to all patients. The actual routes and the schedule are determined by processing the patients one at a time according to the patients’ sequence, and adding them at the end of the respective routes, always at the earliest time. The method considers the

union of three neighborhoods, namely MovePatient, SwapPatients, and InRouteSwap, which move patients in the sequence, changing also their caregivers, swap the global position and caregivers of two patients, and swap positions of two patients within a single route, respectively.

We enhance the previous approach by including all additional features examined in this work. First, we modify the earliest time computation in the extended formulation to incorporate lunch break assignments and the option to shift to a patient's second time window (when available) if the first window cannot be met due to tardiness. Second, a significant modification to the search space addresses optional patients and their potential exclusion from scheduling. While maintaining the global patient permutation concept, we position unscheduled optional patients at the tail of the sequence, effectively excluding them from route construction. The three neighborhood relations (MovePatient, SwapPatients, and InRouteSwap) are applied exclusively to scheduled patients, that is, those positioned in the head portion of the sequence.

To deal with optional patients, we have added a fourth neighborhood called FlipPatient, which handles the insertion and removal of optional patients. A FlipPatient move targets one optional patient and reverses their status—scheduled patients are moved to the tail of the schedule, and caregivers are removed. In contrast, unscheduled patients are positioned at the head and assigned their requested number of caregivers. This novel neighborhood FlipPatient is added to the portfolio of neighborhoods from which we select the move at each stage, and it is assigned a given rate of being extracted.

The solver detects the possible absence of optional patients, and in that case the rate of FlipPatient is set to 0 throughout the entire search. Similarly, if certain features are not used, the corresponding cost component is assigned a weight of 0, and the search procedure does not compute its value. In addition, the auxiliary data structures necessary to speed up their computation are not updated, saving computational time.

In the extended formulation, as for caregivers without the necessary qualification, we exclude a priori incompatible caregivers from the pool of possible assignments. All other additional features contribute to the cost function that guides the local search toward lower-cost solutions.

Building on the work by Ceschia et al. (2026), this solution method has been written in C++ using the optimization framework EasyLocal++.² The code has been compiled using the GNU g++ compiler (v. 13.3) under Ubuntu Linux.

7. Experimental analysis

Experiments on the MILP model ran on an Intel Core i7 12700F machine; running times have been normalized to the machines used by Bazirha et al. Experiments on SA ran on an Ubuntu Linux 24.04 machine with 4 Intel i7-7700 (3.60 GHz) physical cores, hyper-threaded to eight virtual cores. A single virtual core has been dedicated to each experiment.

The SA method has multiple parameters for temperature control and neighborhood selection. We used irace (López-Ibáñez et al., 2016), a state-of-the-art tuning tool, to determine optimal parameter values. Details and results of this process are reported in Appendix B.

²<https://github.com/iolab-uniud/easylocal/releases/tag/v3.3>.

7.1. Comparative results

Comparisons on the datasets of the original formulation of Mankowska et al. (2014) are already reported by Ceschia et al. (2026) for SA and by Montemanni et al. (2025) for the MILP model, respectively. Therefore, we do not repeat them here.

In this section, we compare our solution methods on the formulations and datasets introduced by Bazirha et al. (2023a, 2023b). In Bazirha et al. (2023b), the authors allocated a time limit of four hours; whereas in Bazirha et al. (2023a), they reduced it to two hours on the same machine (Intel i7–7600U 2.80 GHz CPU and 16 GB of RAM).

Bazirha et al. (2023b) proposed a two-stage stochastic model that minimizes transportation cost in the first stage and the expected value of the recourse actions penalizing tardiness and overtime in the second stage. In their deterministic model, the authors focused solely on minimizing transportation costs. Thus, we adopted the same objective function for our comparison. They introduced two datasets: dataset \mathcal{SS} that includes only patients with a single service and dataset \mathcal{MS} with patients requiring multiple services. We excluded from our experimental analysis dataset \mathcal{SS} because MILP models could easily reach all optimal solutions (Bazirha et al., 2023b). Dataset \mathcal{MS} consists of three groups (\mathcal{D} , \mathcal{E} , and \mathcal{F}) with an increasing number of patients (10, 25, 50) and caregivers (3, 5, 10).

Table 8 presents comparative results of the different solution methods. Best values (optimal when available) are highlighted in bold.

We distinguish between the two different MILP models considered—the one from Bazirha et al. (2023b) and the new one described in Section 4—by the name of the solver used, CPLEX IBM (2025) and CP-SAT, respectively. For the MILP models, we report the lower bound (LB), the upper bound (UB), and the CPU time. For CP-SAT, the CPU time has been normalized to the machine used by Bazirha et al. (2023b)

In the case of the stochastic methods, each instance was executed 10 times. We report the average value (Avg), the best value (UB), and the average running time in seconds (t) over 10 executions. Our SA was granted 10M iterations.

Comparing the exact methods, it is evident that CP-SAT outperforms CPLEX in terms of both quality and computational time. In most cases (14 out of 21), CP-SAT was able to find the optimal solution in a shorter time. It also improved the lower bounds for four instances ($\mathcal{F}1$, $\mathcal{F}4$, $\mathcal{F}5$, $\mathcal{F}7$), marked with an asterisk, and certified the optimality for instance $\mathcal{F}3$. CP-SAT provided the best solution even for larger instances of group \mathcal{F} , although it required a long running time.

When focusing on the metaheuristic methods, it is important to note that SA algorithms, both the one by Bazirha et al. (2023a) and ours, clearly outperformed general variable neighborhood search (GVN) and GA by Bazirha et al. (2023b).

Comparing the two versions of SA, we can conclude that their performances are comparable. Specifically, the SA algorithm presented by Bazirha et al. (2023b) found 13 out of 21 best results, while our SA achieved 12 out of 21. The average value of our SA is 2.4% worse than the best solution, compared to a gap of 1.71% for the SA by Bazirha et al. (2023b). However, the average running time of Bazirha et al. (2023b) is 20.4 seconds, while for our SA is 13.8 seconds on comparable machines.

Bazirha et al. (2023a) employed CPLEX and an SA algorithm. Furthermore, the authors generated three datasets, namely *MSMTW*, *Small*, and *Large*. Datasets *MSMTW* have three groups

Table 8
Comparative results on formulation and MS dataset by Bazirha et al. (2023b)

		Bazirha et al. (2023a)						This work																							
		MILP (CPLEX)			GVNS			GA			SA			MILP (CP-SAT)			SA														
		LB	UB	t (seconds)	Avg	UB	LB	UB	t (seconds)	Avg	UB	LB	UB	t (seconds)	Avg	UB	LB	UB	t (seconds)	Avg	UB	LB	UB	t (seconds)	Avg	UB	LB	UB	t (seconds)	Avg	UB
D1	769.0	769	1.6	<1	769.0	769	0.1	769.0	769	0.1	769.0	769	3.1	769	769	0.1	769	769	0.1	769.0	769	769	769	0.1	769.0	769	769	8.6	769.0	769	8.6
D2	872.0	872	1.5	<1	872.0	872	0.1	872.0	872	0.1	872.0	872	2.5	872	872	0.1	872	872	0.1	872.0	872	872	872	0.1	872.0	872	872	8.5	872.0	872	8.5
D3	709.0	709	1.8	<1	709.0	709	0.2	709.8	709	0.2	709.0	709	3.6	709	709	0.2	709	709	0.2	709.0	709	709	709	0.2	709.0	709	709	8.6	709.0	709	8.6
D4	938.0	938	1.6	<1	938.0	938	0.1	938.0	938	0.1	938.0	938	2.3	938	938	0.2	938	938	0.2	938.0	938	938	938	0.2	938.0	938	938	8.6	938.0	938	8.6
D5	777.0	777	1.6	<1	777.0	777	0.2	777.0	777	0.2	777.0	777	3.3	777	777	0.1	777	777	0.1	777.0	777	777	777	0.1	777.0	777	777	8.6	777.0	777	8.6
D6	588.0	588	1.5	<1	588.0	588	0.1	594.2	588	0.1	588.0	588	2.8	588	588	0.1	588	588	0.1	588.0	588	588	588	0.1	588.0	588	588	8.6	588.0	588	8.6
D7	609.0	609	1.5	<1	609.0	609	0.1	609.0	609	0.1	609.0	609	2.9	609	609	0.1	609	609	0.1	609.0	609	609	609	0.1	609.0	609	609	8.6	609.0	609	8.6
E1	1317.0	1317	18.5	5.9	1336.1	1317	6.1	1382.7	1317	6.1	1355.0	1317	12.9	1317	1317	13.8	1317	1317	13.8	1355.0	1317	1317	1317	13.8	1330.9	1317	12.7	1330.9	1317	12.7	
E2	1361.0	1361	8.1	5.8	1391.7	1374	5.8	1407.8	1387	5.8	1385.0	1361	12.5	1361	1361	14.9	1361	1361	14.9	1385.0	1361	1361	1361	14.9	1480.6	1384	12.9	1480.6	1384	12.9	
E3	1338.0	1338	7.2	5.6	1357.5	1338	6.4	1419.1	1353	6.4	1367.5	1338	14.1	1338	1338	4.5	1338	1338	4.5	1367.5	1338	1338	1338	4.5	1362.3	1338	12.7	1362.3	1338	12.7	
E4	1150.0	1150	5.6	4.6	1159.4	1150	5.6	1207.2	1171	5.6	1151.4	1150	15.2	1150	1150	10.1	1150	1150	10.1	1151.4	1150	1150	1150	10.1	1150.3	1150	12.9	1150.3	1150	12.9	
E5	1246.0	1246	3.6	4.8	1268.6	1246	4.8	1286.9	1247	6.4	1250.0	1246	12.0	1246	1246	4.1	1246	1246	4.1	1250.0	1246	1246	1246	4.1	1254.2	1254	13.0	1254.2	1254	13.0	
E6	1251.0	1251	11.3	5.2	1259.3	1251	5.2	1286.7	1251	6.8	1268.6	1251	12.9	1251	1251	10.2	1251	1251	10.2	1268.6	1251	1251	1251	10.2	1253.6	1251	12.9	1253.6	1251	12.9	
E7	1145.0	1145	5.1	3.3	1147.6	1145	3.3	1155.1	1145	4.7	1145.3	1145	12.7	1145	1145	1.7	1145	1145	1.7	1145.3	1145	1145	1145	1.7	1159.1	1145	12.8	1159.1	1145	12.8	
F1	1754.0	1754	1971.0	276.6	1903.0	1808	276.6	2169.4	2069	113.8	1858.3	1780	46.8	1754	1754	760.6	1754	1754	760.6	1858.3	1780	1780	1780	760.6	1856.5	1796	20.1	1856.5	1796	20.1	
F2	1668.7	1864	14400.0	186.3	1933.6	1832	186.3	2166.2	1989	119.7	1876.4	1832	38.6	*1749	1828	14400.0	*1749	1828	14400.0	1876.4	1832	1832	1832	14400.0	1890.3	1841	19.9	1890.3	1841	19.9	
F3	1685.0	1731	14400.0	205.2	1822.7	1813	205.2	2033.5	1940	129.1	1759.2	1731	43.1	1674	1726	14400.0	1674	1726	14400.0	1759.2	1731	1731	1731	14400.0	1769.1	1734	19.9	1769.1	1734	19.9	
F4	1721.7	1962	14400.0	231.2	1955.3	1908	231.2	2193.5	2105	96.9	1937.1	1908	47.4	*1883	1883	14400.0	*1883	1883	14400.0	1937.1	1908	1908	1908	14400.0	1956.3	1930	20.0	1956.3	1930	20.0	
F5	1793.0	2011	14400.0	273.5	2179.0	2040	273.5	2502.5	2296	139.6	2037.1	2015	51.3	*1941	2009	14400.0	*1941	2009	14400.0	2037.1	2015	2015	2015	14400.0	2103.8	2044	19.8	2103.8	2044	19.8	
F6	1808.0	1808	2207.0	209.6	1853.1	1818	209.6	2189.0	2035	148.7	1840.7	1820	42.1	1808	1808	4738.0	1808	1808	4738.0	1840.7	1820	1820	1820	4738.0	1854.5	1835	20.1	1854.5	1835	20.1	
F7	1708.1	1730	14400.0	196.1	1783.1	1738	196.1	2098.8	2012	113.6	1772.5	1740	45.7	*1730	1730	2801.4	*1730	1730	2801.4	1772.5	1740	1740	1740	2801.4	1775.3	1748	20.1	1775.3	1748	20.1	
Avg	1248.0	1282.4	43630.9	6115.3	1314.9	1287.6	6115.3	1417.5	1360.9	443.1	1298.4	1280.8	820.4	1135.6	1276.6	63141.0	1135.6	1276.6	63141.0	1298.4	1280.8	1280.4	1280.4	63141.0	1307.6	1287.1	113.8	1307.6	1287.1	113.8	

(\mathcal{J} , \mathcal{K} , and \mathcal{L}) with 10, 25, and 50 patients, respectively. All patients in this dataset have two time windows. Dataset *Small* consists of the nine instances of group \mathcal{M} , each with 10 patients, several double-service patients ranging from 20% to 60%, and several time windows per patient varying from 1 to 3. Finally, dataset *Large* (group \mathcal{N}) includes only four instances that assign to each patient three time windows and up to three services. We excluded instances $\mathcal{N}100$ s and $\mathcal{N}200$ s because they involve patients requiring three services, and we focused only on instances $\mathcal{N}100$ p and $\mathcal{N}200$ p with 100 and 200 patients, respectively.

Table 9 shows the comparative results between our solution methods and those proposed by Bazirha et al. (2023a) in their formulation and datasets. In their paper, the workload balance deviation s_v is rounded up to $\lceil s_v \rceil$ for each single caregiver; we follow the same approximation for the comparison.

It is worth mentioning that our SA approach is not customized to this specific formulation and its constraints and features. For example, we do not exploit the fact that time windows are strict, and we allow the search not to meet them but penalize this case with a high artificial cost.

Observing the results, all methods find optimal solutions for small instances of groups \mathcal{J} and \mathcal{M} .

In instances of group \mathcal{K} (25 patients), CP-SAT could still identify numerous best solutions (four out of nine) within 2 hours of CPU time. In contrast, CPLEX could find only one best solution. Regarding the SA algorithms, they display similar performances on average.

For instances of group \mathcal{L} (with 50 patients), the exact methods struggle; in particular, CPLEX could not deliver any solution within two hours, while CP-SAT could find a solution for five instances, albeit not of good quality. Comparing the two metaheuristic methods on this set, our SA outperforms, on average, the one by Bazirha et al. (2023a), obtaining several best solutions (i.e., seven out of nine) in a shorter running time (always less than 23.5 seconds).

Finally, our SA, if granted more iterations (100M) to have comparable running times, obtained better results on the larger instances ($\mathcal{N}100$ p and $\mathcal{N}200$ p), outperforming the method of Bazirha et al. (2023a) by 25% in the average value of the objective function.

7.2. Results on the new dataset

We now consider our newly generated dataset. Out of the complete set (465 instances), we randomly select 24 instances for demonstration purposes.

CP-SAT was granted a maximum of 3600 seconds, whereas SA was granted 10M iterations (an average of 67.9 seconds) and configured as described in Appendix B.

Table 10 reports both the main features of the instances and the results we obtained.

Regarding the features, the first five columns report the number of patients (C'), caregivers (C), services ($\sum_{i \in C'} |S_i|$), the percentage of double-service patients ($|C_2|/|C'|$), and optional patients ($|C'_{opt}|/|C'|$). The dataset exhibits significant diversity, comprising instances from 10 to 450 patients. The proportion of patients requiring two services varies from 0 % to 100%, as in the case of the optional patients.

CP-SAT optimally solved instances up to 15 patients (i-116 and i-134) in a short CPU time. For instance, with 25 patients, it still achieves the best values for three out of four cases in one hour; results of SA in these instances are, on average, worse by 13%, but with a much shorter running time

Table 9
Comparative results on formulation and dataset by Bazirha et al. (2023a)

	Bazirha et al. (2023a)					This work					
	MILP (CPLEX)		SA			MILP (CP-SAT)			SA		
	UB	t (seconds)	Avg	UB	t (seconds)	LB	UB	t (seconds)	Avg	UB	t (seconds)
$\mathcal{J}1$	161.0	3.7	161.8	161.0	6.3	161.0	161.0	16.5	161.0	161.0	8.9
$\mathcal{J}2$	130.5	12.9	130.5	130.5	5.2	130.5	130.5	49.6	130.5	130.5	9.2
$\mathcal{J}3$	102.5	6.0	105.3	102.5	6.0	102.5	102.5	10.0	102.5	102.5	9.1
$\mathcal{J}4$	311.0	6.0	311.0	311.0	4.9	311.0	311.0	16.7	311.0	311.0	9.8
$\mathcal{J}5$	160.5	13.6	160.5	160.5	4.7	160.5	160.5	13.9	160.5	160.5	9.4
$\mathcal{J}6$	160.0	10.4	160.0	160.0	4.0	160.0	160.0	32.8	160.0	160.0	9.6
$\mathcal{J}7$	246.5	5.7	246.5	246.5	4.9	246.5	246.5	8.3	246.5	246.5	7.9
$\mathcal{J}8$	57.5	6.6	57.5	57.5	4.8	57.5	57.5	0.7	57.5	57.5	9.5
$\mathcal{J}9$	165.0	32.5	166.1	165.0	4.2	165.0	165.0	144.6	205.0	165.0	9.6
$\mathcal{K}1$	177.0	7200.0	143.4	116.0	25.6	38.0	99.0	7200.0	93.0	77.0	14.4
$\mathcal{K}2$	82.5	7200.0	136.1	110.0	25.5	57.5	64.5	7200.0	98.3	72.0	14.0
$\mathcal{K}3$	118.5	7200.0	154.0	103.5	24.7	52.5	131.5	7200.0	122.6	90.0	14.3
$\mathcal{K}4$	381.0	7200.0	478.0	385.0	20.5	84.0	398.0	7200.0	441.3	409.5	14.7
$\mathcal{K}5$	219.5	7200.0	282.0	216.5	20.9	63.0	191.5	7200.0	225.7	199.5	15.0
$\mathcal{K}6$	286.0	7200.0	303.6	244.0	18.5	82.5	291.0	7200.0	268.8	219.0	15.0
$\mathcal{K}7$	294.0	7200.0	270.0	165.5	18.7	42.5	288.0	7200.0	238.3	182.0	15.1
$\mathcal{K}8$	76.5	7200.0	150.6	94.5	20.1	39.0	61.5	7200.0	114.6	82.5	14.9
$\mathcal{K}9$	201.5	7200.0	158.4	121.5	18.9	40.5	98.0	7200.0	139.5	123.5	14.9
$\mathcal{L}1$	—	7200.0	200.6	152.0	89.0	0.0	335.5	7200.0	188.0	169.0	22.1
$\mathcal{L}2$	—	7200.0	163.2	125.5	80.1	—	—	7200.0	129.0	90.0	22.1
$\mathcal{L}3$	—	7200.0	186.7	147.5	86.9	—	—	7200.0	158.3	121.5	22.2
$\mathcal{L}4$	—	7200.0	434.4	403.0	81.4	—	—	7200.0	520.6	438.0	22.3
$\mathcal{L}5$	—	7200.0	334.1	264.0	85.5	0.0	440.5	7200.0	309.6	202.0	22.4
$\mathcal{L}6$	—	7200.0	574.4	504.0	82.4	—	—	7200.0	598.5	429.0	22.5
$\mathcal{L}7$	—	7200.0	280.8	218.0	87.2	0.0	291.0	7200.0	280.2	183.5	23.3
$\mathcal{L}8$	—	7200.0	216.6	187.0	85.1	0.0	281.0	7200.0	177.2	128.5	23.1
$\mathcal{L}9$	—	7200.0	258.9	183.5	79.5	0.5	217.5	7200.0	227.4	137.5	23.4
$\mathcal{M}1$	223.0	2.8	223.0	223.0	3.0	223.0	223.0	1.5	223.0	223.0	8.7
$\mathcal{M}2$	144.0	6.8	144.0	144.0	4.3	144.0	144.0	3.7	b	144.0	8.8
$\mathcal{M}3$	8.5	8.8	8.5	8.5	4.2	8.5	8.5	4.6	8.5	8.5	8.8
$\mathcal{M}4$	162.5	2.2	162.5	162.5	4.7	162.5	162.5	15.0	162.5	162.5	9.0
$\mathcal{M}5$	146.5	14.9	146.5	146.5	6.8	146.5	146.5	20.0	146.5	146.5	9.3
$\mathcal{M}6$	51.0	25.9	51.7	51.0	7.0	51.0	51.0	26.6	55.0	51.0	9.2
$\mathcal{M}7$	424.5	202.0	428.0	424.5	5.4	424.5	424.5	1504.2	424.8	424.5	10.6
$\mathcal{M}8$	51.0	436.9	70.0	51.0	6.6	51.0	51.0	208.7	51.2	51.0	10.6
$\mathcal{M}9$	29.5	32.0	33.8	29.5	6.6	29.5	29.5	1554.3	29.6	29.5	10.7
$\mathcal{N}100$ p	—	7200.0	99.3	81.0	269.8	—	—	7200.0	77.2	57	673.7
$\mathcal{N}200$ p	—	7200.0	242.8	215.5	1409.1	—	—	7200.0	197.6	138	673.7
Avg		3811.3	206.2	178.2	71.7			3885.0	194.3	165.4	40.1

of ≈ 26 seconds. CP-SAT can provide a feasible solution for up to 75 patients within the time limit; however, SA typically achieves an average improvement of about 13% in approximately 40 seconds.

Comparing the two methods is not possible for larger instances because CP-SAT could not provide any solution in one hour. We can observe that the running time of SA increases proportionally

Table 10
Results on the new dataset

	Features					MILP (CP-SAT)			SA		
	$ C' $	$ V $	$\sum_{i \in C'} S_i $	$ C_2 / C' $	$ C'_{opt} / C' $	LB	UB	t (seconds)	Avg	UB	t (seconds)
i-054	200	30	220	10%	10%	—	—	3600.0	69904.4	65188	70.7
i-077	100	17	120	20%	20%	—	—	3600.0	17114.5	16164	48.6
i-083	70	11	98	40%	0%	7349	15725	3600.9	15828.3	12831	40.3
i-100	25	5	35	40%	0%	14538	14744	3600.0	17916.6	15119	25.1
i-116	10	4	13	30%	60%	17117	17117	0.1	18303.5	17393	13.6
i-126	200	32	260	30%	20%	—	—	3600.0	49607.6	47475	86.9
i-134	15	3	17	13%	60%	15616	15616	6.4	15820.0	15820	16.0
i-164	200	27	240	20%	0%	—	—	3600.0	61539.7	58873	83.9
i-167	100	13	109	9%	100%	—	—	3600.0	16818.2	13805	40.0
i-175	450	63	539	20%	0%	—	—	3600.0	34314.7	34014	186.3
i-185	340	62	442	30%	20%	—	—	3600.0	26839.8	26591	137.6
i-219	100	18	130	30%	20%	—	—	3600.0	11535.0	11086	51.4
i-235	25	6	44	76%	48%	1913	4702	3600.0	7694.3	6021	21.9
i-247	25	5	32	28%	48%	18881	19492	3600.0	21383.0	21186	24.7
i-250	190	26	217	14%	30%	—	—	3600.0	29697.0	28639	85.4
i-263	250	42	324	30%	20%	—	—	3600.0	38049.8	36933	111.8
i-272	170	26	194	14%	50%	—	—	3600.0	27974.6	27672	74.2
i-316	25	5	50	100%	0%	9063	10300	3600.0	11510.3	10196	32.5
i-360	200	35	300	50%	20%	—	—	3600.0	49530.6	46205	89.7
i-369	75	13	97	29%	20%	16402	26346	3600.0	26947.4	25225	41.8
i-406	100	18	130	30%	50%	—	—	3600.0	26197.6	b	43.1
i-414	70	11	98	40%	0%	24660	37633	3600.0	35982.0	33457	38.4
i-416	400	46	400	0%	0%	—	—	3600.0	42150.7	41009	131.6
i-446	75	13	111	48%	20%	15965	28872	3600.0	21768.3	21463	49.3

to the number of patients and that the deviation between average and UB values is more relevant when there are many double-service patients (instances i-083, i-100, i-235, and i-316) and optional patients (instance i-167).

We observe from Tables 8–10 that the running time for both our methods increases significantly with the complexity of the formulation. For example, for SA, for instance, with 25 patients, it goes from 13, to 15 and to 26 seconds, approximately.

7.3. Managerial insights

In this section, we frame the analysis through the lens of value-based purchasing (VBP), a policy instrument increasingly used by Ministries and Departments of Health of different countries to realign incentives across patients, providers, and payers. Value-Based Purchasing (VBP) shifts reimbursements and incentives away from mere activity volume toward outcome- and experience-based metrics, thus rewarding hospitals and other healthcare organizations for delivering high-quality, patient-centered care, while containing cost growth (Khalil et al., 2025). For example, in 2022, the Centers for Medicare & Medicaid Services (USA) launched the expanded home health value-based

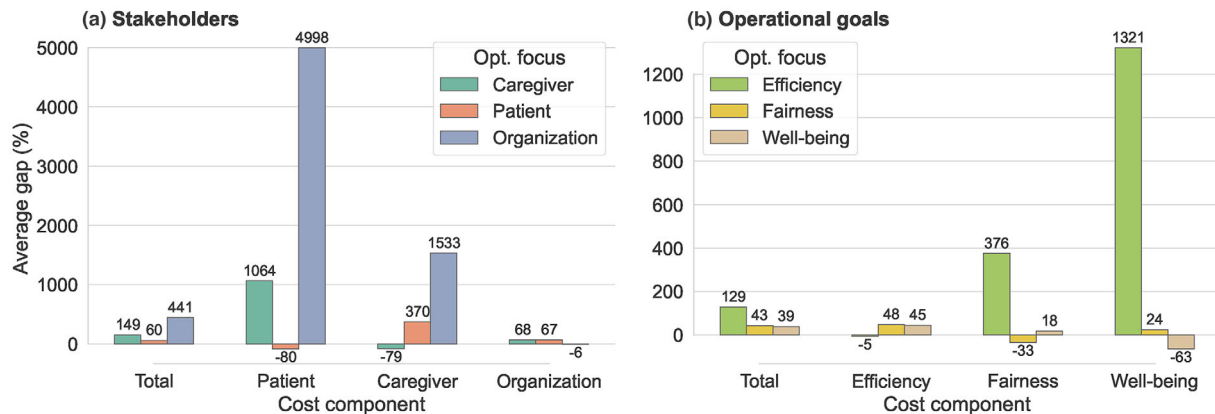


Fig. 3. Average gap (%) relative to the fully optimized solution when the optimization is carried out for different stakeholder groups (left) and operational goals (right)

purchasing model,³ whose primary goal is to “provide incentives for better quality care with greater efficiency.” Comparable initiatives are now emerging in Europe (Telos Management Consulting, 2016) with the adoption of VBP broadly in healthcare-related topics.⁴ These programs underscore the need to analyze the trade-offs among stakeholder objectives that our model addresses.

Therefore, we decided to consider the perspectives of the different stakeholders and investigate how individual cost components behave when the optimization is steered toward a single stakeholder or operational goal. More specifically, we adopted the objective-function breakdown presented in Table 3, which separates cost components by stakeholder (patient, caregiver, organization) and by operational goal (efficiency, fairness, well-being). For every instance, we reran our SA under the six scenarios, each time considering only one cost group (patient, caregiver, organization, efficiency, fairness, well-being) while keeping all other algorithmic parameters identical to the configuration described in Appendix B. We thus conducted $10 \text{ runs} \times 24 \text{ instances} \times 6 \text{ scenarios} = 1440$ experiments in total.

Figure 3 shows the results of these analyses through bar charts. Different colors represent results for various stakeholders (left) or operational goals (right). To compare, we calculated the average percentage gap as the difference between the cost from a single cost group and the total cost, divided by the total cost. A negative value indicates that focusing on a specific group leads to an improvement in the value of the corresponding cost group.

It could be noticed that, when we adopt a patient-centered perspective and ignore the soft constraints related to caregivers and organization, the total cost increases only moderately (60%) compared to other cases, which show increases of 149% for caregivers and 441% for organization. The cost components related to patients decrease by 80%, while costs for the organization and caregivers worsen by 67% and 370%, respectively. Focusing on caregivers, their related cost components decrease by 79%, whereas costs related to the organization increase by 68%, and those related to

³<https://www.cms.gov/priorities/innovation/innovation-models/expanded-home-health-value-based-purchasing-model> (accessed July 22, 2025).

⁴<https://www.sanidad.gob.es/gabinete/notasPrensa.do?id=6475> accessed July 22, 2025).

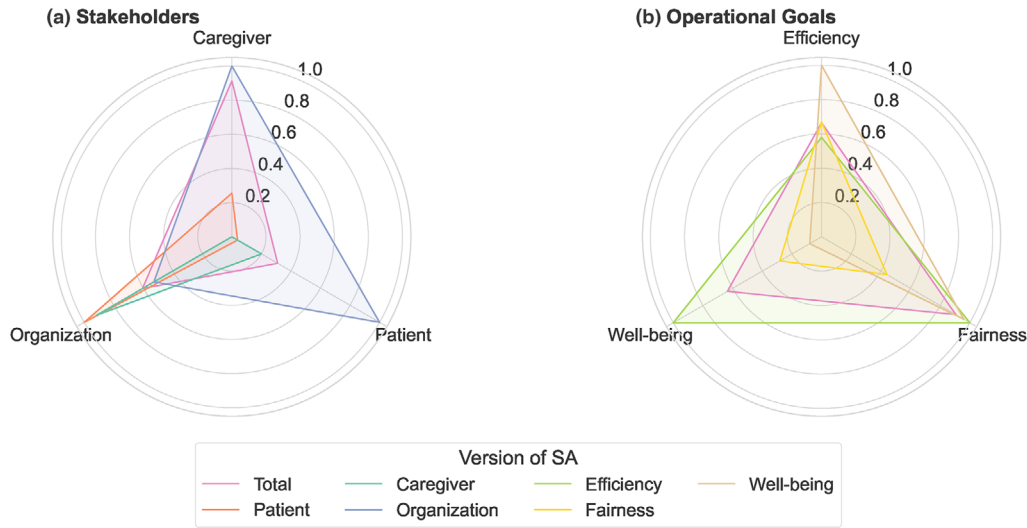


Fig. 4. Radar plot for instance i-083.

patients rise by 1064%. Unexpectedly, the worst results are observed from the organization’s perspective. Specifically, the reduction in organization costs is only 6%, while costs for patients and caregivers increase by 4998% and 1533%, respectively. This cost analysis suggests adopting a patient perspective, in line with recent trends that push healthcare institutions towards the creation of a patient-centered healthcare delivery system (Mohammed et al., 2016).

Figure 4 contrasts the optimization strategies for a representative instance, that is, i-083. Each axis reports the normalized cost for one stakeholder (left) or operational goal (right), where smaller values denote better performance. Considering the stakeholders, the organization-focused optimization attains the lowest organizational cost (≈ 0.5 , as one would expect), but does so at the expense of patients and caregivers (both ≈ 1.0), and eventually has the highest mean cost (≈ 0.8). Caregiver-focused optimization almost eliminates caregiver cost (≈ 0), keeps patient cost low (≈ 0.2), yet pushes the organizational cost up to ≈ 0.95 . These large spreads among objectives signal a strong imbalance. Patient-focused optimization excels in patient metrics (≈ 0.05) and performs well on caregiver metrics (≈ 0.25), but incurs the highest organizational costs (1.0). Joint optimization of the full cost components offers a more balanced compromise across all three criteria with a relatively low standard deviation (≈ 0.25), suggesting it provides a fairer solution that does not significantly sacrifice any single stakeholder’s priorities at the expense of others.

Regarding the operational goals, efficiency and well-being clearly trade off, while fairness holds the lowest average cost (≈ 0.4). As expected, the efficiency goal’s performance aligns with those obtained when focusing on the organization’s perspective. Additionally, it is confirmed that the complete SA achieves the most balanced solutions among all three criteria (with a standard deviation ≈ 0.14).

To isolate the influence of incompatibility and preference constraints on the total cost, we conducted an additional experiment in which these two constraint classes were removed while all other

hard and soft constraints remained active. Across the test set, on average, 41% of patients express preferences (range 0–75%) and 3% of patient–caregiver pairings exhibit incompatibilities (range 0–15%). Eliminating both incompatibilities and preferences yielded a mean improvement of less than 1%, rising to roughly 6% when the focus is restricted to patient–related cost components. This outcome is unsurprising: without incompatibilities and preferences, the solver can form routes with lower tardiness. Because such constraints reflect real-world clinical requirements and stakeholders' wishes, and because their omission delivers only a marginal cost saving, we believe their explicit inclusion in the formulation is well justified.

8. Conclusions and future work

We proposed a unified formulation for the HHCRSP that incorporates most of the features proposed in the literature into a single general framework. We supply all essential infrastructure to facilitate future research on our formulation in a measurable and reproducible manner. Specifically, we provide a new json-based file format, a dataset adequately distributed across the instance space and a toolbox for instance generation and conversions of previously published datasets into this new format. Moreover, the toolbox allows for semantic validation of instances and solutions.

We also adapted to this general problem two search methods that already proved to be at a state-of-the-art level for the basic formulation by Mankowska et al. (2014), namely a compact MILP model and an SA algorithm. We provided results obtained by these methods on both available and novel datasets. Both methods found results at the same level or better than the solution method specifically developed for the given formulation.

Some of the features proposed in the literature (Section 2) have been left out of our current framework. The main ones are as follows:

Multiday: When planning across multiple days, additional challenges emerge, including maintaining continuity of care and balancing caregiver workloads between days. Multiday planning can be approached as a series of sequential single-day problems, adjusting preferences and patient status (optional/mandatory) between consecutive days. However, solving the entire planning horizon simultaneously typically yields superior overall solutions.

Uncertainty/stochasticity: Since real-world healthcare operations have some inherently non-deterministic aspects, stochastic elements will better reflect their nature. For example, variability in service times, travel delays, and patient availability is inevitable in practice. Robust optimization models that accommodate unexpected events and dynamic rescheduling mechanisms would significantly enhance the practical applicability of our solution approaches.

Multi-modal: Currently, we assume that all caregivers use the same means of transportation (i.e., cars) and, consequently, that the traveling time between two locations is a single value. It is generally possible to use alternative means exhibiting different travel times, costs, and environmental impacts (e.g., employing public and private transportation). This extension would add a new set of decision variables representing the modeling choices for the transportation modality of each route.

Perishable biological samples and medicines: When collecting biological samples that could degrade, caregivers must return to the hospital or a clinical analysis lab within a short timeframe after service. Similarly, caregivers must promptly travel from the hospital/pharmacy to the patient when delivering perishable medicines, avoiding unnecessary intermediate stops (Varas et al., 2024).

In the future, we plan to incrementally add all the above features, trying to blend them smoothly within the current framework and without compromising the ability to solve the overall problem with a single general solution method.

Acknowledgments

We thank Mohammed Bazirha for sending us the datasets of his work and answering all our questions about it. This research has been partly funded by the European Union—*NextGenerationEU* under the projects PNR DM 737/21 *Modeling and solving real-world home healthcare routing and scheduling problem*(CUPG25F21003390007) and *FAIR MaPSART* (PE00000013, CUPC63C22000770006)—and by the Italian Ministry of University and Research (MUR)—PRIN-2020 program under the project *Models and algorithms for the optimization of integrated healthcare management* (no. 2020LNEZYC). This publication is of interest to COST Action (CA22137) *Randomised Optimisation Algorithms Research Network* (ROAR-NET), with respect to Working Group 6 (*Benchmarking*).

References

- Ait Haddadene, S.R., Labadie, N., Prodhon, C., 2016. A GRASP \times ILS for the vehicle routing problem with time windows, synchronization and precedence constraints. *Expert Systems with Applications* 66, 274–294.
- Bazirha, M., Benmansour, R., Kadrani, A., 2023a. An efficient two-phase heuristic for the home care routing and scheduling problem. *Computers & Industrial Engineering* 181, 109329.
- Bazirha, M., Kadrani, A., Benmansour, R., 2023b. Stochastic home health care routing and scheduling problem with multiple synchronized services. *Annals of Operations Research* 320, 2, 573–601.
- Begur, S.V., Miller, D.M., Weaver, J.R., 1997. An integrated spatial DSS for scheduling and routing home-health-care nurses. *Interfaces* 27, 4, 35–48.
- Bertels, S., Fahle, T., 2006. A hybrid setup for a hybrid scenario: combining heuristics for the home health care problem. *Computers and Operations Research* 33, 10, 2866–2890.
- Bredström, D., Rönnqvist, M., 2008. Combined vehicle routing and scheduling with temporal precedence and synchronization constraints. *European Journal of Operational Research* 191, 1, 19–31.
- Ceschia, S., Da Ros, F., Di Gaspero, L., Mancini, S., Maniezzo, V., Montemanni, R., Rosati, R.M., Schaerf, A., 2025. UHHC problem repository: unified home healthcare routing and scheduling dataset and toolbox. Zenodo. <https://doi.org/10.5281/zenodo.17713764>.
- 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, 100379.
- Ceschia, S., Di Gaspero, L., Rosati, R.M., Schaerf, A., 2026. Multi-neighborhood simulated annealing for the home healthcare routing and scheduling problem. *International Transactions in Operational Research* 33, 1, 38–67.
- Cheng, E., Rich, J., 1998. A home health care routing and scheduling problem. Technical Report CAAM TR98–04, Rice University.

- Cissé, M., Yalçındağ, S., Kergosien, Y., Şahin, E., Lenté, C., Matta, A., 2017. OR problems related to home health care: a review of relevant routing and scheduling problems. *Operations Research for Health Care* 13–14, 1–22.
- Clapper, Y., Berkhout, J., Bekker, R., Moeke, D., 2023. A model-based evolutionary algorithm for home health care scheduling. *Computers and Operations Research* 150, 106081.
- de Aguiar, A.R.P., Ramos, T.R.P., Gomes, M.I., 2023. Home care routing and scheduling problem with teams' synchronization. *Socio-Economic Planning Sciences* 86, 101503.
- Decerle, J., Grunder, O., El Hassani, A.H., Barakat, O., 2018. A memetic algorithm for a home health care routing and scheduling problem. *Operations Research for Health Care* 16, 59–71.
- Di Gaspero, L., Urli, T., 2014. A CP/LNS approach for multi-day homecare scheduling problems. In Blesa, M.J., Blum, C. and Voß, S. (eds), *Hybrid Metaheuristics*. Springer International Publishing, Berlin, pp. 1–15.
- Di Mascolo, M., Martinez, C., Espinouse, M.L., 2021. Routing and scheduling in home health care: a literature survey and bibliometric analysis. *Computers & Industrial Engineering* 158, 107255.
- Euchi, J., Masmoudi, M., Siarry, P., 2022. Home health care routing and scheduling problems: a literature review. *4OR* 20, 3, 351–389.
- Eurostat, 2024. Eurostat GISCO: grid data. <https://ec.europa.eu/eurostat/web/gisco/geodata/grids> (accessed: 4 August 2025).
- Fikar, C., Hirsch, P., 2017. Home health care routing and scheduling: a review. *Computers and Operations Research* 77, 86–95.
- Genet, N., Boerma, W., Kroneman, M., Hutchinson, A., Saltman, R.B., 2012. *Home Care Across Europe: Current Structure and Future Challenges*. World Health Organization, Regional Office for Europe, Copenhagen, Denmark.
- Gewers, F.L., Ferreira, G.R., Arruda, H.F.D., Silva, F.N., Comin, C.H., Amancio, D.R., Costa, L.D.F., 2021. Principal component analysis: a natural approach to data exploration. *ACM Computing Surveys* 54, 4.
- Grenouilleau, F., Legrain, A., Lahrichi, N., Rousseau, L.M., 2019. A set partitioning heuristic for the home health care routing and scheduling problem. *European Journal of Operational Research* 275, 1, 295–303.
- Grieco, L., Utley, M., Crowe, S., 2021. Operational research applied to decisions in home health care: a systematic literature review. *Journal of the Operational Research Society* 72, 9, 1960–1991.
- Hammersley, J.M., Handscomb, D.C., 1964. *Monte Carlo Methods*. Chapman and Hall, London.
- Hiermann, G., Prandtstetter, M., Rendl, A., Puchinger, J., Raidl, G.R., 2015. Metaheuristics for solving a multimodal home-healthcare scheduling problem. *Central European Journal of Operations Research* 23, 89–113.
- IBM, 2025. IBM CPLEX Optimizer. <https://www.ibm.com/de-de/analytics/cplex-optimizer>. IBM, New York.
- Ivan Chang, Y.-c., 2025. A survey: potential dimensionality reduction methods. arXiv. 2502.11036.
- Khalil, H., Ameen, M., Davies, C., Liu, C., 2025. Implementing value-based healthcare: a scoping review of key elements, outcomes, and challenges for sustainable healthcare systems. *Frontiers in Public Health* 13, 1514098.
- Kordi, G., Divsalar, A., Emami, S., 2023. Multi-objective home health care routing: a variable neighborhood search method. *Optimization Letters* 17, 2257–2298.
- Koskosidis, Y.A., Powell, W.B., Solomon, M.M., 1992. An optimization-based heuristic for vehicle routing and scheduling with soft time window constraints. *Transportation Science* 26, 2, 69–85.
- Kummer, A.F., 2021. A study on the home care routing and scheduling problem. Ph.D. thesis, Universidade Federal do Rio Grande do Sul.
- Kummer, A.F., de Araújo, O.C.B., Buriol, L.S., Resende, M.G.C., 2024. A biased random-key genetic algorithm for the home health care problem. *International Transactions in Operational Research* 31, 3, 1859–1889.
- Kummer, A.F., Buriol, L.S., de Araújo, O.C.B., 2020. A biased random key genetic algorithm applied to the VRPTW with skill requirements and synchronization constraints. In *Proceedings of the 2020 Genetic and Evolutionary Computation Conference*, Association for Computing Machinery, New York, pp. 717–724.
- Lasfargeas, S., Gagné, C., Sioud, A., 2019. Solving the home health care problem with temporal precedence and synchronization. In Talbi, E.G., Nakib, A. (eds) *Bioinspired Heuristics for Optimization*, Studies in Computational Intelligence, vol 774. Springer International Publishing, Cham, pp. 251–267.
- Liu, R., Yuan, B., Jiang, Z., 2017. Mathematical model and exact algorithm for the home care worker scheduling and routing problem with lunch break requirements. *International Journal of Production Research* 55, 2, 558–575.
- Liu, R., Yuan, B., Jiang, Z., 2019. A branch-and-price algorithm for the home-caregiver scheduling and routing problem with stochastic travel and service times. *Flexible Services and Manufacturing Journal* 31, 989–1011.

- Liu, W., Dridi, M., Fei, H., El Hassani, A.H., 2021. Hybrid metaheuristics for solving a home health care routing and scheduling problem with time windows, synchronized visits and lunch breaks. *Expert Systems with Applications* 183, 115307.
- López-Ibáñez, M., Dubois-Lacoste, J., Cáceres, L.P., Birattari, M., Stützle, T., 2016. The irace package: Iterated racing for automatic algorithm configuration. *Operations Research Perspectives* 3, 43–58.
- Luxen, D., Vetter, C., 2011. Real-time routing with OpenStreetMap data. In *Proceedings of the 19th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*. ACM, New York, pp. 513–516.
- Ma, X., Fu, Y., Gao, K., Sadollah, A., Wang, K., 2022. Integration routing and scheduling for multiple home health care centers using a multi-objective cooperation evolutionary algorithm with stochastic simulation. *Swarm and Evolutionary Computation* 75, 101175.
- Mankowska, D.S., Meisel, F., Bierwirth, C., 2014. The home health care routing and scheduling problem with interdependent services. *Health Care Management Science* 17, 1, 15–30.
- Masmoudi, M., Jarboui, B., Borchani, R., 2023. Efficient metaheuristics for the home (health)-care routing and scheduling problem with time windows and synchronized visits. *Optimization Letters* 17, 9, 2135–2167.
- Mohammed, K., Nolan, M.B., Rajjo, T., Shah, N.D., Prokop, L.J., Varkey, P., Murad, M.H., 2016. Creating a patient-centered health care delivery system: a systematic review of health care quality from the patient perspective. *American Journal of Medical Quality* 31, 1, 12–21.
- Montemanni, R., 2025. Solving a home healthcare routing and scheduling problem with real-world features. In *Proceedings of the 9th International Conference on Machine Learning and Soft Computing*. Springer International, Cham, pp. 111–119.
- Montemanni, R., Ceschia, S., Schaerf, A., 2025. A compact model for the home healthcare routing and scheduling problem. *EURO Journal on Computational Optimization* 13, 100101.
- Montemanni, R., Dell'Amico, M., 2023. Solving the parallel drone scheduling traveling salesman problem via constraint programming. *Algorithms* 16, 1, 40.
- Montemanni, R., Dell'Amico, M., Corsini, A., 2024. Parallel drone scheduling vehicle routing problems with collective drones. *Computers & Operations Research* 163, 106514.
- Oladzad-Abbasabady, N., Tavakkoli-Moghaddam, R., Mohammadi, M., Vahedi-Nouri, B., 2023. A bi-objective home care routing and scheduling problem considering patient preference and soft temporal dependency constraints. *Engineering Applications of Artificial Intelligence* 119, 105829.
- OSRM, 2025. Open Source Routing Machine (OSRM). <https://project-osrm.org> (accessed 4 August 2025).
- Perron, L., Didier, F., 2025. Google OR-Tools - CP-SAT. https://developers.google.com/optimization/cp/cp_solver/ (accessed 18 March 2025).
- Pezoa, F., Reutter, J.L., Suarez, F., Ugarte, M., Vrgoč, D., 2016. Foundations of JSON schema. In *Proceedings of the 25th International Conference on World Wide Web*, International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, Switzerland, pp. 263–273.
- Queiroz, M., Lucas, F., Sörensen, K., 2024. Instance generation tool for on-demand transportation problems. *European Journal of Operational Research* 317, 3, 696–717.
- Rasmussen, M.S., Justesen, T., Dohn, A., Larsen, J., 2012. The home care crew scheduling problem: Preference-based visit clustering and temporal dependencies. *European Journal of Operational Research* 219, 3, 598–610.
- Rendl, A., Prandtstetter, M., Hiermann, G., Puchinger, J., Raidl, G., 2012. Hybrid heuristics for multimodal homecare scheduling. In Beldiceanu, N., Jussien, N. and Pinson, É. (eds), *Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems*, Springer, Berlin, Heidelberg, pp. 339–355.
- Shahnejat-Bushehri, S., Tavakkoli-Moghaddam, R., Boronoos, M., Ghasemkhani, A., 2021. A robust home health care routing-scheduling problem with temporal dependencies under uncertainty. *Expert Systems with Applications* 182, 115209.
- Smith-Miles, K., Muñoz, M.A., 2023. Instance space analysis for algorithm testing: Methodology and software tools. *ACM Computing Survey* 55, 12.
- Taskesen, E., 2024. pca: A Python Package for Principal Component Analysis. <https://erdogant.github.io/pca>. (accessed: 4 August 2025).

- Telos Management Consulting, 2016. Value based healthcare, sfide e opportunità per il sistema sanitario italiano. White paper (Italian).
- Varas, M., Baesler, F., Basso, F., Contreras, J.P., Pezoa, R., Rojas-Goldsack, M.F., Ronco, R., 2024. A home hospitalization assignment and routing problem with multiple time windows, mandatory returns and perishable biological samples: a Chilean case study. *Computers & Industrial Engineering* 189, 109951.
- Vieira, B., Demirtas, D., van de Kamer, J.B., Hans, E.W., Jongste, W., van Harten, W., 2021. Radiotherapy treatment scheduling: Implementing operations research into clinical practice. *PLoS ONE* 16, 2, 1–13.
- Xiang, T., Li, Y., Szeto, W.Y., 2021. The daily routing and scheduling problem of home health care: based on costs and participants' preference satisfaction. *International Transactions in Operational Research* 30, 1, 39–69.
- Yadav, N., Tanksale, A., 2022. An integrated routing and scheduling problem for home healthcare delivery with limited person-to-person contact. *European Journal of Operational Research* 303, 3, 1100–1125.
- Zhang, T., Liu, Y., Yang, X., Chen, J., Huang, J., 2023. Home health care routing and scheduling in densely populated communities considering complex human behaviours. *Computers & Industrial Engineering* 182, 109332.

Appendix A: Data format

A.1. Instance representation

The core structure of an HHCRSP instance is shown in Listing 1. This format captures the main components of the problem (Section 3): a list of patients (`patients`), a list of caregivers (`caregivers`), and a list of services (`services`) that caregivers can provide and patients may require. Additionally, the instance includes information about the departure and arrival points for caregiver routes (`terminal_points`).

The instance also provides a distance matrix (`distances`) reporting the travel times in an array of arrays to enhance usability. Although this information is redundant—since patient and terminal point locations are specified using latitude and longitude—it ensures that all solvers operate with the same distances.

Finally, the `metadata` field contains supplementary information that, while not directly part of the instance, is essential for ensuring comparability, facilitating benchmarking, and preserving relevant details. For example, it records the weights of the cost components—either as numerical values for soft constraints or as the string `HARD` otherwise—and includes details about the instance generator when the instance is synthetically produced.

Listing 1: General instance structure.

```
{
  "patients": [...],
  "caregivers": [...],
  "services": [...],
  "terminal_points": [...],
  "distances": [...],
  "metadata": {...},
}
```

Each patient in the `patients` list is represented by a unique identifier and a set of attributes. The example in Listing 2 illustrates these attributes, including the list of required services

(`required_services`), where each item specifies the service and its estimated duration—this may differ from the default duration for that service type (Listing 5). Additionally, the patient's preferred time windows for receiving care are recorded under the `time_windows` field, where time is expressed in units (e.g., minutes) relative to the beginning of the planning horizon. The instance also specifies the patient's preferred and incompatible caregivers through the `preferred_caregivers` and `incompatible_caregivers` fields. The optional attribute indicates whether the patient is mandatory or not. Finally, the patient's geographical coordinates are provided along with the `distance_matrix_index`, which maps the location to an entry in the distance matrix (`distances`).

Listing 1: Patient specification.

```
{
  "id": "p2",
  "required_services": [ { "service": "s3", "duration": 30 } ],
  "time_windows": [ { "start": 60, "end": 300 } ],
  "location": [ 13.3011, 38.1034 ], "distance_matrix_index": 4,
  "incompatible_caregivers": [ "c2" ], "preferred_caregivers": [ "c1", "c4" ],
  "optional": false
}
```

Each caregiver in the `caregivers` list is represented by a unique identifier and a set of attributes. The example in Listing 3 illustrates these attributes. This identifier may also appear in the `preferred_caregivers` and `incompatible_caregivers` fields of the patients. The `abilities` field specifies the services the caregiver is qualified to perform. The departure and arrival points are identified by their respective indices, as defined in the corresponding fields in the related section (`terminal_points`, Listing 4 reports the structure). The `working_shift` field specifies the caregiver's working hours.

Listing 1: Caregiver specification.

```
{
  "id": "c1",
  "abilities": [ "s3" ],
  "departing_point": "d1", "arrival_point": "d1",
  "working_shift": { "start": 0, "end": 600 }
}
```

Listing 1: Terminal point (i.e., departing/arrival point) specification.

```
{ "id": "d0", "distance_matrix_index": 0, "location": [ 13.3468, 38.1109 ] }
```

Each service in the `services` list is represented by a unique identifier, its type, and a default duration (which may differ from the duration required by a specific patient's needs). Listing 5 reports an example. The distinction between the identifier and its type allows for grouping related services under broader categories. For example, under the type “medical care,” specific services such as physiotherapy or wound care can be included, while the type “daily activity” may include tasks like

bathing or house cleaning. The identifier is referenced in the `required_services` list of patients and the `abilities` field of caregivers.

Listing 1: Service specification.

```
{ "id": "s0", "type": "t0", "default_duration": 45 }
```

The generic instance can be customized by adding additional fields, such as specifications for lunch breaks (Listing 6). Similarly, the patient, caregiver, and service components can be adapted to include or exclude specific features depending on the problem formulation. For example, if the formulation accounts for lunch breaks, each caregiver will include a flag indicating whether they are entitled to such a break (`lunch_break`); conversely, if the formulation does not consider optional patients, the corresponding field can be omitted. This flexible structure allows the instance representation to accommodate many problem variants without altering the core format. By selectively including or excluding fields, the same schema can represent basic and extended formulations.

Listing 1: Lunch specification.

```
"lunch_breaks": { "start": 180, "end": 360, "min_duration": 30 }
```

A.2. Solution representation

A solution to an HHCRSP instance consists of the following decisions:

- (i) The route each caregiver follows (i.e., the order in which patients are visited).
- (ii) The arrival time, departure time, and service provided by each caregiver to each patient in their route.
- (iii) If lunch breaks are considered, each caregiver's lunch break location, start time, and end time.

A common solution representation ensures consistent evaluation, making it easier to compare different approaches. It also improves interpretability for practitioners, aiding in debugging, optimization model implementation, and facilitating reproducibility and benchmarking, as solutions can be stored, analyzed, and shared in a standardized manner. Therefore, we encode the solution by employing the `json` format as in the case of the instance (Appendix A.1). The core structure of an HHCRSP solution is shown in Listing 7. This format captures the routes (`routes`), along with optional fields such as cost components and overall cost (both in terms of total cost and the number of violations). While these additional fields are not part of the solution itself, they facilitate comparability, benchmarking, and debugging.

Listing 1: General solution structure.

```
{
  "routes": [...],
  "cost_components": {...},
  "cost": { "cost": 2706, "violations": 0 }
}
```

Each item in the `routes` list (Listing 8) represents a caregiver and includes their identifier (`caregiver_id`) along with the list of locations they visit (`locations`). Each location entry spec-




Fig. A1. Example of solution visualization.

ifies the service performed, the corresponding place (i.e., the patient), and the temporal details of the visit (i.e., the start and end time of the service are required). If a lunch break is scheduled, it is recorded as a service labeled lunch-break.

Listing 1: Route specification.

```
{
  "caregiver_id": "c2",
  "locations": [
    ...
    { "arrival_time": 430, "start_time": 430, "departure_time": 460,
      "patient": "p7", "service": "s1" },
    ...
  ]
}
```

A.3. Solution visualization

While the representation above enables easy benchmarking and explicitly reports all solution characteristics, inspecting solutions visually may also be beneficial. A solution to the HHCSP can be effectively visualized as a Gantt chart, as put forward by Ceschia et al. (2026). In this representation, each item on the y -axis corresponds to either a patient or a terminal point, while the x -axis represents the scheduling horizon. The activities of each caregiver are distinguished by different colors, including the lunch break (abbreviated as LB). Additionally, patient availability times are shown in gray to provide a clear visual reference. Figure A1 reports an example considering the instance and solution employed, also in the solution representation snippets and in Section A.1.

A.4. Existing instances conversion

Converting Mankowska et al. (2014) and Kummer (2021) instances proved easy due to their simple textual data format. This format uses basic placeholders (e.g., nbNodes for the number of nodes)

© 2025 The Author(s).

International Transactions in Operational Research published by John Wiley & Sons Ltd on behalf of International Federation of Operational Research Societies.

followed by either a single scalar value, space-separated values for vectors, or line-separated sequences for matrices.

As for Bazirha et al. (2023b) instances, translation was straightforward regarding problem formulations; however, the main challenge lay in the format itself. Their dataset comprises six groups (A through F) of seven instances each, with each group provided in a separate Excel file. Each instance occupied a distinct tab within its file (seven tabs per file), with properties arranged in layered tables. The Bazirha et al. (2023a) dataset presented similar format challenges. It contains four groups with nine instances each, plus four additional very large instances. It is worth noting that, in the instance translation, the objective weights were multiplied by 2 in order to work exclusively with integer numbers. Consequently, the final objective values must be divided by 2 for a proper comparison with the original results reported in Bazirha et al. (2023a) (this does not apply to Bazirha et al., 2023b, as it considers only one cost component). Translating Ceschia et al. (2026) instances proved straightforward since these were generated using an earlier version of our generator.

Translating Hiermann et al. (2015) instances required several assumptions to align with our specifications (Section 3). Each instance was split across three files: car time distance matrix, public transportation time distance matrix, and entity details (caregivers, patients, services). We standardized timing data by multiplying values by 5, as specified in the original paper's compendium. The four randomly generated instances had no mandatory services, so all were treated as optional. For patients requiring multiple independent services, we represented each service as a separate patient. We explicitly defined incompatibilities through strict relations in patients' descriptions, while marking compatible caregivers as preferred. This approach allows patients to be served by caregivers without all the required qualifications, with appropriate cost penalties. We excluded this dataset from our analysis since our solution methods currently support only one distance matrix (with multiple transportation modes planned for future work; Section 8).

Appendix B: SA parameter tuning

SA relies on a set of parameters that must be carefully tuned. Given that the datasets we analyze deal with different constraints and objectives, and that the instances present different characteristics, we perform three tuning procedures—one for each dataset, that is,

- (i) Bazirha et al. (2023b);
- (ii) Bazirha et al. (2023a);
- (iii) our newly generated dataset.

For the first two datasets, we used the three-modal neighborhood SA, which uses the original neighborhoods by Ceschia et al. (2026) (i.e., MovePatient, SwapPatients, and InRouteSwap). For our new dataset, we employed quadri-modal SA, which also adds the FlipPatient neighborhood to the SA algorithm. This neighborhood was not employed in solving other datasets since they did not consider optional patients.

Table B.1 describes the SA parameters. Note that the MovePatient probability bias is not included as a parameter since it is implicitly calculated as 1 minus the sum of the other probability biases.

The tuning procedures were conducted using irace (López-Ibáñez et al., 2016), with an experimental budget of 10,000 experiments. The precision was set at three significant figures. Based on preliminary experiments, the SA was granted in each case 10M evaluations (parameter “Max evaluations”). Table B.2 presents both the parameter search ranges explored during tuning and the final parameter values.

Table B.1
SA parameter descriptions

Parameter	Description
Start temperature	The initial temperature that controls the probability of accepting worse solutions at the beginning of the search.
Min temperature	The final temperature at which the search terminates.
Cooling rate	Determines how quickly the temperature decreases during the search.
Neighbors accepted ratio	The target ratio of accepted neighbor solutions to evaluated neighbors. Controls the adaptive cooling of the temperature.
Max evaluations	The maximum number of solution evaluations allowed during the search process. Acts as an additional termination criterion.
SwapPatients rate	The probability bias of selecting a swap operation (exchanging two patients) as the neighborhood move.
InRouteSwap rate	The probability bias of selecting an insert-remove operation (removing a patient from one route and inserting into another) as the neighborhood move.
FlipPatient rate	The probability bias of selecting a flip operation, (removing a scheduled optional patient or adding an unscheduled one) as the neighborhood move. Only applicable to our new dataset.

Table B.2
Parameter ranges and selected values for SA

Parameter	Search range	Final value		
	Bazirha et al. (2023a)	New dataset		
Start temperature	[10.0, 30.0]	20.803	90.1175	25.7795
Min temperature	[0.1, 1.0]	0.501	0.8035	0.5442
Cooling rate	[0.985, 0.995]	0.986	0.9932	0.9873
Neighbors accepted ratio	[0.07, 0.15]	0.122	0.0793	0.1181
SwapPatients rate	[0.0, 0.2]	0.101	0.1225	0.1677
InRouteSwap rate	[0.0, 0.2]	0.075	0.0858	0.0205
FlipPatient rate	[0.0, 0.1]	—	—	0.0915

© 2025 The Author(s).

International Transactions in Operational Research published by John Wiley & Sons Ltd on behalf of International Federation of Operational Research Societies.