Louvain School of Management
and Instituto Superior Técnico

A Metaheuristic Approach to Routing and Scheduling Hospital-at-Home Visits

Author:
Maria Ema Furtado Neves

Master's thesis with the view of getting the degrees:
Master in Industrial Engineering and Management, Operations and Logistics +
Master in Business Engineering, Supply Chain Management

Supervisors:
Prof. Daniel Rebelo dos Santos
Prof. Diana Jorge
Prof. Philippe Chevalier

Declaration of the Use of AI Tools

During the preparation of this master's dissertation, I used the following AI-based tools:

1. **ChatGPT / Codex (OpenAI)**- to support the literature review by summarising relevant information; To support the development and debugging of the Python code used to implement the optimisation model and the metaheuristic, and the implementation and refinement of the prototype;
2. **Claude (Anthropic)**- to support the analysis and interpretation of the computational results, the drafting and structuring of parts of the written content, and assistance with LaTeX formatting;
3. **DeepL**- to support the translation of text from Portuguese to English.

After using these tools, I diligently reviewed, verified and edited all resulting content, and I take full responsibility for the final content presented in this dissertation. By signing this declaration, I affirm that the content of this master's dissertation reflects my original work, augmented by the responsible use of AI.

Assinado por: **Maria Ema Furtado Neves**
Num. de Identificação: 14869401
Data: 2026.08.12 14:21:36+01'00'

Acknowledgments

As I reach the end of this chapter and of a journey that has taken me through five years, two universities and two countries, I find myself reflecting not only on what I have learned, but also on the people who made this journey possible.

First and foremost, I would like to express my deepest gratitude to my supervisors, Professor Daniel Rebelo dos Santos and Professor Diana Jorge. Thank you for your guidance, availability and support throughout every stage of this work. I am especially grateful for the trust you placed in me, allowing me the freedom to make my own decisions, explore different ideas and learn through both successes and mistakes. Thank you for your patience, your encouragement and for the knowledge you shared. I would also like to thank Professor Philippe Chevalier for accepting to be the reader of this dissertation and for his support throughout my time at Louvain School of Management.

To my family, thank you for always being my safe place. To my sisters, thank you for your constant support, your encouragement and for celebrating every small achievement as if it were your own.

To my mother, thank you for allowing me to study, to live abroad and to complete this journey. Everything I have achieved is built upon the opportunities you created for me through your strength, your love and your unwavering support. I hope this achievement makes you as proud as I am grateful.

To Bernardo, thank you for standing beside me through every high and every low. Thank you for your patience, for listening, celebrating every milestone and for believing in me even during the moments when I struggled to believe in myself.

To my friends, thank you for making these years unforgettable. A special thank you to my double-degree partner, for sharing every challenge, every deadline and every achievement.

Finally, to my grandmother and to Luna. Although you are no longer here, the love and support you gave me throughout my life stayed with me during every step of this journey. I carry your memory with me always and I dedicate this achievement to you both.

Abstract

The rapid ageing of the population and the growing pressure on hospital capacity have accelerated the adoption of Hospital at Home (HaH) as an alternative to conventional inpatient care. Yet, the daily organisation of teams and visiting routes is still often performed manually. Nursing and doctor teams must visit geographically dispersed patients while satisfying clinical requirements, priority rules and route-duration limits.

This dissertation addresses a multi-depot HaH setting decomposed into depot-level routing and scheduling subproblems under operational constraints. Two complementary solution approaches are developed: a Mixed Integer Linear Programming (MILP) model, used as an exact benchmark, and an Adaptive Large Neighbourhood Search (ALNS) metaheuristic, designed to obtain high-quality solutions within short computational times. The objective combines total travel time with the use of doctor teams, while feasibility is enforced through team-patient compatibility, maximum shift duration constraints and prioritisation rules for urgent patients.

The computational study considers different demand scenarios and instance sizes, allowing the metaheuristic to be compared with the exact method across a broad set of instances. Results show that the ALNS obtains solutions very close to those produced by the exact method, while requiring substantially less computational effort. This benefit becomes more pronounced in larger instances, where the exact method frequently cannot prove optimality within the available time and the ALNS matches or improves upon the incumbent solutions obtained by the exact solver.

Finally, a decision-support prototype is developed to translate the resulting solutions into route visualisations, illustrating the practical applicability of the proposed approach for the daily planning of HaH visits.

Keywords: Hospital at Home; Routing; Scheduling; Adaptive Large Neighbourhood Search.

Contents

1 Introduction

The world's population is ageing rapidly. The World Health Organization (WHO) estimates that the number of people aged 60 and older will nearly double between 2015 and 2050 (World Health Organization, 2025), with Portugal among the two most ageing countries in the European Union (INSA, 2025). The consequences of this demographic shift, compounded by the rising prevalence of chronic diseases and the shortage of professionals and hospital beds, are placing increasing pressure on global health systems. Advances in telemedicine, remote monitoring and digital health have created favourable conditions for shifting part of acute care from hospital wards to patients' homes. Among the emerging solutions, Hospital at Home (HaH) stands out, offering hospital treatment in the patient's home, with continuous medical monitoring and multidisciplinary teams (Edgar et al., 2024; Leong et al., 2021). The aim is not to replace the hospital, but to complement it in a smart, safe and evidence-based way (Mendes, 2026).

Scientific evidence shows that HaH provides clinical results equivalent to or superior to traditional hospitalisation, with direct cost reductions, a lower incidence of hospital infections and greater satisfaction among patients and families (Leff et al., 2005; Saenger et al., 2022). HaH is divided into Early Supported Discharge (ESD), which allows early discharge with continued treatment at home, and Admission Avoidance (AA), which replaces conventional hospitalisation with treatment at home (Leong et al., 2021). However, these advantages depend on the ability to plan and coordinate home visits efficiently, which introduces a set of complex logistical constraints rarely addressed in existing research.

In Portugal, the concept of HaH was institutionalised in 2018 through Order No. 9323-A/2018, which defined the Implementation Strategy for HaH Units in the National Health Service (*Serviço Nacional de Saúde*- SNS) (Inácio, 2024). Since then, the number of units and their capacity have grown significantly, with all the 39 Local Health Units (*Unidade Local de Saúde*- ULS) in mainland Portugal offering this service.

Although the Portuguese HaH network has expanded rapidly, the operational dimension of HaH, particularly the daily organisation of teams and routes, is not sufficiently studied. Daily planning requires deciding which medical team to assign each patient to, managing resources and defining visit routes that comply with various restrictions (Varas et al., 2024). Every morning, each hospital spends around 30 to 45 minutes allocating staff and organising their visit route with no standardised optimisation approach to support these decisions.

This dissertation addresses the lack of decision-support tools for efficient, clinically compliant HaH routing. Existing Home Health Care (HHC) routing research cannot be applied directly to HaH because HaH imposes stricter clinical requirements. The objective is to develop and evaluate a metaheuristic for daily HaH routing in a real-world setting. A prototype decision-support tool is also developed to convert clinical, location and urgency inputs into route visualisations and performance metrics.

The dissertation is structured as follows. Chapter 2 presents the context, while Chapter 3 reviews the literature and identifies the research gap. Chapter 4 defines the problem; Chapter 5 presents the Mixed Integer Linear Programming (MILP) and Adaptive Large Neighbourhood Search (ALNS) approaches; Chapter 6 describes the data, instances and calibration; Chapter 7 reports the results; Chapter 8 presents the prototype; and Chapter 9 concludes the dissertation.

2 Context

2.1 Health System Challenges

The growing imbalance between the increasing demand for healthcare services and the limited hospital capacity has become a pressing global concern. With population ageing accelerating, the WHO estimates that the number of people aged 60 years and older will almost double between 2015 and 2050. Advances in medical technology and healthcare delivery have extended life expectancy, but they have also led to a rise in chronic conditions requiring long-term management. Consequently, the demand for healthcare resources and expenditures continues to escalate. If hospital systems become overwhelmed, this may result in declining patient outcomes and rising healthcare costs (Larsen et al., 2025; World Health Organization, 2025).

Beyond these demographic and clinical pressures, financial projections highlight the growing unsustainability of the traditional hospital-centric model, with health systems facing severe budgetary constraints and rising operational costs (Hosseini et al., 2026). The sector also faces a critical workforce shortage affecting formal and informal care, including declining family-provided care driven by sociological and demographic shifts (Bassoli et al., 2025). Together, these trends reinforce the need for alternative models of healthcare delivery that can alleviate pressure on hospitals while maintaining high-quality clinical outcomes.

2.2 Hospital at Home

In response to these pressures, HaH has emerged as a promising alternative to conventional inpatient care. HaH provides hospital-level care at home under hospital responsibility, supported by safety and monitoring protocols (Cheng et al., 2009; Leff et al., 2005), and it is also known as home-based hospital care, home-based acute care or home hospitalisation (Bonomi et al., 2025). The HaH model was first piloted in the United Kingdom in the late 1970s and subsequently implemented in the United States during the 1990s (Patel & West, 2021). Two major types of HaH programmes are identified: AA, which replaces traditional inpatient hospitalisation for acute episodes such as respiratory or cardiac exacerbations, and ESD, which allows patients to continue recovery at home following a partial hospital stay (Correia Azevedo et al., 2024). By substituting or shortening inpatient stays, these programmes aim to deliver equivalent outcomes while reducing pressure on hospital infrastructure.

In analyses that include the episode + 30 days post-acute, HaH has shown lower total costs (Saenger et al., 2022). In evaluations of large-scale programmes, lower mortality in several diagnostic groups and lower average expenditure are reported, with heterogeneity between hospitals and programme dependence (Levine et al., 2020). A meta-analysis of randomised trials showed even lower mortality, fewer readmissions and lower costs in HaH, with higher satisfaction compared to hospital care (Caplan et al., 2012).

According to Shepperd et al. (2016), the AA model likely yields little to no difference in mortality or hospital readmission rates when compared to acute inpatient care. However, for a select group of elderly patients, HaH may increase the likelihood of living at home (reducing institutionalisation) and result in greater patient satisfaction. A review of reviews shows that AA/ESD achieve outcomes equal to or better than hospitalisation; AA tends to generate greater efficiency (Leong et al., 2021). In addition to cost savings, international sectoral analyses report gains in capacity (beds freed up), reduction in nosocomial infections and better use of resources (De Sousa Vale et al., 2020; Deloitte, 2022).

At the same time, these benefits rely on complex organisational and logistical processes. HaH requires assigning patients to suitable teams, synchronising professionals, respecting clinical windows and managing returns (Bonomi et al., 2025; Varas et al., 2024). Chua et al. (2022) also show that HaH candidate selection remains challenging, especially when social support, housing, caregiver availability and technological readiness are considered.

2.3 Hospital at Home VS Home Health Care

Table 1 summarises the main differences between HaH and HHC.

For this dissertation, the key distinction is that HaH replaces acute inpatient care under hospital responsibility, supported by continuous monitoring, telemedicine and advanced medical equipment (Cheng et al., 2009; Levine et al., 2020). By contrast, HHC mainly provides post-acute, rehabilitative or long-term support, usually under community nursing or home-care organisations (Fikar & Hirsch, 2017; Shahnejat-Bushehri et al., 2021). Recent studies have highlighted that HHC involves multi-period decision making, where patient demands evolve over several consecutive days, requiring dynamic reallocation of caregivers and adaptive route planning (Zhou et al., 2026).

Table 1: Comparison between Home Health Care and Hospital at Home.

Dimension	Home Health Care	Hospital at Home
Purpose	Long-term or post-acute support aimed at recovery and chronic disease management.	Substitute for inpatient hospitalisation, providing hospital-level care at home.
Duration	Continuous or periodic, depending on the patient's condition and rehabilitation needs.	Short-term, limited to the acute treatment or stabilisation phase.
Provider type	Licensed community nurses or caregivers, under home-care organisations.	Hospital doctors and nurses, often the same staff as inpatient wards.
Equipment	Minimal; mostly basic medical/therapeutic tools.	Advanced hospital-grade equipment (monitoring, diagnostics, intravenous therapy).
Supervision	Community-/primary-care-based with periodic visits.	Hospital-based with continuous clinical oversight and telemonitoring.
Proximity to hospital	No strict geographic limitation.	Within hospital coverage area.

2.4 Hospital at Home in Portugal

Portugal has one of the highest ageing indices in Europe, rising from 26.6% in 1960 to 192.4% in 2025 and increasing pressure on healthcare systems (PORDATA, 2025).

In Portugal, Hospital Garcia de Orta, EPE (HGO) pioneered the creation and implementation of the first HaH programme, inspired by the experience of Infanta D. Leonor Hospital in Spain, created in 2015 (Helena, 2021). The success of this pilot programme, combined with the consolidated experience of international models, led to the creation of a national regulatory framework for the expansion of HaH. Thus, Order No. 9323-A/2018 was published in the Official Gazette, defining the strategy for implementing HaH in the SNS. Subsequently, the Directorate-General for Health (DGS), with the support of the pioneering hospital team, issued Standard No. 020/2018, dedicated to HaH in Adulthood (Helena, 2021). Afterwards, in April 2025, the Ministry of Health published an order formalising the extension of this system to all SNS hospitals (DN/Lusa, 2025).

In 2024, 108,923 days of hospitalisation were avoided, representing an 11% increase compared to the previous year. By August 2025, there were 415 home hospitalisation beds nationwide (SNS, 2025) and more than 59,000 people had already been hospitalised at home, involving 721 healthcare workers (Alves, 2025).

Beyond its scale, HaH is recognised as an innovative, patient-centred model of care, ensuring high-quality clinical care in the comfort of the patient's own home (AESE, 2026). Nowadays, HaH in Portugal is considered a medium-sized hospital, with its potential for expansion limited only by the human resources allocated to this model of care (Lopes, 2025).

The private healthcare sector has also recognised the potential of this model, contributing to the overall growth of HaH. In 2020, CUF became the first private hospital to adopt this system across various parts of the country. It operates with specialised doctor and nurse teams and has a capacity of 47 patients (CUF, 2026).

2.5 Routing and Visit-planning Context at Portuguese HaH Units

The service is available to individuals aged 18 and over whose place of residence is within a distance or travel time that meets the safety requirements for a timely response (normally a maximum of 30 minutes). It operates 24 hours a day, every day of the year. The first home visit takes place within 24 hours of admission. On the day of discharge, there should be an interdisciplinary visit (doctor/nurse) (CUF, 2023; SNS, 2018).

Home visits are scheduled at daily clinical meetings according to each patient's needs. At Portuguese HaH units, typically, there are teams composed of physicians and nurses, although not all patients necessarily require the same type of team. Some patients may require a physician visit, while others can be visited by a nursing team only. The duration of each visit may vary depending on the unique needs of each patient, but no visit will be shorter than 20 minutes. Planning must also account for the balance of workload between teams (Martins et al., 2026).

Once teams have been assigned to patients, each team makes its own routing decisions. Each nurse therefore defines the route by considering three main aspects: distance, local traffic and patient status. Patients who require urgent consultations are prioritised and will be visited first. The start and end of the route for doctors and nurses must always be at the depot, i.e., the hospital (Bonomi et al., 2025). This manual planning delays team departure and relies on professionals' experience without systematic decision support (Martins et al., 2026).

3 Literature Review

3.1 Theoretical Overview

This section reviews routing and scheduling literature for home-based healthcare services. Because HaH-specific routing studies are scarce, most insights come from HHC literature. Although HaH and HHC differ in clinical intensity and organisational structure, both require daily caregiver allocation to geographically dispersed patients under temporal, clinical and operational constraints. Thus, HHC routing and scheduling literature provides a foundation for HaH logistics. The key differences between HaH and HHC were discussed in Section 2.3.

3.2 Home Health Care Routing and Scheduling Problem: Definition and Variants

The Home Health Care Routing and Scheduling Problem (HHCRSP) is a specialised variant of the vehicle routing and workforce scheduling problem. When multiple hospital depots are considered, the problem further extends into a Multi-Depot Vehicle Routing Problem with Time Windows (MDVRPTW) (Di Mascolo et al., 2021).

These problems are characterised by constraints inherent to the sector, such as the compatibility of staff and patient schedules, shift boundaries, visiting windows, operational costs and geographical dispersion. Balancing all these factors to allocate resources efficiently and plan optimised routes is a significant challenge (B. Li et al., 2026).

Planning decisions are usually divided into strategic, tactical and operational levels (Hulshof et al., 2012). Strategic decisions define service territories, tactical decisions estimate workforce and capacity needs, and operational decisions assign caregivers to patients and schedule visits (Makboul et al., 2024). At this operational level, the key features discussed below can be grouped into temporal constraints, assignment and geographical considerations, and sources of uncertainty.

Temporal-based Features

The time windows for each visit, which define the intervals during the day when each patient can be seen, form the basis of the time restrictions. More recent articles have begun to consider multiple time windows per patient, i.e., morning or afternoon visits, increasing the flexibility of the service (Bazirha et al., 2024; Du & Wang, 2024). In contexts with uncertainty or capacity limits, models may allow missed visits with objective penalties (Ma et al., 2025; Shahnejat-Bushehri et al., 2021). This flexibility allows algorithms to find solutions even when there are conflicts between schedules and availability.

In addition, violations of time windows may occur as a result of a delay or early arrival, which are penalised by additional costs in the model (Atta et al., 2025). As a rule, each caregiver is required to interrupt their activity for a certain period of the day, for example, for lunch or even for the duration of each visit. This can be modelled as a fictitious service window (Shiri et al., 2021), or even as a return to the hospital base between shifts, especially when professionals need to replenish supplies. Also, the models impose maximum limits on the number of visits per shift, avoiding physical and emotional overload for caregivers and maintaining service quality and worker satisfaction. Finally, the most realistic models also integrate the schedule preferences of patients and caregivers, which are generally treated as soft constraints, represented by mild penalties in the objective function (Atta et al., 2025; Trautsamwieser & Hirsch, 2014).

Assignment and Geographical-based Features

Another characteristic of these models is the compatibility between the skills of professionals and the needs of patients. For example, the Stochastic Districting Staff Dimensioning – Assignment Routing Problem model proposed by Nikzad et al. (2021) accounts for this constraint. Their model ensures that a caregiver can only serve a patient if the patient's minimum required skill level is equal to or lower than the caregiver's skill level. Furthermore, this constraint is integrated into the model to optimise assignment and routing decisions.

Another relevant factor is the preference for keeping the same caregiver visiting the same patient over time. Shahnejat-Bushehri et al. (2021) propose a model that addresses this by considering the dependency between precedence synchronised services and the appropriate level of continuity of care. This approach reduces the loss of patient medical information and, more importantly, raises the level of confidence between the patient and caregiver. However, due to patients' varying needs, achieving full continuity of care may not always be feasible.

Another condition that may also influence assignment is the gender and language of the caregiver. Yadav and Tanksale (2022) incorporate these specific patient preferences as hard constraints, meaning that assignment is only feasible if these restrictions are respected. This requirement is directly linked to the objective of maximising the total revenue generated by the service, as failure to respect these preferences results in the impossibility of carrying out the visit, negatively impacting the financial outcome of the operation.

Some articles consider a single start and finish location for care workers (a depot). However, it is already common for several heuristics to perform multi-depot routing, where each professional starts and ends their route at an associated health centre or even at their own homes. This decision depends on organisational constraints (end-of-shift times, available transport) and the balance between travel costs and operational flexibility. However, the existing literature generally assumes a prior assignment of caregivers to centres. In contrast, the article by Decerle et al. (2018) propose a MILP model that solves the assignment of caregivers to health centres, rather than considering it as a fixed parameter. By avoiding *a priori* assignment, the model obtains better planning combinations and reduces the total travel time of caregivers.

Uncertainty-based Features

Although the models began with planning that assumed all variables were known (deterministic), the reality of operations is governed by significant uncertainty in travel and service times (Tsang & Shehadeh, 2023; Zhan et al., 2021). These random parameters result in delays in arrival and completion of services, incurring penalties that affect both operator utilisation and customer satisfaction (Zhan et al., 2021). One way to address this issue is the solution proposed by Hosseini et al. (2026), which introduces a chance-constrained optimisation framework that guarantees a high probability of on-time arrivals, despite the uncertainty of travel and service times. This ensures reliable and consistent schedules.

While some delays are random (e.g., accidents), others follow predictable patterns. Y. Li et al. (2022) highlight that rush hours create deterministic variations in travel speed throughout the day, requiring models that can adapt to both predictable time-dependent traffic and unpredictable fuzzy disruptions. In addition, adverse weather conditions can also affect the service (Shahnejat-Bushehri et al., 2021).

According to Shi et al. (2017), demand is non-deterministic and difficult to estimate using standard probability distributions due to a lack of historical data. To address this, they propose a fuzzy chance-constrained programming model where patient demand is treated as a fuzzy variable rather than a stochastic one. By applying fuzzy credibility theory, the model balances the minimisation of transport costs with the risk of violating vehicle capacity constraints.

Ziya-Gorabi et al. (2022) address the specific impact of urban logistics, focusing on the uncertainty caused by traffic congestion and its environmental impact. They propose a fuzzy chance-constrained programming model that explicitly treats congestion as a fuzzy parameter.

In addition to time uncertainties, in the HHCRSP, new requests may arrive and some scheduled visits may be cancelled abruptly. Shiri et al. (2021) proposed a three-phase methodology that allows routes to be rescheduled in the face of new emerging requests or cancellations, while maintaining compliance with service windows and minimising operating costs.

3.3 Solution Approaches

The review by Atta et al. (2025) analyses various optimisation methods for the HHCRSP, including exact, heuristic, metaheuristic and hybrid algorithms.

Exact methods typically use integer programming to formulate the HHCRSP, usually extending the Vehicle Routing Problem (VRP) with binary assignment and time variables. For example, R. Liu et al. (2017) proposed a three-dimensional MILP model with lunch interval constraints and solved it by branch-and-price. In this model, a label-correcting algorithm handles lunch windows during column generation. Other exact articles include compact models and branch-and-cut algorithms. Du and Wang (2024) developed a branch-and-price-and-cut algorithm with multiple time windows and service priorities.

However, the computational cost of exact approaches grows rapidly with instance size. To overcome the scalability limitations of exact methods, a variety of heuristic and metaheuristic approaches have been proposed. Several authors adopt constructive or clustering-based heuristics. Erdem et al. (2022) proposed a cluster-first, route-second approach for a multi-period HHCRSP. Grenouilleau et al. (2019) construct subsets of feasible routes and combine them through a set partitioning heuristic. Bazirha et al. (2023) proposed a two-phase heuristic in which routes are first generated using Simulated Annealing (SA) and then specific time windows and synchronised visits are assigned in a second phase. Their later work also considers multiple objectives including service lateness, waiting times and load balancing, solved with evolutionary algorithms (Bazirha et al., 2024).

Population-based metaheuristics have also been widely applied. Kummer et al. (2024) applied a biased random-key genetic algorithm (BRKGA), achieving up to 26.1% improvement over previous local search methods. Belhor et al. (2023) developed a multi-objective genetic algorithm (GA) approach using K-means clustering, and Clapper et al. (2023) proposed a memetic genetic algorithm for multi-objective optimisation. Swarm-based algorithms have also been explored: Y. Li et al. (2022) use a discrete Grey Wolf optimiser, while Fu et al. (2024) apply multi-objective Migrating Birds Optimisation under stochasticity. Hybrid approaches that combine exact and metaheuristic components have shown promising results. W. Liu et al. (2021a) formulate a time-based MILP with synchronised visits and lunch breaks and implement four hybrid metaheuristics. Clapper et al. (2024) propose adaptive simheuristics that allocate simulation budgets within metaheuristics for stochastic HHCRSP instances.

Among the metaheuristic frameworks applied to HHC, ALNS has emerged as particularly effective for Vehicle Routing Problem with Time Windows (VRPTW)-based problems. W. Liu et al. (2024) applied ALNS combined with multi-objective simulated annealing (AMOSA) to a pandemic HHC environment, incorporating constraints such as restrictions on contact between doctors/nurses and patients, lunch breaks at home, multiple trips a day and patients' preferences. More recently, Zhang et al. (2025) proposed an improved ALNS (IALNS) for the HHCRSP with multiple mixed time windows, skill compatibility, simultaneous services and workload balancing, demonstrating its superiority over the exact model.

The healthcare routing and scheduling problem, of which HaH planning is a variant, is an NP-hard combinatorial problem, which justifies the use of metaheuristics rather than exact methods. Metaheuristics are a form of optimisation that explore the solution space without exhaustive enumeration. Metaheuristic performance depends on balancing diversification, which explores the search space, avoiding local minima traps, and intensification, consisting of the ability to refine promising solutions found in order to quickly converge to what is expected to be the global optimum (Al-Sorori & Mohsen, 2020).

Among the various metaheuristics available in the literature, the ALNS proposed by Ropke and Pisinger (2006) has established itself as one of the most successful approaches to routing problems, having originated from Large Neighbourhood Search (LNS). The fundamental principle of LNS consists of iteratively improving a solution through a cycle of destruction and repair, thereby exploring large neighbourhoods that allow escape from local optimum (W. Liu et al., 2021b). ALNS proves to be superior compared to LNS by allowing the simultaneous use of multiple destruction and repair operators, selected dynamically. This learning mechanism allows the algorithm to automatically adapt to the characteristics of the instance being solved (Gendreau & Potvin, 2019).

In VRP problems including additional constraints, for instance multiple depots, heterogeneous teams with different qualifications, time windows, varying levels of urgency, or shift duration constraints, a fast and robust algorithm is required. In this context, ALNS is particularly well suited for tightly constrained problems, in which small neighbourhoods are not sufficient to escape a local minimum (Pisinger & Ropke, 2007). As it also has a great ability to integrate and add new features, it is highly adaptable to various VRP extensions (Zhang et al., 2025).

3.4 HaH Routing and Scheduling

Although few, recent articles focused specifically on routing for HaH have begun to appear. Notably, Varas et al. (2024) formulate a MILP model that minimises total travel time, workload imbalance between teams and waiting time, with multiple time slots for visiting patients on the same day, the requirement to take lunch breaks at the hospital and the need to return immediately to the hospital after collecting perishable biological samples. In addition, the model ensures that each patient is seen by a team with the correct medical specialty and that visits requiring multiple professionals are synchronised in time. According to the computational experiments carried out using CPLEX, it is only feasible with fewer than 80 requirements, as the calculation time grows exponentially. The heuristics developed, on the other hand, were able to generate high-quality solutions in extremely short times. In the practical case of Padre Hurtado Hospital, instances with 141 patients and 155 tasks were tested, including antibiotic treatments, blood tests and visits to risk areas. Comparing the results of the model with the nurse's manual planning, there was a 31.8% reduction in total travel time, an 83.6% reduction in waiting time and a 95.4% reduction in the imbalance between teams.

Similarly, Bonomi et al. (2025) study a multi-hospital HaH scenario in Portugal: they develop a mixed integer model with three objectives, minimising total travelled time (TTT), nurses' workload balance (NWB) and total number of doctors (TND). The process involves two phases. First, Single-Objective Optimisation, in which each objective is optimised individually to establish benchmark values. Second, Bi-Objective Analysis (Pareto Frontiers), in which the objectives are combined in pairs (TTT-NWB, TTT-TND and TND-NWB). The model has been tested using real-world data from Portugal (Lisboa region) and shows that there are clear trade-offs between cost, capacity and resource use, making it possible to inform decisions such as opening new hospitals or redistributing patients more efficiently.

Another key study on this topic is the work by Quintanilla et al. (2020), which addresses a HaH routing problem for a hospital in Valencia, Spain. The article proposes two main models: Current Practice Model and Extended Model. The first one reflects the hospital's rigid policies, where the doctors and nurses are in pre-defined teams, but also the taxis are pre-assigned to a set of workers who travel together for the whole route. The second model is highly complex and explores the benefits of introducing sustainable strategies and more flexible policies. Workers can walk between houses for short distances, the set of workers transported by a taxi may change during the route and a taxi's route can start and/or finish at a home. The core objective of the proposed models is to minimise the transportation costs related to the total time of the taxi ride, including both the travel and waiting costs.

Despite these contributions, significant gaps remain in the current HaH routing literature. Firstly, the reviewed models rely primarily on deterministic frameworks, neglecting the critical impact of real-world uncertainty in travel and service times (Bonomi et al., 2025; Quintanilla et al., 2020; Varas et al., 2024). Secondly, they often impose site-specific rigid constraints, such as mandatory hospital lunch breaks (Varas et al., 2024) or exclusive taxi-based transportation with pre-defined teams (Quintanilla et al., 2020), which limits their transferability to other providers.

Thirdly, no metaheuristic approach has been proposed, nor has any study provided a systematic benchmark comparing exact and metaheuristic methods in this setting.

To address these gaps, the following chapter presents the methodology, including the mathematical formulation of the problem as a MILP and the design of an ALNS metaheuristic.

4 Problem Definition

HaH consists of delivering hospital-level care to patients at their homes through medical teams. For this purpose, every day, a group of patients must be visited by care teams composed of nurses or doctors, subject to clinical compatibility requirements, operational restrictions and working time limitations.

This work addresses this daily routing of care teams to their assigned patients. Given a depot and assigned patients, the problem determines patient-team assignments and visit sequences. Each team performs a maximum of one route per day, starting and ending at the depot, attending to multiple patients within a maximum predefined shift period. Although planning may involve multiple depots, this dissertation decomposes the problem by depot, reflecting the operational and organisational structure of the case-study: each patient belongs to a hospital unit and each depot has dedicated teams. It also reduces computational complexity by transforming the multi-depot setting into a set of smaller single-depot routing and assignment subproblems. Patient-depot assignment is handled in pre-processing, using the input depot or the nearest depot when missing, after which each depot is solved independently.

Each depot has teams of nurses and doctors, the size of which is defined for each depot. Each team carries out a maximum of one route per day, starting and ending at their respective depot. Every patient must be visited exactly once by a team whose clinical expertise satisfies their needs. Patients who require a doctor can only be assigned to doctor teams, while the rest must be assigned to nurse teams. This decision is relevant from an operational standpoint because it allows doctors to be assigned to patients who truly need them.

Additionally, patients are classified as urgent or non-urgent, reflecting different levels of clinical priority. Whenever a route includes both urgent and non-urgent patients, all urgent patients must be visited before any non-urgent patients, imposing a precedence structure within each route. Furthermore, to avoid excessive concentration of urgent cases in a single team and to balance workload distribution, the number of urgent patients assigned to each team is limited to a predefined value M. Throughout this work, instances are assumed to satisfy this condition, so that the total number of urgent patients never exceeds the available urgent-care capacity.

Each team is subject to a maximum shift duration constraint, defined as the sum of travel time and service time at patients' homes. Routes whose total duration exceeds the maximum limit are considered infeasible. Travel times between locations are treated as input data, and service times depend on the specific needs of each patient, being incorporated directly into the route duration constraint. In addition to operational feasibility, the model also considers resource intensity. Activating a team that includes a doctor implies an additional fixed cost, reflecting the higher level of resources associated with the presence of a doctor. The objective function thus balances total travel time with the cost of activating doctor teams, capturing the trade-off between operational efficiency and the use of specialised resources.

To model this problem accurately and provide a benchmark for the metaheuristic approach developed in this work, a MILP formulation is proposed.

5 Methodology

This chapter presents two approaches to the problem defined in Chapter 4: a MILP formulation, used as an exact benchmark, and an ALNS metaheuristic, calibrated before the final experiments. The workflow covers HaH data preparation, depot-level instance generation and travel-matrix construction. The MILP and ALNS approaches are then evaluated through computational experiments, and the resulting routes are used in the prototype visualisation.

5.1 Mixed Integer Linear Programming Model

The problem is formulated as a MILP model since it combines binary assignment and routing variables with continuous auxiliary variables and all objective terms and constraints are linear. Before presenting the formulation, the notation used throughout the model is introduced, covering the sets, parameters and decision variables.

Sets

- $k \in K$: set of care teams, $k = 1,2,\dots,|K|$.
- $i \in N$: set of patients, $i = 1,2,\dots,|N|$.
- $V = \{0\} \cup N$: set of nodes, where node 0 represents the depot.
- $U = \{i \in N : U_i = 1\}$: subset of urgent patients.

Parameters

- t_{ij}: travel time between nodes $i \in V$ and $j \in V$.
- s_i: service time required by patient $i \in N$.
- $U_i \in \{0,1\}$: urgency indicator of patient $i \in N$.
- $D_i \in \{0,1\}$: indicator equal to 1 if patient $i \in N$ requires a doctor.

- $H_k \in \{0,1\}$: indicator equal to 1 if team $k \in K$ includes a doctor.
- $T^{\max}$: maximum shift duration.
- M: maximum number of urgent patients allowed per route.
- W_{travel}: weight associated with total travel time.
- W_{doc}: fixed penalty for using a doctor team.

Decision Variables

- $x_{kij} \in \{0,1\}$: equals 1 if team $k \in K$ travels directly from node $i \in V$ to node $j \in V$.
- $a_{ki} \in \{0,1\}$: equals 1 if patient $i \in N$ is assigned to team $k \in K$.
- $y_k \in \{0,1\}$: equals 1 if team $k \in K$ is used.
- $u_{ki} \geq 0$: continuous auxiliary variables for subtour elimination, representing the position in the route of team k of patient i.

Objective function

Taking the above notation into account, the formulation is as follows.

The objective function minimises the total operational cost, defined as the sum of total travel time and a fixed penalty for each activated doctor team:

$$\min \; W_{\text{travel}} \sum_{k \in K} \sum_{i \in V} \sum_{\substack{j \in V \\ j \neq i}} t_{ij} x_{kij} + W_{\text{doc}} \sum_{k \in K} H_k y_k \tag{1}$$

Assignment and Routing Constraints

Each patient must be assigned to exactly one compatible team. Patients requiring a doctor can only be assigned to doctor teams, while non-doctor patients are assigned to nurse teams.

$$\sum_{\substack{k \in K \\ H_k = D_i}} a_{ki} = 1 \quad \forall i \in N \tag{2}$$

For each team and patient, exactly one arc enters and leaves the patient if the patient is assigned.

$$\sum_{\substack{j \in V \\ j \neq i}} x_{kij} = a_{ki} \quad \forall k \in K, i \in N \tag{3}$$

$$\sum_{\substack{j \in V \\ j \neq i}} x_{kji} = a_{ki} \quad \forall k \in K, i \in N \tag{4}$$

Each active team starts and ends its route at the depot.

$$\sum_{j \in N} x_{k0j} = y_k \quad \forall k \in K \tag{5}$$

$$\sum_{j \in N} x_{kj0} = y_k \quad \forall k \in K \tag{6}$$

A team can only serve patients if it is activated.

$$a_{ki} \leq y_k \quad \forall k \in K, \forall i \in N \tag{7}$$

The total travel and service time of each team must not exceed the maximum shift duration.

$$\sum_{i\in V}\sum_{\substack{j\in V\\ j\neq i}} t_{ij}x_{kij} + \sum_{i\in N} s_i a_{ki} \leq T^{\max} y_k \quad \forall k \in K \tag{8}$$

Urgent-care Policy Constraints

The model enforces a strict urgent-first routing policy composed of two sets of hard constraints.

First, if a route contains both urgent and non-urgent patients, all urgent patients must be visited before any non-urgent patient. This is enforced by forbidding arcs from non-urgent to urgent patients.

$$x_{kij} = 0 \quad \forall k \in K,\, \forall i \in N \setminus U,\, \forall j \in U \tag{9}$$

To prevent excessive concentration of urgent cases in a single team, the number of urgent patients served by each team is limited by M. This constraint directly caps the number of urgent patients assigned to each team while preserving the urgent-first routing structure.

$$\sum_{i\in N} U_i\, a_{ki} \leq M \quad \forall k \in K \tag{10}$$

Subtour Elimination Constraints

Subtours are eliminated using Miller-Tucker-Zemlin (MTZ) constraints,

$$u_{ki} - u_{kj} + |N| x_{kij} \leq |N| - 1 \quad \forall k \in K,\, \forall i \neq j \in N, \tag{11}$$

with bounds

$$a_{ki} \leq u_{ki} \leq |N|\, a_{ki} \quad \forall k \in K,\, \forall i \in N. \tag{12}$$

Finally, the domains of the decision variables are:

$$\begin{aligned} &x_{kij} \in \{0,1\},\ \forall k \in K,\ \forall i,j \in V,\ i \neq j, && y_k \in \{0,1\},\ \forall k \in K,\\ &a_{ki} \in \{0,1\},\ \forall k \in K,\ \forall i \in N, && u_{ki} \in \mathbb{R}_+,\ \forall k \in K,\ \forall i \in N. \end{aligned} \tag{13}$$

5.2 Solution Approach

Knowing that exact methods are not usually able to solve large scale NP-hard problems within reasonable computational time, this dissertation proposes an algorithm based on the ALNS method proposed by Ropke and Pisinger (2006). This metaheuristic is based on the Large Neighbourhood Search and has been successfully adopted in the combinatorial optimisation field in the last decade (W. Liu et al., 2021b).

The ALNS follows an iterative process of solution improvement, in which a current solution is partially destroyed through removal operators and then rebuilt by insertion heuristics. The systematic alternation between destruction and reconstruction phases allows large neighbourhoods to be explored, enabling significant structural changes to routes.

Throughout the iterations, the algorithm seeks to identify which operators produce higher quality solutions. In the proposed method, this adaptation is implemented through a classical Roulette Wheel Selection (RWS) mechanism, which maintains separate weight vectors for destroy and repair operators. At each iteration, each operator is selected with probability proportional to its current weight. Weights are updated at the end of fixed-length segments according to the performance scores accumulated during that segment, rewarding operators that led to new global best solutions and accepted improvements or accepted non-worsening moves.

Additionally, to prevent the algorithm from getting stuck on a local optimum, the acceptance of worse solutions is controlled by SA. SA accepts worse solutions with temperature-dependent probability, encouraging early exploration and later intensification.

The ALNS stops after a fixed number of iterations. An iteration-based stopping criterion was preferred over a time-based one because it ensures that each run performs the same number of destroy-and-repair cycles, making comparisons across instances more fair.

Proposed Solution Algorithm

The ALNS requires a complete feasible initial solution to begin. In this work, this solution is constructed using a greedy repair heuristic, which sequentially inserts patients at the feasible position with the lowest incremental travel time cost. At each insertion, the hard constraints are checked, namely team compatibility, maximum shift duration, urgent-first ordering and the limit on urgent patients per route. Starting from this initial solution, the proposed procedure follows the iterative structure summarised in Algorithm 1 (Appendix A). At each iteration, a candidate solution is generated, evaluated and either accepted or rejected. The search returns the best feasible solution found over the fixed iteration limit.

The following subsections describe the main components of the method: the solution representation, the destroy and repair operators, the acceptance criterion and the adaptive weight adjustment mechanism.

Solution Representation

A solution S is represented by an ordered set of K routes, one for each team available at the depot,

$$S = (r_1, r_2, \ldots, r_K), \tag{14}$$

where each route r_k is an ordered sequence of patients,

$$r_k = (p_1^k, p_2^k, \ldots, p_{n_k}^k), \quad p_i^k \in N, \tag{15}$$

where $n_k \geq 0$ is the number of patients assigned to team k. Empty routes ($n_k = 0$) represent unused teams.

Each team k is either a doctor team ($H_k = 1$) or a nurse team ($H_k = 0$). To ensure compliance with the urgency policy, each route r_k follows an internal two-block structure,

$$r_k = (u_1^k, \ldots, u_{m_k}^k, v_1^k, \ldots, v_{l_k}^k), \tag{16}$$

where $u_1^k, \ldots, u_{m_k}^k$ are urgent patients and $v_1^k, \ldots, v_{l_k}^k$ are non-urgent patients, with $m_k + l_k = n_k$ and $m_k \leq M$. This ordering is preserved throughout the search by a sanitisation routine applied after every destruction operation.

A solution is considered complete when all patients are assigned to exactly one team. The ALNS operates exclusively on complete solutions since the initial solution is complete and each iteration destroys and repairs the solution until it is made complete.

Operators

Each iteration selects a pair (d, r) composed of a destroy operator and a repair operator, with the candidate solution being constructed by the composition $S' = r\big(d(S_{\text{cur}})\big)$, where S_{cur} is the current solution and S' is the candidate solution.

Let $A(S_{\text{cur}}) = \bigcup_{k=1}^{K} r_k$ denote the set of patients assigned in the current complete solution. The destroy operator removes a subset of $A(S_{\text{cur}})$, producing a partial solution $\tilde{S} = d(S_{\text{cur}})$. For this partial solution, $A(\tilde{S})$ denotes the patients that remain assigned to routes, and $U(\tilde{S}) = N \setminus A(\tilde{S})$ denotes the removed, temporarily unassigned patients. The repair operator acts on $U(\tilde{S})$ by reinserting these patients until a complete candidate solution $S' = r(\tilde{S})$ is obtained.

In each iteration, the number of patients removed is not fixed, being controlled by a random fraction $q \sim \mathscr{U}(q_{\min}, q_{\max})$. This range should follow the guidelines established by Ropke and Pisinger (2006), promoting diversification without significantly changing the structure of the solution. Thus, if $|A(S_{\text{cur}})|$ is the number of patients assigned before destruction, the number of removals is $n_{\text{rem}} = \max\{1, \lfloor q\,|A(S_{\text{cur}})|\rfloor\}$.

Destroy Operators

Three destroy operators were implemented: worst removal, shaw removal and random removal. After every destruction, a sanitisation step restores the urgent-first ordering within each route by rearranging patients into the format defined in (Eq. (16)), ensuring that the repair operators can rely on this structure.

The **random removal** operator selects n_{rem} positions uniformly among all visits present in the solution. Let $\Pi(S_{\text{cur}})$ be the set of pairs route k, position i for all assigned visits in the current solution. The operator chooses a subset $C \subseteq \Pi(S_{\text{cur}})$ with $|C| = n_{\text{rem}}$ without replacement and removes the corresponding visits. This operator promotes diversification, as it removes patients regardless of their contribution to the current cost.

The **worst removal** operator preferentially removes visits whose removal yields the largest saving in travel time on the route. For a patient in position i on a route $R = (\dots, p_{i-1}, p_i, p_{i+1}, \dots)$, the local saving (in minutes) is defined as:

$$\Delta_i = t_{p_{i-1},p_i} + t_{p_i,p_{i+1}} - t_{p_{i-1},p_{i+1}}, \tag{17}$$

adopting the depot as predecessor/successor when i is the end of the route. The operator sorts all candidates by decreasing Δ_i and removes n_{rem} patients with the highest marginal contribution.

In the **shaw removal** operator, a patient $p^{(0)}$ is initially selected at random from the current solution. Next, patients related to those already removed are iteratively removed, where the relationship is measured by a similarity function based on the travel time between patients,

$$\text{rel}(p,q) = \phi_{\text{time}} \cdot t_{p,q} - \phi_{\text{urg}} \cdot \mathbb{I}\{U_p = U_q\}, \tag{18}$$

where U_p is the patient's urgency attribute, $\mathbb{I}\{\cdot\}$ is the indicator function and $\phi_{\text{time}} = 1$. The weight $\phi_{\text{urg}} = 2$ reflects the strong coupling between patients with the same urgency label, while geographical proximity remains the main driver of relatedness. Removing them together creates more opportunities for restructuring during the repair. At each step, the remaining patients are sorted in ascending order of proximity and the patient to be removed is selected using a randomised mechanism: $v = L[\lfloor y^p \cdot |L| \rfloor]$, where y is drawn from the uniform distribution $\mathscr{U}(0,1)$ and $p = 6$ controls the degree of determinism. A high value of p will bias the selection towards the most related patients, whilst $p = 1$ reduces the selection to a uniform random selection. The value $p = 6$ follows the recommendation of Ropke and Pisinger (2006).

Repair Operators

Two repair operators were implemented: greedy repair and regret-2 repair. To insert a patient p into a route $R = (p_1, \dots, p_m)$ at position j, the incremental cost is defined by

$$\Delta(p,R,j) = t_{\text{prev},p} + t_{p,\text{next}} - t_{\text{prev},\text{next}}, \tag{19}$$

where prev $= 0$ if $j = 0$ and prev $= p_j$ otherwise and next $= 0$ if $j = m$ and next $= p_{j+1}$ otherwise. This cost allows a quick assessment of the marginal impact of the insertion without recomputing the entire route.

The **greedy repair** operator iteratively reinserts patients unassigned by the destroy phase, choosing for each patient the feasible insertion with the lowest incremental cost Δ. Urgent patients are reinserted first and, in case of ties, a deterministic criterion based on the position and index of the team is applied.

The **regret-2 repair** operator implements a second-order regret heuristic. For each patient p, the best incremental cost $\Delta_1(p)$ and the second-best $\Delta_2(p)$ are calculated among the feasible insertions, defining regret as

$$\text{regret}_2(p) = \Delta_2(p) - \Delta_1(p), \tag{20}$$

with $\Delta_2(p)$ defined as a large value when there is only one feasible insertion. The algorithm then selects the patient with the highest regret_2, reinserting them in their best position.

Both share a common feasibility-checking mechanism which, for each candidate insertion, simultaneously checks all hard constraints:

1. **Skills compatibility**: the patient may only be inserted into a route whose team matches their requirement for a doctor ($H_k = D_i$).
2. **Urgent-first ordering**: an urgent patient may only be inserted before or between the existing urgent block; a non-urgent patient may only be inserted after all urgent patients.
3. **Urgent patient limit**: if inserting an urgent patient causes the number of urgent patients on the route to exceed M, the insertion is rejected.
4. **Shift duration**: the insertion is only accepted if the total duration of the resulting route (travel time + service time) does not exceed $T^{\max}$. This check is performed incrementally.

Acceptance Criterion

At each iteration, the candidate solution S' generated by the pair of operators (d, r) is evaluated and compared with the current solution S_{cur}. The decision to accept or reject S' is evaluated using a SA criterion, which helps the algorithm escape local optima by accepting some degraded solutions during the search process.

A candidate solution S' is accepted if it is better than or equal to the current solution, or, otherwise, with a probability that decreases exponentially with the deterioration of the objective, i.e., if

$$f(S') \leq f(S_{\text{cur}}) \quad \text{or} \quad \mathscr{U}(0,1) < \exp\Big(-\frac{f(S') - f(S_{\text{cur}})}{T}\Big), \tag{21}$$

where $T > 0$ is the current temperature of the system and $\mathscr{U}(0,1)$ represents a uniform random variable.

The initial temperature T_0 is defined in proportion to the value of the objective function of the initial solution S_0

$$T_0 = \max\big(T_{\min},\ \beta \cdot f(S_0)\big), \tag{22}$$

where β is the scaling factor and $T_{\min}$ is the minimum permissible temperature. This approach automatically adapts the temperature scale to the magnitude of the objective.

The temperature is updated at the end of each iteration according to a cooling rule

$$T \leftarrow \max\big(T_{\min},\ \alpha \cdot T\big), \tag{23}$$

with $\alpha \in (0,1)$ as the cooling rate. After I iterations, the temperature is approximately $T_0 \cdot \alpha^I$. As the temperature decreases, the algorithm gradually shifts from broad exploration to intensification, where only improvements are accepted.

Adaptive Weight Adjustment

The selection of destroy and repair operators at each iteration is carried out adaptively using a RWS mechanism. The aim is to prioritise, throughout the execution, the operators that contribute most to a better solution.

Two independent weight vectors are maintained: $\mathbf{w}_d$ for the destroy operators and $\mathbf{w}_r$ for the repair operators. At the start of the ALNS, all methods are initialised with weights equal to 1 and scores equal to 0.

Let j denote a destroy operator in the set $\mathscr{D}$. In each iteration, the destroy operator d is selected with probability proportional to its current weight:

$$P(d = j) = \frac{w_{d,j}}{\sum\limits_{i=1}^{|\mathscr{D}|} w_{d,i}}, \tag{24}$$

where $w_{d,j}$ is the current weight of destroy operator j. The same rule is applied to select the repair operator r from $\mathscr{R}$ using the repair weights $\mathbf{w}_r$. The selection of the two operators is independent.

The entire search is divided into many segments and each one is associated with 100 consecutive iterations, following the standard value adopted in the ALNS literature (W. Liu et al., 2021b; Ropke & Pisinger, 2006). During each segment, a score accumulator π_j and a usage counter θ_j are maintained for each operator j, both initialised to zero at the start of each segment. At the end of each iteration, the selected pair (d, r) receives a score:

$$\sigma = \begin{cases} \sigma_1 & \text{if } S' \text{ is a new global best,} \\ \sigma_2 & \text{if } S' \text{ is accepted and } f(S') < f(S_{\text{cur}}), \\ \sigma_3 & \text{if } S' \text{ is accepted but } f(S') \geq f(S_{\text{cur}}), \\ 0 & \text{if } S' \text{ is rejected.} \end{cases} \tag{25}$$

The score σ is simultaneously added to the accumulators of both operators selected in that iteration, i.e., $\pi_d \leftarrow \pi_d + \sigma$ and $\pi_r \leftarrow \pi_r + \sigma$.

At the end of each segment, the weights are updated based on the average performance of each operator j with $\theta_j > 0$

$$w_j \leftarrow (1-\lambda) \cdot w_j + \lambda \cdot \frac{\pi_j}{\theta_j}, \tag{26}$$

where $\lambda \in (0,1)$ is the reaction factor that controls how quickly the weights adapt to recent performance. Low values of λ make the adaptation more conservative, whilst high values make the weights more responsive to short-term variations. Operators not used during the segment ($\theta_j = 0$) retain their weights unchanged. After the update, the accumulators π_j and θ_j are reset to zero for the next segment.

In the original approach of Ropke and Pisinger (2006), scores in scenarios σ_2 and σ_3 are assigned only to solutions not previously visited, using a hash table to check for duplicates. In the present implementation, the novelty check is omitted for reasons of computational simplicity, given that the SA, which progressively reduces the probability of accepting worse solutions, minimises the probability of revisiting solutions systematically. This simplification was considered acceptable in this implementation.

6 Data Collection

6.1 Instance Information

To evaluate the proposed approaches, instances and travel-time matrices were generated. This subsection also reports the parameter values adopted throughout the experiments.

The starting point is an Excel file with 316 real HaH patients, including ID, coordinates, urgency, doctor requirement, service time and depot. The four depots correspond to existing CUF Saúde hospital units. Table 2 presents their characteristics, including the number of nurse teams (nurse only) and doctor teams (which include both a nurse and a doctor), and the share of patients assigned to each depot. Five instance sizes are considered: **small**, **medium**, **large**, **xlarge** and **xxlarge**, corresponding to 50, 100, 150, 200 and 250 patients. The per-depot breakdown for each size is derived proportionally from the source data. Since proportional targets are generally non-integer, the largest-remainder method is applied to ensure that per-depot allocations sum exactly to N. Table 3 presents the resulting allocations.

The second step is to create different scenarios, which introduces controlled bias into the sampling to create instances with distinct characteristics. Four scenarios are considered:

1. **Balanced scenario**, which maintains the original proportions of the population without

Table 2: Depot characteristics.

Depot	Patients	Share	Nurse teams	Doctor teams
Lisboa	146	46.2%	12	12
Cascais	89	28.2%	8	8
Sintra	56	17.7%	6	6
Odivelas	25	7.9%	4	4
Total	316	100%	30	30

Table 3: Patient allocation per depot and size.

Size	N	Lisboa	Cascais	Sintra	Odivelas
Small	50	23	14	9	4
Medium	100	46	28	18	8
Large	150	69	42	27	12
XLarge	200	92	56	36	16
XXLarge	250	116	70	44	20

any distortion;

2. **Clustered scenario**, which prioritises the selection of patients who are geographically close to one another;
3. **High urgent scenario**, which over-represents urgent patients using weights of 1.8 for urgent and 0.8 for non-urgent patients, making urgent patients 2.25 times more likely to be selected;
4. **High doctor scenario**, which applies the same over-representation logic to patients requiring a doctor.

For each combination of size and scenario, eight replicates are generated with distinct random seeds. Each instance is decomposed by depot, yielding 160 depot-level instances per scenario (four depots $\times$ five dimensions $\times$ eight replicates), and 640 depot-level instances across the four scenarios. Then, the OpenRouteService API is used to pre-calculate the travel time and distance matrices between each depot and patient, and between each pair of patients.

The parameter values used throughout the experiments are as follows. The maximum shift duration is $T^{\max} = 240$ minutes, representing a morning shift (9.00am-1.00pm) for nurses and doctors in a hospital environment. The maximum number of urgent patients per route is $M = 3$, introduced to make the routes more realistic, promoting a faster response and a balanced workload across teams. The travel-time weight is $W_{\text{travel}} = 1$, since the objective is measured in minutes and each unit of travel time contributes directly to the cost. The doctor-team penalty is $W_{\text{doc}} = 200$, chosen to roughly match the duration of a full shift ($T^{\max} = 240$ min), so that activating a doctor team is penalised on a scale comparable to the travel-time component.

6.2 ALNS Parameter Tuning

This work uses the Design of Experiments (DoE) method for parameter tuning, motivated by its successful use in calibrating the parameters of optimisation algorithms. DoE techniques provide a robust way to design experiments efficiently, improving understanding of the interaction between variables and the desired performance characteristics. Three key parameters were identified as significantly influencing the performance of the metaheuristic based on the literature: the removal fraction $[q_{\min}, q_{\max}]$, initial temperature scaling factor β and the geometric cooling rate α. The remaining parameters followed recommendations from Ropke and Pisinger (2006) and R. Liu et al. (2019). Three levels were defined for each tuning factor, as shown in Table 4. Using three levels (low, medium, and high) enables capture of non-linear effects of each parameter on performance. With three factors and three levels, a full factorial design 3^3 was adopted, resulting in 27 experiments.

Table 4: Experimental design and range analysis for ALNS parameter tuning.

Factors and levels				**Range analysis by influence**				
Param.	L1	L2	L3	Param.	L1	L2	L3	Range
$[q_{\min}, q_{\max}]$	[0.10; 0.20]	[0.15; 0.25]	[0.25; 0.35]	$[q_{\min}, q_{\max}]$	26.228	26.143	26.057	**0.171**
β	0.050	0.072	0.100	α	26.168	26.129	26.130	0.039
α	0.995	0.998	0.999	β	26.149	26.136	26.143	0.013

Each configuration was evaluated across the four scenarios, considering **medium** ($N = 100$) and **large** ($N = 150$) instances, replicates 0–2 and the four depots. To account for the stochastic nature of the algorithm, each combination was run with 5 independent seeds, totalling $27 \times 4 \times 2 \times 3 \times 4 \times 5 = \mathbf{12{,}960}$ ALNS runs with $I = 2000$ iterations each, amounting to approximately 25.9×10^6 iterations in total. The number of iterations, $I = 2000$, is justified separately in Section 7.3, based on the convergence behaviour of the ALNS. For configuration c and seed s, the response variable was computed as the objective per patient, with the total objective across all 96 depot-specific runs divided by the total number of patients.

Table B1 shows the aggregated results for each configuration. Configuration C23 achieved the lowest mean value per patient and was therefore selected. The standard deviation was used only as a measure of stochastic variability, not as a classification criterion. This is relevant because C22 had slightly lower stochastic variability and C26–C27 showed similar variability, none improved the mean objective and both C26 and C27 required substantially longer runtimes.

Tables B2 and B3 break the results down by scenario and dimension, grouped by removal fraction, with each value averaged over the three levels of β and α, which have lower influence. The time column reports the mean cumulative time per seed for the 12 corresponding depot-specific runs, averaged across the nine configurations. These tables show that the effect of the destruction fraction depends on the instance size. For **medium** instances, the differences between variants are marginal. For larger instances, the differences become more pronounced. The **high doctor** scenario is particularly notable: $q = [0.25; 0.35]$ yields *Avg* = 28.18 and $q =$

[0.10; 0.20] yields 28.98, creating a gap of 2.8%.

Among the three factors included in the experimental design, the range analysis in Table 4 shows that the removal fraction interval is the most influential. Increasing q from $[0.10, 0.20]$ to $[0.25, 0.35]$ reduced the mean objective per patient from 26.228 to 26.057. In contrast, the ranges associated with α and β were only 0.039 and 0.013, respectively. However, this improvement was accompanied by an increase in computational effort. Configuration C15 provides a relevant runtime reduction compared to C23, from 495.8 to 272.5 seconds, while increasing the mean objective per patient by only 0.178%, from 26.045 to 26.091. Nevertheless, C23 was retained because solution quality was defined as the primary selection criterion.

Table 5 summarises the parameters used in the following experiments.

Table 5: Complete configuration of ALNS parameters.

Parameter	Symbol	Value
Removal fraction	$[q_{\min}, q_{\max}]$	[0.25; 0.35]
Initial temperature scaling factor	β	0.072
Cooling rate	α	0.998
Parameters in adaptive mechanism	$\sigma_1, \sigma_2, \sigma_3$	33, 13, 9
Reaction factor	λ	0.1
Segment length		100
Shaw removal	p	6.0
Shaw relatedness weights	$\phi_{\text{time}}, \phi_{\text{urg}}$	1.0; 2.0
Iterations	I	2000
Minimum temperature	$T_{\min}$	10^{-6}
Max repair attempts		8

7 Numerical Experiments

This chapter evaluates the ALNS in terms of solution quality and computational time. The experiments use the instances described in Chapter 6 to compare the calibrated ALNS with the MILP formulation and to analyse the behaviour of the main ALNS components. The algorithm was implemented in Python 3.11 and the MILP problem was solved using Gurobi 11. All experiments were conducted on a Dell Precision T7500 workstation equipped with two Intel Xeon X5680 CPUs @ 3.33 GHz (12 physical cores, 24 threads in total) and 24 GB of DDR3 RAM, running Windows 10 Pro (64-bit).

7.1 Scalability of the MILP Formulation

This section identifies the largest instance size solvable to optimality within the time limit. The study focuses on the Lisboa depot and considers cohorts with $N = 10$ to 100 patients. A nested sampling strategy is adopted to ensure consistent comparisons across dimensions: the set for $N = 20$ is a strict superset of the set for $N = 10$, etc. This ensures that differences across dimensions reflect problem scale rather than random patient variation. Two solvers are used:

- **Gurobi**, which uses a callback mechanism to measure both the time to the first feasible solution and the time to the best solution;
- **CBC** (open-source solver, via PuLP), included to assess the accessibility of the formulation without commercial licences.

Both solvers have a 30-minute time limit per instance. This reflects the operational context, in which daily routes must be produced within a short planning window, making solution times of several hours impractical. It also defines a common computational budget across instance sizes, allowing the scalability of the MILP formulation to be assessed consistently.

Table 6 reports, for each value of N and both solvers, the active teams, first-feasible-solution statistics and final runtime and optimality gap.

Gurobi proved optimality for instances up to $N = 30$. From $N = 40$ onwards, it reached the imposed 30-minute limit, but the final optimality gaps remained small, generally at or below 6%. The instance with $N = 90$ is an exception, with a final gap of 14.3%, likely attributable to the specific structure of the sample for this dimension, which may have resulted in a more demanding combinatorial analysis for the branch-and-bound process. The time to the first feasible solution increased with instance size, from 0.2 s for $N = 10$ to 46 s for $N = 100$, confirming that Gurobi can quickly identify an initial feasible assignment even for the largest instances tested.

The open-source solver performed significantly worse, hitting the time limit as early as $N = 20$ and subsequently reaching gaps of 100% for $N \geq 70$. For $N = 10$, CBC also proved optimality, matching Gurobi's number of active teams. For $N = 20$, CBC found a solution with the same number of active teams as Gurobi, but with a residual gap of 5%. From $N = 30$ onwards, the number of active teams found by CBC began to diverge from Gurobi's, and residual gaps increased substantially, reaching 35% at $N = 50$.

Two aspects are worth clarifying separately here. First, team activation (y_k) is a decision variable of the model, so any feasible incumbent found before optimality is proven simply reflects the state each solver's search had reached when the time limit was met, not a structural difference between the formulations. Only doctor teams are penalised ($W_{\text{doc}} = 200$), so the model has a clear incentive to use few doctors which is why the two solvers almost always agree on this point. Because nurse teams have no activation cost, many nurse configurations yield equivalent objectives. Solver differences therefore reflect search behaviour rather than model structure, with the commercial Gurobi exploring this region more effectively than the open-source CBC.

The second is that from $N = 70$ onward, CBC simply cannot push its lower bound anywhere close to its best solution within the time limit, so the gap tops out at 100%. This reflects a weak lower bound on CBC's side, not necessarily a poor-quality solution.

Table 6: Scalability results of the MILP formulation.

	Gurobi					CBC (PuLP)		
		First feasible		Best solution			Best solution	
N	Teams	Time (s)	Gap (%)	Time (s)	Gap (%)	Teams	Time (s)	Gap (%)
10	2	0.2	81.4	0.3	0.0	2	40.2	0.0
20	4	1.4	35.3	187.7	0.0	4	1799.5	5.0
30	5	2.6	41.0	149.9	0.0	6	1798.8	13.0
40	6	3.8	100.0	1800.0	0.2	7	1798.9	17.0
50	8	5.4	100.0	1800.0	0.6	10	1798.6	35.0
60	8	7.4	66.3	1800.0	1.4	9	1798.1	16.0
70	10	11.4	62.9	1800.0	1.4	8	1797.0	100.0
80	11	20.4	42.2	1800.0	3.8	10	1797.1	100.0
90	13	20.9	52.6	1800.0	14.3	11	1800.0	100.0
100	15	46.1	54.3	1800.0	6.0	12	1800.0	100.0

7.2 Solution Quality and Computational Performance of ALNS

Comparison with exact method

This section compares ALNS solution quality with the exact model. Each instance was solved in two ways: using the MILP model with Gurobi 11 with a time limit of 30 minutes per instance, justified by the results in Section 7.1, and using the ALNS with the calibrated configuration shown in Table 5. Since the ALNS is stochastic, with randomness present in the choice of the destroy and repair operators and in the acceptance criterion, each instance was solved using five independent seeds. In turn, its quality is summarised by the average objective across seeds.

For each instance, the relative gap between the two methods is defined as

$$\text{gap} \;=\; 100 \times \frac{\bar{z}_{\text{ALNS}} - z_{\text{MILP}}}{z_{\text{MILP}}} \quad (\%), \tag{27}$$

where $\bar{z}_{\text{ALNS}}$ is the average objective of the ALNS across the five seeds and z_{MILP} is the objective of the best solution returned by Gurobi within the time limit.

There may be two outcomes: a positive gap, meaning that the ALNS solution is more expensive than the model, and a negative gap, meaning that the ALNS achieves a lower-cost solution.

When Gurobi finishes proving optimality, z_{MILP} is the lowest cost amongst all feasible solutions to that instance and, by the definition of optimality, the gap cannot be negative. When Gurobi reaches the time limit without proving optimality, z_{MILP} is merely the best incumbent found, that is, an upper bound on the optimum; therefore, the ALNS may legitimately achieve a lower-cost solution, resulting in a negative gap.

The comparison is carried out independently for each of the four instance-generation scenarios described in Chapter 6: **balanced**, **clustered**, **high urgent** and **high doctor**, throughout the following chapters. The **balanced** scenario, which preserves the original proportions of patients (Table 2), constitutes the baseline case and is therefore analysed first.

As described in Chapter 6, each scenario contributes 160 depot-level instances, so the benchmark covers 640 depot-level instances across the four scenarios. This corresponds to 3,200 ALNS runs, each of the 640 instances being evaluated with five seeds.

Balanced scenario

The **balanced** scenario consists of 160 instances (four depots $\times$ five dimensions $\times$ eight replicates). Within the 30-minute time limit, Gurobi certified optimality in 112 of the 160 instances (70.0%) and reached the time limit in the remaining 48 (30.0%), with the latter becoming increasingly frequent as the dimension increases. Table 7 summarises, by dimension, the average MILP and ALNS objective values, the number of certified optima, the gap statistics according to (Eq. (27)) and the average running time per instance for both methods.

Across dimensions, the objective values show that the ALNS remains very close to the Gurobi reference. In the **small** and **medium** dimensions, the average MILP and ALNS objectives are almost identical, increasing from 426.3 and 426.4 to 613.0 and 613.3, respectively, with all 64 instances remaining within 0.6% of the Gurobi benchmark. From the **large** dimension onwards, the number of optima certified by Gurobi decreases and the minimum gap becomes negative, indicating that the ALNS sometimes improves upon the incumbent found by Gurobi within the time limit. This is particularly visible in the **xxlarge** dimension, where the average ALNS objective is 1295.5, below the average Gurobi objective of 1299.1.

Table 7: ALNS VS Gurobi in the balanced scenario, by instance dimension.

		Optimal	Objective		Gap (%)			Time (s)
Instance size	N_{inst}	certificates	z_{MILP}	$\bar{z}_{\text{ALNS}}$	Mean	Min	Max	ALNS / Gurobi
Small	32	32	426.3	426.4	+0.027	0.000	+0.511	1.9 / 6.4
Medium	32	29	613.0	613.3	+0.029	0.000	+0.471	6.7 / 255.3
Large	32	21	827.9	827.6	+0.007	−2.581	+1.173	16.0 / 765.4
XLarge	32	16	1112.8	1112.2	−0.012	−0.658	+0.287	31.9 / 1054.9
XXLarge	32	14	1299.1	1295.5	−0.094	−1.954	+0.679	55.4 / 1118.4

Across all instances, the ALNS matches the Gurobi reference objective within a margin of 0.01% in 76 cases (47.5%) and is within less than 1% in 154 cases (96.2%). The overall average gap is -0.008%.

In the 112 instances where Gurobi proved optimality, the average gap is just 0.053% and in 73 of these instances (65.2%), the ALNS solution is within 0.01% of the certified optimum. This indicates that whenever there is a proven optimal solution, the ALNS either achieves it or comes within a very small margin of it. Gurobi solved all 32 **small** instances and 29 of the 32 **medium** instances, whilst its success rate dropped to 21, 16 and 14 out of 32 in the **large**, **xlarge** and **xxlarge** dimensions, respectively, as the combinatorial search could not be closed within the 30-minute time limit.

In the 48 instances where Gurobi reached the time limit, the average ALNS solution matched or improved the incumbent in 20 cases. When considering the best solution obtained across the five independent ALNS seeds, this number increased to 34 cases, corresponding to 70.8%. For these 48 time-limited instances, Gurobi's own remaining optimality gap averaged 3.59%. The advantage in solution quality of the ALNS is therefore concentrated in the larger and computationally more difficult instances where the MILP solver does not close the optimality gap within the imposed time limit.

A particularly relevant feature is the contrast in computational time at the instance level. In the balanced scenario, the ALNS requires from 1.9 s in **small** instances to 55.4 s in **xxlarge** instances on average. This remains well below the average runtime of Gurobi, which increases from 6.4 s in **small** instances to 1118.4 s in **xxlarge** instances.

The depot-level summary reported in Table C1 shows that the largest differences occur in Lisboa, where the average ALNS objective is 1427.1 compared with 1432.3 for Gurobi. This is consistent with Lisboa being the largest and most computationally demanding depot, with only 13 of the 40 Lisboa instances certified as optimal. In addition, the largest average improvement occurs in a **large** Lisboa instance, where the average ALNS gap is -2.581% and the best seed reaches -12.234%. By contrast, Odivelas was solved to certified optimality in all 40 instances and the ALNS reproduced the same objective in most cases, with an average gap of only 0.012%. In addition, the runtime results reinforce this conclusion. The average runtime of Gurobi ranges from a fraction of a second in Odivelas to more than 21 minutes in Lisboa, whereas the ALNS stays close to one minute for Lisboa. This is operationally significant, since Lisboa concentrates the highest share of patients and is therefore the depot where manual morning planning is expected to be most time-consuming.

Finally, the ALNS results are stable across seeds. All 160 instances produced a feasible solution for all five seeds and the average coefficient of variation of the objective across seeds is only 0.11%. The complete per-instance results supporting the scenario-level comparisons are provided in Table D1 in Appendix.

Clustered scenario

The **clustered** scenario, which selects geographically close patients, tests how the presence of many near-equivalent routing alternatives affects the comparison. It also consists of 160 instances. Within the 30-minute time limit, Gurobi certified optimality in 105 of the 160 instances (65.6%) and reached the time limit in the remaining 55 (34.4%). Table 8 summarises the results by dimension.

Table 8: ALNS VS Gurobi in the clustered scenario, by instance dimension.

		Optimal	Objective		Gap (%)			Time (s)
Instance size	N_{inst}	certificates	z_{MILP}	$\bar{z}_{\text{ALNS}}$	Mean	Min	Max	ALNS / Gurobi
Small	32	31	367.8	367.8	+0.002	0.000	+0.054	1.9 / 66.0
Medium	32	26	577.6	577.8	+0.022	−0.130	+0.424	6.9 / 405.8
Large	32	20	772.2	771.4	−0.042	−2.891	+1.196	16.3 / 829.5
XLarge	32	14	1069.1	1065.4	−0.187	−7.140	+0.388	31.9 / 1055.8
XXLarge	32	14	1250.6	1247.7	−0.096	−1.807	+0.702	55.5 / 1166.0

Across dimensions, the objective values confirm that the ALNS stays very close to the Gurobi reference. In the **small** and **medium** dimensions the two methods are practically identical (367.8 VS 367.8 and 577.6 VS 577.8). The average ALNS objective falls below the Gurobi reference already from the **large** dimension onwards (771.4 against 772.2), and the largest absolute difference occurs in the **xlarge** dimension (1065.4 against 1069.1), precisely the dimension that contains the most favourable outlier of the scenario.

The overall average gap is −0.060%. In the 105 instances where Gurobi certifies optimality, the average ALNS gap relative to the certified optimum is only 0.022%. In 83 of these instances (79.0%), the ALNS solution is within 0.01% of the certified optimum, including 75 instances where the gap is exactly zero. As in the **balanced** scenario, certified optima are concentrated in the smaller sizes, with the certification rate dropping from 31 out of 32 in the **small** dimension to 14 out of 32 in both the **xlarge** and **xxlarge** dimensions.

In the 55 instances that reached the time limit, the average ALNS gap across the five seeds was −0.217%. When considering the best solution obtained from the five independent seeds, the ALNS matched or outperformed the incumbent in 35 cases (63.6%), with an average best-seed gap of −0.618%. For these 55 time-limited instances, Gurobi's own remaining optimality gap averaged 3.99%. It is worth noting that, in this scenario, Gurobi certifies fewer instances from Lisboa (9 out of 40, compared to 13 in the **balanced** scenario). One possible explanation is that the presence of several near-equivalent routing alternatives makes it harder for the branch-and-bound process to close the optimality gap within the time limit.

The largest advantage occurs in an **xlarge** instance in Lisboa. The average ALNS gap is -7.140% and the best ALNS seed improves the Gurobi incumbent by **11.626%**. There is also an outlier in a **large** instance in the same depot, with 69 patients, where the average gap is -2.891% and the best seed reaches -12.041%. The depot-level summary in Table C1 confirms that Lisboa is the main source of computational difficulty in this scenario, while the smaller depots remain essentially tied.

The ALNS requires from 1.9 s in **small** instances to 55.5 s in **xxlarge** instances on average. By comparison, Gurobi requires 66.0 s on average in **small** instances and 1166.0 s in **xxlarge** instances. All 160 instances produced feasible solutions for all five seeds, with an average coefficient of variation for the objective of only 0.13%. The complete per-instance results supporting the scenario-level comparisons are provided in Table D2 in Appendix.

High urgent scenario

In the **high urgent** scenario, the aim is to test the behaviour of the ALNS under greater pressure from the priority policy, where the order is urgent first and the limit is M urgent patients per route. It comprises 160 instances. Within the 30-minute time limit, Gurobi certified optimality in 125 of the 160 instances (78.1%), the highest rate of the four scenarios, and reached the time limit only in the remaining 35 (21.9%). Table 9 summarises the results by dimension.

Table 9: ALNS VS Gurobi in the high urgent scenario, by instance dimension.

		Optimal	Objective		Gap (%)			Time (s)
Instance size	N_{inst}	certificates	z_{MILP}	$\bar{z}_{\text{ALNS}}$	Mean	Min	Max	ALNS / Gurobi
Small	32	32	435.5	435.5	+0.004	0.000	+0.076	1.9 / 2.4
Medium	32	32	740.2	740.4	+0.020	0.000	+0.181	6.6 / 93.3
Large	32	26	1121.0	1122.8	+0.108	0.000	+0.560	15.1 / 579.8
XLarge	32	24	1402.5	1404.2	+0.185	−0.095	+1.461	29.0 / 526.3
XXLarge	32	11	1465.9	1461.0	−0.170	−1.568	+0.298	55.7 / 1298.7

Across all instances, the ALNS matches the Gurobi reference objective within a margin of 0.01% in 83 cases (51.9%) and remains within 1% in 154 cases (96.2%). The overall average gap is $+0.030\%$. The objective values in Table 9 confirm this close agreement. In the **small** and **medium** dimensions, the average MILP and ALNS objectives are practically identical, while the differences remain small in the **large** and **xlarge** dimensions. The only dimension where the average ALNS objective is lower than the Gurobi reference is the **xxlarge** dimension, with 1461.0 compared with 1465.9.

What sets this scenario apart is the higher certification rate obtained by Gurobi. The MILP certified optimality in 125 out of 160 instances, which is the highest value among the four scenarios. One possible explanation is that, although the urgent-first rule applies in all scenarios, in this scenario it affects a larger proportion of patients. As a result, more routing decisions are constrained by the priority order of urgent patients and by the maximum limit on the number of urgent patients per route. This may reduce the effective sequencing flexibility and help Gurobi certify optimality in a larger proportion of instances.

In the 35 instances where Gurobi reached the time limit, the average ALNS solution was equal to or better than the Gurobi incumbent in 12 cases. When considering the best solution obtained from the five independent seeds, this number rose to 21 cases. The average gap in this subset is -0.101%, with the largest average improvement reaching 1.568%. For these 35 time-limited instances, Gurobi's own remaining optimality gap averaged 2.32%.

In line with the general pattern observed in the other scenarios, the largest improvements achieved by the metaheuristic occur in the **xxlarge** instances for **Lisboa**, where Gurobi reached the time limit. In these cases, the average ALNS gaps range from approximately -1.304% to -1.568%. The depot-level results in Table C1 show the same structure: Odivelas was solved to optimality in all 40 instances, whereas Lisboa was again the only depot large enough for the exact method's certification to fail frequently.

In terms of runtime, the ALNS ranges from around 1.9 s to 55.7 s on average across instance sizes. Gurobi, by contrast, increases from 2.4 s in **small** instances to 1298.7 s in **xxlarge** instances. The results are stable across seeds, with an average coefficient of variation for the objective of just 0.05%, which is the lowest of the four scenarios. The complete per-instance results supporting the scenario-level comparisons are provided in Table D3 in Appendix.

High doctor scenario

The **high doctor** scenario was defined to evaluate the behaviour of the methods when there is a higher proportion of patients requiring teams with a doctor and, consequently, to place greater emphasis on the penalty for activating teams of doctors, W_{doc}, in the objective function.

Table 10 presents the aggregated results by instance size. Across the 160 instances, the ALNS matches the reference within 0.01% in 88 instances and remains within 1% in 152 instances (95.0%). The overall average gap is $+0.027\%$, indicating that the average performance of ALNS is almost identical to the Gurobi reference.

The main difficulty for Gurobi appears in the larger dimensions. All **small** instances were solved to optimality, but only 16 of the 32 **xlarge** instances and 14 of the 32 **xxlarge** instances were certified. Therefore, as in the previous scenarios, the certification ability of the MILP decreases as instance size increases.

Table 10: ALNS VS Gurobi in the high doctor scenario, by instance dimension.

		Optimal	Objective		Gap (%)			Time (s)
Instance size	N_{inst}	certificates	z_{MILP}	$\bar{z}_{\text{ALNS}}$	Mean	Min	Max	ALNS / Gurobi
Small	32	32	475.7	475.8	+0.011	0.000	+0.361	1.9 / 4.0
Medium	32	26	775.4	775.6	+0.019	0.000	+0.151	6.6 / 351.7
Large	32	24	1067.9	1065.5	−0.064	−10.234	+11.021	15.5 / 520.9
XLarge	32	16	1372.2	1377.5	+0.360	−0.881	+9.105	30.5 / 909.4
XXLarge	32	14	1500.2	1494.8	−0.191	−3.155	+0.225	54.2 / 1122.5

Among the 112 instances where Gurobi proved optimality, the average ALNS gap is +0.032%, with 82 cases (73.2%) within 0.01% of the certified optimum. In the 48 instances where Gurobi reached the time limit, the average ALNS gap is +0.016%. When considering the best solution obtained from the five independent seeds, the ALNS matched or improved the Gurobi incumbent in 31 of these 48 cases. For these 48 time-limited instances, Gurobi's own remaining optimality gap averaged 7.04%.

However, this scenario exhibits greater variability, especially in the **large** and **xlarge** instances. The minimum and maximum gaps reach −10.234% and +11.021%, both in the **large** dimension. One reason is that, whenever the ALNS stabilises on a solution that uses one more doctor than the exact method, the gap adds around 200 cost units, regardless of the efficiency of the underlying routes.

The depot-level summary in Table C1 shows that this additional variability is mainly associated with Lisboa and Cascais. Lisboa is the only depot with a negative average gap, with an average ALNS objective of 1770.8 compared with 1776.9 for Gurobi. Cascais contains the largest positive outlier (+9.105%), associated with the maximum gap observed in the **xlarge** instances. This supports the interpretation that the activation of one additional doctor team can have a visible effect on the gap, especially in larger depots.

In computational terms, the ALNS remains faster on average on a per-instance basis. It requires from 1.9 s in **small** instances to 54.2 s in **xxlarge** instances on average. Gurobi, in contrast, requires 4.0 s on average in **small** instances and 1122.5 s in **xxlarge** instances. Across seeds, despite the outliers, the average coefficient of variation remains low, at 0.17%, although it is the highest among the four scenarios.

Complete per-instance results are provided in Table D4 in Appendix. Because this scenario had the strongest outliers, an additional robustness analysis increased the ALNS seeds from 5 to 20 for the **large**, **xlarge** and **xxlarge** dimensions with the results in Table 11.

Table 11: ALNS gap statistics for 5 and 20 seeds (high doctor scenario).

	5 seeds			20 seeds		
Instance size	Mean	Min	Max	Mean	Min	Max
Large	−0.064%	−10.234%	+11.021%	−0.113%	−10.110%	+11.027%
XLarge	+0.360%	−0.881%	+9.105%	+0.266%	−1.064%	+7.762%
XXLarge	−0.191%	−3.155%	+0.225%	−0.191%	−3.203%	+0.288%
Total	+0.035%	−10.234%	+11.021%	−0.013%	−10.110%	+11.027%

The average gap in the larger dimension instances decreased from 0.035% with five seeds to -0.013% with 20 seeds.

In the Lisboa **large** instance of replicate 7 (69 patients; Table D4), the twenty seeds converge to solutions with eleven active teams and the best gap remains at approximately +10.9%, suggesting a structural difficulty in this specific instance.

In contrast, other cases of high gaps prove to be primarily the result of stochastic variability. In the Cascais **xlarge** instance of replicate 0 (56 patients; Table D4), where the average gap with five seeds was +9.1%, the best of the twenty seeds achieves a gap of 0.000%. In this instance, the solutions obtained range from eight to nine active teams.

7.3 ALNS Search Behaviour

Whilst Section 7.2 evaluated the quality of the solutions produced by the ALNS, this section characterises how the adaptive operator selection and the SA acceptance criterion presented in Section 5.2 behave.

When tested on the same instances as in the previous section, the number of failed repairs and skipped iterations was zero in all runs.

Operator selection

The RWS weights are uniformly initialised to 1.0 and are updated at the end of each segment of 100 iterations towards the average score obtained per use, rewarding moves that produce a new best solution (σ_1), an accepted improvement over the current solution (σ_2) or an accepted non-improving move (σ_3). Table 12 shows, for each operator, the average selection frequency and the average final weight across the same 3,200 ALNS runs used in the scenario-level comparison of Section 7.2.

The three destroy operators are selected with almost identical frequencies (32–34%) and their final weights range between 6.53 and 7.20. Both repair operators are equally close.

Table 12: Selection frequency and final weight of operators (average over 3,200 runs of the ALNS).

Operator	Selection frequency (%)	Final weight
Random removal	34.0	7.20
Worst removal	33.7	6.72
Shaw removal	32.3	6.53
Greedy repair	48.2	6.31
Regret-2 repair	51.8	7.33

At the end of each segment, an operator's weight is shifted towards the average score it achieved per use (Eq. (26)), and this score is determined by the result with σ_1, σ_2, σ_3 or zero (rejection). As the final weights are so close to one another, all operators necessarily accumulate a similar average reward, so none is consistently more productive than the others. Thus, similar weights result in an almost uniform selection.

Regret-2 receives a higher weight than greedy repair in 87.6% of runs, while shaw removal is the least rewarded destroy operator.

It is also observed that as N increases, regret-2 rises from 50.5% in the smallest dimension to 52.9% in the largest, whilst worst removal increases from 33.3% to 34.4% and shaw removal decreases from 33.1% to 31.4%. Finally, the selection pattern and the ordering of the weights remain essentially unchanged across the four scenarios, regardless of the instance type.

Acceptance behaviour

Table 13 shows acceptance behaviour by instance size, aggregated across all runs.

Table 13: Acceptance behaviour by instance dimension (average across ALNS runs, $I = 2000$ per run).

Size	Accepted (%)	Improved (%)	Same (%)	Worse (%)	Rejected (%)	New best solutions
Small	96.1	4.8	87.3	3.9	3.9	3.6
Medium	88.3	14.9	61.5	12.0	11.7	9.6
Large	81.0	20.4	44.1	16.5	19.0	14.7
XLarge	78.4	25.3	32.3	20.7	21.6	17.3
XXLarge	75.0	26.4	26.6	22.0	25.0	19.3
Global	83.8	18.3	50.4	15.0	16.2	12.9

Intensification dominates the search. Overall, 83.8% of candidates are accepted, but only 15.0% of all decisions correspond to accepted worse moves. Around half of all decisions, 50.4%, are changes that do not change the objective value. In the smallest instances, cost-neutral moves account for 87.3% of decisions, partly because several teams are interchangeable. For example, assigning two equivalent nurse teams to swapped routes changes the solution representation but not the objective value. In larger instances, this proportion falls to 26.6%, as the search space contains more genuinely different routing alternatives. The same pattern runs throughout Table 13: as N increases, moves of equal cost give way to real improvements

(4.8% $\rightarrow$ 26.4%) and to accepted worse moves (3.9% $\rightarrow$ 22.0%), and the number of new local optima rises from 3.6 to 19.3 per run.

The data suggest that the simulated annealing acceptance rule is appropriately scaled across instance sizes. Regardless of the instance size, the SA accepts roughly half of the worse moves it encounters. This is because the initial temperature is proportional to the objective of the initial solution, $T_0 = \max(T_{\min}, \beta \cdot f(S_0))$ (Eq. (22)).

The same pattern reappears when looking at each depot (Table 14), given that they differ mainly in the number of patients.

Table 14: Acceptance behaviour by depot (average across ALNS runs).

Depot	N_d	Accepted (%)	Worse accepted (%)	New best solutions
Lisboa	23-116	73.0	22.6	20.8
Cascais	14-70	80.2	19.9	16.7
Sintra	9-44	85.8	13.7	10.8
Odivelas	4-20	96.0	3.9	3.2

Lisboa is where the search explores the most, accepting 22.6% of worse moves and finding the highest number of new best solutions (around 21 per run). Odivelas, the smallest, behaves with almost no exploration.

Ablation study

In order to assess the individual contribution of the main operators and components of the ALNS, an ablation study was carried out on the same set of four scenarios. Unlike the benchmark using Gurobi, where each ALNS solution is compared directly with the exact model's solution for the same instance using the total objective, the ablation study aggregates instances of different dimensions. Thus, the results are presented in terms of objective per patient and travel time per patient, to prevent larger instances from dominating the aggregated averages.

The seven tests cover: the full ALNS (ALNS full); the ALNS without worst removal (ALNS no worst), without shaw removal (ALNS no shaw), without simulated annealing (ALNS no SA) and without regret-2 (ALNS no regret-2); the random-only variant that keeps only random removal and greedy repair (ALNS random-only); and the constructive greedy heuristic (Greedy). The random-only variant represents the minimal configuration of the ALNS, as the algorithm always requires at least one destruction operator and one repair operator to generate new candidate solutions. Accordingly, 4 scenarios, 5 dimensions, 5 replicates, 4 depots and 5 random seeds were tested using the 7 methods. All runs concluded with feasible solutions. Table 15 presents the average results by method.

Table 15: Ablation study results by method.

Method	Runs	Obj./patient	Doctors used	Travel/patient	Runtime (s)
ALNS full	2000	29.490	3.219	8.646	22.001
ALNS no worst	2000	29.493	3.219	8.648	22.754
ALNS no shaw	2000	29.534	3.228	8.649	21.206
ALNS no SA	2000	29.535	3.226	8.661	26.727
ALNS no regret-2	2000	29.539	3.222	8.679	4.780
ALNS random-only	2000	29.564	3.232	8.670	4.014
Greedy	2000	31.779	3.338	10.438	0.004

The full ALNS achieves the best average objective per patient, at 29.490, although the other ALNS variants remain close. Removing worst removal has almost no effect, increasing the objective by only 0.003 units. The variants without simulated annealing, without shaw removal and without regret-2 show slightly higher losses, but these are still small. From a computational perspective, the regret-2 operator is the most expensive component: removing it reduces the average runtime from 22.001 s to 4.780 s, while increasing the objective from 29.490 to 29.539. This suggests that omitting regret-2 may be reasonable when computation time is prioritised over a very small loss in solution quality. The random-only variant is faster still, but gives the poorest solution quality among the ALNS configurations.

The ablation suggests that no single component drives most of the performance. This conclusion is consistent with the previous analysis of operator selection, in which selection frequencies and final weights remained relatively balanced. The removal of individual operators leads only to small average losses, indicating that the operators are partially complementary. However, the full ALNS continues to deliver the best average performance and the random-only configuration, whilst faster, is the one with the poorest quality among the ALNS variants.

The greedy heuristic performs substantially worse. Its objective per patient is 31.78, compared with 29.49 for the full ALNS, corresponding to a 7.20% reduction with ALNS. Travel time per patient falls from 10.44 to 8.65 minutes, a 17.17% reduction. In terms of runtime, the greedy heuristic is practically instantaneous, as it merely constructs an initial solution and does not carry out any iterative optimisation process.

Table 16: Improvement of the full ALNS over the greedy heuristic by scenario.

Scenario	Greedy obj./patient	ALNS obj./patient	Obj. reduction	Greedy travel/patient	ALNS travel/patient	Travel reduction
Balanced	30.056	27.565	8.29%	10.633	8.693	18.24%
Clustered	28.365	26.106	7.97%	9.197	7.607	17.29%
High doctor	34.687	32.168	7.26%	10.722	8.970	16.34%
High urgent	34.009	32.121	5.55%	11.201	9.314	16.85%

The improvement is consistent across scenarios: the full ALNS reduces the average objective by 8.29% in the balanced scenario, 7.97% in the clustered scenario, 7.26% in the high doctor scenario and 5.55% in the high urgent scenario. Travel-time reductions range from 16.34% to 18.24%, as shown in Table 16.

Convergence behaviour

Figure E1 in Appendix shows the average convergence curves for the six ALNS variants considered in the ablation study, with the objective normalised by the value of the initial solution.

Initially, the objective decreases sharply, with most improvement achieved within the first few hundred iterations. From then on, only marginal improvements occur. The choice of $I = 2000$ aligns this behaviour with SA cooling. According to the geometric cooling rule (Eq. (23)), the temperature after I iterations is approximately $T_0 \cdot \alpha^I$. With $\alpha = 0.998$, this corresponds to 1.82% of T_0 at iteration 2000. Terminating the search much earlier would leave the SA in a still exploratory phase, whilst extending it beyond 2000 would merely increase the computational cost without any noticeable gain. In addition, with a segment length of 100 iterations, the search performs 20 adaptive operator-weight updates.

8 Prototype

To demonstrate the practical applicability of the ALNS algorithm, a prototype decision-support tool was developed. It serves as a functional proof of concept, illustrating how the algorithm's results can be communicated to operational teams intuitively and immediately.

8.1 Prototype Workflow

The prototype follows a data-to-dashboard workflow.

1. **Input data preparation:** Patient and depot data are first read from Excel files. The patient file contains the information required to build the optimisation instances: patient identifier, coordinates, service time, urgency indicator, doctor requirement and assigned depot. The depot file contains each depot's coordinates and team composition.
2. **Travel-time matrix construction:** Travel-time and distance matrices are computed using OpenRouteService. A preliminary depot-to-patient matrix identifies the nearest depot for unassigned patients. Then, for each depot, a local matrix is built with the depot and the patients assigned to it.
3. **Optimisation:** Each depot instance is solved independently with the ALNS metaheuristic, producing the assignment of patients to teams and the visiting sequence of each route.
4. **Output generation and dashboard visualisation:** After the optimisation is completed, the best solution found for each depot is exported to a JSON file. This file contains the

depot coordinates, objective value, active teams, route sequences and patient attributes. The dashboard then reads this output and generates an interactive map with the routes, depot information and patient-level details.

8.2 Dashboard Structure and Features

The prototype offers several features. The side panel is organised into two vertical sections:

- Depot selection, displaying four Key Performance Indicators (KPIs): patients, active teams, urgent patients and patients requiring a doctor. When selected, the map centres on the corresponding area and plots the routes.
- List of active teams in the selected depot. Each row shows the team name, type, route colour and number of patients. Clicking on a team focuses the view.

On the map, the depot is represented by a circular marker, labelled D. Clicking on a patient marker displays a pop-up with detailed information: patient ID, assigned team, order number on the route, service time (in minutes) and visual labels for urgency and need for a doctor.

8.3 Route Visualisation Example

To illustrate how the prototype works, the Lisboa depot was selected as the main example, due to its larger number of patients.

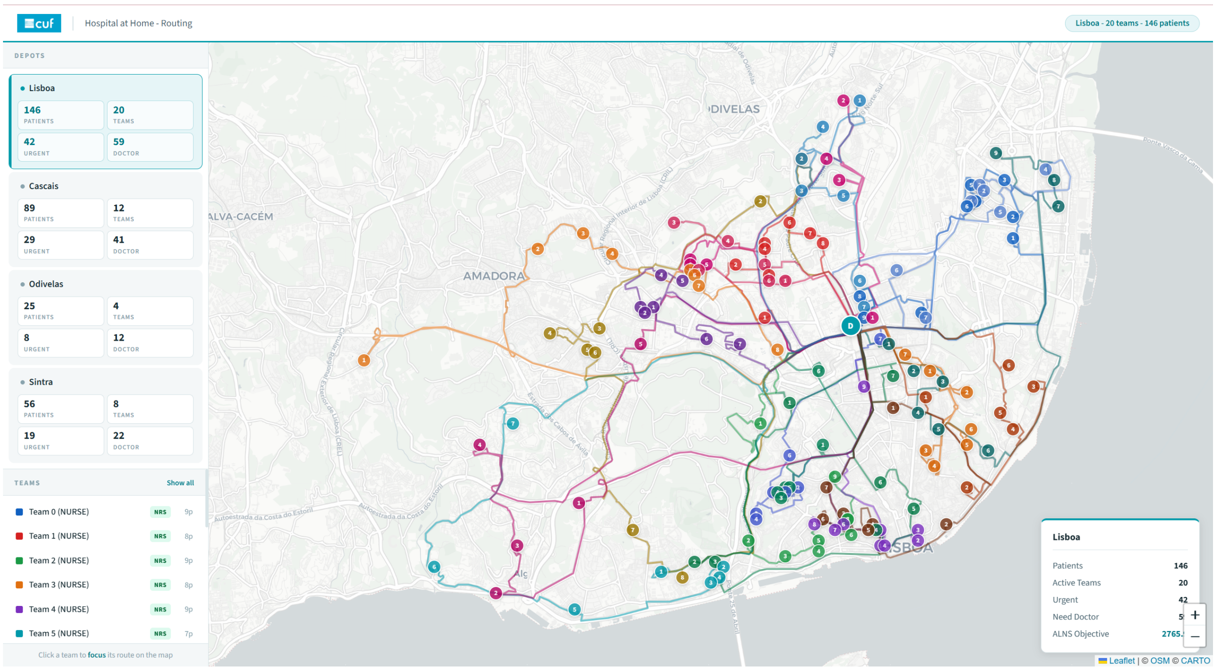

Figure 1: Lisboa routes.

Figure 1 shows an overview of the dashboard for the Lisboa depot. It displays the depot's KPIs, the list of active teams and the geographic distribution of the generated routes. Using different colours for each team makes it easier to visualise the data, allowing the user to assess whether the routes are geographically consistent with the assigned patients.

For a more detailed view, each team can be selected to inspect its visit sequence and spatial layout. This functionality is illustrated with Team 5 (nurse team) in the Lisboa depot, whose route is reported in Table F1 and Figure F1. The example shows that the two urgent patients are visited first and that all patients are assigned to a compatible nurse team, confirming that the dashboard can be used not only to visualise routes but also to verify relevant constraints.

Another feature is the pop-up associated with each patient. By clicking on each number on the route, the dashboard displays the patient ID and the corresponding requests. The first urgent patient on this route is shown in Figure F2.

Overall, the prototype acts not only as a route visualisation tool but also as an interpretation layer between the optimisation algorithm and the operational user. Additional dashboard outputs for the remaining depots are presented in Appendix F, showing that the same visualisation structure can be applied consistently across all depots.

9 Conclusion, limitations and future work

Population ageing and hospital capacity pressure make HaH a promising alternative to conventional hospitalisation, but daily team and route planning remains manual and time-consuming.

The existing HaH routing and scheduling literature does not include a metaheuristic approach or a systematic comparison between exact and approximate methods. In response to this gap, this dissertation formulated the problem by first assigning patients to depots, with each depot then solved independently, incorporating constraints such as the exclusivity of doctor and nurse teams, an urgent-first priority with a limit per route, maximum shift duration and a cost associated with deploying teams with a doctor. Two complementary approaches were developed: a MILP model and an ALNS metaheuristic, supported by a decision-support prototype for visualising solutions.

Based on a set of 640 instances derived from 316 patients, results show that the ALNS closely matches the exact method with overall average gaps between -0.060% and $+0.030\%$. It also demonstrated that the advantage of the ALNS is concentrated in the larger instances associated with the Lisboa depot, in which the exact method mostly reaches the 30-minute limit without proving optimality. Moreover, it does so at a fraction of the computational cost. Even when the five independent seeds are run sequentially, the ALNS requires only a few seconds to a few minutes, up to roughly 4.6 minutes in the largest instances, whereas Gurobi requires around 18-22 minutes on average in the largest dimensions and frequently reaches the 30-minute limit.

The ablation study revealed that the individual contribution of each operator is less than 0.5%. Even so, the full ALNS reduces the objective per patient by around 7% and the travel time by around 17% relative to the greedy construction, confirming that its value lies in the repeated reorganisation of routes rather than in any single operator.

These results directly address the proposed objective, showing that the exact approach is sufficient for smaller-scale instances and that the metaheuristic ensures the necessary scalability for larger-scale instances. In practical terms, this manual morning planning can instead be produced automatically in seconds to a few minutes, in a systematic and reproducible way, while guaranteeing by construction that the urgent-first policy and the cap on urgent patients per team are always respected.

The study has several limitations. Although the patients have real coordinates, these are still not the same as those in the actual hospital case. Furthermore, only a single morning shift is considered, failing to account for continuity of care across days or the dynamic arrival of new requests throughout the day. It should also be noted that, in instances where the ALNS outperforms the exact method, the comparison is made against the best incumbent result obtained within the time limit, whilst the optimal value remains unknown for both methods. The approach developed is, moreover, deterministic. Travel and service times are treated as known, although in practice they are subject to uncertainty. Nevertheless, the results obtained provide a solid basis for future extensions that explicitly incorporate uncertainty.

These limitations also indicate several directions for future work. Future extensions could incorporate multi-period planning, allowing continuity of care to be considered. This would make the approach closer to real HaH operations, where clinical follow-up is a relevant dimension of service quality. Second, the approach could be extended to account for uncertainty, either through stochastic travel and service times or through robust optimisation mechanisms that protect route feasibility under adverse conditions. Third, a dynamic version of the approach could update routes as new information becomes available, while preserving as much of the original plan as possible to avoid operational disruption. Finally, future work could further investigate the effect of equivalent teams on the search process. Since teams of the same type and depot are interchangeable, different mathematical solutions may correspond to the same operational plan and objective value. Addressing this symmetry, either in the exact formulation or in the solution representation used by the ALNS, could reduce redundant alternatives and improve computational efficiency.

References

AESE. (2026). Hospitalização domiciliária e o futuro dos cuidados de saúde em Portugal. Retrieved May 2, 2026, from https://www.aese.pt/noticias/hospitalizacao-domiciliaria-e-o-futuro-dos-cuidados-de-saude-em-portugal/

Al-Sorori, W., & Mohsen, A. M. (2020). New Caledonian crow learning algorithm: A new metaheuristic algorithm for solving continuous optimization problems. *Applied Soft Computing*, *92*, 106325. https://doi.org/10.1016/j.asoc.2020.106325

Alves, A. M. (2025, August). Estar no hospital "era uma escuridão". Em casa, há mais ânimo e profissionais de saúde cuidam "com gosto". Retrieved May 2, 2026, from https://www.publico.pt/2025/08/11/sociedade/reportagem/estar-hospital-escuridao-casa-ha-animo-profissionais-saude-cuidam-gosto-2142597

Atta, S., Basto-Fernandes, V., & Emmerich, M. (2025). A concise review of the home health care routing and scheduling problem. *Operations Research Perspectives*, *15*, 100347. https://doi.org/10.1016/j.orp.2025.100347

Bassoli, E., Lefebvre, M., & Schoenmaeckers, J. (2025). Home vs. nursing care: Unpacking the impact on health and well-being. *Social Science & Medicine*, *385*, 118533. https://doi.org/10.1016/j.socscimed.2025.118533

Bazirha, M., Benmansour, R., & Kadrani, A. (2023). An efficient two-phase heuristic for the home care routing and scheduling problem. *Computers & Industrial Engineering*, *181*, 109329. https://doi.org/10.1016/j.cie.2023.109329

Bazirha, M., Kadrani, A., & Benmansour, R. (2024). Pareto and decomposition based approaches for the multi-objective home health care routing and scheduling problem with lunch breaks. *Engineering Applications of Artificial Intelligence*, *128*, 107502. https://doi.org/10.1016/j.engappai.2023.107502

Belhor, M., El-Amraoui, A., Jemai, A., & Delmotte, F. (2023). Multi-objective evolutionary approach based on K-means clustering for home health care routing and scheduling problem. *Expert Systems with Applications*, *213*, 119035. https://doi.org/10.1016/j.eswa.2022.119035

Bonomi, V., Lopes, M., Santos, D. R., Mansini, R., & Barbosa-Póvoa, A. P. (2025). Allocation-routing problem in a multi-hospital home hospitalization system: The case of a healthcare provider in Portugal. *International Transactions in Operational Research*, itor.70025. https://doi.org/10.1111/itor.70025

Caplan, G. A., Sulaiman, N. S., Mangin, D. A., Aimonino Ricauda, N., Wilson, A. D., & Barclay, L. (2012). A meta-analysis of "hospital in the home". *The Medical Journal of Australia*, *197*(9), 512–519. https://doi.org/10.5694/mja12.10480

Cheng, J., Montalto, M., & Leff, B. (2009). Hospital at home. *Clinics in Geriatric Medicine*, *25*(1), 79–91, vi. https://doi.org/10.1016/j.cger.2008.10.002

Chua, C. M. S., Ko, S. Q., Lai, Y. F., Lim, Y. W., & Shorey, S. (2022). Perceptions of hospital-at-home among stakeholders: A meta-synthesis. *Journal of General Internal Medicine*, *37*(3), 637–650. https://doi.org/10.1007/s11606-021-07065-0

Clapper, Y., Berkhout, J., & Bekker, R. (2024). Adaptive budget allocation in simheuristics applied to stochastic home healthcare routing and scheduling. *Computers & Industrial Engineering*, *198*, 110651. https://doi.org/10.1016/j.cie.2024.110651

Clapper, Y., Berkhout, J., Bekker, R., & Moeke, D. (2023). A model-based evolutionary algorithm for home health care scheduling. *Computers & Operations Research*, *150*, 106081. https://doi.org/10.1016/j.cor.2022.106081

Correia Azevedo, P., Rei, C., Grande, R., Saraiva, M., Guede-Fernández, F., Oliosi, E., & Londral, A. (2024). Assessment of the impact of home-based hospitalization on health outcomes: An observational study. *Acta Médica Portuguesa*, *37*(6), 445–454. https://doi.org/10.20344/amp.20474

CUF. (2023). Internamento em casa, tal como acontece no hospital — CUF. Retrieved July 31, 2026, from https://www.cuf.pt/mais-saude/internamento-em-casa-tal-como-acontece-no-hospital

CUF. (2026). Unidade de Hospitalização Domiciliária — CUF. Retrieved May 3, 2026, from https://www.cuf.pt/servicos-cuf/servicos-domiciliarios-cuf/unidade-de-hospitalizacao-domiciliaria-cuf

De Sousa Vale, J., Franco, A. I., Oliveira, C. V., Araújo, I., & Sousa, D. (2020). Hospital at home: An overview of literature. *Home Health Care Management & Practice*, *32*(2), 118–123. https://doi.org/10.1177/1084822319880930

Decerle, J., Grunder, O., El Hassani, A. H., & Barakat, O. (2018). A matheuristic for a multi-depot home health care problem. *IFAC-PapersOnLine*, *51*(11), 340–345. https://doi.org/10.1016/j.ifacol.2018.08.306

Deloitte. (2022). Hospital at home: A model with a future. Retrieved May 2, 2026, from https://www.deloitte.com/ch/en/Industries/government-public/perspectives/hospital-at-home.html

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. https://doi.org/10.1016/j.cie.2021.107255

DN/Lusa. (2025, April). Governo alarga hospitalização domiciliária a todos os hospitais do SNS. Retrieved May 2, 2026, from https://www.dn.pt/sociedade/governo-alarga-hospitaliza%C3%A7%C3%A3o-domicili%C3%A1ria-a-todos-os-hospitais-do-sns

Du, J., & Wang, X. (2024). A branch-and-price-and-cut algorithm for the home health care routing and scheduling problem with multiple prioritized time windows. *Computers & Operations Research*, *170*, 106749. https://doi.org/10.1016/j.cor.2024.106749

Edgar, K., Iliffe, S., Doll, H. A., Clarke, M. J., Gonçalves-Bradley, D. C., Wong, E., & Shepperd, S. (2024). Admission avoidance hospital at home (Cochrane Effective Practice and Organisation of Care Group, Ed.). *Cochrane Database of Systematic Reviews*, *2024*(3). https://doi.org/10.1002/14651858.CD007491.pub3

Erdem, M., Koç, Ç., & Yücel, E. (2022). The electric home health care routing and scheduling problem with time windows and fast chargers. *Computers & Industrial Engineering*, *172*, 108580. https://doi.org/10.1016/j.cie.2022.108580

Fikar, C., & Hirsch, P. (2017). Home health care routing and scheduling: A review. *Computers & Operations Research*, *77*, 86–95. https://doi.org/10.1016/j.cor.2016.07.019

Fu, Y., Ma, X., Gao, K., Wang, H., Sadollah, A., & Chen, L. (2024). Multi-objective migrating birds optimization for solving stochastic home health care routing and scheduling problems considering caregiver working time constraints. *Swarm and Evolutionary Computation*, *85*, 101484. https://doi.org/10.1016/j.swevo.2024.101484

Gendreau, M., & Potvin, J.-Y. (Eds.). (2019). *Handbook of Metaheuristics* (Vol. 272). Springer International Publishing. https://doi.org/10.1007/978-3-319-91086-4

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. https://doi.org/10.1016/j.ejor.2018.11.025

Helena, I. (2021, June). A Evolução da Hospitalização Domiciliária em Portugal. *SPMI*. Retrieved May 2, 2026, from https://www.spmi.pt/a-evolucao-da-hospitalizacao-domiciliaria-em-portugal/

Hosseini, S.-S., Adulyasak, Y., & Rousseau, L.-M. (2026). Consistent home health care routing and scheduling problem under time uncertainty. *Transportation Research Part E: Logistics and Transportation Review*, *205*, 104509. https://doi.org/10.1016/j.tre.2025.104509

Hulshof, P. J. H., Kortbeek, N., Boucherie, R. J., Hans, E. W., & Bakker, P. J. M. (2012). Taxonomic classification of planning decisions in health care: A structured review of the state of the art in OR/MS. *Health Systems*, *1*(2), 129–175. https://doi.org/10.1057/hs.2012.18

Inácio, A. M. (2024, March). Doentes tratados em casa já equivalem a um hospital de média dimensão. E vão aumentar. Retrieved May 2, 2026, from https://www.dn.pt/sociedade/doentes-tratados-em-casa-ja-equivalem-a-um-hospital-de-media-dimensao-e-vao-aumentar

INSA. (2025). Infográfico INSA – Envelhecimento e Saúde em Portugal – Quem Somos e Como Envelhecemos - INSA. Retrieved May 2, 2026, from https://www.insa.min-saude.pt/infografico-insa-envelhecimento-e-saude-em-portugal-quem-somos-e-como-envelhecemos/

Kummer, A. F., De Araújo, O. C., Buriol, L. S., & Resende, M. G. (2024). A biased random-key genetic algorithm for the home health care problem. *International Transactions in Operational Research*, *31*(3), 1859–1889. https://doi.org/10.1111/itor.13221

Larsen, M. N., Dreisig, T. S., Rasmussen, M. K., Schultz, A. N., Nielsen, T. L., & Fischer, T. K. (2025). Virtual health care in hospital-at-home models for patients with acute infections: A scoping review. *Mayo Clinic Proceedings: Digital Health*, *3*(3), 100250. https://doi.org/10.1016/j.mcpdig.2025.100250

Leff, B., Burton, L., Mader, S. L., Naughton, B., Burl, J., Inouye, S. K., Greenough, W. B., Guido, S., Langston, C., Frick, K. D., Steinwachs, D., & Burton, J. R. (2005). Hospital at home: Feasibility and outcomes of a program to provide hospital-level care at home for acutely ill older patients. *Annals of Internal Medicine*, *143*(11), 798–808. https://doi.org/10.7326/0003-4819-143-11-200512060-00008

Leong, M. Q., Lim, C. W., & Lai, Y. F. (2021). Comparison of Hospital-at-Home models: A systematic review of reviews. *BMJ open*, *11*(1), e043285. https://doi.org/10.1136/bmjopen-2020-043285

Levine, D. M., Ouchi, K., Blanchfield, B., Saenz, A., Burke, K., Paz, M., Diamond, K., Pu, C. T., & Schnipper, J. L. (2020). Hospital-level care at home for acutely ill adults: A randomized controlled trial. *Annals of Internal Medicine*, *172*(2), 77–85. https://doi.org/10.7326/M19-0600

Li, B., Pan, J., & Wei, Q. (2026). An adaptive large neighborhood search algorithm with customized search operators for solving home healthcare scheduling and routing problems with skill matching and periodic visits. *Applied Soft Computing*, *197*, 115152. https://doi.org/10.1016/j.asoc.2026.115152

Li, Y., Ye, C., Wang, H., Wang, F., & Xu, X. (2022). A discrete multi-objective grey wolf optimizer for the home health care routing and scheduling problem with priorities and uncertainty. *Computers & Industrial Engineering*, *169*, 108256. https://doi.org/10.1016/j.cie.2022.108256

Liu, R., Tao, Y., & Xie, X. (2019). An adaptive large neighborhood search heuristic for the vehicle routing problem with time windows and synchronized visits. *Computers & Operations Research*, *101*, 250–262. https://doi.org/10.1016/j.cor.2018.08.002

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. https://doi.org/10.1080/00207543.2016.1213917

Liu, W., Dridi, M., Fathollahi-Fard, A. M., & El Hassani, A. H. (2024). A customized adaptive large neighborhood search algorithm for solving a multi-objective home health care problem in a pandemic environment. *Swarm and Evolutionary Computation*, *86*, 101507. https://doi.org/10.1016/j.swevo.2024.101507

Liu, W., Dridi, M., Fei, H., & El Hassani, A. H. (2021a). 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. https://doi.org/10.1016/j.eswa.2021.115307

Liu, W., Dridi, M., Fei, H., & El Hassani, A. H. (2021b). Solving a multi-period home health care routing and scheduling problem using an efficient matheuristic. *Computers & Industrial Engineering*, *162*, 107721. https://doi.org/10.1016/j.cie.2021.107721

Lopes, P. (2025). Hospitalização Domiciliária: um hospital seguro em cada casa. *SPMI*. Retrieved May 2, 2026, from https://www.spmi.pt/hospitalizacao-domiciliaria-um-hospital-seguro-em-cada-casa/

Ma, X., Pu, X., Fu, Y., Gao, K., & Xu, Y. (2025). Home health care routing and scheduling problems considering patient classification and outsourcing: Modeling and a solution algorithm. *Computers & Operations Research*, *182*, 107143. https://doi.org/10.1016/j.cor.2025.107143

Makboul, S., Kharraja, S., Abbassi, A., & El Hilali Alaoui, A. (2024). A multiobjective approach for weekly Green Home Health Care routing and scheduling problem with care continuity and synchronized services. *Operations Research Perspectives*, *12*, 100302. https://doi.org/10.1016/j.orp.2024.100302

Martins, M., Bernardino, R., Santos, D., & Marques, I. (2026). A bi-objective hospital at home routing and scheduling problem. *International Transactions in Operational Research*, itor.70220. https://doi.org/10.1111/itor.70220

Mendes, J. J. (2026, March). Estamos preparados para levar o hospital a casa? Retrieved May 2, 2026, from https://expresso.pt/opiniao/2026-03-26-estamos-preparados-para-levar-o-hospital-a-casa--6bbdbaa0

Nikzad, E., Bashiri, M., & Abbasi, B. (2021). A matheuristic algorithm for stochastic home health care planning. *European Journal of Operational Research*, *288*(3), 753–774. https://doi.org/10.1016/j.ejor.2020.06.040

Patel, H. Y., & West, D. J. (2021). Hospital at home: An evolving model for comprehensive healthcare. *Global Journal on Quality and Safety in Healthcare*, *4*(4), 141–146. https://doi.org/10.36401/JQSH-21-4

Pisinger, D., & Ropke, S. (2007). A general heuristic for vehicle routing problems. *Computers & Operations Research*, *34*(8), 2403–2435. https://doi.org/10.1016/j.cor.2005.09.012

PORDATA. (2025). Ageing index. Retrieved May 2, 2026, from https://www.pordata.pt/en/europe/ageing+index-1609

Quintanilla, S., Ballestín, F., & Pérez, Á. (2020). Mathematical models to improve the current practice in a Home Healthcare Unit. *OR Spectrum*, *42*(1), 43–74. https://doi.org/10.1007/s00291-019-00565-w

Ropke, S., & Pisinger, D. (2006). An adaptive large neighborhood search heuristic for the pickup and delivery problem with time windows. *Transportation Science*, *40*(4), 455–472. https://doi.org/10.1287/trsc.1050.0135

Saenger, P. M., Ornstein, K. A., Garrido, M. M., Lubetsky, S., Bollens-Lund, E., DeCherrie, L. V., Leff, B., Siu, A. L., & Federman, A. D. (2022). Cost of home hospitalization versus inpatient hospitalization inclusive of a 30-day post-acute period. *Journal of the American Geriatrics Society*, *70*(5), 1374–1383. https://doi.org/10.1111/jgs.17706

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. https://doi.org/10.1016/j.eswa.2021.115209

Shepperd, S., Iliffe, S., Doll, H. A., Clarke, M. J., Kalra, L., Wilson, A. D., & Gonçalves-Bradley, D. C. (2016). Admission avoidance hospital at home (Cochrane Effective Practice and Organisation of Care Group, Ed.). *Cochrane Database of Systematic Reviews*, *2016*(9). https://doi.org/10.1002/14651858.CD007491.pub2

Shi, Y., Boudouh, T., & Grunder, O. (2017). A hybrid genetic algorithm for a home health care routing problem with time window and fuzzy demand. *Expert Systems with Applications*, *72*, 160–176. https://doi.org/10.1016/j.eswa.2016.12.013

Shiri, M., Ahmadizar, F., & Mahmoudzadeh, H. (2021). A three-phase methodology for home healthcare routing and scheduling under uncertainty. *Computers & Industrial Engineering*, *158*, 107416. https://doi.org/10.1016/j.cie.2021.107416

SNS. (2018). Hospitalização domiciliária em idade adulta. Retrieved May 2, 2026, from https://normas.dgs.min-saude.pt/2018/12/20/hospitalizacao-domiciliaria-em-idade-adulta/

SNS. (2025). 100% de cobertura em hospitalização domiciliária. Retrieved May 2, 2026, from https://www.sns.gov.pt/noticias/2025/08/11/100-de-cobertura-em-hospitalizacao-domiciliaria/

Trautsamwieser, A., & Hirsch, P. (2014). A Branch-Price-and-Cut approach for solving the medium-term home health care planning problem. *Networks*, *64*(3), 143–159. https://doi.org/10.1002/net.21566

Tsang, M. Y., & Shehadeh, K. S. (2023). Stochastic optimization models for a home service routing and appointment scheduling problem with random travel and service times. *European Journal of Operational Research*, *307*(1), 48–63. https://doi.org/10.1016/j.ejor.2022.09.020

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. https://doi.org/10.1016/j.cie.2024.109951

World Health Organization. (2025). Ageing and health. Retrieved May 2, 2026, from https://www.who.int/news-room/fact-sheets/detail/ageing-and-health

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. https://doi.org/10.1016/j.ejor.2022.03.022

Zhan, Y., Wang, Z., & Wan, G. (2021). Home service routing and appointment scheduling with stochastic service times. *European Journal of Operational Research*, *288*(1), 98–110. https://doi.org/10.1016/j.ejor.2020.05.037

Zhang, W., Ma, W., Yang, S., Chen, S., & Zhang, J. (2025). An improved adaptive large neighborhood search for the home health care routing and scheduling problem with multiple mixed time windows. *Information Sciences*, *719*, 122458. https://doi.org/10.1016/j.ins.2025.122458

Zhou, Y., Pan, Z., Zhao, Y., & Li, Z. (2026). Optimizing multi-period multi-mode multi-time window home health care scheduling with an improved tabu search algorithm. *Swarm and Evolutionary Computation*, *100*, 102263. https://doi.org/10.1016/j.swevo.2025.102263

Ziya-Gorabi, F., Ghodratnama, A., Tavakkoli-Moghaddam, R., & Asadi-Lari, M. S. (2022). A new fuzzy tri-objective model for a home health care problem with green ambulance routing and congestion under uncertainty. *Expert Systems with Applications*, *201*, 117093. https://doi.org/10.1016/j.eswa.2022.117093

Appendices

Appendix A: ALNS pseudo-code

Algorithm 1 The pseudo-code of ALNS

Input: Initial solution S_0; iterations I; max repair attempts A; destroy ops $\mathscr{D} = \{\text{random, worst, shaw}\}$; repair ops $\mathscr{R} = \{\text{greedy, regret-2}\}$; SA params β, α, $T_{\min}$; RWS params $\sigma_1, \sigma_2, \sigma_3, \lambda$, seg

Output: Best feasible solution S_{best}

$S_{\text{cur}} \leftarrow S_0$; $S_{\text{best}} \leftarrow S_0$; $T_0 \leftarrow \max(T_{\min},\ \beta \cdot f(S_0))$; $T \leftarrow T_0$
$\mathbf{w}_d, \mathbf{w}_r \leftarrow \mathbf{1}$ ▷ *uniform initial weights*
for iter $= 1$ **to** I **do**
  Select $d \in \mathscr{D}$ and $r \in \mathscr{R}$ via RWS on $\mathbf{w}_d$, $\mathbf{w}_r$
  $f_{\text{cur}} \leftarrow f(S_{\text{cur}})$; $S_{\text{new}} \leftarrow$ null
  for attempt $= 1$ **to** A **do** ▷ *retry candidate generation until repair succeeds*
    try $S_{\text{new}} \leftarrow r(d(S_{\text{cur}}))$; **break**
    on failure: $S_{\text{new}} \leftarrow$ null; **continue**
  end for
  if $S_{\text{new}} =$ null **then** ▷ *no feasible candidate this iteration*
    $T \leftarrow \max(T_{\min}, T \cdot \alpha)$; **continue**
  end if
  $\Delta f \leftarrow f(S_{\text{new}}) - f_{\text{cur}}$
  if $\Delta f \leq 0$ **or** $\mathscr{U}(0,1) < \exp(-\Delta f / T)$ **then**
    $S_{\text{cur}} \leftarrow S_{\text{new}}$; accepted $\leftarrow$ **true**
  else
    accepted $\leftarrow$ **false**
  end if
  if accepted **and** $f(S_{\text{new}}) < f(S_{\text{best}})$ **then**
    $S_{\text{best}} \leftarrow S_{\text{new}}$; score $\leftarrow \sigma_1$
  else if accepted **and** $f(S_{\text{new}}) < f_{\text{cur}}$ **then**
    score $\leftarrow \sigma_2$
  else if accepted **then**
    score $\leftarrow \sigma_3$
  else
    score $\leftarrow 0$
  end if
  Add score to the segment accumulators of d and r
  if iter mod seg $= 0$ **then**
    for all operators i used in the segment (uses$_i > 0$) **do**
      $\mathbf{w}[i] \leftarrow (1-\lambda)\,\mathbf{w}[i] + \lambda\,(\text{score}_i / \text{uses}_i)$
    end for
    Reset segment accumulators
  end if
  $T \leftarrow \max(T_{\min}, T \cdot \alpha)$
end for
return S_{best}

Appendix B: Full Parameter-Tuning Results

Table B1: Results of the 3^3 full-factorial experimental design.

Config.	$[q_{\min}, q_{\max}]$	β	α	Obj./patient	Time/seed (s)
C1	$[0.10, 0.20]$	0.050	0.995	26.292 ± 0.035	202.3
C2	$[0.10, 0.20]$	0.050	0.998	26.201 ± 0.049	174.6
C3	$[0.10, 0.20]$	0.050	0.999	26.221 ± 0.051	179.9
C4	$[0.10, 0.20]$	0.072	0.995	26.267 ± 0.034	179.4
C5	$[0.10, 0.20]$	0.072	0.998	26.205 ± 0.034	172.9
C6	$[0.10, 0.20]$	0.072	0.999	26.212 ± 0.032	171.7
C7	$[0.10, 0.20]$	0.100	0.995	26.235 ± 0.058	172.0
C8	$[0.10, 0.20]$	0.100	0.998	26.217 ± 0.051	169.7
C9	$[0.10, 0.20]$	0.100	0.999	26.203 ± 0.067	172.8
C10	$[0.15, 0.25]$	0.050	0.995	26.189 ± 0.064	264.3
C11	$[0.15, 0.25]$	0.050	0.998	26.118 ± 0.054	257.9
C12	$[0.15, 0.25]$	0.050	0.999	26.129 ± 0.053	257.3
C13	$[0.15, 0.25]$	0.072	0.995	26.162 ± 0.048	260.0
C14	$[0.15, 0.25]$	0.072	0.998	26.121 ± 0.055	272.5
C15	$[0.15, 0.25]$	0.072	0.999	26.091 ± 0.002	272.5
C16	$[0.15, 0.25]$	0.100	0.995	26.186 ± 0.031	260.8
C17	$[0.15, 0.25]$	0.100	0.998	26.145 ± 0.042	257.7
C18	$[0.15, 0.25]$	0.100	0.999	26.147 ± 0.048	258.6
C19	$[0.25, 0.35]$	0.050	0.995	26.052 ± 0.023	521.1
C20	$[0.25, 0.35]$	0.050	0.998	26.066 ± 0.041	502.2
C21	$[0.25, 0.35]$	0.050	0.999	26.071 ± 0.028	494.2
C22	$[0.25, 0.35]$	0.072	0.995	26.066 ± 0.003	518.1
C23	$\mathbf{[0.25, 0.35]}$	**0.072**	**0.998**	$\mathbf{26.045 \pm 0.025}$	**495.8**
C24	$[0.25, 0.35]$	0.072	0.999	26.054 ± 0.031	491.0
C25	$[0.25, 0.35]$	0.100	0.995	26.066 ± 0.039	552.1
C26	$[0.25, 0.35]$	0.100	0.998	26.046 ± 0.026	923.5
C27	$[0.25, 0.35]$	0.100	0.999	26.047 ± 0.025	939.1

Table B2: Results by scenario for **medium** instances.

	$q = [0.10;\ 0.20]$			$q = [0.15;\ 0.25]$			$q = [0.25;\ 0.35]$		
Scenario	Best	Avg	$Time(s)$	Best	Avg	$Time(s)$	Best	Avg	$Time(s)$
Balanced	24.81	24.83	14.9	24.80	24.81	20.9	24.79	24.80	45.3
Clustered	23.23	23.23	14.3	23.21	23.22	20.9	23.20	23.21	46.5
High urgent	29.44	29.46	13.5	29.43	29.44	19.0	29.42	29.43	42.2
High doctor	31.02	31.03	14.4	30.99	31.00	20.4	30.87	30.98	44.9
Average	27.12	27.14	14.3	27.11	27.12	20.3	27.07	27.10	44.7

Table B3: Results by scenario for **large** instances.

	$q = [0.10;\ 0.20]$			$q = [0.15;\ 0.25]$			$q = [0.25;\ 0.35]$		
Scenario	Best	Avg	$Time(s)$	Best	Avg	$Time(s)$	Best	Avg	$Time(s)$
Balanced	22.58	22.59	31.1	22.50	22.55	47.4	22.35	22.42	109.8
Clustered	20.77	20.80	30.8	20.76	20.77	46.6	20.68	20.75	111.2
High urgent	30.09	30.11	27.6	30.07	30.09	41.5	30.05	30.08	97.7
High doctor	28.81	28.98	30.7	28.24	28.57	45.7	28.12	28.18	106.5
Average	25.56	25.62	30.0	25.39	25.49	45.3	25.30	25.36	106.3

Appendix C: Depot-Level Scenario Results

Table C1: ALNS VS Gurobi by scenario and depot.

Scenario	Depot	N_d	Optimal certificates	Objective z_{MILP}	Objective $\bar{z}_{\text{ALNS}}$	Gap (%) Mean	Gap (%) Min	Gap (%) Max	Time (s) ALNS / Gurobi
Balanced	Lisboa	23–116	13/40	1432.3	1427.1	−0.242	−2.581	+1.173	65.9 / 1286.0
Balanced	Cascais	14–70	24/40	932.5	933.5	+0.099	−0.469	+0.670	16.8 / 912.2
Balanced	Sintra	9–44	35/40	669.9	670.8	+0.097	0.000	+0.679	5.7 / 361.8
Balanced	Odivelas	4–20	40/40	388.6	388.6	+0.012	0.000	+0.183	1.1 / 0.3
Clustered	Lisboa	23–116	9/40	1352.0	1345.1	−0.354	−7.140	+1.196	66.2 / 1452.3
Clustered	Cascais	14–70	23/40	869.7	870.5	+0.066	−0.046	+0.702	16.9 / 924.9
Clustered	Sintra	9–44	33/40	631.3	631.7	+0.045	0.000	+0.305	5.7 / 440.9
Clustered	Odivelas	4–20	40/40	376.7	376.8	+0.002	0.000	+0.090	1.1 / 0.4
High urgent	Lisboa	23–116	19/40	1678.0	1674.8	−0.103	−1.568	+0.560	63.4 / 1126.4
High urgent	Cascais	14–70	31/40	1200.8	1202.0	+0.081	−0.111	+0.298	16.4 / 539.5
High urgent	Sintra	9–44	35/40	822.2	822.7	+0.046	0.000	+0.259	5.7 / 333.8
High urgent	Odivelas	4–20	40/40	431.2	431.7	+0.096	0.000	+1.461	1.2 / 0.7
High doctor	Lisboa	23–116	10/40	1776.9	1770.8	−0.218	−10.234	+11.021	63.6 / 1359.9
High doctor	Cascais	14–70	24/40	1131.2	1135.0	+0.284	−0.257	+9.105	16.7 / 780.4
High doctor	Sintra	9–44	38/40	817.3	817.7	+0.042	0.000	+0.200	5.5 / 185.9
High doctor	Odivelas	4–20	40/40	427.8	427.8	+0.001	0.000	+0.022	1.1 / 0.6

Appendix D: Per-Instance ALNS and Gurobi Results

Table D1: Per-instance ALNS VS Gurobi results for the *balanced* scenario (160 depot instances).

Size	Depot	Rep	N_d	z_{MILP}	St.	$\bar{z}_{\text{ALNS}}$	Gap_{mean}	Gap_{best}
Small	Lisboa	0	23	641.1	Opt	641.1	+0.000	+0.000
Small	Cascais	0	14	303.5	Opt	305.0	+0.511	+0.511
Small	Odivelas	0	4	247.4	Opt	247.4	+0.000	+0.000
Small	Sintra	0	9	292.6	Opt	292.6	+0.000	+0.000
Small	Lisboa	1	23	616.1	Opt	616.1	+0.000	+0.000
Small	Cascais	1	14	526.2	Opt	526.2	+0.000	+0.000
Small	Odivelas	1	4	253.0	Opt	253.0	+0.000	+0.000
Small	Sintra	1	9	312.6	Opt	312.6	+0.000	+0.000
Small	Lisboa	2	23	666.6	Opt	666.6	+0.000	+0.000
Small	Cascais	2	14	537.9	Opt	537.9	+0.000	+0.000
Small	Odivelas	2	4	241.1	Opt	241.1	+0.000	+0.000
Small	Sintra	2	9	290.9	Opt	290.9	+0.000	+0.000
Small	Lisboa	3	23	655.9	Opt	655.9	+0.000	+0.000
Small	Cascais	3	14	570.6	Opt	572.7	+0.357	+0.000
Small	Odivelas	3	4	257.0	Opt	257.0	+0.000	+0.000
Small	Sintra	3	9	286.8	Opt	286.8	+0.000	+0.000

continued on next page

(Table D1 continued)

Size	Depot	Rep	N_d	z_{MILP}	St.	$\bar{z}_{\text{ALNS}}$	Gap_{mean}	Gap_{best}
Small	Lisboa	4	23	656.1	Opt	656.1	+0.000	+0.000
Small	Cascais	4	14	504.8	Opt	504.8	+0.000	+0.000
Small	Odivelas	4	4	258.4	Opt	258.4	+0.000	+0.000
Small	Sintra	4	9	312.2	Opt	312.2	+0.000	+0.000
Small	Lisboa	5	23	623.6	Opt	623.6	+0.000	+0.000
Small	Cascais	5	14	544.1	Opt	544.1	+0.000	+0.000
Small	Odivelas	5	4	256.3	Opt	256.3	+0.000	+0.000
Small	Sintra	5	9	326.6	Opt	326.6	+0.000	+0.000
Small	Lisboa	6	23	611.6	Opt	611.6	+0.000	+0.000
Small	Cascais	6	14	532.0	Opt	532.0	+0.000	+0.000
Small	Odivelas	6	4	266.8	Opt	266.8	+0.000	+0.000
Small	Sintra	6	9	308.7	Opt	308.7	+0.000	+0.000
Small	Lisboa	7	23	649.9	Opt	649.9	+0.000	+0.000
Small	Cascais	7	14	523.6	Opt	523.6	+0.000	+0.000
Small	Odivelas	7	4	235.6	Opt	235.6	+0.000	+0.000
Small	Sintra	7	9	332.2	Opt	332.2	+0.000	+0.000
Medium	Lisboa	0	46	1002.4	TL	1007.1	+0.471	+0.045
Medium	Cascais	0	28	614.8	Opt	614.8	+0.000	+0.000
Medium	Odivelas	0	8	272.2	Opt	272.2	+0.000	+0.000
Medium	Sintra	0	18	589.8	Opt	589.8	+0.000	+0.000
Medium	Lisboa	1	46	987.9	Opt	987.9	+0.000	+0.000
Medium	Cascais	1	28	589.2	Opt	589.2	+0.000	+0.000
Medium	Odivelas	1	8	295.3	Opt	295.3	+0.000	+0.000
Medium	Sintra	1	18	599.2	Opt	599.2	+0.000	+0.000
Medium	Lisboa	2	46	997.7	Opt	997.9	+0.024	+0.000
Medium	Cascais	2	28	602.8	Opt	602.8	+0.000	+0.000
Medium	Odivelas	2	8	290.9	Opt	290.9	+0.000	+0.000
Medium	Sintra	2	18	602.2	Opt	602.3	+0.029	+0.000
Medium	Lisboa	3	46	1002.3	Opt	1002.9	+0.058	+0.000
Medium	Cascais	3	28	630.3	Opt	630.3	+0.000	+0.000
Medium	Odivelas	3	8	289.6	Opt	289.6	+0.000	+0.000
Medium	Sintra	3	18	582.9	Opt	582.9	+0.000	+0.000
Medium	Lisboa	4	46	980.0	TL	980.0	+0.000	+0.000
Medium	Cascais	4	28	595.8	Opt	595.8	+0.000	+0.000
Medium	Odivelas	4	8	274.2	Opt	274.2	+0.000	+0.000

continued on next page

(Table D1 continued)

Size	Depot	Rep	N_d	z_{MILP}	St.	$\bar{z}_{\text{ALNS}}$	Gap$_{\text{mean}}$	Gap$_{\text{best}}$
Medium	Sintra	4	18	578.6	Opt	578.6	+0.000	+0.000
Medium	Lisboa	5	46	986.9	Opt	986.9	+0.000	+0.000
Medium	Cascais	5	28	616.8	Opt	617.4	+0.085	+0.000
Medium	Odivelas	5	8	285.7	Opt	285.7	+0.000	+0.000
Medium	Sintra	5	18	582.4	Opt	582.4	+0.000	+0.000
Medium	Lisboa	6	46	996.7	TL	998.0	+0.132	+0.000
Medium	Cascais	6	28	627.3	Opt	627.3	+0.000	+0.000
Medium	Odivelas	6	8	292.3	Opt	292.3	+0.000	+0.000
Medium	Sintra	6	18	609.6	Opt	610.1	+0.085	+0.000
Medium	Lisboa	7	46	979.1	Opt	979.5	+0.046	+0.000
Medium	Cascais	7	28	613.6	Opt	613.6	+0.000	+0.000
Medium	Odivelas	7	8	298.4	Opt	298.4	+0.000	+0.000
Medium	Sintra	7	18	350.2	Opt	350.2	+0.000	+0.000
Large	Lisboa	0	69	1524.0	TL	1484.7	-2.581	-12.234
Large	Cascais	0	42	911.3	TL	912.6	+0.144	+0.000
Large	Odivelas	0	12	298.2	Opt	298.2	+0.000	+0.000
Large	Sintra	0	27	635.6	Opt	635.7	+0.015	+0.000
Large	Lisboa	1	69	1543.0	TL	1541.7	-0.088	-0.134
Large	Cascais	1	42	894.1	Opt	895.3	+0.141	+0.000
Large	Odivelas	1	12	309.9	Opt	310.3	+0.102	+0.000
Large	Sintra	1	27	626.9	Opt	627.1	+0.029	+0.000
Large	Lisboa	2	69	1552.7	TL	1552.9	+0.013	-0.104
Large	Cascais	2	42	908.5	Opt	908.6	+0.005	+0.000
Large	Odivelas	2	12	297.8	Opt	297.8	+0.000	+0.000
Large	Sintra	2	27	655.1	Opt	655.2	+0.010	+0.000
Large	Lisboa	3	69	1523.1	TL	1524.4	+0.086	-0.012
Large	Cascais	3	42	909.5	Opt	912.6	+0.346	+0.131
Large	Odivelas	3	12	310.0	Opt	310.0	+0.000	+0.000
Large	Sintra	3	27	616.3	Opt	616.6	+0.048	+0.000
Large	Lisboa	4	69	1303.3	TL	1303.1	-0.012	-0.081
Large	Cascais	4	42	925.0	Opt	927.4	+0.253	+0.251
Large	Odivelas	4	12	326.2	Opt	326.2	+0.000	+0.000
Large	Sintra	4	27	630.2	Opt	630.6	+0.070	+0.000
Large	Lisboa	5	69	1517.6	TL	1517.9	+0.019	-0.004
Large	Cascais	5	42	901.9	TL	903.7	+0.200	+0.008

continued on next page

(Table D1 continued)

Size	Depot	Rep	N_d	z_{MILP}	St.	$\bar{z}_{\text{ALNS}}$	Gap_{mean}	Gap_{best}
Large	Odivelas	5	12	328.6	Opt	328.8	+0.059	+0.000
Large	Sintra	5	27	657.3	Opt	657.3	+0.000	+0.000
Large	Lisboa	6	69	1328.1	TL	1343.7	+1.173	+0.346
Large	Cascais	6	42	878.4	TL	878.4	+0.000	+0.000
Large	Odivelas	6	12	307.5	Opt	307.5	+0.000	+0.000
Large	Sintra	6	27	635.8	Opt	635.8	+0.000	+0.000
Large	Lisboa	7	69	1340.5	TL	1340.7	+0.012	-0.304
Large	Cascais	7	42	928.5	Opt	928.9	+0.036	+0.000
Large	Odivelas	7	12	321.5	Opt	321.5	+0.000	+0.000
Large	Sintra	7	27	647.1	Opt	648.1	+0.156	+0.080
XLarge	Lisboa	0	92	1852.4	TL	1842.6	-0.529	-0.732
XLarge	Cascais	0	56	1179.3	TL	1179.9	+0.046	-0.043
XLarge	Odivelas	0	16	521.6	Opt	521.6	+0.000	+0.000
XLarge	Sintra	0	36	891.7	Opt	892.7	+0.106	+0.000
XLarge	Lisboa	1	92	1864.7	TL	1868.0	+0.178	+0.002
XLarge	Cascais	1	56	1184.8	Opt	1188.2	+0.287	+0.098
XLarge	Odivelas	1	16	521.3	Opt	521.3	+0.000	+0.000
XLarge	Sintra	1	36	905.3	Opt	905.5	+0.023	+0.000
XLarge	Lisboa	2	92	1860.8	TL	1850.0	-0.577	-0.626
XLarge	Cascais	2	56	1170.0	Opt	1171.5	+0.129	+0.000
XLarge	Odivelas	2	16	541.5	Opt	541.5	+0.000	+0.000
XLarge	Sintra	2	36	923.3	TL	923.6	+0.038	+0.000
XLarge	Lisboa	3	92	1891.3	TL	1878.9	-0.658	-0.771
XLarge	Cascais	3	56	1161.9	TL	1161.7	-0.024	-0.055
XLarge	Odivelas	3	16	528.7	Opt	528.7	+0.000	+0.000
XLarge	Sintra	3	36	668.9	Opt	669.4	+0.078	+0.000
XLarge	Lisboa	4	92	1889.7	TL	1882.7	-0.374	-0.564
XLarge	Cascais	4	56	1190.4	Opt	1193.8	+0.286	+0.206
XLarge	Odivelas	4	16	546.4	Opt	547.4	+0.183	+0.183
XLarge	Sintra	4	36	912.8	Opt	913.2	+0.042	+0.000
XLarge	Lisboa	5	92	1863.6	TL	1863.9	+0.019	-0.249
XLarge	Cascais	5	56	1174.9	TL	1175.4	+0.044	+0.000
XLarge	Odivelas	5	16	525.1	Opt	525.1	+0.000	+0.000
XLarge	Sintra	5	36	887.6	TL	887.6	+0.000	+0.000
XLarge	Lisboa	6	92	1861.9	TL	1858.8	-0.167	-0.279

continued on next page

(Table D1 continued)

Size	Depot	Rep	N_d	z_{MILP}	St.	$\bar{z}_{\text{ALNS}}$	Gap_{mean}	Gap_{best}
XLarge	Cascais	6	56	1199.1	TL	1201.4	+0.188	+0.118
XLarge	Odivelas	6	16	551.0	Opt	551.0	+0.000	+0.000
XLarge	Sintra	6	36	907.9	Opt	907.9	+0.002	+0.000
XLarge	Lisboa	7	92	1858.2	TL	1862.7	+0.245	-0.003
XLarge	Cascais	7	56	1152.8	TL	1153.2	+0.033	+0.000
XLarge	Odivelas	7	16	538.8	Opt	538.8	+0.000	+0.000
XLarge	Sintra	7	36	881.1	TL	881.4	+0.035	+0.000
XXLarge	Lisboa	0	116	2189.2	TL	2170.4	-0.859	-1.059
XXLarge	Cascais	0	70	1458.4	TL	1454.9	-0.245	-0.362
XXLarge	Odivelas	0	20	549.9	Opt	549.9	+0.000	+0.000
XXLarge	Sintra	0	44	953.5	Opt	957.9	+0.455	+0.000
XXLarge	Lisboa	1	116	2211.3	TL	2181.3	-1.358	-1.755
XXLarge	Cascais	1	70	1488.4	TL	1488.7	+0.016	-0.040
XXLarge	Odivelas	1	20	554.7	Opt	554.7	+0.013	+0.000
XXLarge	Sintra	1	44	973.3	TL	979.3	+0.617	+0.457
XXLarge	Lisboa	2	116	2226.6	TL	2202.4	-1.083	-1.310
XXLarge	Cascais	2	70	1442.1	TL	1435.3	-0.469	-0.556
XXLarge	Odivelas	2	20	553.6	Opt	553.7	+0.019	+0.000
XXLarge	Sintra	2	44	946.7	Opt	949.2	+0.261	+0.000
XXLarge	Lisboa	3	116	2216.9	TL	2198.3	-0.840	-1.169
XXLarge	Cascais	3	70	1467.2	TL	1469.2	+0.132	+0.006
XXLarge	Odivelas	3	20	560.4	Opt	560.4	+0.000	+0.000
XXLarge	Sintra	3	44	965.9	Opt	966.4	+0.050	+0.000
XXLarge	Lisboa	4	116	2196.4	TL	2196.8	+0.017	-0.317
XXLarge	Cascais	4	70	1448.4	TL	1456.4	+0.553	+0.512
XXLarge	Odivelas	4	20	559.8	Opt	559.9	+0.011	+0.000
XXLarge	Sintra	4	44	974.5	Opt	981.1	+0.679	+0.124
XXLarge	Lisboa	5	116	2242.4	TL	2198.6	-1.954	-2.189
XXLarge	Cascais	5	70	1471.5	TL	1469.8	-0.120	-0.372
XXLarge	Odivelas	5	20	564.6	Opt	564.9	+0.049	+0.000
XXLarge	Sintra	5	44	977.7	Opt	981.9	+0.424	+0.259
XXLarge	Lisboa	6	116	2184.1	TL	2160.8	-1.064	-1.258
XXLarge	Cascais	6	70	1461.5	TL	1471.3	+0.670	+0.435
XXLarge	Odivelas	6	20	549.8	Opt	549.8	+0.000	+0.000
XXLarge	Sintra	6	44	985.9	Opt	989.4	+0.358	+0.078

continued on next page

(Table D1 continued)

Size	Depot	Rep	N_d	z_{MILP}	St.	$\bar{z}_{\text{ALNS}}$	Gap$_{\text{mean}}$	Gap$_{\text{best}}$
XXLarge	Lisboa	7	116	2196.9	TL	2196.7	-0.010	-0.164
XXLarge	Cascais	7	70	1458.9	TL	1464.0	+0.346	+0.148
XXLarge	Odivelas	7	20	561.2	Opt	561.4	+0.044	+0.000
XXLarge	Sintra	7	44	978.0	TL	980.8	+0.290	+0.179

Table D2: Per-instance ALNS VS Gurobi results for the *clustered* scenario (160 depot instances).

Size	Depot	Rep	N_d	z_{MILP}	St.	$\bar{z}_{\text{ALNS}}$	Gap$_{\text{mean}}$	Gap$_{\text{best}}$
Small	Lisboa	0	23	603.9	Opt	603.9	+0.000	+0.000
Small	Cascais	0	14	298.4	Opt	298.4	+0.000	+0.000
Small	Odivelas	0	4	250.7	Opt	250.7	+0.000	+0.000
Small	Sintra	0	9	288.4	Opt	288.4	+0.000	+0.000
Small	Lisboa	1	23	560.7	Opt	560.7	+0.000	+0.000
Small	Cascais	1	14	501.7	Opt	501.7	+0.000	+0.000
Small	Odivelas	1	4	262.4	Opt	262.4	+0.000	+0.000
Small	Sintra	1	9	304.6	Opt	304.6	+0.000	+0.000
Small	Lisboa	2	23	589.2	Opt	589.3	+0.013	+0.000
Small	Cascais	2	14	525.3	Opt	525.3	+0.000	+0.000
Small	Odivelas	2	4	245.7	Opt	245.7	+0.000	+0.000
Small	Sintra	2	9	280.2	Opt	280.2	+0.000	+0.000
Small	Lisboa	3	23	584.2	Opt	584.5	+0.054	+0.000
Small	Cascais	3	14	300.1	Opt	300.1	+0.000	+0.000
Small	Odivelas	3	4	243.1	Opt	243.1	+0.000	+0.000
Small	Sintra	3	9	293.7	Opt	293.7	+0.000	+0.000
Small	Lisboa	4	23	572.2	Opt	572.2	+0.000	+0.000
Small	Cascais	4	14	326.8	Opt	326.8	+0.000	+0.000
Small	Odivelas	4	4	259.8	Opt	259.8	+0.000	+0.000
Small	Sintra	4	9	266.8	Opt	266.8	+0.000	+0.000
Small	Lisboa	5	23	594.3	TL	594.4	+0.013	+0.000
Small	Cascais	5	14	282.5	Opt	282.5	+0.000	+0.000
Small	Odivelas	5	4	258.4	Opt	258.4	+0.000	+0.000
Small	Sintra	5	9	296.7	Opt	296.7	+0.000	+0.000
Small	Lisboa	6	23	580.3	Opt	580.3	+0.000	+0.000
Small	Cascais	6	14	278.1	Opt	278.1	+0.000	+0.000
Small	Odivelas	6	4	247.1	Opt	247.1	+0.000	+0.000
Small	Sintra	6	9	284.8	Opt	284.8	+0.000	+0.000

continued on next page

(Table D2 continued)

Size	Depot	Rep	N_d	z_{MILP}	St.	$\bar{z}_{\text{ALNS}}$	Gap_{mean}	Gap_{best}
Small	Lisboa	7	23	558.9	Opt	558.9	+0.000	+0.000
Small	Cascais	7	14	296.8	Opt	296.8	+0.000	+0.000
Small	Odivelas	7	4	248.8	Opt	248.8	+0.000	+0.000
Small	Sintra	7	9	283.6	Opt	283.6	+0.000	+0.000
Medium	Lisboa	0	46	928.8	Opt	929.6	+0.081	+0.000
Medium	Cascais	0	28	609.9	Opt	609.9	+0.000	+0.000
Medium	Odivelas	0	8	292.9	Opt	292.9	+0.000	+0.000
Medium	Sintra	0	18	550.4	Opt	550.4	+0.000	+0.000
Medium	Lisboa	1	46	905.7	TL	905.9	+0.023	+0.000
Medium	Cascais	1	28	566.3	Opt	566.7	+0.071	+0.000
Medium	Odivelas	1	8	269.9	Opt	269.9	+0.000	+0.000
Medium	Sintra	1	18	545.0	Opt	545.0	+0.000	+0.000
Medium	Lisboa	2	46	909.4	TL	910.9	+0.164	+0.159
Medium	Cascais	2	28	593.9	Opt	593.9	+0.000	+0.000
Medium	Odivelas	2	8	259.8	Opt	259.8	+0.000	+0.000
Medium	Sintra	2	18	534.8	Opt	534.8	+0.000	+0.000
Medium	Lisboa	3	46	910.4	TL	914.3	+0.424	+0.293
Medium	Cascais	3	28	588.2	Opt	588.2	+0.000	+0.000
Medium	Odivelas	3	8	275.9	Opt	275.9	+0.000	+0.000
Medium	Sintra	3	18	560.7	Opt	560.7	+0.000	+0.000
Medium	Lisboa	4	46	919.7	TL	919.8	+0.019	+0.000
Medium	Cascais	4	28	575.0	Opt	575.0	+0.000	+0.000
Medium	Odivelas	4	8	267.1	Opt	267.1	+0.000	+0.000
Medium	Sintra	4	18	559.6	Opt	559.6	+0.000	+0.000
Medium	Lisboa	5	46	914.0	Opt	914.3	+0.029	+0.000
Medium	Cascais	5	28	610.6	Opt	610.6	+0.000	+0.000
Medium	Odivelas	5	8	265.5	Opt	265.5	+0.000	+0.000
Medium	Sintra	5	18	350.8	Opt	350.8	+0.000	+0.000
Medium	Lisboa	6	46	904.3	TL	904.5	+0.021	+0.000
Medium	Cascais	6	28	604.3	Opt	604.3	+0.000	+0.000
Medium	Odivelas	6	8	280.8	Opt	280.8	+0.000	+0.000
Medium	Sintra	6	18	569.5	Opt	569.5	+0.007	+0.000
Medium	Lisboa	7	46	908.1	TL	906.9	-0.130	-0.232
Medium	Cascais	7	28	599.8	Opt	599.8	+0.000	+0.000
Medium	Odivelas	7	8	268.7	Opt	268.7	+0.000	+0.000

continued on next page

(Table D2 continued)

Size	Depot	Rep	N_d	z_{MILP}	St.	$\bar{z}_{\text{ALNS}}$	Gap_{mean}	Gap_{best}
Medium	Sintra	7	18	583.0	Opt	583.0	+0.000	+0.000
Large	Lisboa	0	69	1457.5	TL	1461.1	+0.252	+0.054
Large	Cascais	0	42	871.8	TL	872.4	+0.074	+0.000
Large	Odivelas	0	12	297.0	Opt	297.0	+0.000	+0.000
Large	Sintra	0	27	646.3	Opt	646.5	+0.028	+0.000
Large	Lisboa	1	69	1264.7	TL	1258.5	-0.493	-0.894
Large	Cascais	1	42	873.7	Opt	875.2	+0.165	+0.086
Large	Odivelas	1	12	311.6	Opt	311.6	+0.000	+0.000
Large	Sintra	1	27	594.4	Opt	595.0	+0.100	+0.000
Large	Lisboa	2	69	1281.1	TL	1283.6	+0.193	-0.156
Large	Cascais	2	42	865.1	TL	865.6	+0.049	+0.021
Large	Odivelas	2	12	293.3	Opt	293.3	+0.000	+0.000
Large	Sintra	2	27	600.0	Opt	600.0	+0.000	+0.000
Large	Lisboa	3	69	1228.8	TL	1230.8	+0.162	+0.070
Large	Cascais	3	42	855.7	Opt	856.6	+0.098	+0.076
Large	Odivelas	3	12	311.6	Opt	311.6	+0.000	+0.000
Large	Sintra	3	27	613.0	Opt	613.0	+0.000	+0.000
Large	Lisboa	4	69	1308.6	TL	1305.5	-0.238	-0.345
Large	Cascais	4	42	850.9	Opt	850.9	+0.005	+0.000
Large	Odivelas	4	12	309.0	Opt	309.0	+0.000	+0.000
Large	Sintra	4	27	620.0	Opt	620.0	+0.000	+0.000
Large	Lisboa	5	69	1257.8	TL	1272.9	+1.196	+0.206
Large	Cascais	5	42	903.8	TL	903.9	+0.007	+0.000
Large	Odivelas	5	12	280.0	Opt	280.0	+0.000	+0.000
Large	Sintra	5	27	594.9	Opt	594.9	+0.000	+0.000
Large	Lisboa	6	69	1491.1	TL	1448.0	-2.891	-12.041
Large	Cascais	6	42	850.9	Opt	851.5	+0.066	+0.000
Large	Odivelas	6	12	274.3	Opt	274.3	+0.000	+0.000
Large	Sintra	6	27	587.3	Opt	587.3	+0.000	+0.000
Large	Lisboa	7	69	1248.4	TL	1245.3	-0.247	-0.353
Large	Cascais	7	42	857.2	TL	857.7	+0.061	+0.000
Large	Odivelas	7	12	289.6	Opt	289.6	+0.000	+0.000
Large	Sintra	7	27	622.3	Opt	622.6	+0.056	+0.000
XLarge	Lisboa	0	92	1782.1	TL	1781.2	-0.050	-0.123
XLarge	Cascais	0	56	1138.6	TL	1139.9	+0.117	+0.000

continued on next page

(Table D2 continued)

Size	Depot	Rep	N_d	z_{MILP}	St.	$\bar{z}_{\text{ALNS}}$	Gap_{mean}	Gap_{best}
XLarge	Odivelas	0	16	523.3	Opt	523.3	+0.000	+0.000
XLarge	Sintra	0	36	914.5	Opt	915.4	+0.099	+0.000
XLarge	Lisboa	1	92	1844.8	TL	1844.6	-0.012	-0.208
XLarge	Cascais	1	56	1153.0	Opt	1155.9	+0.255	+0.000
XLarge	Odivelas	1	16	505.7	Opt	505.7	+0.000	+0.000
XLarge	Sintra	1	36	834.7	Opt	836.3	+0.186	+0.034
XLarge	Lisboa	2	92	1805.6	TL	1803.3	-0.130	-0.284
XLarge	Cascais	2	56	1113.0	TL	1112.5	-0.046	-0.141
XLarge	Odivelas	2	16	515.3	Opt	515.3	+0.000	+0.000
XLarge	Sintra	2	36	860.7	TL	861.2	+0.065	+0.000
XLarge	Lisboa	3	92	1802.1	TL	1807.0	+0.271	+0.191
XLarge	Cascais	3	56	1118.5	Opt	1120.2	+0.148	+0.138
XLarge	Odivelas	3	16	524.4	Opt	524.4	+0.000	+0.000
XLarge	Sintra	3	36	832.5	TL	833.0	+0.054	-0.009
XLarge	Lisboa	4	92	1784.8	TL	1774.7	-0.565	-0.815
XLarge	Cascais	4	56	1119.8	TL	1122.6	+0.248	+0.128
XLarge	Odivelas	4	16	500.3	Opt	500.3	+0.000	+0.000
XLarge	Sintra	4	36	628.5	TL	628.5	+0.000	+0.000
XLarge	Lisboa	5	92	1769.4	TL	1776.2	+0.388	+0.102
XLarge	Cascais	5	56	1189.3	TL	1188.9	-0.029	-0.102
XLarge	Odivelas	5	16	516.8	Opt	516.8	+0.000	+0.000
XLarge	Sintra	5	36	839.0	TL	839.0	+0.000	+0.000
XLarge	Lisboa	6	92	1765.8	TL	1765.6	-0.014	-0.207
XLarge	Cascais	6	56	1108.7	TL	1109.3	+0.053	-0.110
XLarge	Odivelas	6	16	544.5	Opt	544.5	+0.000	+0.000
XLarge	Sintra	6	36	841.9	TL	842.4	+0.056	+0.000
XLarge	Lisboa	7	92	1825.7	TL	1695.3	-7.140	-11.626
XLarge	Cascais	7	56	1126.2	Opt	1126.9	+0.066	+0.016
XLarge	Odivelas	7	16	511.6	Opt	511.6	+0.000	+0.000
XLarge	Sintra	7	36	869.4	Opt	869.4	+0.000	+0.000
XXLarge	Lisboa	0	116	2129.1	TL	2111.1	-0.848	-1.134
XXLarge	Cascais	0	70	1393.8	TL	1394.7	+0.067	-0.006
XXLarge	Odivelas	0	20	542.6	Opt	542.6	+0.000	+0.000
XXLarge	Sintra	0	44	896.8	Opt	897.5	+0.073	+0.000
XXLarge	Lisboa	1	116	2136.8	TL	2119.4	-0.816	-1.079

continued on next page

(Table D2 continued)

Size	Depot	Rep	N_d	z_{MILP}	St.	$\bar{z}_{\text{ALNS}}$	Gap$_{\text{mean}}$	Gap$_{\text{best}}$
XXLarge	Cascais	1	70	1397.0	TL	1406.8	+0.702	+0.515
XXLarge	Odivelas	1	20	542.4	Opt	542.4	+0.000	+0.000
XXLarge	Sintra	1	44	921.7	TL	923.1	+0.160	+0.065
XXLarge	Lisboa	2	116	2189.9	TL	2169.1	-0.946	-1.322
XXLarge	Cascais	2	70	1411.1	TL	1411.3	+0.012	-0.068
XXLarge	Odivelas	2	20	541.3	Opt	541.3	+0.000	+0.000
XXLarge	Sintra	2	44	885.9	Opt	886.4	+0.047	+0.000
XXLarge	Lisboa	3	116	2116.7	TL	2118.7	+0.094	-0.215
XXLarge	Cascais	3	70	1387.0	TL	1386.6	-0.026	-0.117
XXLarge	Odivelas	3	20	563.4	Opt	563.4	+0.000	+0.000
XXLarge	Sintra	3	44	900.1	TL	902.8	+0.305	+0.031
XXLarge	Lisboa	4	116	2162.8	TL	2160.1	-0.124	-0.395
XXLarge	Cascais	4	70	1401.4	TL	1400.9	-0.035	-0.096
XXLarge	Odivelas	4	20	535.0	Opt	535.0	+0.000	+0.000
XXLarge	Sintra	4	44	903.6	Opt	905.8	+0.242	+0.084
XXLarge	Lisboa	5	116	2184.9	TL	2145.4	-1.807	-2.179
XXLarge	Cascais	5	70	1435.8	TL	1435.6	-0.015	-0.084
XXLarge	Odivelas	5	20	537.3	Opt	537.8	+0.090	+0.000
XXLarge	Sintra	5	44	885.1	Opt	886.3	+0.137	+0.066
XXLarge	Lisboa	6	116	2156.5	TL	2140.5	-0.742	-1.064
XXLarge	Cascais	6	70	1381.9	TL	1385.4	+0.255	+0.099
XXLarge	Odivelas	6	20	542.4	Opt	542.4	+0.000	+0.000
XXLarge	Sintra	6	44	906.5	Opt	906.5	+0.000	+0.000
XXLarge	Lisboa	7	116	2142.5	TL	2134.7	-0.363	-0.559
XXLarge	Cascais	7	70	1427.5	TL	1431.3	+0.267	+0.064
XXLarge	Odivelas	7	20	560.5	Opt	560.5	+0.000	+0.000
XXLarge	Sintra	7	44	900.7	Opt	902.5	+0.199	+0.114

Table D3: Per-instance ALNS VS Gurobi results for the **high urgent** scenario (160 depot instances).

Size	Depot	Rep	N_d	z_{MILP}	St.	$\bar{z}_{\text{ALNS}}$	Gap$_{\text{mean}}$	Gap$_{\text{best}}$
Small	Lisboa	0	23	641.7	Opt	641.7	+0.000	+0.000
Small	Cascais	0	14	555.4	Opt	555.4	+0.000	+0.000
Small	Odivelas	0	4	247.0	Opt	247.0	+0.000	+0.000
Small	Sintra	0	9	351.0	Opt	351.0	+0.000	+0.000
Small	Lisboa	1	23	668.8	Opt	668.8	+0.000	+0.000

continued on next page

(Table D3 continued)

Size	Depot	Rep	N_d	z_{MILP}	St.	$\bar{z}_{\text{ALNS}}$	Gap_{mean}	Gap_{best}
Small	Cascais	1	14	546.1	Opt	546.1	+0.000	+0.000
Small	Odivelas	1	4	244.2	Opt	244.2	+0.000	+0.000
Small	Sintra	1	9	300.0	Opt	300.0	+0.000	+0.000
Small	Lisboa	2	23	640.6	Opt	640.6	+0.000	+0.000
Small	Cascais	2	14	532.6	Opt	532.6	+0.000	+0.000
Small	Odivelas	2	4	235.4	Opt	235.4	+0.000	+0.000
Small	Sintra	2	9	300.6	Opt	300.6	+0.000	+0.000
Small	Lisboa	3	23	621.2	Opt	621.2	+0.000	+0.000
Small	Cascais	3	14	514.6	Opt	514.6	+0.000	+0.000
Small	Odivelas	3	4	252.6	Opt	252.6	+0.000	+0.000
Small	Sintra	3	9	330.0	Opt	330.0	+0.000	+0.000
Small	Lisboa	4	23	669.1	Opt	669.4	+0.046	+0.000
Small	Cascais	4	14	549.9	Opt	549.9	+0.000	+0.000
Small	Odivelas	4	4	243.5	Opt	243.5	+0.000	+0.000
Small	Sintra	4	9	300.0	Opt	300.0	+0.000	+0.000
Small	Lisboa	5	23	636.8	Opt	637.3	+0.076	+0.000
Small	Cascais	5	14	524.7	Opt	524.7	+0.000	+0.000
Small	Odivelas	5	4	255.7	Opt	255.7	+0.000	+0.000
Small	Sintra	5	9	317.1	Opt	317.1	+0.000	+0.000
Small	Lisboa	6	23	626.9	Opt	626.9	+0.000	+0.000
Small	Cascais	6	14	534.0	Opt	534.0	+0.000	+0.000
Small	Odivelas	6	4	245.2	Opt	245.2	+0.000	+0.000
Small	Sintra	6	9	310.4	Opt	310.4	+0.000	+0.000
Small	Lisboa	7	23	646.7	Opt	646.7	+0.000	+0.000
Small	Cascais	7	14	538.5	Opt	538.5	+0.000	+0.000
Small	Odivelas	7	4	259.3	Opt	259.3	+0.000	+0.000
Small	Sintra	7	9	297.0	Opt	297.0	+0.000	+0.000
Medium	Lisboa	0	46	1205.3	Opt	1205.3	+0.000	+0.000
Medium	Cascais	0	28	839.5	Opt	839.5	+0.000	+0.000
Medium	Odivelas	0	8	306.8	Opt	306.8	+0.000	+0.000
Medium	Sintra	0	18	582.1	Opt	582.1	+0.000	+0.000
Medium	Lisboa	1	46	1199.7	Opt	1200.9	+0.098	+0.050
Medium	Cascais	1	28	850.2	Opt	851.7	+0.181	+0.148
Medium	Odivelas	1	8	292.6	Opt	292.6	+0.000	+0.000
Medium	Sintra	1	18	597.6	Opt	597.6	+0.000	+0.000

continued on next page

(Table D3 continued)

Size	Depot	Rep	N_d	z_{MILP}	St.	$\bar{z}_{\text{ALNS}}$	Gap_{mean}	Gap_{best}
Medium	Lisboa	2	46	1209.5	Opt	1209.6	+0.004	+0.000
Medium	Cascais	2	28	868.4	Opt	868.4	+0.000	+0.000
Medium	Odivelas	2	8	273.5	Opt	273.5	+0.000	+0.000
Medium	Sintra	2	18	606.9	Opt	606.9	+0.000	+0.000
Medium	Lisboa	3	46	1220.1	Opt	1220.1	+0.000	+0.000
Medium	Cascais	3	28	853.1	Opt	853.4	+0.028	+0.000
Medium	Odivelas	3	8	283.2	Opt	283.2	+0.000	+0.000
Medium	Sintra	3	18	618.7	Opt	618.7	+0.000	+0.000
Medium	Lisboa	4	46	1222.2	Opt	1222.3	+0.008	+0.006
Medium	Cascais	4	28	836.2	Opt	836.2	+0.000	+0.000
Medium	Odivelas	4	8	291.2	Opt	291.2	+0.000	+0.000
Medium	Sintra	4	18	616.6	Opt	616.6	+0.000	+0.000
Medium	Lisboa	5	46	1215.5	Opt	1216.5	+0.087	+0.000
Medium	Cascais	5	28	837.3	Opt	837.6	+0.031	+0.000
Medium	Odivelas	5	8	278.2	Opt	278.2	+0.000	+0.000
Medium	Sintra	5	18	620.5	Opt	620.5	+0.000	+0.000
Medium	Lisboa	6	46	1213.1	Opt	1214.9	+0.146	+0.054
Medium	Cascais	6	28	849.7	Opt	849.7	+0.000	+0.000
Medium	Odivelas	6	8	288.6	Opt	288.6	+0.000	+0.000
Medium	Sintra	6	18	631.8	Opt	631.8	+0.000	+0.000
Medium	Lisboa	7	46	1229.8	Opt	1230.2	+0.036	+0.000
Medium	Cascais	7	28	851.6	Opt	851.9	+0.033	+0.000
Medium	Odivelas	7	8	291.2	Opt	291.2	+0.000	+0.000
Medium	Sintra	7	18	604.8	Opt	604.8	+0.000	+0.000
Large	Lisboa	0	69	1768.7	TL	1771.6	+0.162	+0.005
Large	Cascais	0	42	1355.7	TL	1359.0	+0.244	+0.072
Large	Odivelas	0	12	520.5	Opt	520.5	+0.000	+0.000
Large	Sintra	0	27	842.0	Opt	842.0	+0.000	+0.000
Large	Lisboa	1	69	1764.5	Opt	1767.9	+0.190	+0.078
Large	Cascais	1	42	1354.2	Opt	1357.8	+0.264	+0.125
Large	Odivelas	1	12	511.7	Opt	511.7	+0.000	+0.000
Large	Sintra	1	27	905.1	Opt	905.1	+0.000	+0.000
Large	Lisboa	2	69	1771.2	Opt	1775.9	+0.267	+0.003
Large	Cascais	2	42	1348.3	Opt	1349.4	+0.081	+0.000
Large	Odivelas	2	12	520.4	Opt	520.4	+0.000	+0.000

continued on next page

(Table D3 continued)

Size	Depot	Rep	N_d	z_{MILP}	St.	$\bar{z}_{\text{ALNS}}$	Gap_{mean}	Gap_{best}
Large	Sintra	2	27	876.3	Opt	876.5	+0.016	+0.000
Large	Lisboa	3	69	1738.9	Opt	1746.5	+0.436	+0.176
Large	Cascais	3	42	1316.6	Opt	1316.9	+0.023	+0.000
Large	Odivelas	3	12	534.6	Opt	534.6	+0.000	+0.000
Large	Sintra	3	27	881.2	Opt	881.2	+0.000	+0.000
Large	Lisboa	4	69	1788.3	TL	1791.9	+0.205	+0.169
Large	Cascais	4	42	1336.4	Opt	1338.6	+0.162	+0.037
Large	Odivelas	4	12	509.8	Opt	509.8	+0.000	+0.000
Large	Sintra	4	27	863.4	Opt	863.4	+0.000	+0.000
Large	Lisboa	5	69	1789.2	TL	1793.0	+0.211	+0.041
Large	Cascais	5	42	1348.9	Opt	1349.6	+0.052	+0.000
Large	Odivelas	5	12	494.2	Opt	494.2	+0.000	+0.000
Large	Sintra	5	27	856.6	Opt	856.6	+0.004	+0.000
Large	Lisboa	6	69	1764.5	TL	1773.5	+0.510	+0.312
Large	Cascais	6	42	1333.2	Opt	1333.2	+0.000	+0.000
Large	Odivelas	6	12	515.5	Opt	515.5	+0.000	+0.000
Large	Sintra	6	27	841.6	Opt	841.6	+0.000	+0.000
Large	Lisboa	7	69	1739.2	TL	1748.9	+0.560	+0.394
Large	Cascais	7	42	1330.6	Opt	1331.6	+0.078	+0.000
Large	Odivelas	7	12	503.4	Opt	503.4	+0.000	+0.000
Large	Sintra	7	27	848.5	Opt	848.5	+0.000	+0.000
XLarge	Lisboa	0	92	2328.5	TL	2326.3	-0.095	-0.367
XLarge	Cascais	0	56	1621.9	Opt	1623.9	+0.127	+0.056
XLarge	Odivelas	0	16	543.3	Opt	547.9	+0.857	+0.857
XLarge	Sintra	0	36	1168.1	Opt	1171.1	+0.259	+0.026
XLarge	Lisboa	1	92	2315.7	TL	2315.6	-0.006	-0.112
XLarge	Cascais	1	56	1622.7	Opt	1623.3	+0.034	+0.000
XLarge	Odivelas	1	16	540.1	Opt	540.1	+0.000	+0.000
XLarge	Sintra	1	36	1123.4	Opt	1123.6	+0.015	+0.000
XLarge	Lisboa	2	92	2306.0	TL	2306.4	+0.018	-0.121
XLarge	Cascais	2	56	1626.4	Opt	1628.6	+0.137	+0.016
XLarge	Odivelas	2	16	545.2	Opt	545.2	+0.000	+0.000
XLarge	Sintra	2	36	1130.3	Opt	1131.7	+0.128	+0.000
XLarge	Lisboa	3	92	2298.3	TL	2298.9	+0.026	-0.167
XLarge	Cascais	3	56	1619.9	Opt	1622.9	+0.187	+0.010

continued on next page

(Table D3 continued)

Size	Depot	Rep	N_d	z_{MILP}	St.	$\bar{z}_{\text{ALNS}}$	Gap$_{\text{mean}}$	Gap$_{\text{best}}$
XLarge	Odivelas	3	16	545.1	Opt	545.1	+0.000	+0.000
XLarge	Sintra	3	36	1150.6	Opt	1151.0	+0.031	+0.000
XLarge	Lisboa	4	92	2313.9	TL	2314.5	+0.025	-0.187
XLarge	Cascais	4	56	1589.5	Opt	1589.8	+0.022	+0.000
XLarge	Odivelas	4	16	540.5	Opt	548.4	+1.461	+1.461
XLarge	Sintra	4	36	1146.4	Opt	1148.5	+0.180	+0.000
XLarge	Lisboa	5	92	2292.3	TL	2300.0	+0.334	-0.007
XLarge	Cascais	5	56	1593.1	Opt	1595.2	+0.133	+0.017
XLarge	Odivelas	5	16	551.0	Opt	551.0	+0.000	+0.000
XLarge	Sintra	5	36	1133.7	Opt	1134.0	+0.030	+0.000
XLarge	Lisboa	6	92	2337.1	TL	2335.6	-0.065	-0.409
XLarge	Cascais	6	56	1613.6	Opt	1615.4	+0.111	+0.077
XLarge	Odivelas	6	16	548.4	Opt	548.4	+0.000	+0.000
XLarge	Sintra	6	36	1126.3	Opt	1126.9	+0.049	+0.000
XLarge	Lisboa	7	92	2315.2	TL	2319.1	+0.171	+0.084
XLarge	Cascais	7	56	1603.4	Opt	1606.7	+0.206	+0.036
XLarge	Odivelas	7	16	548.5	Opt	556.2	+1.404	+1.404
XLarge	Sintra	7	36	1141.4	Opt	1143.1	+0.150	+0.000
XXLarge	Lisboa	0	116	2450.3	TL	2418.3	-1.310	-1.746
XXLarge	Cascais	0	70	1661.7	TL	1664.3	+0.156	+0.022
XXLarge	Odivelas	0	20	569.3	Opt	569.3	+0.000	+0.000
XXLarge	Sintra	0	44	1184.1	TL	1185.6	+0.126	+0.000
XXLarge	Lisboa	1	116	2444.5	TL	2410.9	-1.374	-1.452
XXLarge	Cascais	1	70	1667.7	TL	1669.4	+0.098	-0.057
XXLarge	Odivelas	1	20	559.4	Opt	559.4	+0.000	+0.000
XXLarge	Sintra	1	44	1165.0	TL	1167.0	+0.172	+0.139
XXLarge	Lisboa	2	116	2481.3	TL	2442.4	-1.568	-1.815
XXLarge	Cascais	2	70	1676.2	TL	1679.4	+0.192	+0.159
XXLarge	Odivelas	2	20	551.8	Opt	551.8	+0.000	+0.000
XXLarge	Sintra	2	44	1183.0	TL	1184.6	+0.140	+0.000
XXLarge	Lisboa	3	116	2446.4	TL	2434.7	-0.477	-0.677
XXLarge	Cascais	3	70	1652.4	TL	1652.1	-0.017	-0.164
XXLarge	Odivelas	3	20	559.8	Opt	559.8	+0.000	+0.000
XXLarge	Sintra	3	44	1167.5	Opt	1168.4	+0.074	+0.000
XXLarge	Lisboa	4	116	2452.3	TL	2435.3	-0.695	-1.088

continued on next page

(Table D3 continued)

Size	Depot	Rep	N_d	z_{MILP}	St.	$\bar{z}_{\text{ALNS}}$	Gap_{mean}	Gap_{best}
XXLarge	Cascais	4	70	1674.0	TL	1677.8	+0.232	+0.008
XXLarge	Odivelas	4	20	570.2	Opt	570.8	+0.112	+0.000
XXLarge	Sintra	4	44	1201.0	TL	1202.7	+0.137	-0.058
XXLarge	Lisboa	5	116	2462.9	TL	2438.7	-0.983	-1.130
XXLarge	Cascais	5	70	1651.5	TL	1655.5	+0.240	+0.107
XXLarge	Odivelas	5	20	557.2	Opt	557.2	+0.000	+0.000
XXLarge	Sintra	5	44	1180.0	Opt	1182.1	+0.180	+0.000
XXLarge	Lisboa	6	116	2448.1	TL	2416.2	-1.304	-1.650
XXLarge	Cascais	6	70	1669.3	TL	1674.3	+0.298	+0.192
XXLarge	Odivelas	6	20	563.3	Opt	563.3	+0.000	+0.000
XXLarge	Sintra	6	44	1187.8	TL	1188.8	+0.084	+0.000
XXLarge	Lisboa	7	116	2435.8	TL	2438.8	+0.123	-0.011
XXLarge	Cascais	7	70	1681.9	TL	1680.0	-0.111	-0.281
XXLarge	Odivelas	7	20	556.0	Opt	556.0	+0.000	+0.000
XXLarge	Sintra	7	44	1198.0	Opt	1198.5	+0.045	+0.014

Table D4: Per-instance ALNS VS Gurobi results for the **high doctor** scenario (160 depot instances).

Size	Depot	Rep	N_d	z_{MILP}	St.	$\bar{z}_{\text{ALNS}}$	Gap_{mean}	Gap_{best}
Small	Lisboa	0	23	843.8	Opt	843.8	+0.000	+0.000
Small	Cascais	0	14	531.0	Opt	531.0	+0.000	+0.000
Small	Odivelas	0	4	248.9	Opt	248.9	+0.000	+0.000
Small	Sintra	0	9	315.8	Opt	315.8	+0.000	+0.000
Small	Lisboa	1	23	828.8	Opt	828.8	+0.000	+0.000
Small	Cascais	1	14	510.4	Opt	510.4	+0.000	+0.000
Small	Odivelas	1	4	248.6	Opt	248.6	+0.000	+0.000
Small	Sintra	1	9	306.1	Opt	306.1	+0.000	+0.000
Small	Lisboa	2	23	850.8	Opt	850.8	+0.000	+0.000
Small	Cascais	2	14	541.1	Opt	541.1	+0.000	+0.000
Small	Odivelas	2	4	247.6	Opt	247.6	+0.000	+0.000
Small	Sintra	2	9	321.1	Opt	321.1	+0.000	+0.000
Small	Lisboa	3	23	628.1	Opt	630.4	+0.361	+0.000
Small	Cascais	3	14	532.3	Opt	532.3	+0.000	+0.000
Small	Odivelas	3	4	257.1	Opt	257.1	+0.000	+0.000
Small	Sintra	3	9	321.2	Opt	321.2	+0.000	+0.000
Small	Lisboa	4	23	863.3	Opt	863.3	+0.000	+0.000

continued on next page

(Table D4 continued)

Size	Depot	Rep	N_d	z_{MILP}	St.	$\bar{z}_{\text{ALNS}}$	Gap_{mean}	Gap_{best}
Small	Cascais	4	14	512.3	Opt	512.3	+0.000	+0.000
Small	Odivelas	4	4	260.8	Opt	260.8	+0.000	+0.000
Small	Sintra	4	9	287.5	Opt	287.5	+0.000	+0.000
Small	Lisboa	5	23	826.1	Opt	826.1	+0.000	+0.000
Small	Cascais	5	14	528.3	Opt	528.3	+0.000	+0.000
Small	Odivelas	5	4	257.6	Opt	257.6	+0.000	+0.000
Small	Sintra	5	9	308.8	Opt	308.8	+0.000	+0.000
Small	Lisboa	6	23	841.6	Opt	841.6	+0.000	+0.000
Small	Cascais	6	14	531.4	Opt	531.4	+0.000	+0.000
Small	Odivelas	6	4	246.9	Opt	246.9	+0.000	+0.000
Small	Sintra	6	9	301.1	Opt	301.1	+0.000	+0.000
Small	Lisboa	7	23	838.2	Opt	838.2	+0.000	+0.000
Small	Cascais	7	14	525.0	Opt	525.0	+0.000	+0.000
Small	Odivelas	7	4	247.2	Opt	247.2	+0.000	+0.000
Small	Sintra	7	9	313.1	Opt	313.1	+0.000	+0.000
Medium	Lisboa	0	46	1402.0	TL	1403.3	+0.092	+0.000
Medium	Cascais	0	28	842.7	Opt	842.7	+0.000	+0.000
Medium	Odivelas	0	8	288.2	Opt	288.2	+0.000	+0.000
Medium	Sintra	0	18	594.8	Opt	595.5	+0.108	+0.000
Medium	Lisboa	1	46	1402.6	Opt	1404.7	+0.151	+0.097
Medium	Cascais	1	28	824.1	Opt	824.1	+0.000	+0.000
Medium	Odivelas	1	8	283.2	Opt	283.2	+0.000	+0.000
Medium	Sintra	1	18	617.3	Opt	617.3	+0.000	+0.000
Medium	Lisboa	2	46	1365.1	TL	1365.1	+0.000	+0.000
Medium	Cascais	2	28	816.7	Opt	816.7	+0.000	+0.000
Medium	Odivelas	2	8	281.1	Opt	281.1	+0.000	+0.000
Medium	Sintra	2	18	581.4	Opt	581.4	+0.000	+0.000
Medium	Lisboa	3	46	1408.1	Opt	1408.1	+0.000	+0.000
Medium	Cascais	3	28	837.7	Opt	838.2	+0.064	+0.000
Medium	Odivelas	3	8	308.4	Opt	308.4	+0.000	+0.000
Medium	Sintra	3	18	570.9	Opt	571.5	+0.102	+0.000
Medium	Lisboa	4	46	1385.3	TL	1385.4	+0.009	+0.000
Medium	Cascais	4	28	839.5	Opt	839.5	+0.000	+0.000
Medium	Odivelas	4	8	280.7	Opt	280.7	+0.000	+0.000
Medium	Sintra	4	18	591.9	Opt	591.9	+0.000	+0.000

continued on next page

(Table D4 continued)

Size	Depot	Rep	N_d	z_{MILP}	St.	$\bar{z}_{\text{ALNS}}$	Gap_{mean}	Gap_{best}
Medium	Lisboa	5	46	1399.8	TL	1399.8	+0.000	+0.000
Medium	Cascais	5	28	813.4	Opt	814.1	+0.075	+0.000
Medium	Odivelas	5	8	275.0	Opt	275.0	+0.000	+0.000
Medium	Sintra	5	18	631.6	Opt	631.6	+0.000	+0.000
Medium	Lisboa	6	46	1399.9	TL	1399.9	+0.000	+0.000
Medium	Cascais	6	28	809.2	Opt	809.2	+0.000	+0.000
Medium	Odivelas	6	8	282.9	Opt	282.9	+0.000	+0.000
Medium	Sintra	6	18	593.4	Opt	593.4	+0.000	+0.000
Medium	Lisboa	7	46	1390.5	TL	1390.5	+0.000	+0.000
Medium	Cascais	7	28	820.0	Opt	820.1	+0.011	+0.000
Medium	Odivelas	7	8	280.2	Opt	280.2	+0.000	+0.000
Medium	Sintra	7	18	594.7	Opt	594.7	+0.000	+0.000
Large	Lisboa	0	69	1920.3	TL	1723.7	-10.234	-10.391
Large	Cascais	0	42	1098.5	Opt	1099.4	+0.079	+0.030
Large	Odivelas	0	12	514.5	Opt	514.5	+0.000	+0.000
Large	Sintra	0	27	855.1	Opt	855.1	+0.000	+0.000
Large	Lisboa	1	69	1753.1	TL	1760.7	+0.434	+0.154
Large	Cascais	1	42	1078.5	Opt	1080.3	+0.170	+0.000
Large	Odivelas	1	12	522.2	Opt	522.2	+0.000	+0.000
Large	Sintra	1	27	873.7	Opt	873.7	+0.000	+0.000
Large	Lisboa	2	69	1753.0	TL	1755.3	+0.133	-0.098
Large	Cascais	2	42	1113.8	Opt	1116.7	+0.258	+0.102
Large	Odivelas	2	12	522.4	Opt	522.5	+0.000	+0.000
Large	Sintra	2	27	853.9	Opt	853.9	+0.000	+0.000
Large	Lisboa	3	69	1938.9	TL	1822.9	-5.982	-10.284
Large	Cascais	3	42	1109.7	Opt	1110.2	+0.045	+0.000
Large	Odivelas	3	12	496.6	Opt	496.6	+0.000	+0.000
Large	Sintra	3	27	893.8	Opt	893.8	+0.000	+0.000
Large	Lisboa	4	69	1738.5	TL	1748.3	+0.567	+0.327
Large	Cascais	4	42	1068.8	Opt	1070.3	+0.138	+0.000
Large	Odivelas	4	12	535.8	Opt	535.8	+0.000	+0.000
Large	Sintra	4	27	859.9	Opt	859.9	+0.000	+0.000
Large	Lisboa	5	69	1728.6	TL	1726.9	-0.100	-0.249
Large	Cascais	5	42	1096.3	Opt	1101.5	+0.476	+0.000
Large	Odivelas	5	12	514.7	Opt	514.7	+0.000	+0.000

continued on next page

(Table D4 continued)

Size	Depot	Rep	N_d	z_{MILP}	St.	$\bar{z}_{\text{ALNS}}$	Gap_{mean}	Gap_{best}
Large	Sintra	5	27	874.8	Opt	875.4	+0.072	+0.000
Large	Lisboa	6	69	1740.9	TL	1750.2	+0.534	+0.251
Large	Cascais	6	42	1113.0	Opt	1114.5	+0.137	+0.000
Large	Odivelas	6	12	509.6	Opt	509.6	+0.000	+0.000
Large	Sintra	6	27	867.2	Opt	867.2	+0.000	+0.000
Large	Lisboa	7	69	1732.5	TL	1923.5	+11.021	+10.998
Large	Cascais	7	42	1090.4	Opt	1092.7	+0.205	+0.157
Large	Odivelas	7	12	524.7	Opt	524.7	+0.000	+0.000
Large	Sintra	7	27	877.9	Opt	877.9	+0.000	+0.000
XLarge	Lisboa	0	92	2265.4	TL	2261.9	-0.155	-0.368
XLarge	Cascais	0	56	1386.9	TL	1513.2	+9.105	+0.279
XLarge	Odivelas	0	16	525.2	Opt	525.2	+0.000	+0.000
XLarge	Sintra	0	36	1115.8	Opt	1116.5	+0.067	+0.000
XLarge	Lisboa	1	92	2279.5	TL	2271.5	-0.352	-0.521
XLarge	Cascais	1	56	1599.5	TL	1602.0	+0.158	+0.018
XLarge	Odivelas	1	16	535.0	Opt	535.0	+0.000	+0.000
XLarge	Sintra	1	36	1135.1	Opt	1136.1	+0.095	+0.005
XLarge	Lisboa	2	92	2269.5	TL	2316.9	+2.090	-0.132
XLarge	Cascais	2	56	1600.4	TL	1604.7	+0.272	+0.170
XLarge	Odivelas	2	16	535.7	Opt	535.7	+0.000	+0.000
XLarge	Sintra	2	36	1144.6	Opt	1145.3	+0.061	+0.056
XLarge	Lisboa	3	92	2280.7	TL	2260.6	-0.881	-1.249
XLarge	Cascais	3	56	1568.3	TL	1568.8	+0.029	+0.000
XLarge	Odivelas	3	16	547.7	Opt	547.7	+0.000	+0.000
XLarge	Sintra	3	36	1135.7	Opt	1135.9	+0.014	+0.000
XLarge	Lisboa	4	92	2275.0	TL	2275.3	+0.013	-0.156
XLarge	Cascais	4	56	1556.8	TL	1558.5	+0.110	+0.042
XLarge	Odivelas	4	16	536.7	Opt	536.8	+0.022	+0.000
XLarge	Sintra	4	36	1101.3	Opt	1101.3	+0.000	+0.000
XLarge	Lisboa	5	92	2300.8	TL	2286.6	-0.620	-0.992
XLarge	Cascais	5	56	1572.8	TL	1571.1	-0.112	-0.190
XLarge	Odivelas	5	16	530.4	Opt	530.4	+0.000	+0.000
XLarge	Sintra	5	36	1131.3	Opt	1133.5	+0.200	+0.000
XLarge	Lisboa	6	92	2260.6	TL	2284.3	+1.048	+0.604
XLarge	Cascais	6	56	1585.5	TL	1585.9	+0.019	-0.015

continued on next page

(Table D4 continued)

Size	Depot	Rep	N_d	z_{MILP}	St.	$\bar{z}_{\text{ALNS}}$	Gap_{mean}	Gap_{best}
XLarge	Odivelas	6	16	515.5	Opt	515.5	+0.000	+0.000
XLarge	Sintra	6	36	1122.9	Opt	1123.4	+0.042	+0.000
XLarge	Lisboa	7	92	2267.7	TL	2270.7	+0.133	-0.618
XLarge	Cascais	7	56	1558.4	TL	1560.7	+0.152	+0.091
XLarge	Odivelas	7	16	539.8	Opt	539.8	+0.000	+0.000
XLarge	Sintra	7	36	1129.2	Opt	1129.5	+0.025	+0.000
XXLarge	Lisboa	0	116	2589.5	TL	2573.8	-0.606	-0.811
XXLarge	Cascais	0	70	1633.8	TL	1633.8	+0.000	-0.150
XXLarge	Odivelas	0	20	553.4	Opt	553.4	+0.000	+0.000
XXLarge	Sintra	0	44	1163.4	TL	1165.7	+0.196	+0.088
XXLarge	Lisboa	1	116	2619.2	TL	2605.3	-0.533	-1.083
XXLarge	Cascais	1	70	1666.4	TL	1667.1	+0.043	-0.055
XXLarge	Odivelas	1	20	556.7	Opt	556.7	+0.000	+0.000
XXLarge	Sintra	1	44	1189.5	Opt	1191.2	+0.142	+0.003
XXLarge	Lisboa	2	116	2602.4	TL	2581.9	-0.786	-1.116
XXLarge	Cascais	2	70	1635.6	TL	1638.8	+0.193	+0.037
XXLarge	Odivelas	2	20	564.2	Opt	564.2	+0.000	+0.000
XXLarge	Sintra	2	44	1188.6	Opt	1190.4	+0.146	+0.000
XXLarge	Lisboa	3	116	2612.6	TL	2608.8	-0.142	-0.421
XXLarge	Cascais	3	70	1657.4	TL	1655.8	-0.097	-0.229
XXLarge	Odivelas	3	20	540.6	Opt	540.6	+0.000	+0.000
XXLarge	Sintra	3	44	1192.9	Opt	1193.5	+0.052	+0.000
XXLarge	Lisboa	4	116	2572.5	TL	2576.9	+0.172	-0.094
XXLarge	Cascais	4	70	1669.2	TL	1672.9	+0.225	+0.062
XXLarge	Odivelas	4	20	534.9	Opt	534.9	+0.000	+0.000
XXLarge	Sintra	4	44	1176.1	Opt	1176.9	+0.065	+0.000
XXLarge	Lisboa	5	116	2582.5	TL	2576.3	-0.240	-0.629
XXLarge	Cascais	5	70	1663.8	TL	1663.1	-0.044	-0.104
XXLarge	Odivelas	5	20	560.2	Opt	560.2	+0.000	+0.000
XXLarge	Sintra	5	44	1187.0	Opt	1187.8	+0.070	+0.000
XXLarge	Lisboa	6	116	2636.6	TL	2592.5	-1.673	-1.985
XXLarge	Cascais	6	70	1656.6	TL	1652.4	-0.257	-0.314
XXLarge	Odivelas	6	20	548.1	Opt	548.1	+0.000	+0.000
XXLarge	Sintra	6	44	1184.3	TL	1185.0	+0.057	+0.000
XXLarge	Lisboa	7	116	2681.4	TL	2596.8	-3.155	-3.287

continued on next page

(Table D4 continued)

Size	Depot	Rep	N_d	z_{MILP}	St.	$\bar{z}_{\text{ALNS}}$	Gap_{mean}	Gap_{best}
XXLarge	Cascais	7	70	1651.6	TL	1650.3	-0.078	-0.229
XXLarge	Odivelas	7	20	551.3	Opt	551.3	+0.000	+0.000
XXLarge	Sintra	7	44	1185.1	Opt	1186.8	+0.144	+0.087

Appendix E: Convergence behaviour

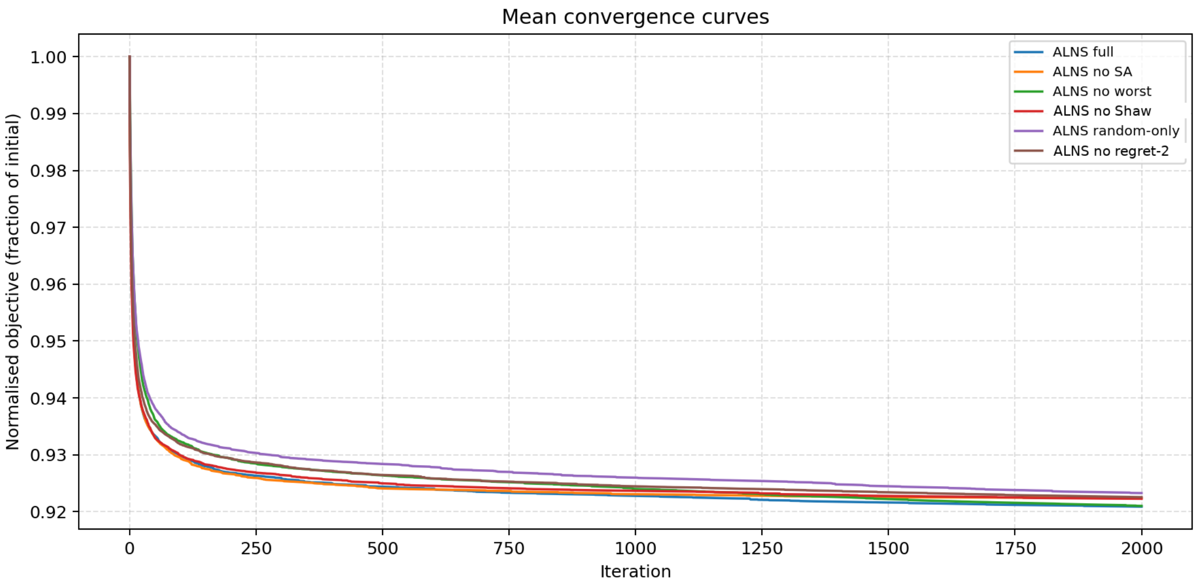

Figure E1: Average convergence curves for the ALNS ablation variants (objective normalised by the initial solution).

Appendix F: Prototype Outputs

Table F1: Team 5's visit sequence at Lisboa depot.

Visit Order	Patient ID	Urgent	Need a Doctor	Service Time (min)
1	24	1	0	25
2	103	1	0	35
3	204	0	0	20
4	62	0	0	20
5	25	0	0	20
6	27	0	0	20
7	57	0	0	20

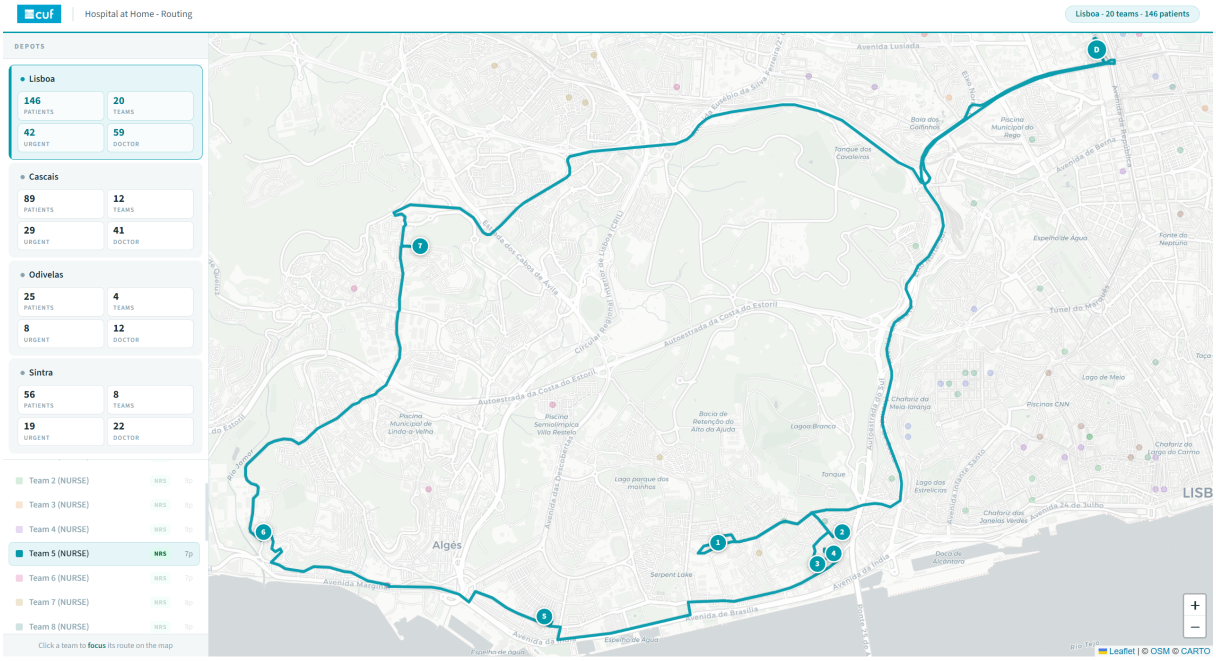

Figure F1: Focused view of Team 5 route at Lisboa depot.

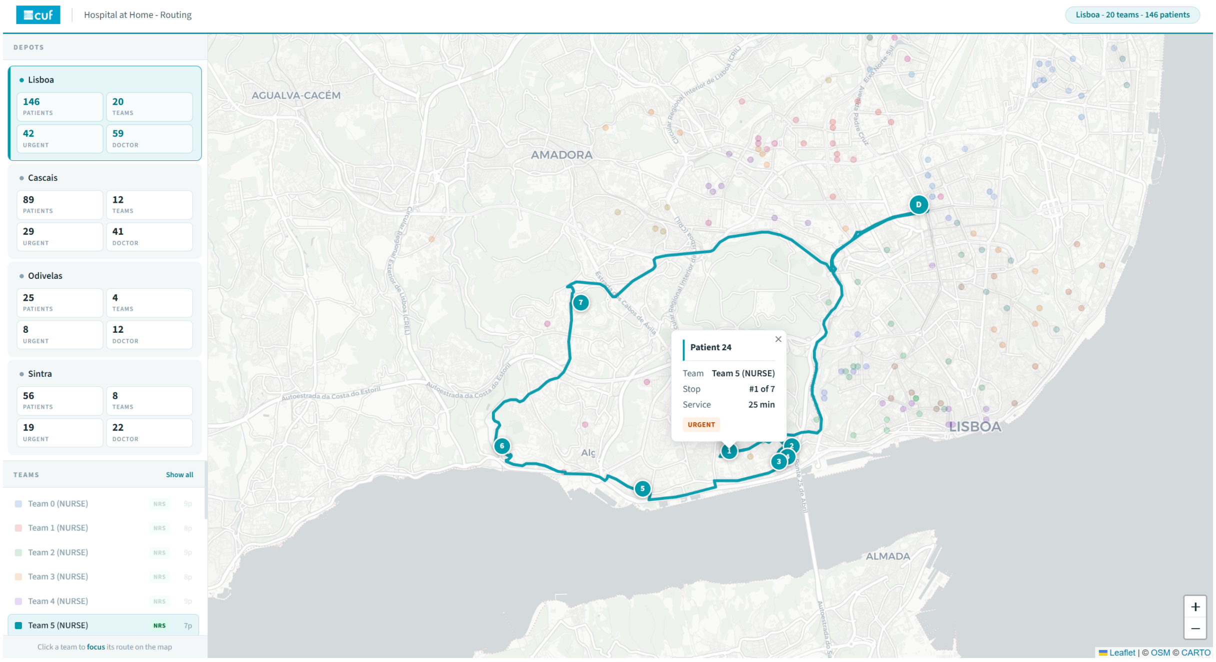

Figure F2: Patient pop-up.

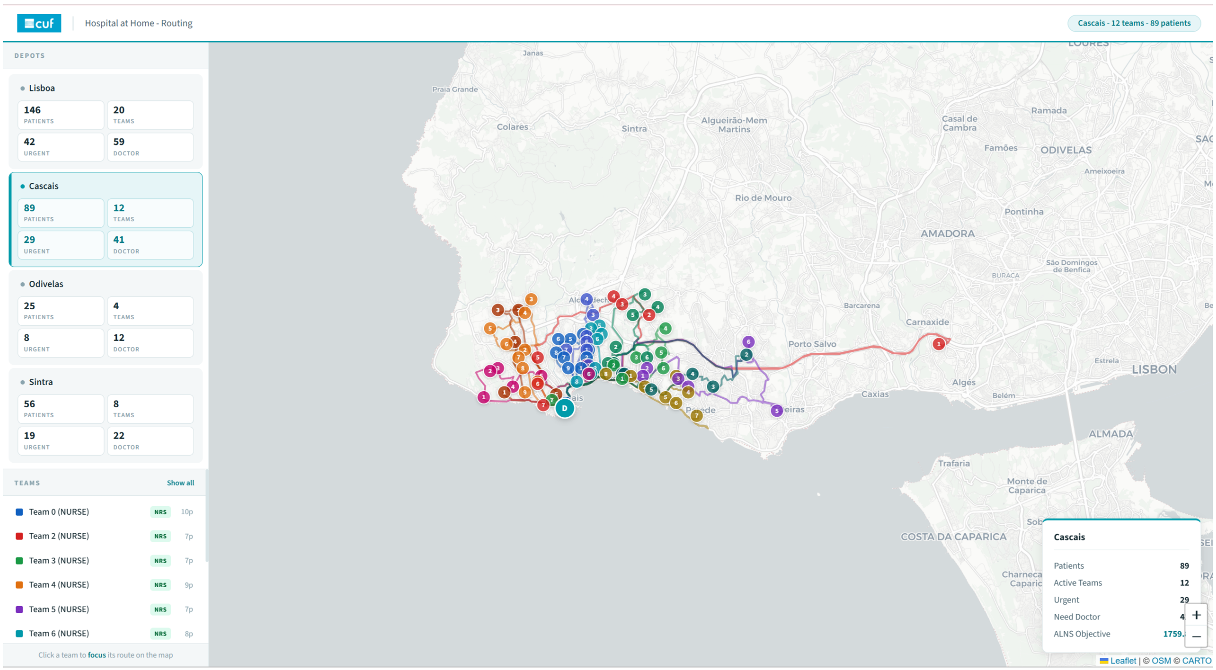

Figure F3: Cascais routes.

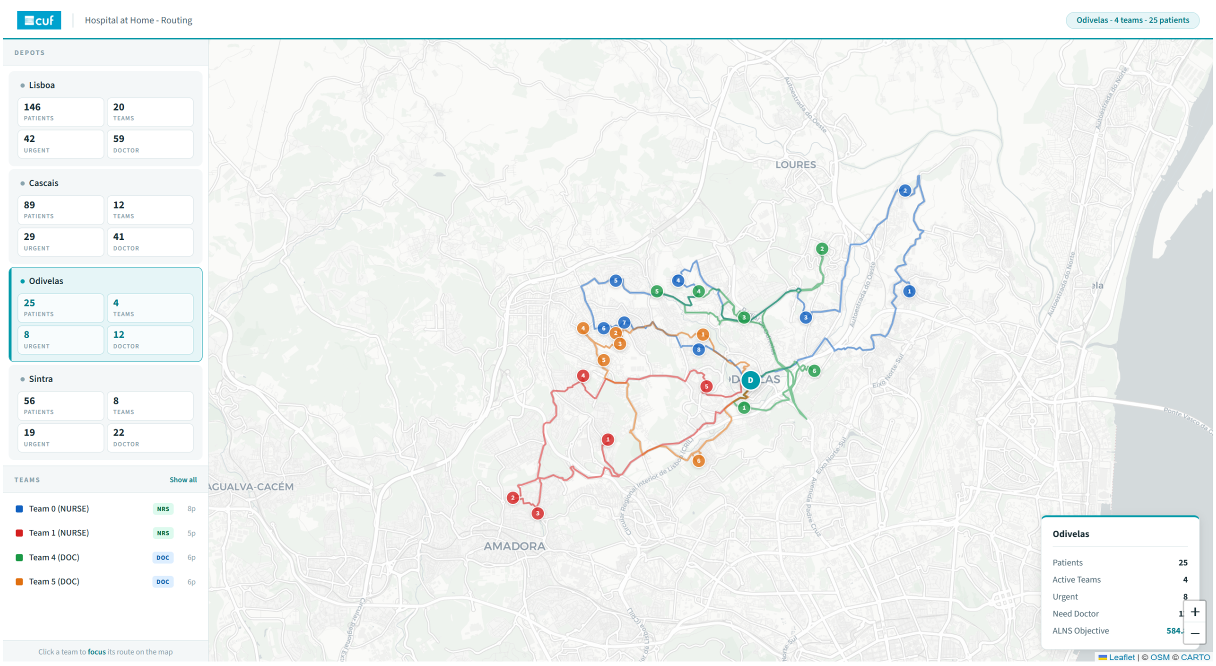

Figure F4: Odivelas routes.

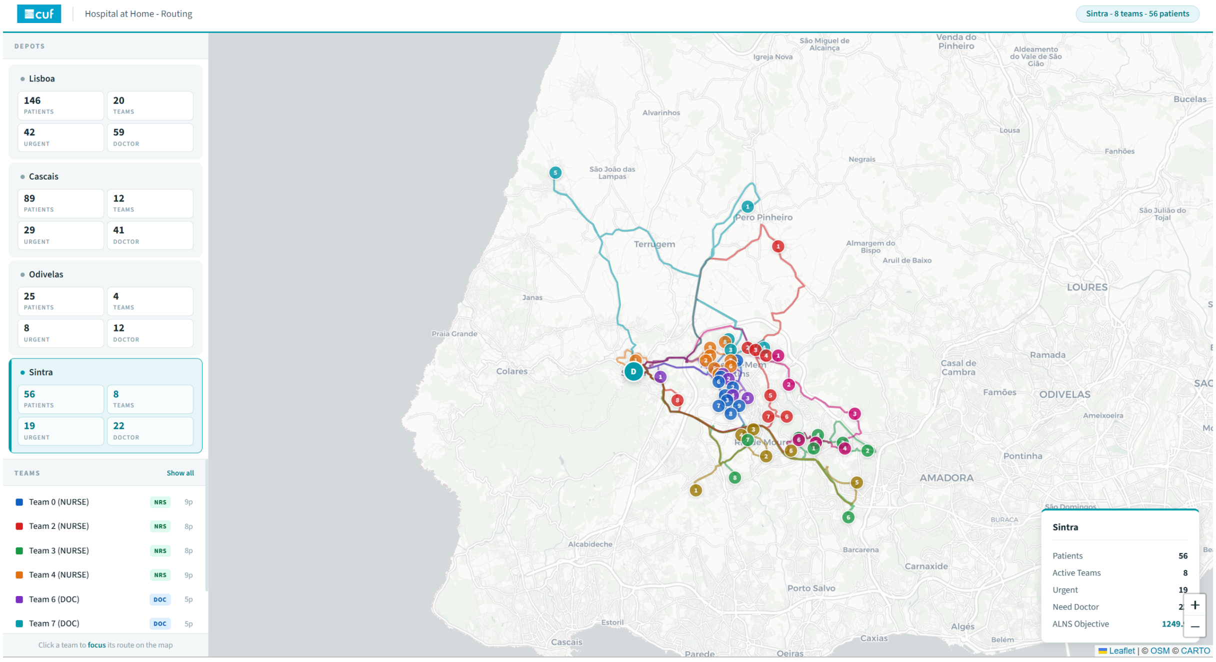

Figure F5: Sintra routes.

UNIVERSITÉ CATHOLIQUE DE LOUVAIN
Louvain School of Management
Place des Doyens, 1 bte L2.01.01, 1348 Louvain-la-Neuve, Belgique | **www.uclouvain.be/lsm**