Logistics Strategies for Emergency Medical Services

Changle Song

A thesis presented for the degree of
Doctor of Philosophy (Engineering)

School of Civil Engineering
Faculty of Engineering
The University of Sydney
Australia
November 10, 2025

Supervisor: Dr. Emily Moylan
Prof. David Levinson

Abstract

This thesis advances Emergency Medical Services (EMS) logistics by developing innovative delivery strategies, optimizing facility locations, and designing adaptive dispatching policies to enhance response efficiency and resource allocation. Recognizing limitations in traditional EMS models, this research explores alternative logistics strategies—including rendezvous and pre-hospital models—that offer greater flexibility in high-demand, resource-constrained environments, laying the groundwork for adaptive solutions in EMS.

A comprehensive literature review identifies and critiques current EMS strategies, revealing critical gaps in facility location planning, dispatch policies, and resource management in real-time. These findings establish a foundation for the research's models, prioritizing patient outcomes, and advancing EMS logistics to meet different healthcare needs.

To address spatial disparities in EMS access, the thesis employs facility location optimization, determining optimal resource placement for rendezvous delivery strategies. Analyzing geographic and demographic factors, these methods enhance response times by ensuring EMS resources are strategically positioned to maximize accessibility and patient medical outcomes.

Dispatch policies are explored through agent-based simulation, with a focus on rule-based assignment methods that adapt dynamically to resource types, incident severity, and varying performance metrics across diverse resources and demand types. This approach enables a nuanced allocation of resources, tailoring responses to the specific needs of each incident while optimizing overall system performance.

In conclusion, this research offers actionable insights into EMS system design and implementation, with broader implications for emergency logistics and urban health policy.

Acknowledgements

I would like to express my deepest gratitude to my primary advisor, Dr. Emily Moylan, and my second supervisor, Prof. David Levinson, for their unwavering mentorship and generous support throughout my doctoral studies. Their expert guidance, insightful feedback, and persistent encouragement have been instrumental in shaping both my research and my growth as a scholar. I am also sincerely grateful to Mark and Brian for their timely and thoughtful assistance, and to Dr. Hao Wu for his help in processing travel time data.

This research was supported by the NSW Health Translational Research Grant Scheme (EVIDENCE Study) and the National Heart Foundation. It was partially funded by iMOVE CRC and received support from the Cooperative Research Centres program, an Australian Government initiative. I gratefully acknowledge the data and advice provided by NSW Ambulance, particularly Sophie Dyson, Sandra Ware, and Harry Lai. The use of historical cardiac arrest data was reviewed and approved by the NSW Population and Health Services Research Ethics Committee (2021/ETH01302). Consent requirements were waived due to the retrospective nature of the study and the inability of cardiac arrest patients to provide informed consent. The National Heart Foundation had no role in the study design, data collection, analysis, interpretation of the data, or the writing and submission of this thesis.

Finally, I owe my greatest appreciation to my family, whose unwavering support, patience, and belief in my abilities sustained me throughout the challenges of this PhD journey. Their constant encouragement and love provided the foundation upon which this work was built.

Statement of Originality

This is to certify that to the best of my knowledge, the content of this thesis is my own work. This thesis has not been submitted for any degree or other purposes.

I certify that the intellectual content of this thesis is the product of my own work and that all the assistance received in preparing this thesis and sources have been acknowledged.

Changle Song
January 28, 2025

Declaration

I hereby declare that the work presented in this thesis is my own and has not been submitted for a degree or other qualification at any other institution. To the best of my knowledge, all sources of information and contributions by others have been properly cited.

Two chapters of this thesis have been published in peer-reviewed journals. Specifically, Chapter 3 was published in the Scandinavian Journal of Trauma, Resuscitation and Emergency Medicine under the title Improving Access to Extracorporeal Membrane Oxygenation for Out-of-Hospital Cardiac Arrest: Pre-hospital ECPR and Alternate Delivery Strategies. Chapter 4 was published in the Journal of Transport Geography with the title Improved Spatial Equity in Healthcare Access Through Novel Logistics Strategies.

Changle Song
October 31, 2025

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Introduction to EMS and the Role of Logistics

1.1.1 Overview of Emergency Medical Services

Emergency Medical Services (EMS) are a vital part of public healthcare systems (Aringhieri et al., 2017). They provide immediate medical care and transportation for people experiencing sudden health emergencies. The main goal of EMS is to offer essential medical intervention before patients receive definitive treatment in a hospital, helping to stabilize their condition, relieve their suffering and improve their chances of recovery. Prompt action by EMS teams is especially important in cases like injuries and infections, heart attacks and strokes, asthma and acute complications of pregnancy (World Health Organization, n.d.a).

EMS systems usually include different levels of service, from Basic Life Support (BLS) to Advanced Life Support (ALS) (Su and Shih, 2003), which vary in different regions. They involve various professionals such as paramedics, pre-hospital doctors and Emergency Medical Technicians. The vehicles and their equipment can also differ. The EMS system relies on a hierarchical system of ambulance bases and hospitals that support tiered treatment through the allocation of resources and personnel. These elements work together through coordinated dispatch systems to respond quickly and effectively to emergencies.

Various societal factors contribute to increased pressure on EMS— for example, a growing population leads to more potential patients and wealth polarization exacerbates disparities in healthcare access (Andrew et al., 2020). This increase in demand exists within the context of congestion and urban sprawl and evolving healthcare treatments and technologies. This underscores the importance of efficient EMS systems that can respond rapidly to different types of medical emergencies with better outcomes.

The development of modern EMS can be traced back to medical practices on battlefields in the early 20th century, where the focus was on quickly evacuating injured soldiers and providing basic first aid. Over time, these practices were adapted for civilian use, especially during the mid-20th century (Committee on the Future of Emergency Care in the United States Health System, 2007). A significant moment in EMS history was the publication of the report *Accidental Death and Disability: The Neglected Disease of Modern Society* in 1966 (National Research Council, 1966). This report highlighted serious shortcomings in emergency medical care in the United States and led to major reforms, including the creation of paramedic programs, standardized emergency communication systems, and specialized emergency vehicles. Emergency Medical Services in African countries also face significant challenges, including a lack of trained professionals, program immaturity, and limited stakeholder engagement (Mosadeghrad et al., 2019). The availability of facilities and equipment in both BLS and ALS ambulances was inadequate in Sri Lanka (Nandasena and Abeysena, 2019).

Advancements in technology throughout the 20th and 21st centuries have greatly transformed EMS. Some technologies, such as automated external defibrillators (AEDs) (Committee on the Future of Emergency Care in the United States Health System, 2007) improve treatment outcomes. Other tools, like real-time dispatch systems and Geographic Information Systems (GIS), help EMS providers respond more quickly and efficiently to emergencies. These innovative applications have been particularly important in urban areas, where heavy traffic congestion and high population density make it more challenging to reach patients promptly.

EMS system is already an essential component of public health and safety infrastructure in most countries and regions. It plays a crucial role in improving patient outcomes during medical emergencies (MacFarlane and Benn, 2003). As the demand for EMS continues to rise, addressing management challenges and finding ways to make EMS operations more efficient remain important goals. Effective logistics that assist the coordination and management of EMS resources are crucial for reducing response times and improving service delivery.

1.1.2 Demonstrated importance of EMS logistics in public health

The logistics of EMS not only affect individual patient outcomes but also have significant implications for population health and community resilience.

In times of public health emergencies, natural disasters, or mass casualty incidents, the capacity of EMS logistics to adapt and respond effectively is crucial. During large-scale emergencies like earthquakes, hurricanes, or pandemics such as COVID-19, EMS systems face immense pressure due to sudden increases in demand and potential resource limitations (World Health Organization, n.d.b). The ability to rapidly mobilize ambulances, reassign personnel, and manage supply chains determines how effectively EMS can respond to these crises. For instance, during the COVID-19 pandemic, EMS providers need to quickly adjust their operations to handle increased demand volumes, implement new safety protocols, and manage shortages of critical equipment to meet the increasing demand (Li et al., 2023).

Longer response times can lead to worse health in life-threatening situations, which is a particular concern for rural areas and underserved urban neighborhoods (Luo et al., 2022). EMS logistics play a vital role in addressing health disparities because they address a number of strategies for minimising response time in a resource constrained system and can be complemented with strategies like treatments, community education programs and advanced communication methods (Farcas et al., 2023). Optimizing logistics to improve EMS coverage in these areas is essential for ensuring equitable access to emergency care in a population-health context.

Effective EMS logistics contribute to the overall resilience of healthcare systems by enhancing their ability to respond to both everyday emergencies and extraordinary events. By strengthening logistical frameworks, EMS can reduce response times, allocate resources more efficiently and address inequitable distribution of service. These interventions will in turn better support public health objectives, improve population health outcomes, and increase community preparedness for future crises.

1.1.3 Applications of logistics in EMS

Logistics are crucial to the effective operation of EMS because most stages of EMS responses in Fig. 1.1 require movement of paramedics, services, and patients. The main elements include facility locations, personnel and resource types, delivery methods, and dispatching systems. Each of these components must be carefully managed to ensure timely and efficient emergency care.

Facility Locations: The placement of EMS facilities, such as ambulance stations and medical centers, is vital for reducing response times. Strategically located facilities enable ambulances to reach patients quickly and transport them to appropriate care centers. Geographic Information Systems (GIS) are often used to analyze data on population density, incident locations, and road networks to determine optimal facility locations (Jia et al., 2007).

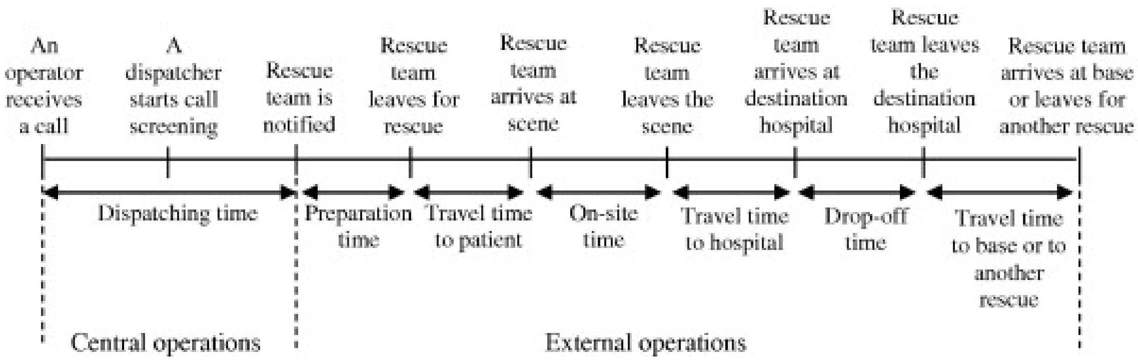

Figure 1.1: EMS Response Process (Aboueljinane et al., 2013)

Personnel and Resources: EMS personnel, including paramedics and Emergency Medical Technicians, must possess the necessary skills and qualifications to provide appropriate care. Their training levels determine the types of medical interventions they can perform, such as BLS or ALS. Ensuring that the EMS workforce has sufficient clinical exposure (Bray et al., 2020) and a mix of skills allows for better response to various medical emergencies. Additionally, the types of resources and equipment available—like defibrillators, ventilators, and medications—affect the level of care that can be provided on the scene (Nitzschke et al., 2017). Although more on-scene time will take longer out-of-hospital time, advanced pre-hospital interventions can improve patients' chances of survival (Ter Avest et al., 2022).

Delivery Methods:We can use the elements from graph theory to represent logistics or delivery strategies, where nodes are locations (e.g., suppliers, warehouses, customer delivery points) and edges are connections (e.g., transportation routes or data flows) between them. For example, a supplier node connects to a warehouse node via a directed edge representing the movement. Edges can have weights to indicate distance, cost, or delivery time.

For instance, in Fig. 1.2, the three nodes represent the delivery rider, the restaurant, and the customer. The links illustrate the movement directions within the food delivery strategy: dashed lines indicate trips with food, while solid lines denote trips without food. Fig. 1.2a depicts the scenario where the customer dines directly at the restaurant. In contrast, Fig. 1.2b illustrates an alternative where the customer picks up the food from the restaurant and subsequently returns. Two types of delivery services are highlighted: Fig. 1.2c represents in-house delivery managed directly by the restaurant, while Fig. 1.2d demonstrates third-party logistics provided by an external platform. Adapting a suitable delivery strategy for a certain scenario to meet a specific target can reform the logistics system. The same concept of designing and comparing delivery strategies underlies the EMS context.

Dispatch Systems: Dispatch centers are responsible for sending the appropriate

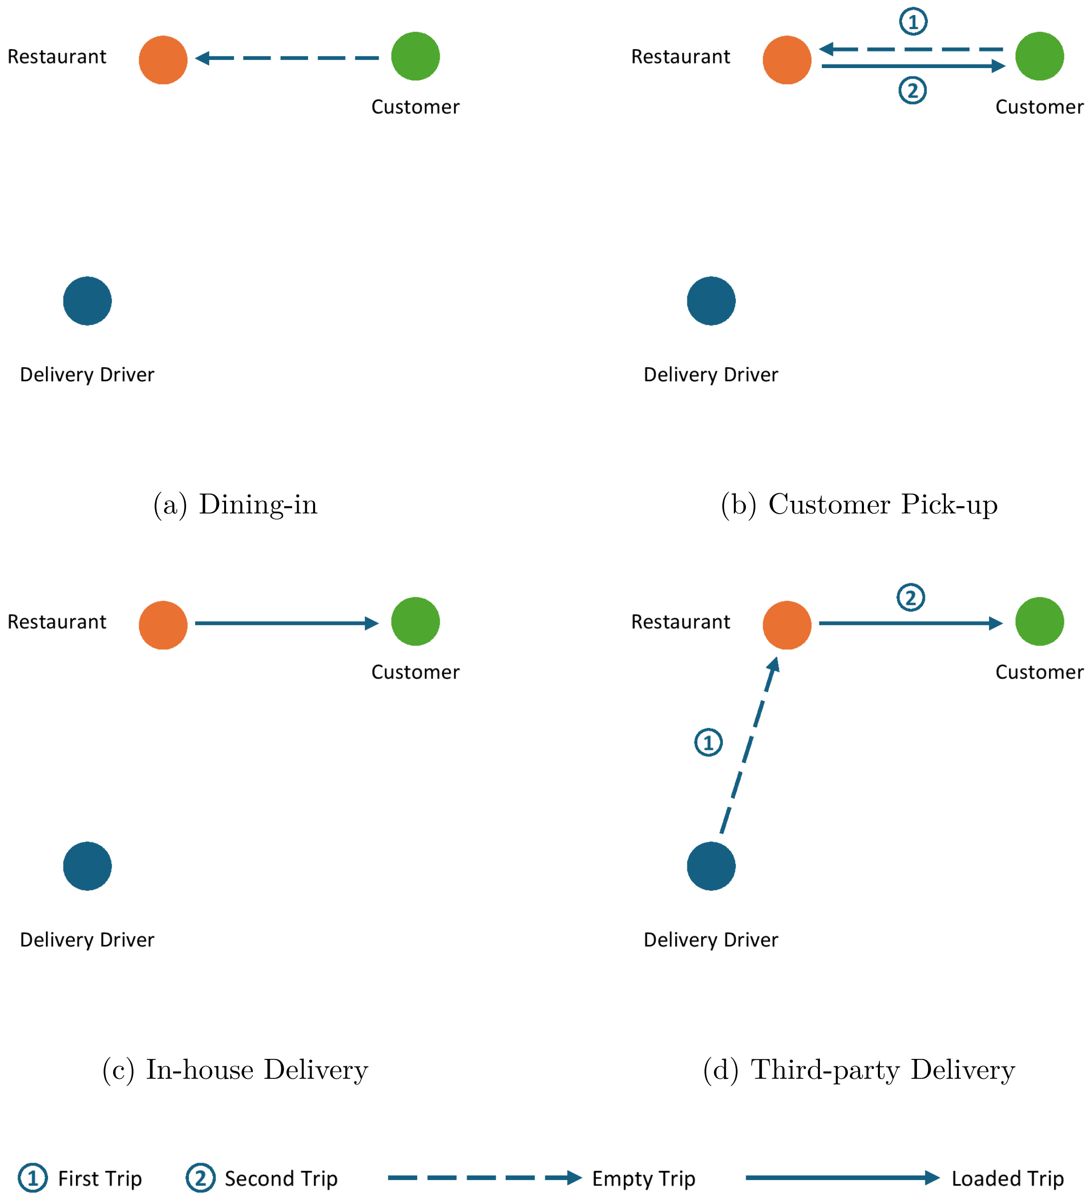

Figure 1.2: Graphs for Meal Service Options

EMS resources to emergency locations. They evaluate emergency calls and determine the best response (Peyravi et al., 2013). The assignment problem is assigning adequate ambulances to patients based on their conditions (Drabecki et al., 2023), when reassignment is redirecting an ambulance en route to one emergency to another(Guigues et al., 2022). A well-functioning dispatch system ensures that EMS units are sent quickly and efficiently, improving response times.

Overall, effective logistics in EMS—covering the placement of facilities, the suitability of staff and resources, the choice of delivery methods, and the efficiency of dispatch systems—are all essential for ensuring timely care. Improving these logistical aspects can lead to better use of resources and faster responses. With growing demand for EMS, finding ways to improve these operations remains a critical focus for policymakers and health service planners.

1.1.4 Transportation Engineering in EMS Logistics

Transportation engineering provides a practical and analytical foundation for improving EMS. While the previous sections outlined how EMS operates and the challenges it faces, transportation research contributes tools to describe and optimize the system's spatial and temporal performance. During the past decade, methods originally developed for transport planning and logistics have been adapted to medical emergency settings, linking the dynamics of the road network with patient outcomes.

Many works have examined where to locate emergency resources. Facility location models determine the placement of ambulance stations and emergency departments to maximize coverage and reduce inequality in accessibility (Brotcorne et al., 2003). Building on this, studies in vehicle routing and dispatching use dynamic optimisation to shorten travel times and coordinate mixed fleets under uncertain demand (Pillac et al., 2013). The idea of accessibility, long central to transport planning, has also been extended to EMS to measure the population that can be reached within clinically critical time limits (Mansourihanis et al., 2024).

Beyond these, demand forecasting models help anticipate call volumes and spatial-temporal patterns using historical data, machine learning, and demographic projections (Altarawneh, 2023). Recent research in traffic flow modelling and network reliability brings further realism by representing congestion, temporal variation and road capacity when estimating response performance (Bimpou and Ferguson, 2020). In parallel, simulation approaches such as agent-based and discrete-event models make it possible to test dispatch and rendezvous strategies in complex, uncertain environments (Cabrera et al., 2011; Lam et al.).

Emerging areas link EMS with resilience planning, intermodal integration, and real-time sensing, showing the breadth of transportation perspectives applied to EMS (Cimellaro et al., 2013). Together, these contributions show that transportation engineering provides more than static optimisation—- it offers a coherent framework for analysing EMS as a dynamic transport system. This framework underpins the methodological approach adopted in the following chapters, where accessibility modelling, spatial optimisation and simulation are combined to evaluate alternative delivery strategies.

1.2 Logistical Challenges in Optimizing EMS Systems

1.2.1 Evaluating EMS Delivery Models

The choice of delivery model in EMS significantly influences the overall effectiveness and efficiency of emergency response systems. Traditional models often emphasize transporting patients to centralized facilities equipped with specialized resources. However, these models are frequently hindered by delays caused by travel distance, limited resource availability, and traffic congestion.

Referring to the delivery strategy definition in Section 1.1.3, we can modify the logistics strategy by changing the nodes and links. The alternative logistics strategy, such as the rendezvous and pre-hospital strategies have the ability to reduce the delay to advanced care. These models each come with their own logistical complexities, as they attempt to balance the trade-offs between accessibility, patient outcomes, and resource utilization.

Effective evaluation of these EMS delivery models requires detailed consideration of patient catchment areas, travel times, and the capacity of each strategy to provide timely care. The need to assess spatial and temporal factors makes the optimization of these delivery models a critical challenge in EMS logistics. Transportation engineering principles, such as accessibility metrics and travel time analysis, are vital in understanding and improving the performance of these diverse EMS delivery strategies.

1.2.2 Strategic Resource Allocation in EMS System

One of the central challenges in optimizing EMS systems is the strategic allocation of resources—primarily ambulances, medical personnel, and specialized equipment, which we introduced as personnel and resources in Section 1.1.3. The geographic distribution of these resources directly influences the system's ability to respond to emergencies within critical time windows. The complexities of urban environments, with their high population densities and traffic congestion, contrast sharply with rural areas, where low population densities and long distances create different logistical problems Alanazy et al. (2019). The facility location planning part in Section 1.1.4 give us a direction to solve the problems.

Strategic resource allocation must consider not only static demand patterns but also temporal fluctuations and unpredictable spikes in emergency incidents. Facility location models such as the MCP provide frameworks for optimizing ambulance station locations to maximize coverage and minimize response times. These models must be adapted to account for the spatial distribution of demand and the variability in emergency incident rates over time.

However, static models alone are insufficient for optimizing EMS logistics. The highly dynamic nature of emergency demand requires the integration of predictive analytics to adjust for peak periods, as well as real-time factors such as traffic conditions and resource availability. By combining static optimization with dynamic adjustments, EMS systems can achieve more effective resource distribution, ensuring that high-priority cases are handled with minimal delay while maintaining sufficient coverage for lower-priority emergencies. Spatial optimization, coupled with the use of real-time data, can enhance EMS network performance across a wide range of geographic and population settings.

1.2.3 Real-time Operations: Dispatching and Adaptation

The static optimization of facility locations and resource distribution provides a critical foundation for EMS logistics, but the real-time operation of dispatch systems is where the practical challenges of EMS delivery manifest. Dispatching systems must dynamically allocate resources in response to the immediate demands of emergency incidents (Lam et al., 2015). This includes adjusting for geographic proximity, traffic conditions, and the prioritization of high-severity cases. The complexity of this task is particularly acute in multi-priority systems, where ambulances may need to be reassigned mid-response to higher-priority emergencies, potentially delaying lower-priority cases (Silva and Serra, 2008).

The challenge in EMS dispatching lies in balancing speed and accuracy. A system that reacts too slowly may result in delayed care for critical patients, while an overly reactive system may create inefficiencies by constantly reallocating resources and disrupting coverage for non-critical incidents. Optimization algorithms, such as those based on agent-based modeling or heuristic methods, have been developed to manage this balance. These algorithms simulate a wide variety of demand scenarios, helping to identify dispatching strategies that minimize delays for critical incidents while maintaining overall system efficiency.

Technological advancements have also improved the operational capacity of EMS systems by integrating real-time traffic data and Geographic Information Systems (GIS). Real-time routing adjustments, based on current traffic patterns and incident locations, allow ambulances to reach patients more quickly, particularly in congested urban areas. However, despite these technological improvements, many EMS systems lack fully integrated platforms that link dispatching operations with real-time transportation data, creating opportunities for further research and development in optimizing dynamic EMS operations (Aboueljinane et al., 2013). The modeling and simulation methods mentioned in Section 1.1.4 could help us examine different dispatching strategies.

1.3 Modeling and Simulation Framework in EMS Operations

1.3.1 Static Optimization for Delivery Strategies

This section formalises EMS delivery strategies using static optimisation models to provide a foundation for the analytical and simulation work that follows. The use of static optimization models is central to solving long-term planning problems, particularly in the context of EMS delivery strategies such as in-hospital, rendezvous, and pre-hospital models. These models are essential for determining the optimal locations for facilities and resources, such as ambulance stations and specialized care centers, to minimize response times and maximize population coverage (Jánošíková et al., 2021).

Facility Location Optimization: One of the most common static optimization techniques used in EMS planning is the MCP or Maximal Covering Location Problem (MCLP). The model identifies the best locations for bases and facilities within predefined clinical time limits—typically the threshold associated with patient survival in time-critical conditions such as OHCA (Coskun and Erol, 2010). Integrating these clinical thresholds connects spatial optimisation directly with patient outcomes.

Static optimization models are especially useful at the planning scale for long term decision making. By optimizing the location of facilities based on real-world travel data and population distributions, these models can ensure efficient resource allocation that accounts for spatial factors. At this scale, planners must also consider the trade-off between dense urban demand and sparsely populated regions where travel is faster but coverage is limited.

Application to Delivery Strategies: Delivery strategies in EMS, such as the in-hospital, rendezvous, and pre-hospital models, each present unique challenges for static optimization. For in-hospital models, the key challenge lies in ensuring that specialized care facilities are accessible to the largest possible population within critical time limits. The rendezvous model introduces an additional layer of complexity, as it involves identifying optimal intermediate locations where specialized teams can meet patients for quicker interventions. Pre-hospital models, where advanced care is initiated at the scene, require a broader distribution of resources, necessitating a more widespread network of ambulances equipped with specialized medical technologies (Song et al., 2022).

By applying static optimization techniques to each of these strategies, EMS systems can balance the trade-offs between resource availability, population coverage, and response times. These models ensure that EMS providers can deliver care efficiently, even in geographically diverse regions with varying infrastructure and demand patterns.

1.3.2 Dynamic Simulation for Dispatching

While static optimization is vital for strategic planning, the real-time operation of EMS systems demands dynamic simulation models that can handle the unpredictability of emergency incidents. Dynamic simulation allows EMS systems to adapt to real-time factors such as incident severity, vehicle availability, and traffic conditions, which are critical for effective dispatching and resource management in fast-paced environments (Wang et al., 2021).

A common approach to dynamic simulation in EMS is agent-based modeling (ABM), where individual agents represent ambulances, emergency incidents, and other system components (Aboueljinane et al., 2013). In this thesis, agents (vehicles, patients, and facilities) follow rule-based behaviours under given dispatch policies. This structure captures coordination among simultaneous incidents while remaining consistent with the operational logic of EMS systems.

ABM also enables the testing of different dispatch policies in various scenarios, providing EMS operators with insights into how different strategies impact system performance. For instance, in a multi-priority system, higher-priority emergencies (e.g., cardiac arrests, trauma, stroke) may necessitate the diversion of ambulances that were originally assigned to lower-priority incidents. By simulating these scenarios, ABM allows EMS planners to evaluate the trade-offs between response times for critical incidents and the potential delays for non-critical cases. This information is invaluable for designing dispatch policies that maximize overall system efficiency while ensuring that life-threatening emergencies are prioritized.

This approach highlights the operational adaptability of EMS systems—balancing response to critical cases with sustained service availability across the network.

1.4 Research Objectives

The primary aim of this thesis is to develop and evaluate advanced modeling and simulation frameworks for optimizing Emergency Medical Services (EMS) operations. This research focuses on two key areas: static optimization for strategic planning of delivery strategies and dynamic simulation for real-time dispatching and resource allocation. The overarching goal is to improve the effectiveness and efficiency of EMS systems, particularly in minimizing response times and maximizing patient outcomes. The specific objectives of the research are as follows:

1. **Optimize EMS Delivery Strategies Using Static Models**

   This objective focuses on developing and applying static optimization models to evaluate and enhance EMS delivery strategies, including in-hospital, rendezvous, and pre-hospital models. Static optimization will be used to determine the optimal location of resources such as ambulance stations and specialized care units to maximize coverage and minimize response times across diverse geographic areas.

   The model is performed on the existing network of hospitals and ambulance bases across Greater Sydney.

   - Develop static optimization models, including the MCP, to address facility location challenges in EMS systems.
   - Evaluate the effectiveness of different EMS delivery strategies, such as the rendezvous strategy, by optimizing the placement of intermediary care units.
   - Analyze how static optimization techniques can be applied to large-scale metropolitan environments to ensure optimal access to emergency care.

2. **Examine Algorithms for EMS Optimization**

   This objective focuses on examining algorithms that enhance EMS efficiency. The aim is to improve decision-making in strategic planning decisions. The optimisation problem compares Genetic algorithm and greedy algorithm with the solver from Gurobi (Gurobi Optimization, LLC, 2023). The GA's performance is validated through repeated runs, and its setting details are discussed in Chapter 4.

   - Select algorithms suited to optimizing EMS resources and analyze their strengths.
   - Formulate adaptive algorithms that can incorporate technological inputs

or have mathematical proofs in approximation.

- Implement the algorithms to solve specific problems and compare the results between algorithms and existing problem solvers.

3. **Improve Real-Time Dispatching Using Dynamic Simulation**

   This objective involves the application of dynamic simulation models to manage real-time operations in EMS, particularly focusing on dispatching policies and vehicle reallocation in multi-priority systems. The goal is to develop simulation frameworks that can adapt to real-time conditions, such as incident severity, traffic congestion, and vehicle availability.

   - Use agent-based simulation models to simulate real-time EMS operations to dispatch ambulances based on incident severity and proximity.

   - Test and refine dynamic dispatch policies, including the reassignment of ambulances in response to changing priorities, and assess their impact on overall system efficiency and response times.

   - Integrate real-world traffic data into dynamic simulation models to optimize dispatching decisions, reduce response times and improve patient outcomes in complex traffic environments.

4. **Analyze the Implications of Resource Allocation Decisions on EMS System Performance**

   This objective focuses on understanding the broader implications of resource allocation decisions within EMS systems, particularly how strategic and operational decisions influence system performance and patient outcomes.

   - Analyze the performance of EMS systems under different resource allocation strategies, including both static optimization and dynamic simulation approaches.

   - Assess the impact of geographic and temporal factors, such as population density, traffic conditions, and incident frequency, on the effectiveness of resource allocation.

   - Provide recommendations for optimizing resource allocation in EMS systems to improve overall system responsiveness and ensure equitable access to emergency care.

Table 1.1: Technical Pathway

Chapter	Stage	Main Tasks	Methods
3	Accessibility and Logistics Strategy	Evaluate logistics strategies by analysing travel-time and survival-time relationships to quantify spatial and temporal accessibility.	Accessibility modelling using fitted survival-time functions as logistics performance indicators.
4	Facility Location	Optimise the spatial configuration of emergency medical facilities to improve coverage and equity.	Mathematical optimisation models and heuristic solving algorithms.
5	Dispatch Policy	Assess dispatch and rendezvous policies to enhance operational efficiency and patient outcomes.	Simulation-based evaluation using agent-based and discrete-event models.

In summary, this research aims to increase EMS system performance through improving and optimizing logistics strategies. Table 1.1 summarizes the problems and the using technologies. I enhance the performance of EMS systems with the integration of advanced static optimization and dynamic simulation models. By optimizing delivery strategies, improving real-time dispatching, and incorporating cutting-edge technologies, the research seeks to provide a comprehensive framework for improving EMS operations in complex network environments. The results of this study will contribute to the development of more efficient, and responsive EMS systems, ultimately improving patient outcomes and the quality of emergency care services.

1.5 Thesis Contributions

This thesis makes several significant contributions to the field of Emergency Medical Services (EMS) logistics, focusing on the development of logistics strategies, facility location optimization, and dynamic operational policies. The key contributions are enumerated as follows:

1. **Proposing Diverse EMS Logistics Strategies**

This research proposes and evaluates several logistics strategies tailored to EMS operations. It compares the in-hospital, rendezvous, and pre-hospital models, identifying strengths and weaknesses under different operational conditions. These strategies provide EMS systems with practical tools to improve response times and optimize resource allocation for various types of emergencies. This contribution will be mainly proposed in Chapter 4.

2. **Recasting the Facility Location Problem and Comparing Optimization Algorithms**

   A significant contribution of this thesis is the reformulation of the facility location problem to better suit EMS needs. By adapting the Maximal Covering Problem (MCP) to the specific requirements of EMS resource allocation to reflect the time dependence of patient outcomes, the research develops a more applicable model for real-world settings. Furthermore, Chapter 3 compares different algorithms for solving the adapted facility location problem, evaluating their performance and providing valuable insights into algorithm selection based on EMS-specific operational goals.

3. **Development of a Comprehensive Simulation Framework for EMS Operations**

   This thesis develops a detailed simulation framework, which is introduced in Chapter 5, that models EMS operations. The framework integrates static optimization for long-term planning and dynamic simulation for real-time decision-making. It enables the evaluation of dispatching policies, vehicle allocation, and resource management, considering factors such as vehicle types, incident types, and incident priorities.

4. **Evaluation of Dispatching Policies Considering Vehicle Types, Incident Types, and Priorities**

   The research evaluates complex dispatching policies by taking into account vehicle types, incident types, and priority levels. Through the simulation framework, the thesis explores how dispatch policies can be optimized to ensure prompt responses to critical incidents while balancing the needs of lower-priority emergencies. Chapter 5 compares several rule-based assignment policies, and the findings help design dispatch strategies that improve system efficiency without compromising the quality of patient care.

5. **Confirmation of the Effectiveness of Proposed Logistics Methods**

   The research confirms the effectiveness of the proposed logistics strategies through testing and validation on real-world EMS data across Greater Sydney. The static optimization models, dynamic simulations, and algorithm comparisons demonstrate measurable improvements in EMS system performance, including reduced response times and better resource distribution. The validation confirms that the proposed methods significantly enhance EMS operations and improve patient outcomes.

This thesis synthesizes and extends core methodologies in EMS planning—facility location modeling, dynamic resource allocation, and multi-objective optimization—into a unified framework that shifts the focus from merely meeting time thresholds to maximally improving patient outcomes. By integrating accessibility, simulation, and advanced survival-based metrics, the research offers a holistic blueprint for EMS policymakers to design more responsive, equitable, and outcome-oriented emergency care networks.

Chapter 2

Literature Review

Emergency Medical Services (EMS) are a critical component of healthcare systems worldwide, providing immediate medical assistance to individuals experiencing acute health emergencies. The primary objective of EMS is to deliver timely and effective care to improve patient outcomes, especially in life-threatening situations where minutes can make the difference between life and death (Blackwell and Kaufman, 2002).

2.1 Current Research Practices in EMS Logistics

Research on EMS logistics can be interpreted through a problem-dimension hierarchy that reflects the sequential nature of EMS decision-making (Reuter-Oppermann et al., 2017). At the strategic level, studies examine the design of logistics structures through various delivery strategies that determine how patients, vehicles, and facilities interact within the EMS system. The planning level focuses on facility location, where spatial optimisation models are used to determine the placement of resources and ensure adequate coverage to support the chosen strategy. The operational level concerns real-time dispatching, which defines how available vehicles are allocated to incidents under dynamic demand and uncertain conditions. This hierarchical view highlights how EMS logistics research progresses from long-term system design to spatial planning and operational control. The following sections discuss representative studies at each level, focusing respectively on delivery strategies, facility location, and dispatching policies, which together capture the core problems addressed in EMS logistics research.

2.1.1 Facility Location

The Purpose of Facility Location Problems

The Facility Location Problem (FLP) is central to optimizing resource placement in emergency logistics, with the overarching goal of ensuring timely and effective responses to emergencies. The purpose of these models can generally be divided into two main categories: *resource minimization* and *coverage maximization* (Wang et al., 2021).

Resource minimization models, such as the Location Set Covering Problem (LSCP), focus on determining the minimum number of facilities required to cover all demand points. This approach is particularly useful for decision-makers aiming to minimize the infrastructure required while maintaining service standards (Toregas et al., 1971). On the other hand, coverage maximization models, such as the Maximal Covering Location Problem (MCLP) or Maximum Coverage Problem (MCP), aim to maximize the population or demand points covered within a given set of resources. These models are typically employed in scenarios where resources are limited, and the goal is to optimize their use to cover as much demand as possible (Church and ReVelle, 1974). More recent work extends these models, considering additional constraints like budget limitations and different service levels for diverse population segments, enhancing their practical relevance in emergency logistics (Wang et al., 2021). From the operational research (OR) perspective, they are *single-objective optimizations*. On the other hand, there are many studies target on *different multi-objective functions* (Chen and Lai, 2022; Naji et al., 2024; Wang et al., 2023).

Coverage-Based Classification

FLP models are often classified based on how they handle coverage, with a key distinction between *binary* and *gradual coverage* models. Traditional binary coverage models, such as the LSCP and MCP, operate on a strict threshold principle: a demand point is fully covered if it falls within a specified distance from a facility, and receives no coverage if it lies beyond this threshold (Church and ReVelle, 1974). This all-or-nothing framework is appropriate for situations where service is either entirely available or not available at all—such as a house being connected to water mains.

In response to this limitation, Gradual Covering Location Problem (GCLP) models allow for a more realistic depiction of service delivery. These models assign partial coverage to demand points that are further away, with coverage progressively declining as distance increases. Berman et al. (2003) were among the first to propose a step function for modeling gradual coverage, providing a more nuanced framework for real-world applications like emergency medical services or telecommunications. This approach aligns with the accessibility framework of Hansen (1959), which has been applied across diverse fields, and more recent research has refined it further, utilizing linear or non-linear decay functions to model varying levels of coverage more accurately (Karatas and Eriskin, 2021). These gradual coverage models are particularly valuable in applications where service quality diminishes over distance, such as mobile phone signal strength or response time in emergency logistics.

2.1.2 Transportation Logistics

Transportation logistics are essential for efficiently managing EMS resources and ensuring timely responses to critical emergencies. Several logistics strategies are used to optimize the allocation and flow of these resources, including basic response, multi-echelon logistics, rendezvous, and reverse logistics. Each strategy has its advantages and limitations, which can affect system performance depending on the operational environment and the type of emergency.

Basic Response Strategy The basic response strategy is the simplest form of EMS logistics, where ambulances are dispatched directly to an incident and return to the station after completing their task (Dick, 2003). This approach works best in smaller, low-demand environments, where response times can remain minimal without needing more complex systems. However, it lacks flexibility for handling high-demand situations or specialized services and can lead to inefficiencies if multiple incidents occur simultaneously.

Multi-Echelon Strategy The multi-echelon strategy introduces a decentralized approach, where resources move through multiple layers or stages. For example, vaccines may first be delivered to a central hub before being distributed to GP clinics (Kar and Jenamani, 2024), or patients may be initially stabilized at smaller clinics before being transferred to more specialized medical centers. This system allows for more efficient resource allocation across larger geographic areas and can handle both high demand and diverse service levels, so it is more often shown in humanitarian or disaster logistics (Döyen et al., 2012; Lohrasbpoor et al., 2023). Compared to the basic response, the multi-echelon strategy is more flexible but requires careful coordination between layers to avoid delays.

Rendezvous Strategy The rendezvous strategy involves synchronizing the movements of EMS units so that specialized equipment or personnel meet the primary response team at an intermediate point (Parker et al., 2020). This is especially useful when advanced care is needed during transport. While the basic response relies on direct dispatch, the rendezvous strategy adds a layer of flexibility by allowing for specialized care without having to return to a centralized station or hospital. However, this requires precise coordination, which may be challenging in fast-paced, time-sensitive environments.

Reverse Logistics Reverse logistics is the practice of making return trips more efficient by carrying additional tasks, such as waste removal or resource collection, on the way back (Ribeiro et al., 2021). In contrast to the basic response strategy, where vehicles typically return empty, reverse logistics ensures that no trip is wasted. This strategy is especially useful for optimizing resource use in rural areas or after large-scale medical deliveries, where materials (e.g., unused medical supplies or waste) can be transported back to a central location (Qi et al., 2023; Zhu et al., 2024). Reverse logistics increases operational efficiency but requires careful planning to align return trips with new tasks.

Pre-Hospital Care Strategy The pre-hospital care strategy focuses on providing life-saving treatments at the scene before transporting the patient to a hospital (Dick, 2003). For critical emergencies, such as cardiac arrest, this approach offers a significant survival advantage over the basic response, where treatment only begins after the patient reaches a medical facility. By enabling EMS teams to perform advanced medical procedures in the field, the pre-hospital strategy can reduce the need for rapid hospital transport, making it a key element of advanced EMS systems.

Each logistics strategy offers unique benefits and limitations that affect system efficiency, resource utilization, and patient outcomes. The basic response strategy, while simple and effective in low-demand settings, can become inefficient in more complex or high-demand environments. It provides no flexibility for specialized services or incident escalation, unlike the multi-echelon strategy, which can better manage complex resource distribution across multiple layers but requires high coordination to avoid bottlenecks.

2.1.3 Dispatch Systems and Policies

Dispatch systems are the operational hubs of EMS logistics, responsible for receiving emergency calls, assessing the nature of incidents, and coordinating the deployment of resources. Traditional dispatch methods often rely on rule-based systems, where dispatchers follow predefined protocols to assign resources based on incident type and resource availability. For instance, studies by Amorim et al. (2018) demonstrated that integrating spatial data into dispatch decisions enhances the accuracy of travel time estimations and improves overall EMS performance.

Dispatch for Different Metrics

Time-based dispatch policies focus on minimizing response time to critical incidents, particularly life-threatening emergencies where faster intervention can significantly improve patient outcomes. Response times could be impacted by many factors such as technological advancements, geographic influence, and humans (Alshammari et al., 2024). These systems typically prioritize high-severity incidents by utilizing the closest idle vehicle policy, ensuring the quickest possible response. This approach is especially vital for conditions like cardiac arrest, where seconds can make a critical difference. Most relevant research focuses on routing algorithms instead of dispatch policies (Khoshgehbari et al., 2023; Rabbani et al., 2022; Sutherland and Chakrabortty, 2023). Research by Nasrollahzadeh et al. (2018) shows that dynamic dispatch policies could reduce response time for high-priority incidents by 30.6%, while also improving system efficiency. The balance of location and dispatch for response time should also be considered (Bélanger et al., 2020; Jánošíková et al., 2021). Prioritizing response alone can result in resource saturation in high-demand areas, leaving other regions underserved.

In addition to response time, *equity-based* dispatch policies aim to ensure fair and equitable distribution of EMS services across different geographic regions, including underserved rural areas. It is an important indicator of EMS to the society (Farcas et al., 2024; Newton et al., 2021). McLay and Mayorga (2013) develops a model for optimally dispatching distinguishable servers to prioritized customers in EMS systems, incorporating four equity constraints. The study of Enayati et al. (2019) also proposes a multicriteria optimization approach for EMS location and dispatch decisions that addresses the trade-off between minimizing system-wide operational cost and reducing disparities in regional service level. VK and Roy (2017) focus on equitable workload distribution among ambulance crews to ensure fairness and experience balance. These policies are often guided by geographic demand models and socio-economic factors.

The effectiveness of emergency medical dispatch system can improve the survival (Ageron et al., 2016), so *medical output* could be a straightforward metric. When proposing new performance measures for EMS system design, McLay (2010) also emphasize the need to focus on patient outcomes. Wajid and Nezamuddin (2022) proposes a model to optimize ambulance utilization and maximize survival rates for high-priority calls under uncertain travel times. Lim et al. (2011) finds that adopting appropriate dispatch policies, such as reroute-enabled dispatch and free-ambulance exploitation, can improve EMS performance.

Dispatch for Multi-Class Resources and Demands

EMS systems rely on a variety of resources, such as Basic Life Support (BLS) units and Advanced Life Support (ALS) units, each serving different roles based on the type and severity of emergencies (Sporer et al., 2007). The challenge is ensuring that these multi-class resources are allocated efficiently. ALS units, for example, should be reserved for severe cases, while BLS units handle lower-priority incidents to prevent resource depletion (Gamberini et al., 2021). Misallocating ALS units can lead to shortages when they are needed for critical emergencies, making proper resource balancing crucial.

Addressing multi-type demand—such as medical emergencies, basic healthcare, and transportation spills—adds complexity. Each type of emergency requires specific combinations of resources, and dispatch systems must adapt to ensure timely and coordinated responses. Various models have been proposed to manage these challenges. Dynamic dispatch models, like the one developed by Gendreau et al. (2001), adjust resource allocation in real time based on fluctuating demand, ensuring that multi-class resources are deployed optimally. Additionally, multi-objective optimization approaches, such as the one proposed by Nasrollahzadeh et al. (2018), balance equity and efficiency by ensuring fair resource distribution across regions while maintaining rapid response times. These models help EMS systems adapt to the variability in resources and emergency types, improving both response times and resource utilization.

Dispatch Principles

Various dispatch principles guide EMS systems in making decisions on which vehicle responds to an incident. These principles are designed to balance speed, resource efficiency, and incident severity.

First-In, First-Out (FIFO) and Closest Dispatch: FIFO dispatch policies serve calls in the order they are received, but they do not account for the urgency of incidents. Lim et al. (2011) highlighted that FIFO dispatching often leads to inefficiencies, especially in high-demand periods, as it does not prioritize more critical cases. The closest-vehicle dispatch policy is more commonly used, but it can result in poor preparedness for future calls when resources are not allocated considering future demand.

Reroute-Enabled Dispatch: Rerouting ambulances to higher-priority incidents while en route to lower-priority calls is a common strategy in modern EMS systems. Gendreau et al. (2001) explored rerouting policies that allow for dynamic reassignment of vehicles, improving system's responsiveness to urgent calls while maintaining overall response time.

Multi-Tier Dispatching: Many EMS systems operate on a tiered basis, where BLS vehicles handle routine incidents and ALS vehicles are reserved for high-severity calls. Boujemaa et al. (2020) found that tiered systems improve response times by sparing ALS units for critical cases while allowing BLS units to handle most calls. This approach not only improves sytem efficiency but also the output robustness

(Boujemaa et al., 2018).

Priority-Based Dispatch: Priority-based dispatching systems such as the Advanced Medical Priority Dispatch System and Criteria-Based Dispatch assign priority levels to incidents based on urgency (Bandara et al., 2014). These systems ensure that high-priority incidents receive immediate attention, even if it requires reallocating resources en route to lower-priority incidents (Schmid, 2012). Lim et al. (2011) showed that rerouting ambulances to more urgent calls can reduce response times for critical cases, improving overall EMS performance.

Preparedness and Impact on Future Demands: Some dispatch strategies, such as those proposed by Schmid (2012), consider not just the immediate response but also the system's preparedness to handle future incidents. These models seek to minimize degradation in the system's ability to respond to future calls by dispatching vehicles in a way that preserves overall coverage and availability.

2.2 Gaps and Opportunities in EMS Logistics Research

This section outlines the primary challenges in EMS logistics, including managing demand variability, integrating complex resource types, bridging long-term planning with real-time operations, ensuring accessibility and equity, and addressing limitations of current delivery strategies.

2.2.1 Demand Variability and Complex Resource Types

One of the most pressing challenges in EMS logistics is managing demand variability, especially during unpredictable events like large-scale accidents, natural disasters, or public health crises. Traditional EMS optimization models often rely on static resource allocation, assuming stable demand conditions. However, during high-demand periods, these models become inadequate, leading to delays in emergency response times and inefficient resource utilization. For instance, Grot (2024) highlighted that during extreme weather conditions or mass casualty events, the limitations of static models become evident, resulting in strained resources and compromised patient care.

To compound this issue, modern EMS systems deploy a variety of complex resource types, including BLS ambulances, ALS response teams. Each resource type has distinct capabilities and operational constraints, which must be considered in resource allocation decisions. While multi-tiered dispatch systems offer frameworks for coordinating different resource types, many existing models do not fully account for these complexities.

Gap: There is a need for dynamic models that can adapt to fluctuating demand and effectively integrate diverse resource types. Such models should account for the stochastic nature of emergency incidents and the operational constraints of different EMS assets.

2.2.2 Bridging Long-Term Planning with Real-Time Operations

Another significant challenge is the disconnection between long-term strategic planning and real-time operational needs. Long-term planning involves optimizing facility locations to maximize coverage and minimize response time. In contrast, real-time operations demand immediate decision-making based on current conditions, such as traffic congestion, resource availability, and incident severity.

Many EMS systems rely on static models for facility location and resource allocation, which are not equipped to adapt to the dynamic nature of emergencies. While dynamic dispatch algorithms offer a promising solution, their implementation is often hindered by organizational inertia and technical barriers, such as integrating real-time data into existing systems. The studies by Abou-Ali and Shouman (2004); Borozdukhin et al. (2016) noted that static models follow fixed schedules, while dynamic models adjust in real-time to system changes, improving adaptability and performance. In flexible manufacturing, dynamic models are more efficient, handling fluctuations and optimizing throughput better than static models. The system calculates time-optimal routes for garbage collection trucks based on real-time data, such as the fullness of garbage containers and current traffic conditions. This differs from static systems where routes are predefined, leading to inefficiencies when containers are only partially filled. However, the paper by Ridler et al. (2022) includes a case study demonstrating the use of a Julia (programming language) package for Emergency Medical Services Simulation to optimize static ambulance deployment in Auckland, New Zealand, with the objective of maximizing the number of emergency calls reached within a target response time.

Gap: There is a need for research that bridge the gap between long-term planning and real-time operations. These models should incorporate real-world data, predictive analytics, and dynamic dispatching to enhance operational flexibility and responsiveness.

2.2.3 Limitations of Current Delivery Strategies

Most EMS systems continue to rely on traditional delivery strategies, primarily involving transporting patients from the scene of an emergency directly to a hospital. While this approach is effective in many scenarios, it has limitations, particularly in urban areas with high traffic congestion or in situations where rapid advanced medical intervention is critical.

Alternative strategies, such as the rendezvous model and pre-hospital care, are not sufficient in the literature. The terms of rendezvous model and pre-hospital care are all for field care currently not the definitive care (Reuter-Oppermann et al., 2017). The rendezvous model should involve coordinating different EMS units to meet at predetermined locations, allowing specialized care to be provided en route to the hospital. Pre-hospital care strategies focus on delivering advanced medical interventions for definitive care at the scene before transportation. Study by Dell'Amico and Hadjidimitriou (2012) have highlighted the potential benefits of innovative delivery strategy but also note the lack of comprehensive models to optimize their implementation.

Gap: There is a need for research that explores and optimizes alternative delivery strategies. Developing models that evaluate the effectiveness of these strategies under various conditions can lead to improved patient outcomes and more efficient EMS operations.

2.3 Modeling Methods and Problem-Solving Techniques

To overcome the complex challenges identified in Emergency Medical Services (EMS) logistics—including managing demand variability, integrating complex resource types, bridging long-term planning with real-time operations, ensuring accessibility and equity, and addressing limitations of current delivery strategies—advanced modeling methods and problem-solving techniques are essential. This section explores various modeling approaches designed to address these challenges, focusing on dynamic facility location models, optimization algorithms, simulation techniques, and their application to EMS delivery strategies.

2.3.1 Facility Location Models for Alternative Delivery Strategies

The MCP has long been utilized in EMS logistics to optimize facility locations, aiming to maximize the population covered within a critical response time (Toregas et al., 1971). Traditional MCP models, however, are static and fail to account for performance for different distances, resource availability, and traffic conditions. This limitation hinders the EMS system's ability to respond effectively to emergencies, particularly when the output is sensitive to distance or travel time.

To address the challenge of managing demand variability and integrating complex resource types, the MCP can be enhanced through dynamic facility location models. These models incorporate comprehensive data, such as demand patterns, traffic congestion, and resource statuses, allowing for continuous adjustment of facility locations and resource allocations.

By integrating comprehensive data, MCP models enable EMS systems to adapt to different delivery strategies, ensuring that resources are positioned optimally to meet different needs. This approach not only enhances converage but also improve response times against demand surges and resource constraints.

2.3.2 Optimization Techniques: Greedy and Genetic Algorithms

Effective resource allocation and dispatching in EMS logistics require optimization methods capable of handling complex, multi-objective problems. Greedy algorithms and genetic algorithms are two such techniques that can address these challenges.

Greedy algorithms make locally optimal choices at each step with the hope of finding a global optimum. They are computationally efficient and useful for problems where quick, approximate solutions are acceptable (Malik et al., 2018). In EMS logistics, greedy algorithms can rapidly assign resources to incidents based on immediate criteria like proximity and availability, which is beneficial in time-sensitive situations. Marla et al. (2017) deploy a greedy algorithm for both static ambulance fleet allocation and dynamic redeployment to maximize system service coverage. Erdemir et al. (2010) utilize a greedy heuristic to simultaneously locate ground and air ambulances, along with landing zones, to maximize coverage based on both response time and total service time.

However, greedy algorithms may not always yield the best overall solution, especially when dealing with multiple constraints and objectives. Genetic algorithms, inspired by the process of natural selection, offer a more robust approach by exploring a wider solution space through operations like selection, crossover, and mutation (Atta et al., 2018). GAs are particularly effective in handling large scale problems, like FLP that is NP-hard on compute solving complexity.

For example, Sasaki et al. (2010) predict EMS demand through 2050 based on demographic factors, using a modified grouping genetic algorithm to optimize ambulance locations. Comber et al. (2011) present a modified grouping genetic algorithm applied to identify optimal ambulance locations, improving average EMS response times by 1 minute 14 seconds, and evaluating the impact of varying ambulance numbers at current and new locations using census areas as spatial units.

2.3.3 Simulation Approaches: Agent-Based Modeling (ABM)

Agent-based models (ABM) have been increasingly adopted in emergency medical services (EMS) to simulate complex, dynamic, and heterogeneous environments. ABMs allow the detailed modeling of individual entities (agents) such as ambulances, patients, and dispatch centers, capturing their interactions and decision-making processes in a simulated environment. This makes them particularly suitable for EMS, where multiple agents with different roles and objectives must coordinate in real-time to optimize outcomes.

Hawe et al. (2015) highlighted that ABMs provide a natural way to simulate EMS operations because they enable each agent (e.g., ambulance, patient) to be modeled with individual behaviors and decision-making capabilities. The agents in these simulations can interact with each other and their environment, which is crucial for modeling real-world complexities such as traffic conditions, patient severity, and dispatcher decision-making. The ability of agents to adapt their behaviors based on changes in the environment makes ABMs particularly effective for EMS systems, which operate in dynamic and uncertain settings.

Moreover, ABMs are used to test and evaluate different EMS strategies, such as ambulance dispatching and resource allocation. For example, Wang et al. (2012) developed an agent-based model to simulate the emergency response to a mass casualty incident, focusing on the allocation of ambulances to incident sites. Their results showed that ABMs could effectively evaluate the trade-offs between different dispatching strategies, providing insights into how best to allocate resources in real-time to maximize patient survival.

Coordination between different emergency responders, such as ambulances, hospitals, and dispatch centers, is another area where ABMs excel. As noted by Hawe et al. (2015), ABMs enable researchers to study the coordination between these entities, helping to identify bottlenecks and inefficiencies in the system. This is particularly important in mass-casualty incidents (MCIs), where the effective coordination of resources can significantly impact survival rates.

Additionally, ABMs provide the flexibility to simulate various scenarios, such as changes in traffic conditions, availability of resources, and variations in patient severity, which traditional models like discrete event simulation and system dynamics struggle to capture. Alotaibi and Ibrahim (2018) emphasized that ABMs allow the simulation of realistic agent behaviors, such as ambulances rerouting based on traffic or the severity of calls, making them a powerful tool for EMS system optimization.

In summary, ABMs offer a robust framework for simulating EMS operations, allowing for the detailed modeling of individual agents and their interactions. This enables the evaluation of different dispatching policies, resource allocation strategies, and coordination mechanisms, ultimately helping to improve EMS performance in real-world scenarios.

2.4 Conclusion

This literature review has explored the critical role of EMS within healthcare systems, emphasizing the necessity of efficient logistics to ensure timely medical responses during emergencies. The examination of current practices has provided insights into various aspects of EMS logistics, including facility location, transportation logistics, and dispatch systems.

The analysis of facility location problems has underscored the importance of balancing resource minimization and coverage maximization, highlighting advancements in gradual coverage models that offer more nuanced solutions to real-world challenges (Church and ReVelle, 1974; Wang et al., 2021). Despite progress, many existing models remain static, failing to address the dynamic nature of emergency demands, which presents an ongoing challenge for EMS operations (Grot, 2024).

Transportation logistics strategies, including basic response, multi-echelon systems, and pre-hospital care strategies, have demonstrated their effectiveness in improving response times and patient outcomes (Dick, 2003; Kar and Jenamani, 2024). However, the limitations of traditional delivery strategies, particularly in urban settings with high traffic congestion, necessitate further exploration of alternative strategies, such as the rendezvous model and reverse logistics (Qi et al., 2023; Reuter-Oppermann et al., 2017). These strategies hold promise for enhancing operational efficiency but lack comprehensive models for optimization in EMS contexts (Dell'Amico and Hadjidimitriou, 2012).

Furthermore, the review of dispatch systems has revealed the significance of integrating equity considerations and medical output metrics into performance measures. Research by McLay (2010) emphasizes the need to focus on patient outcomes, while studies on dynamic dispatch policies show potential improvements in efficiency (Lim et al., 2011; Nasrollahzadeh et al., 2018). As EMS systems increasingly rely on diverse, multi-class resources, the challenge of optimizing resource allocation remains pertinent (Gamberini et al., 2021). Dispatch principles that prioritize urgency and preparedness must also adapt to the realities of fluctuating demand and resource availability (Lim et al., 2011; Schmid, 2012).

Despite the advancements in EMS logistics research, notable gaps persist. The need for dynamic models that bridge long-term planning with real-time operations is evident, as current static frameworks struggle to meet the complexities of modern EMS demands (Abou-Ali and Shouman, 2004; Borozdukhin et al., 2016). Additionally, the exploration of alternative delivery strategies is underdeveloped, with a clear need for research that evaluates and optimizes these approaches in various scenarios (Dell'Amico and Hadjidimitriou, 2012; Reuter-Oppermann et al., 2017).

In summary, while significant strides have been made in understanding and optimizing EMS logistics, ongoing research is essential to address the identified gaps and enhance the effectiveness of EMS systems. Future studies should focus on developing integrated, dynamic models that account for the variability of emergency demands, ensuring efficient resource allocation and improving patient outcomes in diverse contexts.

Chapter 3

Alternate Delivery Strategies for EMS

3.1 Introduction

Emergency medical services (EMS) remain a critical public concern due to their direct impact on human lives. Advances in medical and clinical technology have improved treatment outcomes across various conditions. However, due to the complexity of interventions and the constraints of resource availability (Reynolds et al., 2017), only a limited number of patients can access advanced treatments. From a logistics perspective, optimizing resource allocation and utilization can enhance both effectiveness and efficiency.

Developing effective logistics strategies is essential for EMS to deliver advanced interventions within the critical early minutes of a severe incident. Out-of-hospital cardiac arrest (OHCA) exemplifies the time-sensitive nature of medical emergencies, where even slight delays in reaching advanced care can drastically reduce survival rates (Bartos et al., 2018; Wengenmayer et al., 2017). Survival probabilities often drop to less than 2% if treatment is not initiated within the critical one-hour window (Reynolds et al., 2013).

In recent years, extracorporeal membrane oxygenation (ECMO) during cardiopulmonary resuscitation—commonly known as ECPR—has been introduced to address refractory cardiac arrest (Hadaya et al., 2020; Yannopoulos et al., 2020a). This advanced technique can significantly improve survival rates, provided ECMO flows are established quickly to minimize "low-flow" duration (Richardson et al., 2021; Wengenmayer et al., 2017). Traditional in-hospital ECPR involves transporting patients to specialized centers. While logistically straightforward, this approach often limits access to patients located close to such facilities.

This chapter utilizes transport accessibility and time-threshold methodologies (Wu and Levinson, 2020) alongside empirical data—including travel times, historical arrest locations, and population distributions—to evaluate how in-hospital, rendezvous, and pre-hospital ECPR strategies differ in terms of coverage and predicted outcomes. By emphasizing logistics and spatial planning, this chapter seeks to identify the delivery model that best expands timely access to advanced interventions.

3.2 Method

We defined three strategies for ECPR delivery for refractory OHCA within Sydney, Australia and then applied transport accessibility metric analysis methods to determine the effective patient catchment of each strategy. The tested strategies are illustrated in Figure 3.1. The basic ambulance, which provides initial treatment and patient transport, typically remains at the ambulance station. The ECMO team operates an independent ambulance equipped with ECMO devices and staffed by trained ECMO technicians.

In-hospital ECPR Figure 3.1a: The basic ambulance responds to the patient, stabilises on-scene and transports the patient to an ECPR-capable hospital. ECMO cannulation and ECPR are delivered at the hospital. This delivery strategy is the status quo in most places. Assuming that time of arrest coincides with time of call, arrest to flow time is the sum of response time, on-scene time, travel time and cannulation time.

Rendezvous ECPR Figure 3.1b: The patient is attended by an ambulance and stabilised on-scene then transferred to an emergency department, which may not be ECPR-capable, in order to rendezvous with the ECPR team. The rendezvous hospital is selected based on minimisation of the greater of: ECMO team travel time to the rendezvous hospital equals to the sum of paramedic response time, on-scene time and travel time to rendezvous hospital. Arrest to flow time is the sum of the maximum of these two intervals plus cannulation time. The ECPR team rendezvous with the patient at that hospital, establishes ECMO support and transfers the patient on ECMO to an intensive care unit for definitive treatment. Modelling assumes the ECPR team is notified of the OHCA at the time of the initial cardiac arrest call and begins movement to the rendezvous hospital.

Pre-hospital ECPR Figure 3.1c: A pre-hospital ECMO team is dispatched and ECMO cannulation and flow is completed at the scene of cardiac arrest, with subsequent transfer of the patient back to an ECMO-capable hospital. Unlike the

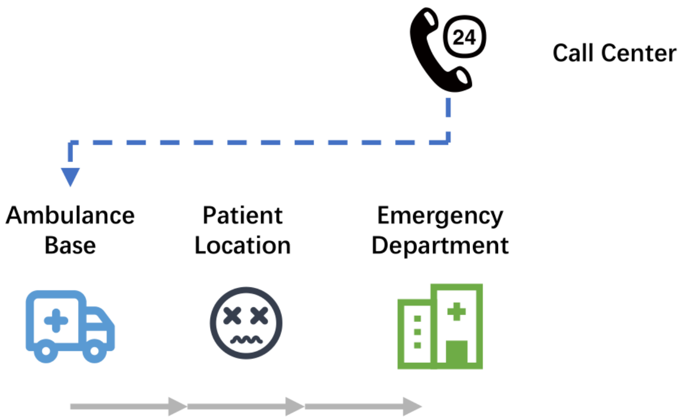

(a) In-hospital Strategy

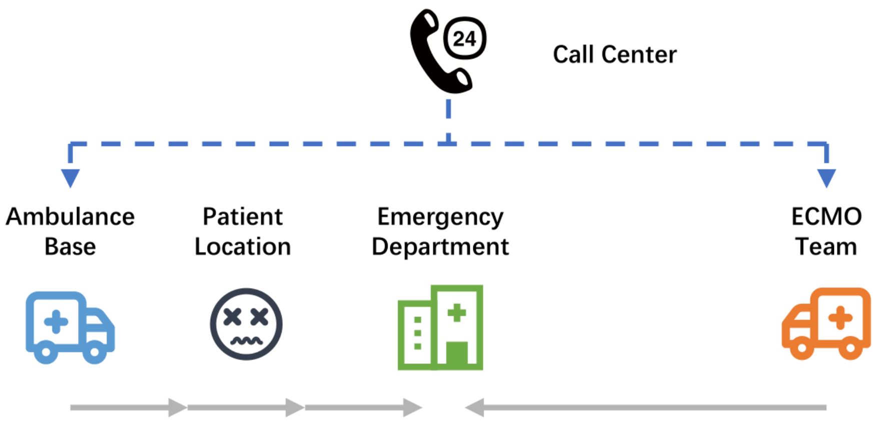

(b) Rendezvous Strategy

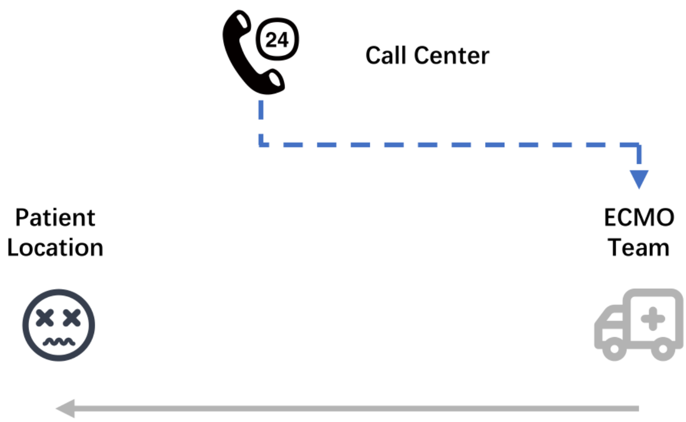

(c) Pre-hospital Strategy

Figure 3.1: Three Types of Delivery Strategies

ECMO team in the Rendezvous ECPR strategy, this ECMO team uses a larger van that allows paramedics to provide ECPR treatment directly at the patient's location. Modelling assumed that the pre-hospital ECPR team is dispatched at the same time of the cardiac arrest emergency call. Arrest to flow time is calculated by: response time + cannulation time.

The immediate activation point for the mobile ECMO teams in Rendezvous and Pre-hospital ECPR of time of initial EMS call, was chosen for two reasons. First, a number of current trials (Lamhaut et al., 2017; Marinaro et al., 2020; Petermichl et al., 2021; Singer et al., 2020), currently utilise this approach. Secondly , previous studies have reported that a majority of OHCA are recognised by emergency dispatchers between 50 s to approximately 2 min from the start of the call, (Berdowski et al., 2009; Culley et al., 1991; Dami et al., 2015; Lewis et al., 2013). This justifies the rapid dispatch of the ECMO team to patients suspected of eligibility for ECPR.

3.2.1 Transport accessibility metrics analyses

The comparison of the three cardiac arrest strategies was addressed using transport accessibility metrics (Wu and Levinson, 2020), as shown in Table 3.1. In this approach, the study area is divided into zones with a known number of potential patients, n_i. ECPR facilities (in-hospital, rendezvous, or pre-hospital) can be allocated to each zone, and y_j represents the number of facilities in zone j. Usually, y_j would be zero or one. The ability of a patient in zone i to access the ECPR facilities in zone j requires knowledge of the complete travel time matrix, t_{ij}. Since the success of ECPR in zone i depends on the time from arrest to ECMO flow, T_i, the travel time is added to other relevant time intervals for each ECPR delivery strategy. The components are defined as:

Table 3.1: Components of the time from arrest to ECMO flow under each ECPR delivery strategy. In the rendezvous strategy, the patient and the ECPR team meet at an intermediate location k, where $y_k = 1$ indicates that zone k contains a suitable emergency department. The patient is in zone i and accessing ECPR locates in zone j. The bottom row shows the total time from arrest to flow.

Component	**In-hospital**	**Rendezvous**	**Pre-hospital**
Ambulance to patient	t_{hi}	t_{hi}	t_{ji}
On-scene and loading	t_s	t_s	–
Ambulance to hospital	t_{ij}	t_{ik}	–
ECMO team to hospital	–	t_{jk}	–
Cannulation	t_c	t_c	t_c
Arrest to flow time, T_i	$\min(t_{hi} + t_s + t_{ij} + t_c)$	$\min_{j,y_j \neq 0} \min_{k,y_k \neq 0} (\max(t_{hi} + t_s + t_{ik}) + T_c)$	$\min_{h,y_h \neq 0} (t_{ji} + t_c + t_p)$

Response time (t_{hi}): The time from the location of the ambulance in zone h to the location of the patient in zone i. This is the time between the call to emergency medical services (EMS) and the arrival of EMS paramedics at the scene of cardiac arrest. For pre-hospital ECPR, the response time is from the location of the mobile ECMO unit in zone j (t_{ji}).

On-scene time (t_s): The time interval between the arrival of paramedics on scene and patient departure to the hospital, including patient access, treatment, and extrication.

Travel time: The transfer time from the location of the cardiac arrest to the ECPR-capable hospital (t_{ij} for in-hospital ECPR) or to an intermediate emergency department (t_{ik} for rendezvous ECPR).

Cannulation time (t_c): The time from the arrival of the ECPR team at the patient to the establishment of ECMO flows.

Penalty time (t_p): The additional time required for the pre-hospital ECPR team to establish ECMO flow at the patient location, relative to in-hospital cannulation.

For an arrest occurring in zone i with ECPR facilities in zone j and a suitable emergency department in zone k, this interval is the minimum across all facilities of the sum of the time components, as shown in Table 3.1.

3.2.2 Determining population coverage by ECPR strategy

To reflect the golden hour cut-off of eligibility (Marini et al., 2024), the time from arrest to ECMO flow for patients in zone i, T_i, is compared to the threshold $\tau = 60$ and that zone is indicated to be either above or below the threshold with a binary variable:

$$b_i = \begin{cases} 1, & \text{if } T_i \leq \tau \\ 0, & \text{if } T_i > \tau \end{cases} \tag{3.1}$$

If a zone is covered by any facility (i.e., $b_i = 1$), then the potential patients in that zone, n_i contribute to the total coverage, A_c:

$$A_c = \sum_{i=1}^{N} n_i b_i \tag{3.2}$$

Higher values of A_c indicate that the strategy or ECPR-facility location offers an advantage in the number of potential patients that can receive ECPR.

3.2.3 Survival benefit modelling

Patients who are commenced on ECMO flow earlier after arrest are more likely to survive (Bartos et al., 2020a; Wengenmayer et al., 2017), therefore we supplemented the population coverage metric with another measure that weights each covered patient by their probability of survival. Probability of survival in zone i, p_i, ranges from 0 to 1, and is estimated by evaluating a decreasing survival function at T_i. As defines the population-weighted survival probability below.

$$A_s = \frac{1}{\sum_{i=1}^{N} n_i} \sum_{i=1}^{N} n_i p_i \tag{3.3}$$

For estimating population-weighted survival probability, we modelled the relationship between survival probability and resuscitation time using the aggregated data reported by Bartos et al. (2020b). The average survival outcomes from that paper are fitted with a logistic curve with time to resuscitation (low-flow time) as the only predictor using the statsmodels package in Python. Based on the principle of selecting the best-performing strategy at each location, the survival probability p_i at meshblock i is defined as:

$$p_i = \max_m s^m(T_i^m) \tag{3.4}$$

where $s^m(\cdot)$ is the survival function for method m and T_i^m is the total time to ECMO flow at location i using method m. This ensures that the population, n_i, at each location is attributed the highest survival probability achievable from all available ECPR strategies, based on the corresponding total response and treatment times.

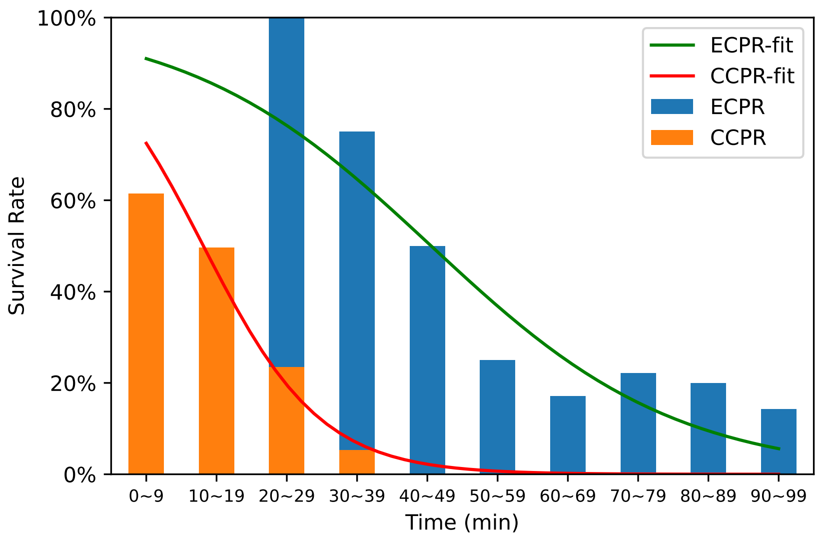

Figure 3.2: Fitted logistic survival rate functions for ECPR (Bartos et al., 2020b) and conventional cardio-pulmonary resuscitation (CCPR) using data from the Amiodarone, Lidocaine or Placebo study (ALPS) (Kudenchuk et al., 2017)

The survival rate is the logistic fit to the data shown in Figure 3.2. Newton's method and a 95% confidence level is used in the regression. The survival rate function of ECPR is $s(T) = 1/(1 + e^{-(2.6004-0.0571*t)})$, while CCPR is $s(T) = 1/(1 + e^{-(1.5635-0.1191*t)})$. P-values of all parameters are less than 0.0005.

Unlike the coverage metric, the population weight survival probability varies from 0 to 1 and gives the overall probability of survival for ECPR-eligible arrests in each delivery scenario. Therefore, it can distinguish between two strategies that reach the same number of potential patients, but one reaches them faster and provides better survival outcomes. Furthermore, the population-weighted survival probability allows us to relax the 60 min threshold and quantify the benefit to patients who sit just outside this coverage boundary.

3.2.4 Modelling location and data levels

Modelling was completed for Greater Metropolitan Sydney, Australia with a 2016 census population of 4.8 million and an area of 12,368 km^2. Hospital-based ECPR services exist at 5 hospitals (Figure 3.3), and rendezvous and pre-hospital ECPR were not offered within Sydney at the start of the modeling. The analysis zones are Greater Sydney's approximately 58,000 meshblocks, the finest spatial resolution available in the Australian census data (Australian Bureau of Statistics, 2017). Patients are assumed to be distributed proportionately to the meshblock resident populations from the 2016 census counts. Historical cardiac arrest cases (NSW Ambulance, 2019) from 2017 to mid-2021 aggregated to the statistical area level 2 were used to calculate localised ambulance response times. The distribution of on-scene treatment times for CCPR were obtained from the NSW OHCA registry (NSW Ambulance, 2019). Meshblock-to-meshblock travel times on the road network were calculated from Compass IoT's connected vehicle data averaging speeds for every link in the Sydney network from one week in November 2019. Travel times were validated against realised ambulance travel times from the cardiac arrest registry and shown to be consistent to within 2% (Figure 3.4). These data comprise the necessary inputs for calculating A_c as described above.

Base case modelling for the status quo ECPR delivery strategy—Figure 3.5, summarised in Table 3.1, uses locally-appropriate response times, a 27 min on-scene treatment time, travel time to the nearest ECPR-capable hospital, and 15 min of cannulation time. The on-scene treatment time of 27 min was chosen based on published data (Belohlavek et al., 2022; Yannopoulos et al., 2020a) of expedited transfer of patients from the scene, as until 2021, mechanical CPR devices were not available in Sydney, NSW. Preliminary data from the recently completed EVIDENCE study (ACTRN12621000668808), comparing expedited transfer to more extended on-scene resuscitation, has reported a median on-scene time of 26 min in the expedited arm in the study area.

Rendezvous modelling assumed that ECPR teams are located at the five current hospitals but can mobilise to move to any of the 26 emergency departments in the

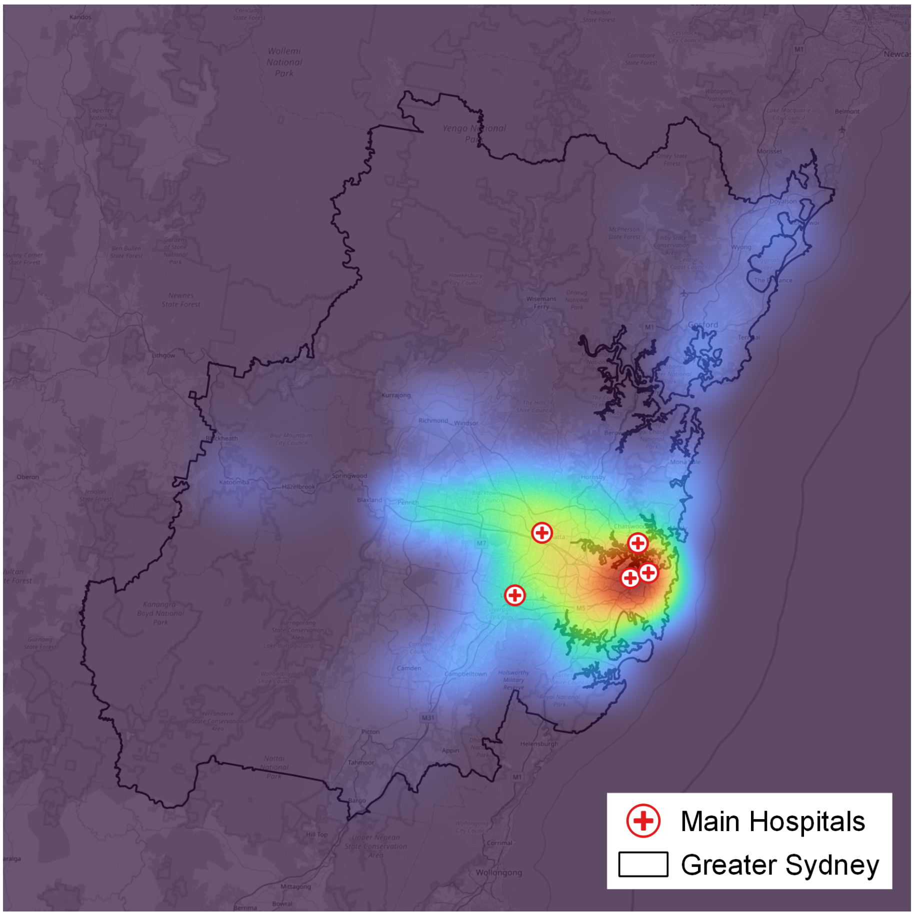

Figure 3.3: Heatmap of Historical Incidents Distribution. The location of the five current ECPR hospitals within Greater Sydney. The heatmap shows the spatial distribution of 4.5 years of out-of-hospital cardiac arrests from the NSW Ambulance OHCA registry.

study area. The time from arrest to treatment includes the locally-appropriate response time, 27 min on-scene time, travel time to the rendezvous hospital emergency department, and 15 min of cannulation time.

The pre-hospital strategy assumes one optimally-positioned mobile ECPR team providing the minimum benefit of a pre-hospital ECPR service. The methodology for identifying the optimal position is described below. The time from arrest to ECMO flow is the sum of the travel time from the optimal location to the patient plus 22 min for cannulation, where the additional 7 min accounts for the difficult cannulation context and has been based on published experience thus far (Lamhaut et al., 2017).

Additional sensitivity analyses (Table 3.2 and Figures 3.8 to 3.12) included variation of key variables to determine changes in the outcome. On-scene treatment was tested at 22 min, 27 min, and 32 min to reflect aspirational, reported (Yannopoulos et al., 2020a), and historical values (Dennis et al., 2022; NSW Ambulance, 2019)

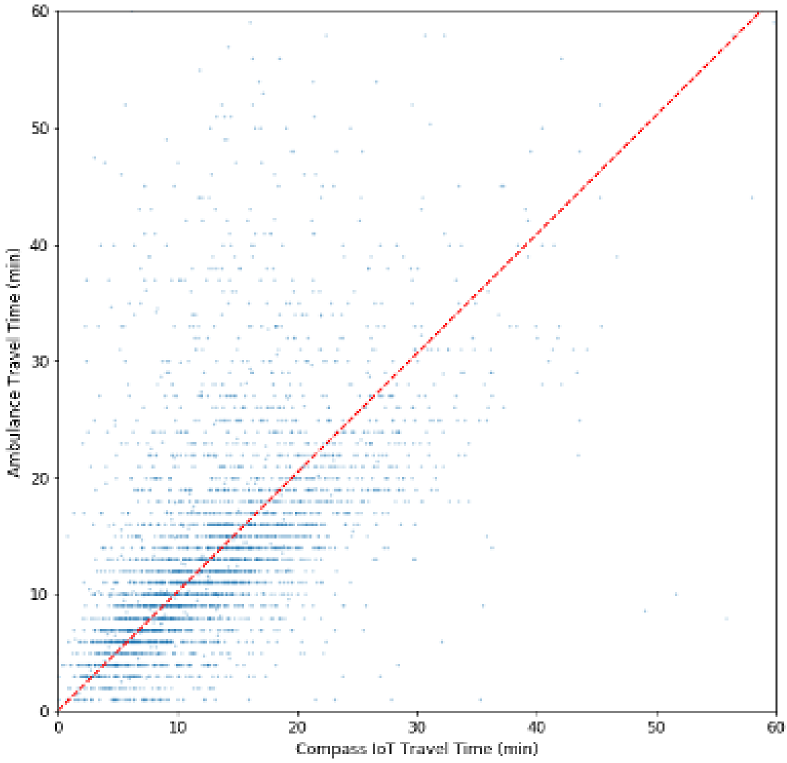

Figure 3.4: Validation of the Compass IoT travel times used to build the complete travel time matrix t_{ij} against the ambulance travel times in the NSW OHCA registry. As expected, the ambulance travel times show more variation. The red line shows a fit indicating consistency within 2%.

respectively. Pre-hospital cannulation time was tested at 22 and 27 min to take into account delays owing to the difficulty of the pre-hospital environment. Additional mobile ECMO teams were also tested (Figure 3.12).

3.2.5 Optimal location for basing pre-hospital ECPR team

The optimal location for placement of a pre-hospital ECPR team was determined by enumerating all possible locations, j, and identifying the one with the highest population-weighted survival probability, A_s. One additional model was completed assuming a mobile ECPR team located at an existing aeromedical base for practicality regarding staffing and restocking—Figure 3.13.

Sensitivity of the timepoint of when mobile ECMO team is dispatched was also assessed. The base case assumes activation of the pre-hospital team at time of EMS call as described in methods above. Modelling was completed assuming the activation of the mobile ECMO team, two minutes after the arrival of the first ambulance to allow for additional review of suitability for the pre-hospital ECPR—Figure 3.14.

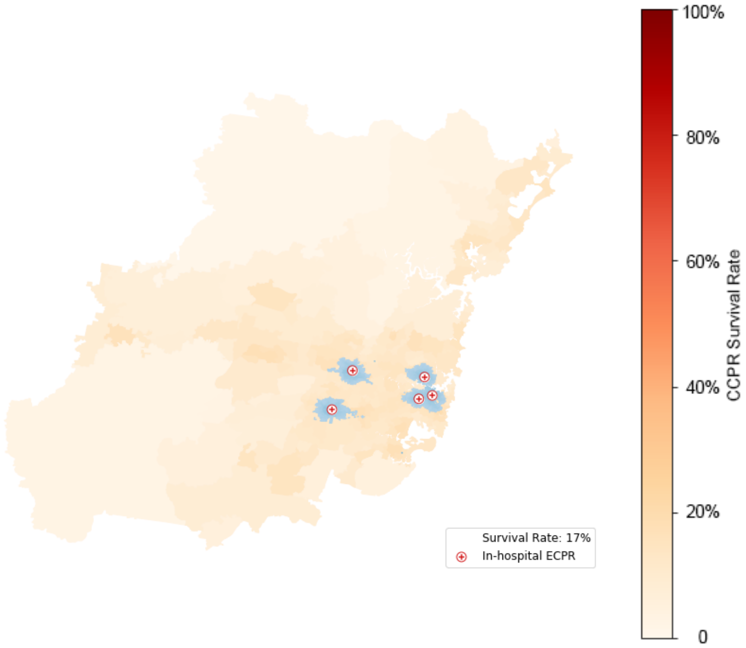

Figure 3.5: Survival rates subject to the current status quo of in-hospital ECPR offered at 5 hospitals assuming local response times, 27 min on-scene time, and 15 min cannulation. Eligibility is subject to a one-hour threshold. Blue shades indicate areas that are both within the 1 h threshold and offer higher survival from ECPR than CCPR.

3.2.6 Population coverage of ECPR delivery strategies

A basic summary of the ECPR delivery strategies is shown in Table 3.2. Results are quantified by three metrics: (1) the number of residents that can access ECPR and establish ECMO flows within 1 h from arrest, (2) the population-weighted average survival probability assuming a 1 h cut-off for eligibility and (3) the number of expected survivors based on the area-wide incidence of ECPR-eligible OHCA (NSW Ambulance, 2019).

Table 3.2: Summary of the performance of each delivery strategy. The two measures are: (1) Population that has access to ECPR within 60 min (2) the average survival probability assuming a threshold of 60 min from arrest to ECMO flow (NSW Ambulance, 2019)

Strategy	**Key model assumptions**	**Paramedic on-scene time (mins)**	**Cannulation time (mins)**	**Population able to access ECPR < 1 h**	**Survival probability across the city (%)**
In-hospital ECPR	5 current ECPR capable hospitals	27	15	811,091	16.56
Rendezvous ECPR	5 mobile teams leaving from current ECPR capable hospitals	27	15	2,175,096	22.42
Pre-hospital ECPR	1 mobile team + 5 current ECPR capable hospitals	N/A	22	3,851,727	42.71

Additional sensitivity analyses assuming variations in length on-scene treatment and pre-hospital cannulation times did not reveal significant changes in the coverage between the different ECPR delivery strategies—(Table 3.2).

3.2.7 Expected survival benefit

The expected survival benefit of different ECPR delivery strategies (blue and green) versus background survival rates from CCPR (orange) are represented in Figures 3.5 to 3.7.

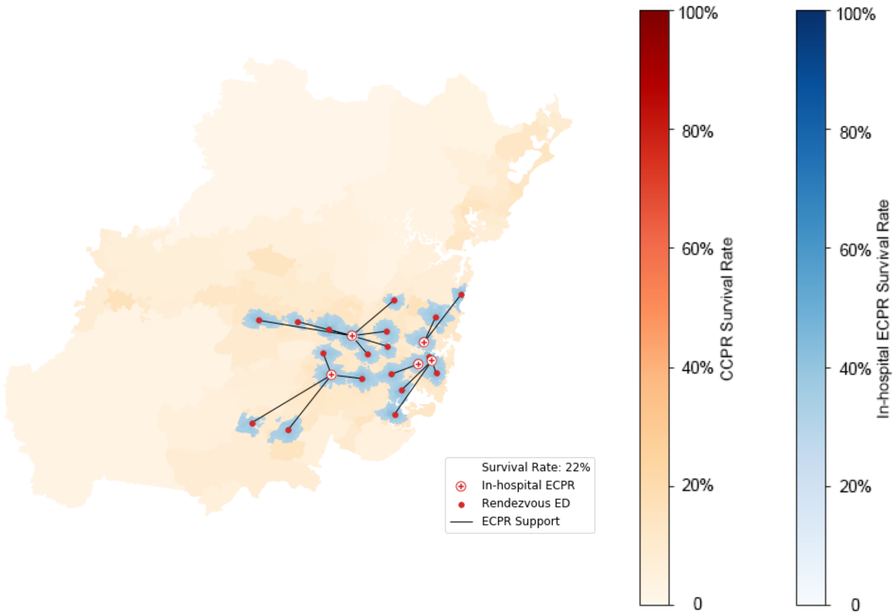

Figure 3.6: Survival rates subject to ECPR teams stationed at 5 hospitals using a rendezvous strategy with local response times, 27 min on-scene time, 15 min cannulation and a 1 h eligibility threshold. The lines connect each emergency department (red dots) with the ECPR team (red crosses) that will provide ECPR. Because coverage is limited by the on-scene time, the number of ECPR teams could be reduced without increasing resuscitation time.

3.2.8 In-hospital ECPR

Due to the time taken for ambulance response, on-scene treatment, transfer and cannulation, meshblocks that can reach in-hospital ECPR within 1 h are limited to patients within approximately 10 min travel time to the five ECPR-capable hospitals, hence only 17% of the population could be established on ECPR within an hour of arrest. These patients, indicated by the blue area in Figure 3.5, have a higher probability of survival compared to those outside the coverage area shown by the lighter shade of orange compared to the shade of blue.

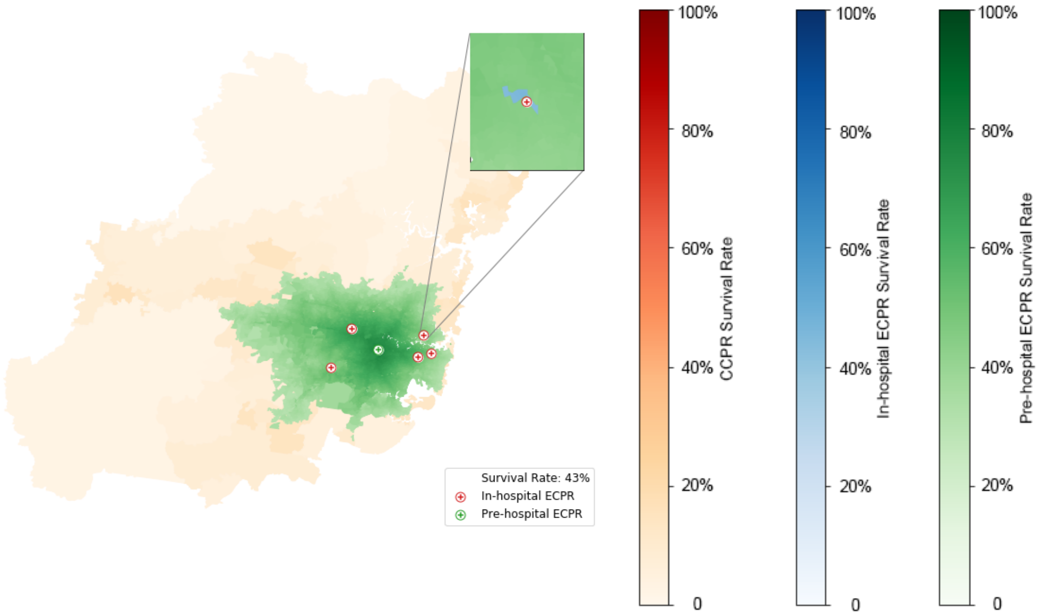

Figure 3.7: Survival rates subject to 1 pre-hospital ECPR team stationed at the optimal location and supported by 5 ECPR facilities. Time to resuscitation is based on the response time of the mobile team, cannulation time of 22 min to reflect the challenging environment, and a 1 h eligibility threshold. The inset shows that pre-hospital ECPR outperforms in-hospital ECPR at almost all locations.

3.2.9 Rendezvous strategy

The rendezvous delivery strategy—Figure 3.6 increases the spatial coverage of the ECPR facilities when compared to in-hospital ECPR by decreasing the hospital transfer time. This increased the population coverage to 2,175,096 and estimated survival to 22.4%. As in the case of the in-hospital strategy, on-scene time limits the coverage of rendezvous ECPR, and we observe discontinuous islands around the emergency departments where ECPR can occur. In the rendezvous strategy, the ECPR teams and patients move towards the emergency department (ED) simultaneously—whichever arrives last is the constraint on the coverage. The black lines connecting the EDs to the ECPR teams show how many EDs can be reached by the team, and the shaded area indicates the population that can receive rendezvous ECPR. The discontinuity between the shaded areas indicates that the patients' timeline is limiting the coverage, so the same level of service might be achieved with fewer ECPR teams because each team could travel further to an ED without impacting the number of patients that could reach that ED.

3.2.10 Pre-hospital ECPR

A single mobile ECPR team stationed at the optimal location, in addition to the existing 5 ECPR facilities, increases the population coverage to 3,851,727 with a potential weighted average survival rate increase to $42 \cdot 7\%$ as shown in Table 3.1 and Figure 3.7. This strategy substantially increases spatial coverage as well as population coverage, by reducing the times to establish ECPR. As shown by the dominance of green over blue, predicted survival from pre-hospital ECPR was higher than in-hospital ECPR, unless the patient arrested adjacent to the hospital (see inset of Figure 3.7).

Basing the pre-hospital ECPR team in the non-optimal location did result in small areas where existing in-hospital ECPR provided additional coverage over pre-hospital ECPR (Figure 3.13). When even more conservative assumptions are made e.g., delaying until the first ambulance arrives before activating the mobile pre-hospital ECMO team (at 2 min after arrival), the pre-hospital coverage is marginally smaller—Table 3.3 and Figure 3.14, with some small areas receiving better coverage from in-hospital ECPR—Figure 3.14, and the effective population able to reach ECPR reducing from 3,851,727 to 2,644,243 and a reduction in modelled survival of 14%.

3.2.11 Sensitivity Analysis

Table 3.3: Sensitivity Analysis for On-scene Time Variations. Each row represents a scenario (a strategy combined with timing assumptions). There are four measures: (1) population coverage, (2) expected survivors, (3) population-weighted survival probability with a 1-hour eligibility limit, and (4) population-weighted survival probability with no time limit.

Strategy	Key Model Assumptions	Additional delay from first ambulance arrival to dispatching the mobile team (mins)	Paramedic On-scene time (mins)	Cannulation time (mins)	Population able to access ECPR within 1 hour	Expected survivors	Survival probability across the city	
							1 Hour	No Limit
In-hospital ECPR	5 current ECPR capable hospitals	NA	22	15	1,572,036	42	20.74%	26.05%
			27	15	811,091	34	16.56%	21.67%
			32	15	276,671	29	14.14%	18.08%
Rendezvous ECPR	5 mobile teams leaving from current ECPR capable hospitals	0	22	15	3,496,950	63	31.23%	33.95%
		0	27	15	2,175,096	46	22.42%	28.35%
		2	27	15	2,134,042	45	22.15%	28.14%
		0	32	15	691,430	32	15.60%	23.31%
Pre-hospital ECPR	5 mobile teams	0	22	15	4,286,352	108	53.14%	53.62%
		2	22	15	3,684,974	78	38.54%	40.24%
		0	27	15	4,119,090	95	46.61%	47.49%
	1 optimal mobile team + 5 current ECPR hospitals	0	22	15	3,851,727	87	42.71%	44.36%
		2	22	15	2,644,243	58	28.58%	31.78%
		0	27	15	3,358,611	73	36.11%	38.46%
	1 alternative mobile + 5 fixed	0	22	15	3,068,491	68	33.25%	35.94%
		2	22	15	1,982,293	47	23.38%	27.33%
		0	27	15	2,543,245	57	28.10%	31.54%

Sensitivity Analysis – On-scene time variations The results presented above assume a 27-min on-scene time for the in-hospital and rendezvous strategies, which is consistent with published trials (Belohlavek et al., 2022; Yannopoulos et al., 2020a). For the in-hospital strategy, the 16.6% population-weighted average survival probability associated with 27 min can be compared to 20.7% with an aspirational 22-min on-scene treatment time or 14.1% with a 32-min on-scene treatment time that is similar to typical values from the cardiac arrest registry (NSW Ambulance, 2019). These differences are illustrated in Figures 3.8 and 3.9 for the in-hospital and rendezvous strategies respectively.

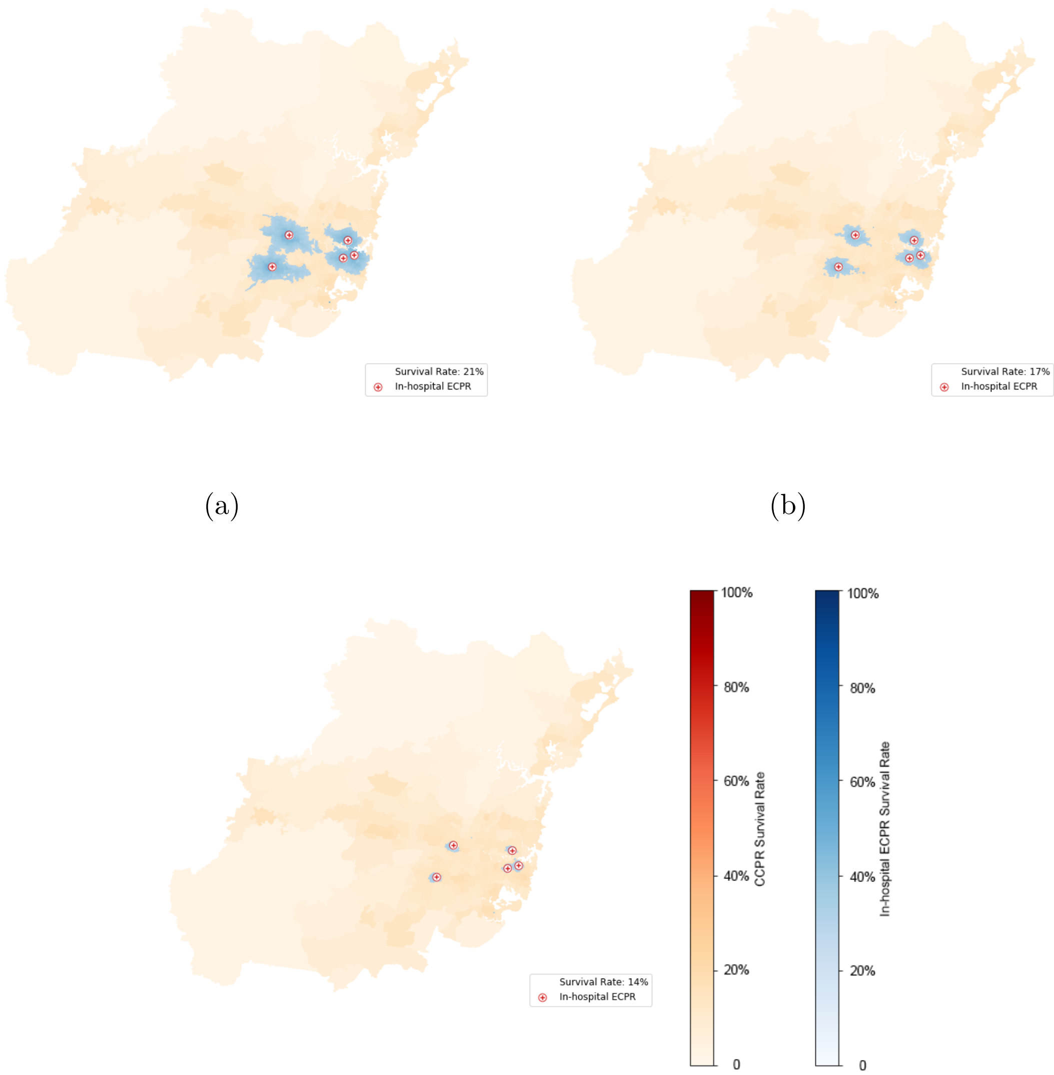

(a) (b) (c)

Figure 3.8: Sensitivity to the On-Scene Time for the In-Hospital Strategy. As the on-scene time increases from 22 min (a) to 27 min (b) to 32 min (c), the coverage area diminishes. Both the coverage and the population weighted survival rate (see legends) are sensitive to the assumptions about on-scene time.

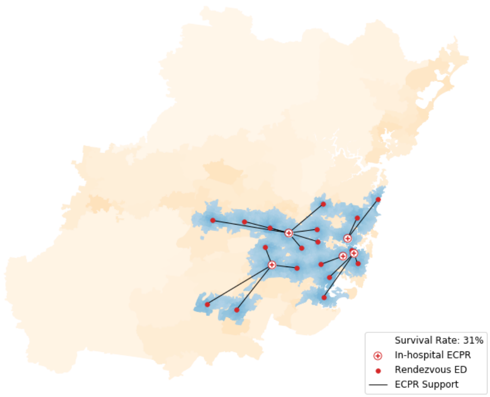

(a)

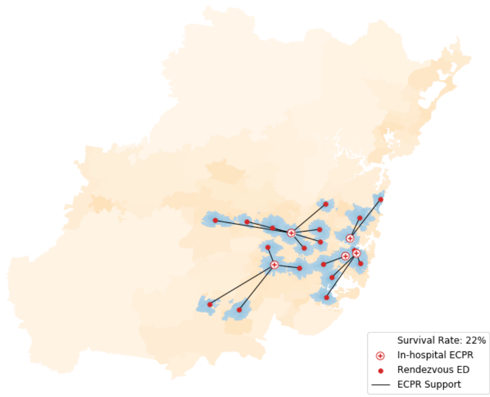

(b)

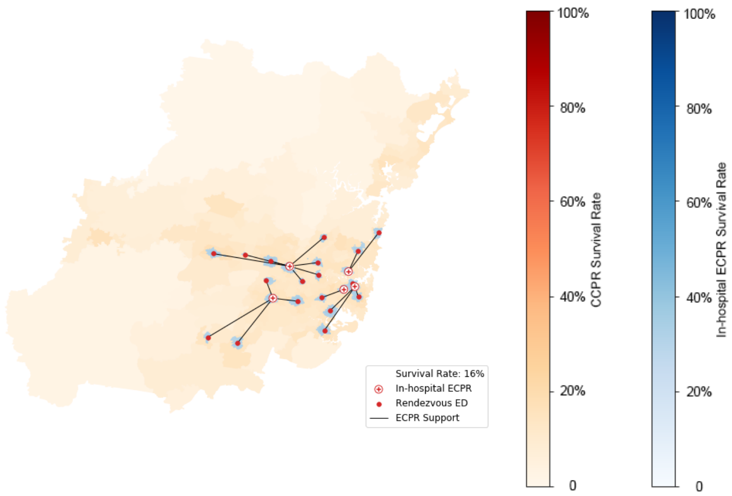

(c)

Figure 3.9: Sensitivity to the On-Scene Time for the Rendezvous Strategy. As the on-scene time increases from 22 min (a) to 27 min (b) to 32 min (c), the coverage area diminishes. Similar to the in-hospital strategy, the results are sensitive to the assumptions about on-scene time. Moreover, the merging of the blue areas for the 22-min on-scene time indicates that for aspirational on-scene times, the rendezvous strategy may start to be constrained by the locations of the five mobile ECPR teams.

Sensitivity Analysis – Cannulation time variations The pre-hospital model requires cannulation in an unpredictable context, which may involve moving the patient. To address this challenge, the modelling assumes an additional seven minutes for cannulation compared to in-hospital cannulation (15 min). The pre-hospital ECPR part of Table 3.3 presents comparisons between this extended cannulation time (22 min) and an additional five minutes (27 min). The results are mapped in Figure 3.10 below.

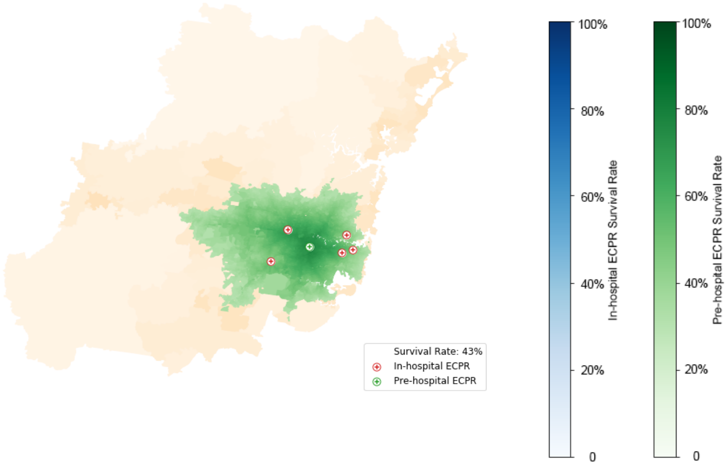

(a)

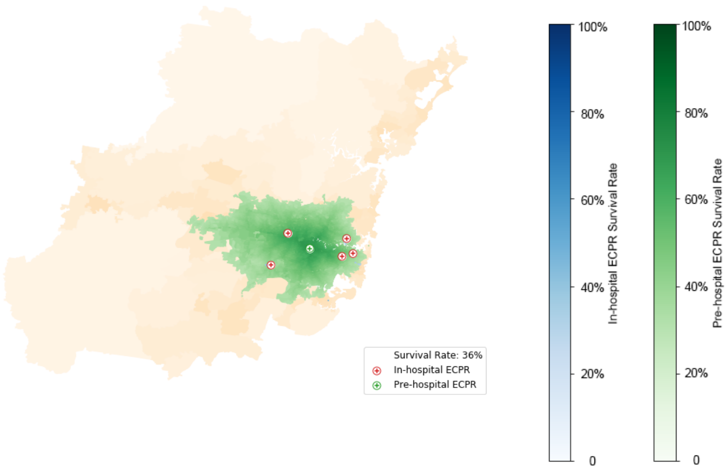

(b)

Figure 3.10: Sensitivity to the assumptions around cannulation time for pre-hospital ECPR. The extended cannulation time of 22 min (a) provides better coverage than the more conservative 27 min cannulation time (b). However, the difference in performance between these two cannulation times is smaller than the difference between the strategies.

Sensitivity Analysis - Relaxing the 1-hour constraint for arrest to ECMO flows Current guidelines suggests that ECPR should be initiated within 60 min of cardiac arrest, which lends itself to the cumulative opportunities approach to accessibility (Wachs and Kumagai, 1973). Even when the resuscitation time exceeds 1 hour, the survival rate of cardiac arrest patients is not zero with ECPR (Scquizzato et al., 2022). Moreover, the time decay of survival with resuscitation time suggests that, even within the coverage area, patients further from ECPR facilities should be weighted proportionally to their probability of survival. Relaxing the eligibility threshold will increase the number of patients that can receive ECPR but it will lower the average survival rate of ECPR patients by including more remote patients with low survival rates who would otherwise have received CCPR.

Sensitivity Analysis – Relaxing the 1-hour Constraint

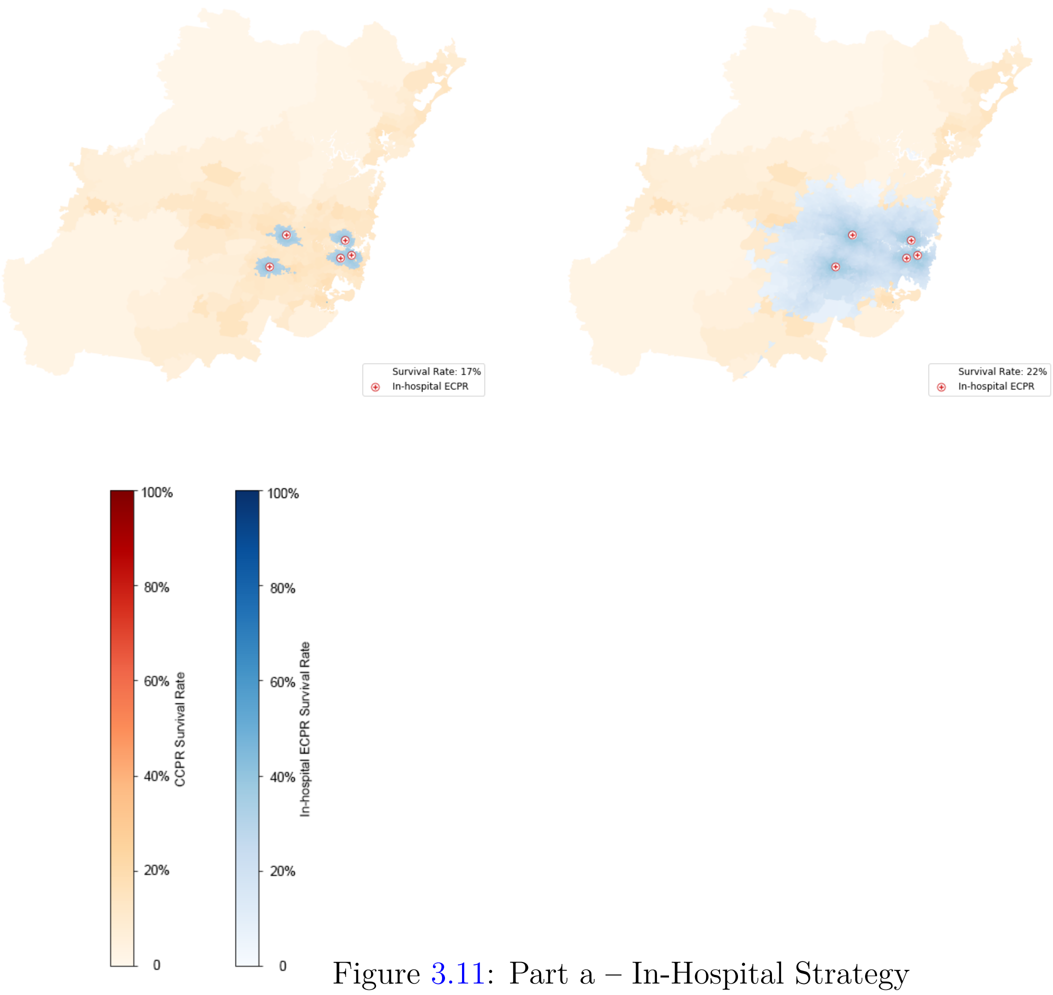

Figure 3.11: Part a – In-Hospital Strategy

Comparing the left and right of Figure 3.11, relaxing the one-hour time limit offers improved survival probability to large areas in all three strategies. Notably, the discontinuous areas merge together when there is no time threshold for eligibility—this removes inequitable edge effects caused by a firm coverage boundary. The improvement is strongest in the rendezvous case where the relaxation allows the addition of three EDs and surrounding catchments. Comparing the population-weighted survival probability (see legends), the relaxation of the 60-min limit has only a small impact because the additional coverage areas tend to have low population and only marginally improved expected outcomes from ECPR compared to CCPR. Indeed, some areas which are far away from ECPR facilities still show higher survival with CCPR because ambulance response is fast—in these regions, reducing ambulance

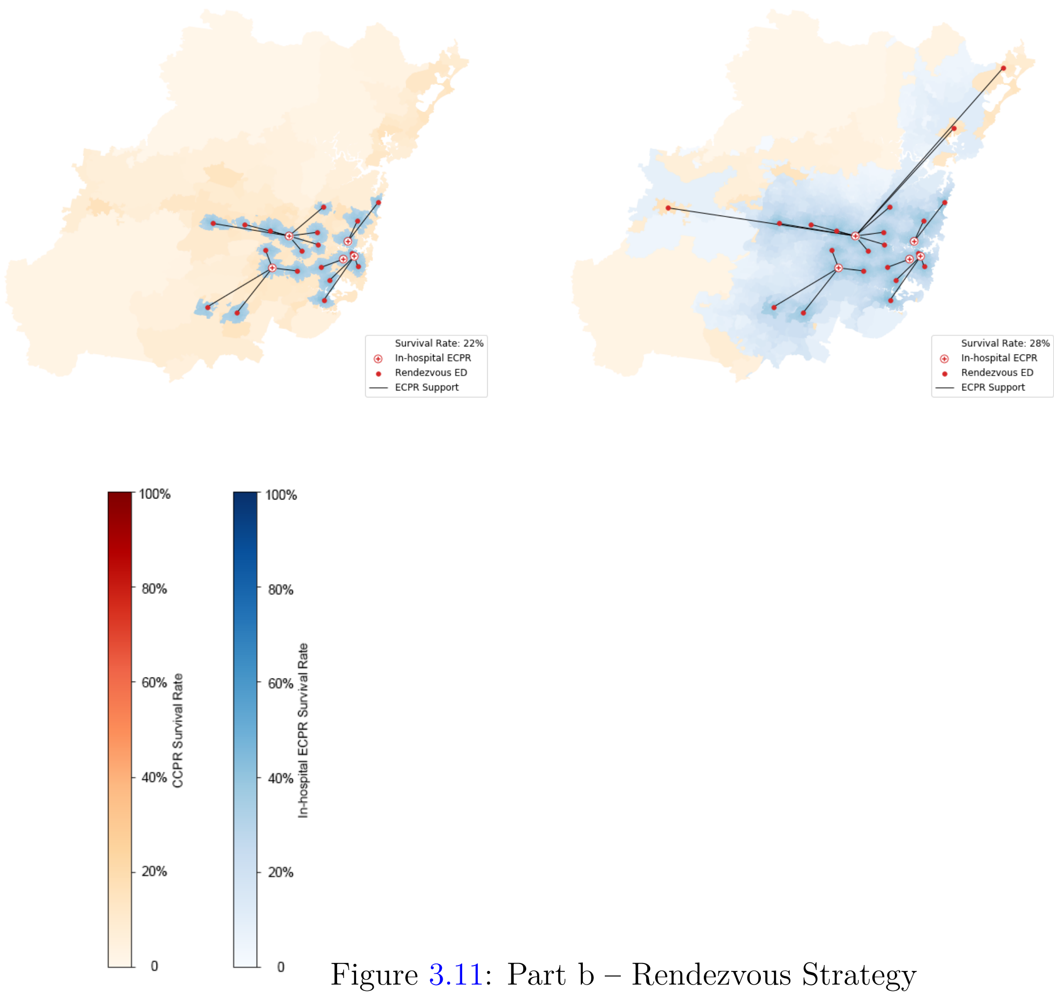

Figure 3.11: Part b – Rendezvous Strategy

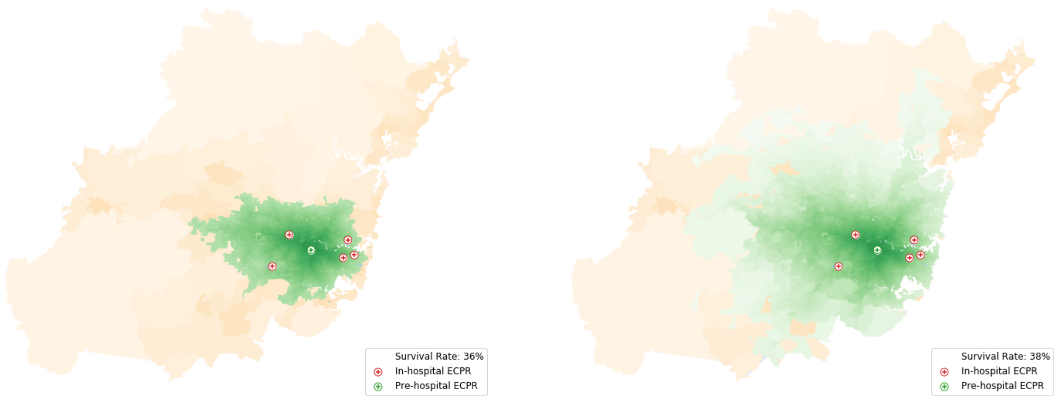

Figure 3.11: Part c – Pre-Hospital Strategy

response time may be more effective than expanding ECPR coverage. Further evidence is needed to explore how relaxing the 60-min cut-off might improve the overall

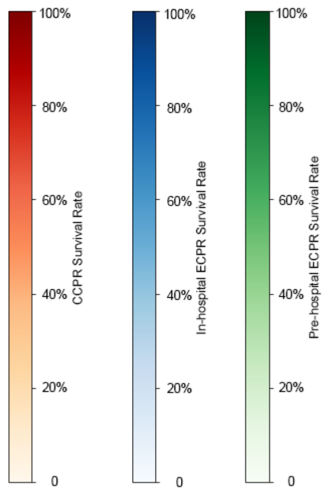

Figure 3.11: Comparison of the Strategies without the 60-Min Cut-Off. The coverage area is reduced substantially in the in-hospital (a), rendezvous (b) and pre-hospital (c) models when you impose the threshold (left) compared to when there is no limit (right). In the no-limit cases, the green and blue shading indicates areas where ECPR at the relevant time offers an improvement over CCPR (orange). The population-weighted survival rate (see legends) shows that the additional areas tend to be low population with marginal improvement over CCPR, so the increase in average survival rate is small.

survival and where CCPR treatment is advantageous to patients. This will depend on whether ECPR teams are unavailable e.g., serving one patient will come at the expense of another, or if the system has significant redundancy. In this analysis we assume ECPR resources and personnel are unconstrained.

Sensitivity Analysis - Number of pre-hospital teams The number of pre-hospital ECPR teams has a substantial impact on system performance. As illustrated in Figure 3.12, deploying five pre-hospital teams—each based at an existing ECPR hospital—significantly increases the coverage area compared to the in-hospital strategy using the same five locations. Survival probabilities in densely populated areas are also improved, as patients receive ECMO flow more rapidly due to the reduced response times provided by mobile teams. These results highlight the operational advantage of delivering ECPR to the patient, rather than relying on transport to fixed facilities.

Sensitivity Analysis - Location of pre-hospital team In the pre-hospital strategy illustrated in Figure 3.13 the pre-hospital team waits for patients at the optimal location (left image). This location is identified by enumerating all possible locations, calculating the population-weighted survival probability and choosing the one with the best outcome. This identified a base in the vicinity of the suburb of Homebush as the optimal location. To understand the sensitivity to the location a second

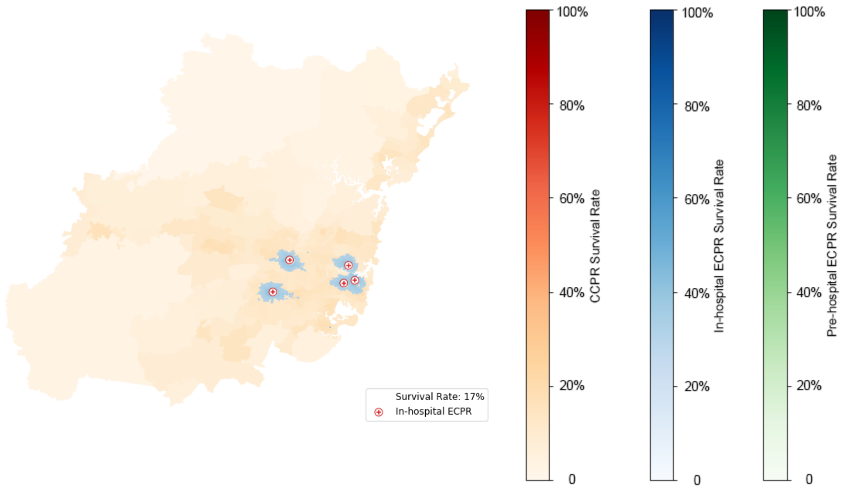

(a)

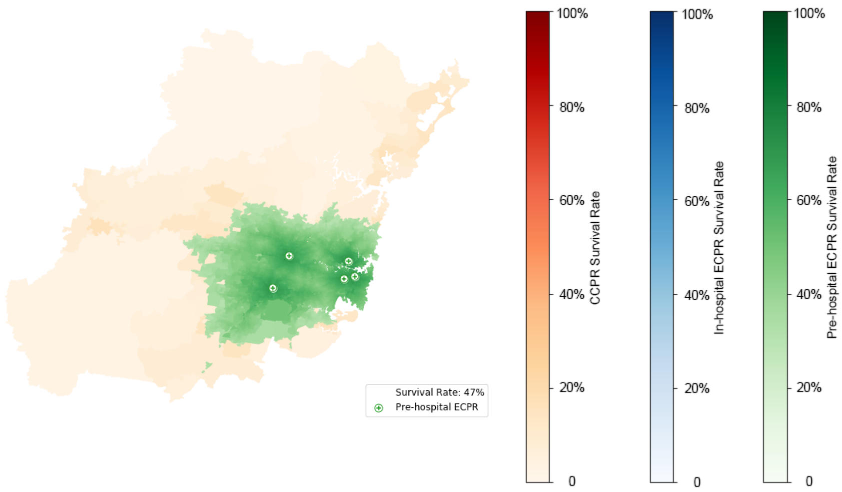

(b)

Figure 3.12: 1-hr coverage from five in-hospital ECPR facilities (a) compared to five pre-hospital ECPR teams at the same locations (b). Bringing ECPR to the patient instead of bringing the patient to ECPR dramatically expands the number of potential patients and the expected survival rate. The bottom maps shows that many of the most populous areas have survival probability of nearly 85% as time to ECMO flow times are substantially shorter.

location where the mobile ECPR team is located at an existing major aeromedical base (right picture) is provided.

The reduced spatial coverage and population-weighted survival rate show that the results are sensitive to the specific location of the ECPR team. Additionally, under the alternative location, the most populous areas of the metropolitan region have better expected outcomes from in-hospital ECPR than pre-hospital ECPR. In the bottom map of Figure 3.13, CCPR, in-hospital ECPR or pre-hospital ECPR are all

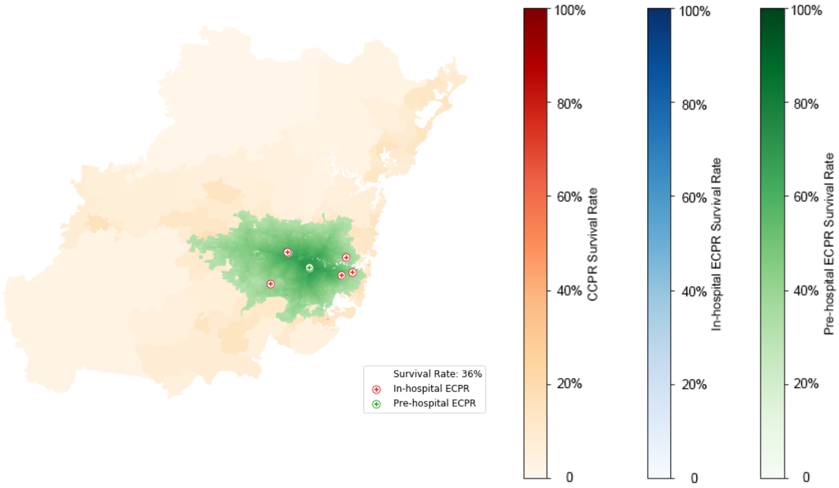

(a)

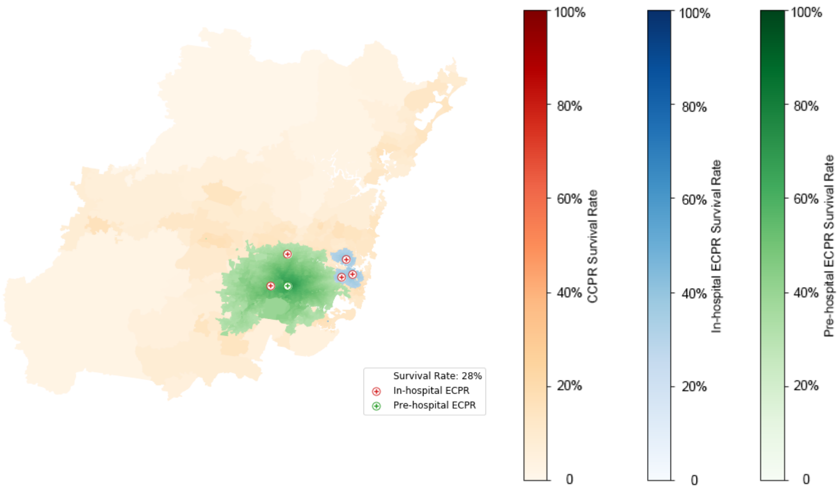

(b)

Figure 3.13: Pre-hospital ECPR coverage with the optimal location for the mobile team (a) compared to another location supported by historical and logistical reasons (b). Although the pre-hospital teams are nearby on the scale of the map, the non-optimal location has less spatial coverage and 10% lower population-weighted survival probability. Moreover, with this location, some areas have been survival probability from in-hospital ECPR (blue) because the response time from the single pre-hospital unit is so long.

used depending on the location and timings of the case. This may have implications for maintaining clinical exposure and equity of care.

Sensitivity Analysis – Dispatching the mobile ECMO team after initial ambulance arrival In our baseline model for the rendezvous and pre-hospital strategies, we assume the mobile ECMO team is dispatched at the time of the call. A more conservative, alternate, approach suggests that the dispatch is delayed until a responding ambulance assesses the patient and determines ECPR eligibility. This is assumed in the model to occur after 2 minutes to allow for brief additional assessment of ECPR eligibility. Delaying pre-hospital ECPR activation until arrival of the first ambulance on scene and initial assessment has important implications for how successful the alternative strategies will be because delaying dispatch decreases probability of patient survival, but does decrease the “false ECPR call outs”.

A delayed dispatch, owing to waiting for the initial ambulance to arrive and assess the patient before activation of the pre-hospital team causes an increase low flow time. This is reflected in lower coverage populations, survival probabilities and expected survivors as detailed in Table 3.3 and Figure 3.14 below. The reduced coverage areas result in in-hospital ECPR being competitive in some areas, which has implications for the resourcing of the best ECPR strategy. The maps shown in Figure 3.14 show that waiting for the responding ambulance to assess the patient will result in a decrease in the modelled survival probability of ECMO-eligible patients of 14%.

3.3 Discussion

This chapter employs transport accessibility frameworks to demonstrate how the choice of delivery strategy in logistics profoundly affects the performance of emergency medical services (EMS). In-hospital ECPR serves as the foundational model, but its geographic reach is limited, as many areas fall outside feasible transfer times. Despite these limitations, in-hospital ECPR contributes to improved survival rates due to its established efficacy in managing cardiac arrest. Alternative strategies, such as the rendezvous and pre-hospital models, significantly enhance accessibility, enabling a greater proportion of the urban population to receive care within the critical 60-minute window. Notably, the pre-hospital strategy exhibits the most pronounced benefits.

These findings affirm core principles in transportation and logistics: a customized delivery strategy can enhance service capacity, provided resource constraints are addressed. The rendezvous model reduces the overall time to advanced care by dividing the journey between the patient and the ECMO team. The pre-hospital model extends this concept by effectively “relocating” the advanced facility to the patient’s location. However, scaling these alternative models requires meticulous planning, inter-agency coordination, and substantial resource investments, including specialized vehicles, trained personnel, and well-defined operational protocols.

Furthermore, the results underscore the necessity of incorporating local contextual factors—such as traffic patterns, road infrastructure, and population density—into

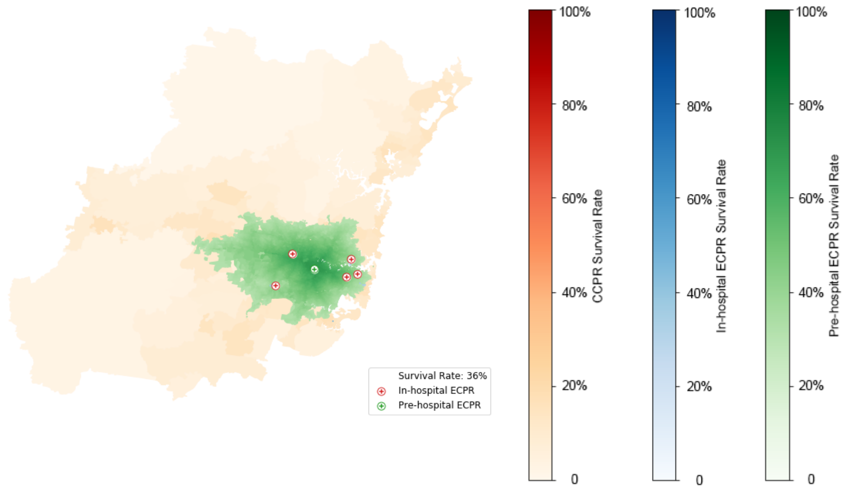

(a)

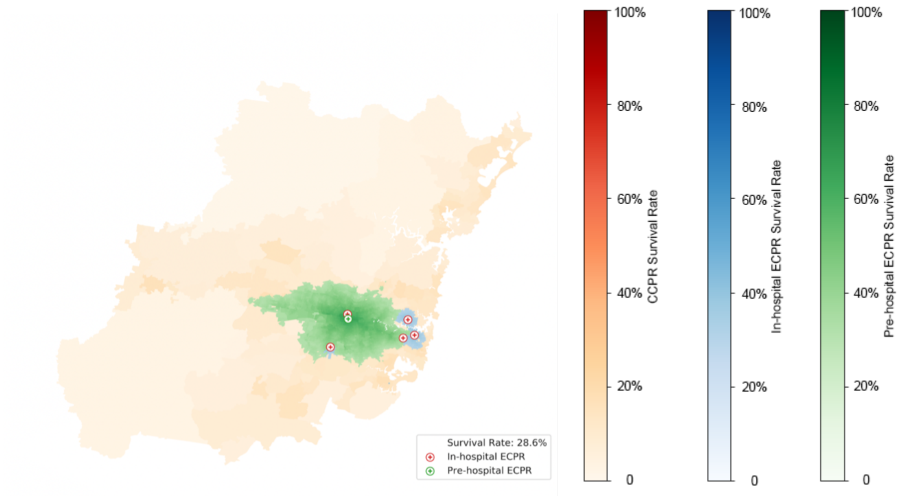

(b)

Figure 3.14: Comparison of pre-hospital ECPR coverage when the mobile team is dispatched immediately versus 2 minutes after the ambulance arrives. Some blue (figure 9 b) areas appear showing that in-hospital ECPR becomes relevant in this more conservative policy.

ECPR strategy development. In certain regions, a rendezvous system may be sufficient to meet demand with fewer mobile teams, while in others, a comprehensive pre-hospital program may be essential to reach the majority of patients. These contextual considerations are pivotal for optimizing resource allocation and ensuring equitable access to advanced interventions across diverse settings.

Chapter 4

Optimizing Access to Medical Resources Through Spatial Modeling

4.1 Introduction

This chapter presents methods for optimising the location of medical facilities to provide more benefits to residents. Building upon past work of Facility Location Problems and accessibility measures, we propose a rendezvous strategy to extend the coverage of facilities and use a survival rate function as the impedance of accessibility. This chapter further develops a problem-specific model that links facility configuration with survival-based accessibility and spatial equity, providing new insights into how location decisions influence the performance of the rendezvous strategy in EMS. The details of the methodology will be explained in Section 4.2. Section 4.3 introduces a test case of ECPR cardiac arrest care in New South Wales Australia, and the results stemming from this application will be analyzed in Section 4.4. The implications for emergency logistics and healthcare planning will be discussed in Section 4.5.

4.2 Methodology

In this study, we propose a rendezvous strategy to increase the coverage of emergency health care facilities. Because the patient distribution is not homogeneous and the survival rate depends on the interval between the incident and treatment, the objective function in the MCP will be represented by accessibility downweighted by patient survival expectation.

4.2.1 Rendezvous Strategy

This problem overall aims to optimise the benefit of a resource-constrained healthcare service with limited equipment and personnel. The baseline scenario is to use ambulances to transfer patients from the scene of the incident to elite hospitals for treatment. In the proposed rendezvous model, a mobile ECPR team consists of trained clinicians capable of providing extracorporeal resuscitation outside hospital facilities. These teams travel in dedicated ambulances that carry portable ECMO equipment or can be transported to a meeting point by other emergency vehicles when required. To make better use of these scarce resources, the mobile team and its equipment can be mobilised to meet patients at an intermediate emergency department. The patient and the mobile team converge at the rendezvous point, where advanced life-support procedures are initiated before definitive transfer to an intensive care unit. Depending on operational feasibility, such rendezvous points may include general emergency departments, ambulance bases, or other predefined meeting sites within the service network. This arrangement shortens the time between the patient departing the scene and receiving treatment without increasing the number of personnel or equipment. The strategy is illustrated in a schematic in Figure 4.1.

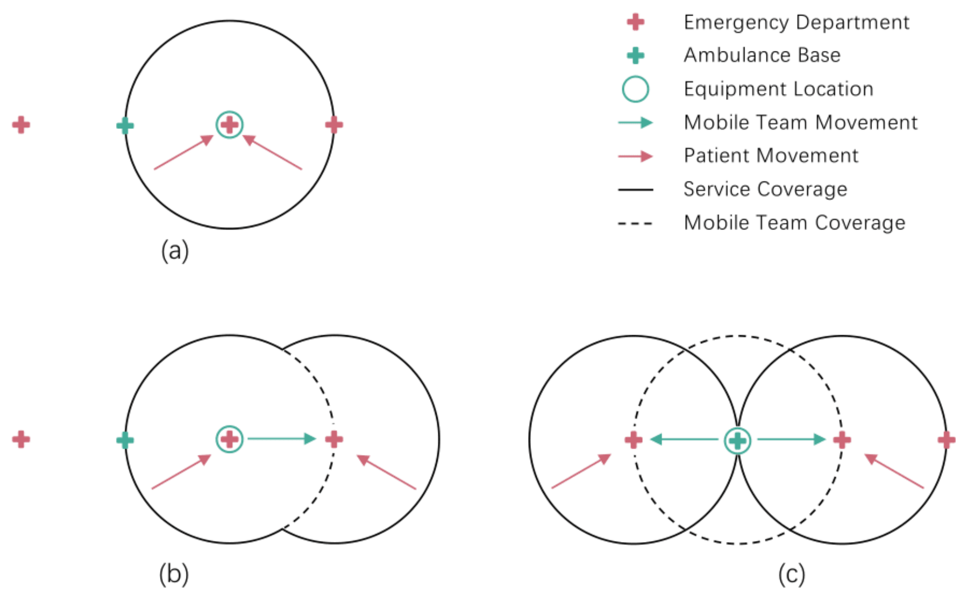

Figure 4.1: Schematic diagram of delivery strategy. (a) In the baseline strategy, only the elite emergency departments provide the service and the catchment is limited to a single black circle. (b) In the rendezvous strategy, the general EDs could also support treatment at a rendezvous point which increases the size of the catchment of each mobile team. (c) The mobile teams can also be on standby at non-resourced locations (i.e. an ambulance base), in which case they can only support other EDs. The total coverage area increases from (a) to (b) to (c).

4.2.2 Generalized Maximum Coverage Model

In general, MCPs maximise the coverage of a facility, whereas a weighted MCP places different values on each area, for example, valuing coverage of some areas more than others. In this methodology, the weight parameter will be generalized to an impedance function of patient survival rate. Compared to a general Facility Location Problem, we constrain the potential rendezvous points to the subset of locations with existing emergency departments and more specifically, those that are within the coverage of the mobile teams.

There are some assumptions for this model.

- Patients will experience a baseline survival expectation associated with their location regardless of whether they are within the coverage area.
- This research is on planning scale, so the travel times are assumed static.
- Optional facility locations and rendezvous locations are emergency departments or ambulance bases only, so they are finite.
- Under the rendezvous strategy for a specific patient location and facility location, the patient could be served by multiple viable rendezvous locations. Treatment can only occur when both the patient and the mobile team arrive at the same location. Hence, in the assumption, the patient will be taken to the best rendezvous point, which is represented by the maximum survival (i.e. minimum travel time) associated with the larger of a) the time for the personnel and equipment travelling to the rendezvous point and b) the patient travelling to the rendezvous point.
- The patients' survival probability is monotonically decreasing with time.

In the following exposition, various variables and parameters have been introduced to streamline the analysis. To begin, I symbolizes the set of demand locations with i representing the specific location in the location set. The collection of potential rendezvous locations that may be activated by the supply or equipment site is denoted as J, while j serves as their corresponding elements. For the positions where the supply or equipment might be allocated, the term K is employed. The element of these optional supply or equipment locations K is represented by k.

The binary variable x_i is deemed to be 1 if any rendezvous location j falls within the expected threshold r of location i, otherwise it stands at 0. Similarly, y_j, another binary variable, equals 1 when the rendezvous location j is within the expected threshold r of any supply location k. z_k is a binary variable that equals 1 only if there is supply or equipment allocated at site k.

Furthermore, $s_i(\mathbf{z},\mathbf{y}, x_i)$ is utilized to denote the survival rate of patients situated at location i. It takes the greater value of the baseline survival from conventional treatment at that location (s'_i) and the survival resulting from the equipment at $\mathbf{k}$ and rendezvous locations at $\mathbf{j}$. f describes the monotonically decreasing survival function. The variable n_i expresses the number of patients at location i. t_{ij} and t_{kj} indicate the total time from the patient location i and the supply/equipment location k to the rendezvous location j, respectively. p represents the total count of equipment or teams.

Although x_i and $\mathbf{y}$ are defined as variables in the formulation, they are not independent decision variables; their values are logically determined by the coverage constraints, whereas $\mathbf{z}$ represents the true decision variable optimized in the model.

Reflecting these requirements for the alternative delivery strategy, the facility location optimisation is formulated as:

$$\max_{\mathbf{z},\mathbf{y},\mathbf{x}} \sum\nolimits_{i\in U} s_i(\mathbf{z},\mathbf{y},x_i)n_i \tag{4.1a}$$

$$\text{s.t.} \sum\nolimits_{k\in K} z_k \leq p \tag{4.1b}$$

$$\sum\nolimits_{j\in J:\ t_{ij}\leq r} y_j \geq x_i \forall\ i \in I, \tag{4.1c}$$

$$\sum\nolimits_{k\in K:\ t_{kj}\leq r} z_k \geq y_j \forall\ j \in J, \tag{4.1d}$$

$$s_i(\mathbf{z},\mathbf{y},x_i) = \max(s'_i, f(\min_{j,k}(\max(t_{ij},\ t_{kj})))x_i y_j z_k) \tag{4.1e}$$

$$x_i \in \{0,1\} \text{ for } \forall\ i \in I \tag{4.1f}$$

$$y_j \in \{0,1\} \text{ for } \forall\ j \in J \tag{4.1g}$$

$$z_k \in \{0,1\} \text{ for } \forall\ k \in K \tag{4.1h}$$

$$\tag{4.1i}$$

I	the set of demand locations
i	the location in location set I
x_i	binary variable, equal to 1 if and only if rendezvous location j is within the expected threshold r of location i
J	the set of potential rendezvous locations that could be activated by supply/equipment location
j	the location in potential rendezvous location set J
y_j	binary variable, equal to 1 if and only if supply location k is in the expected threshold
K	the set of locations where the supply/equipment could be allocated

k	the location in optional supply/equipment location set K
z_k	binary variable, equal to 1 if and only if there is supply/equipment allocated at site k
$\mathbf{z}$	the combination of a group of supply/equipment locations z_k
n_i	the number of patients at location i
s_i	the rendezvous survival rate of patients in location i
s'_i	the baseline survival rate of patients in location i
f	the survival function associated with the rendezvous strategy
$t_{i,j}$	the total time from demand location i to arrival at rendezvous location j
$t_{k,j}$	the total time from supply/equipment location k to arrival at rendezvous location j
p	the total number of equipment/teams
r	the maximum threshold for coverage

This formulation is adapted from the maximal covering location problems of Ahmadi-Javid et al. (2017) in order to accommodate constraints on the coverage of the equipment as well as the coverage of the hospital as presented in the original. Eq. (4.1a) presents the objective function that maximises the sum of the expected survival probability of each location weighted by the population at that location. The first constraint is that the number of locations with equipment should be less than or equal to the total number of equipment (Eq. (4.1b)). The second constraint (Eq. (4.1c)) is if a patient location is covered then there must be at least one ED within the travel time threshold. The third constraint (Eq. (4.1d)) is that if an ED is covered then there must be at least one mobile team within the travel time threshold.

In Eq. (4.1f), we define x_i to be a binary variable indicating whether a patient location, i is either served (1) or not served (0) by a rendezvous location. In Eq. (4.1g) we define y_j to be a binary variable indicating whether potential rendezvous location j is either supported (1) or not supported (0) by a supply/equipment. In Eq. (4.1h) we define binary variable z_k to be a binary variable indicating whether a facility location k is either allocated (1) or not allocated (0) a supply/equipment—the set of z_k are the fundamental decision variables of the facility location problem.

The impedance function, (4.1e), referenced as s_i in Eq. (4.1a), expresses how the quality of the coverage degrades with time. If the function equals to 1 inside the threshold r and 0 elsewhere, the formula will be a classic MCP. Setting the survival rate function as the impedance function in Eq. (4.1a), it expresses a discrete gradual coverage location problem or a special case of uncapacitated facility location model with zero fixed costs (Berman et al., 2010; Cornuéjols et al., 1983). Using both the survival rate function and a threshold r, this formulation represents a maximal survival problem using the rendezvous strategy.

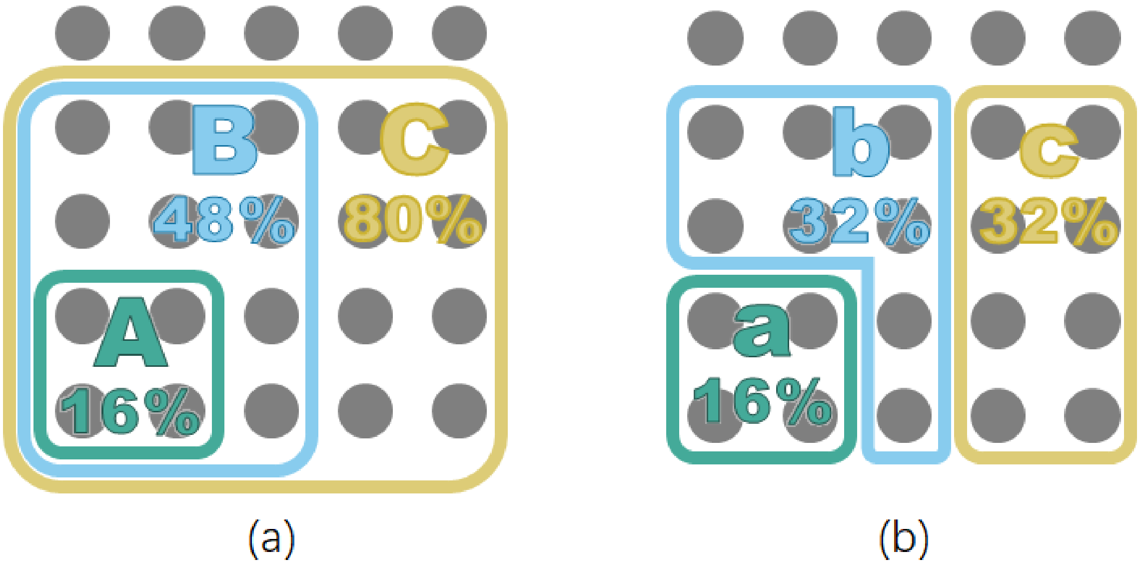

Figure 4.2: Representation of the recasting of the Maximal Survival Problem (a) to the Maximum Coverage Problem (b). Each dot is a person at location i showing that additional facilities offers an incremental improvement in coverage within a single location. In the MSP, each location i has layered expected survival from multiple facilities. In the recasted problem, we have a simple MCP over a bigger space because each location i is mapped to multiple subpopulations.

Notably, this formulation differs from a classic MCP with non-linear aspects in the time-dependent impedance function. This time is also determined by the maximum value between patient and equipment arrival at the rendezvous point. However, the problem features limited optional facility locations and rendezvous points. And the travel times should be static on planning scale. Hence, through some preprocessing, it is feasible to recast the problem as a maximum coverage problem.

For a given facility location, all rendezvous points within the threshold can be set as available. Subsequently, the survival probability of each demand location at each rendezvous point can be computed. For a certain demand location, the maximum of its survival probabilities at all rendezvous points is the survival rate at that facility location. Multiplying this by the demand at that location yields the expected survival count.

Figure 4.2(a) illustrates a scenario where the grey dots represent all demand at a single location. Note that 100% of the patients experiencing x% survival is mathematically equivalent to x% of patients experiencing 100% survival (aka full coverage). Facilities A, B, and C contribute survival rates of 16%, 48%, and 80% at this location, respectively. Since survival rates are not cumulative, the maximum value is selected whenever multiple facilities are deployed. Through preprocessing, this scenario can be transformed into the one depicted in Figure 4.2(b). The preprocessing sorts survival rates at this demand location from low to high, then splits the demand into several subelements based on the differences between adjacent values. For instance, Facility A only covers the demand of subelement a, Facility B covers the demands of subelements a and b, while Facility C can cover all three subelements a, b, and c. If Facility A is present, the addition of Facility B covers an extra subelement b. However, if Facility C is already present, the addition of either Facility A or B does not cover any new subelement.

Therefore, for limited facilities and rendezvous locations and static travel times, this problem can be recast into a maximal coverage problem through the preprocessing described above. The total preprocessing time complexity can be expressed as $O(|I| \cdot |J| \cdot |K|) + O(|I| \cdot |K| \log(|K|))$. The first term, $O(|I| \cdot |J| \cdot |K|)$, represents the survival rate calculations for all demand nodes ($i \in I$), rendezvous locations ($j \in J$), and equipment locations ($k \in K$). The second term, $O(|I| \cdot |K| \log(|K|))$, corresponds to the sorting operations performed over the equipment number dimension for each demand node. The overall preprocessing complexity is determined by the relative sizes of $|J|$ and $\log(|K|)$.

In this recast problem, we define $\tilde{I}$ as the set of potential survival population groups, where each group is indexed by $\tilde{i}$ and $|\tilde{I}| = p \cdot |I|$ are needed. For each group, a binary variable $\tilde{x}_{\tilde{i}}$ is associated, which equals to 1 if there exists at least one supply or piece of equipment that can achieve the survival rate through a rendezvous location within the expected threshold r of that group.

The supply or equipment can be allocated at different locations that constitute the set K, with k denoting the index of these optional supply or equipment locations. A binary variable z_k is utilized, which equals 1 if and only if there is a supply or equipment allocated at site k. The demand at location $\tilde{i}$ is represented by $\tilde{n}_{\tilde{i}}$.

In addition, $\tilde{t}_{\tilde{i},k}$ is defined as the minimum of the maximum total time from demand location $\tilde{i}$ to the rendezvous location and from supply or equipment location k to the rendezvous location.

Lastly, p represents the total number of available supply or equipment, and r is the maximum threshold that is available for coverage.

The transformed formula, which corresponds to a classic MCP, is shown as follows:

$$\max \sum_{\tilde{i}\in\tilde{I}} \tilde{n}_{\tilde{i}}\tilde{x}_{\tilde{i}} \tag{4.2a}$$

$$\text{s.t.} \sum_{k\in K} z_k \leq p \tag{4.2b}$$

$$\sum_{k\in K;\tilde{t}_{\tilde{i},k}\leq\tilde{r}} z_k \geq \tilde{x}_{\tilde{i}};\ \tilde{i}\in\tilde{I} \tag{4.2c}$$

$$\tilde{x}_{\tilde{i}} \in \{0,1\},\ \tilde{i}\in\tilde{I} \tag{4.2d}$$

$$z_k \in \{0,1\},\ k\in K. \tag{4.2e}$$

$\tilde{I}$	the set of potential survival population group
$\tilde{i}$	the index of potential survival population group $\tilde{I}$
$\tilde{x}_{\tilde{i}}$	binary variable, equal to 1 if and only if there is at least one supply/equipment can achieve the survival rate via a rendezvous location in the expected threshold r of group $\tilde{i}$
K	the set of locations where the supply/equipment could be allocated
k	the index of the optional supply/equipment locations K
z_k	binary variable, equal to 1 if and only if there is supply/equipment allocated at site k
$\tilde{n}_{\tilde{i}}$	the demand at location i
$\tilde{t}_{\tilde{i},k}$	the minimum of the maximum total time from demand location $\tilde{i}$ to rendezvous location and from supply/equipment location k to rendezvous location
p	the total number of available supply/equipment
r	the maximum available threshold

After recasting the problem, (4.2a) is the objective function to maximize survival across the population. (4.2b) represents the limited number of equipment/teams. (4.2c) is the coverage constraint— if the subpopulation in $\tilde{i}$ is covered ($\tilde{x}_{\tilde{i}} = 1$), the total number of equipment z_k within the threshold r should be at least 1. If not, both z_k and $\tilde{x}_{\tilde{i}}$ should be 0 . Eq. (4.2d) represents a binary variable indicating whether the location $\tilde{i}$ is served (1) or not served (0) via any rendezvous points. Eq. (4.2e) represents a binary variable which indicates whether the optional facility location k is allocated (1) or not allocated (0) equipment.

The equivalence of the maximal survival and maximal coverage problems demonstrated in the recasting indicates that maximal survival problems can be solved with the same algorithms and receive the same convenient guarantees from the heuristics used for maximal coverage problems.

4.2.3 Solution Approach

We solve the recast problem as a linear program using Gurobi. We also enumerate the solutions to the original formulation to ensure their equivalence. Since recasting the problem increases the size of the problem and MCPs have high time complexity to find the optimal result, we consider approximate approaches. The greedy algorithm works in polynomial time and guarantees a solution with an approximation ratio of $1 - \frac{1}{e}$ (i.e. the solution will be at least ≈ 0.632 as good as the optimal). Building on the recasting presented above, a greedy algorithm, shown in Algorithm 1, picks the facility location in each step that increases the population-weighted average survival rate the most.

```
Data: ED list, potential location list, patient location list, equipment number
Result: equipment location list
Create an empty equipment location list;
Initialise temp survival rate number as zero;
Initialise equipment location as null;
while equipment location list less than equipment number do
    Create a place holder to save temp equipment location
    for each potential equipment location in potential location list do
        Find all covered EDs for the equipment location;
        Find all covered patient locations for these EDs;
        Calculate population weighted average survival rate for all patient
          locations;
        if calculated survival rate larger than temp survival rate then
            Update the temp survival rate and equipment location;
        end
    end
    Add the temp equipment location to location list and remove it from
      potential location list;
end
```

Algorithm 1: Greedy Algorithm

The result of the greedy algorithm is at least locally optimal with a minimum approximation guarantee. However, the process of greedy algorithm is deterministic so that the result will never change unless the data or the algorithm itself changed.

In contrast, a heuristic such as a genetic algorithm introduces randomness to create possibilities to jump out of a local optimum and converge to another. It may take longer to get a result as good as that from the greedy algorithm but, it has potential to improve the result.

In this case, the gene is the vector of locations for the equipment/team and each gene produces a given survival rate. The crossover step in Algorithm 2 uses the times that a facility location appears in two parent combinations as the weight to sample the first generation gene of next generation combination without replacement, and

Data: Iteration, Population, Crossover rate, Mutation rate
Result: Z, r
Initialise sets of facility locations;
while *Optimal result changed or the maximum iteration is not met* **do**
  for *each facility location combination in population* **do**
    Calculate average survival rate of metropolitan area;
    Select combinations of locations with best survival rate;
    Crossover between successful combinations of locations;
    if *There are identical sets of locations in the pool* **then**
      Mutation (Changing some locations randomly);
    end
  end
  Assign best combinations of locations to be the new population;
  Update the best average survival rate;
  Increase iteration number by 1
end
$r \leftarrow$ best average survival rate;
$Z \leftarrow$ equipment locations that produce the best average survival rate

Algorithm 2: Genetic Algorithm

the rest as the gene of another next generation. In the mutation step, the activated gene will replace an allocated location with a randomly selected other location. In this study, we set the genetic algorithm crossover probability as 70% and mutation rate as 5% (Yang, 2020).

4.3 Data and Experiment

The methodology described above is tested for the case of the deployment of ECPR in Sydney, Australia. OHCA is a leading cause of death in adults who are otherwise healthy (Benjamin et al., 2019). OHCAs are a primary motivation for funding metropolitan ambulance services and a key aspect in the training of first responders. Despite the logistical efforts related to these health emergencies, OHCAs have typical survival rates less than 12% (NSW Health, 2021). ECPR provides heart and lung bypass while awaiting definitive treatment for OHCA (Lamhaut et al., 2017). Compared to conventional cardiopulmonary resuscitation (CCPR) treatment, ECPR is highly effective and less time sensitive (Bartos et al., 2020b). ECPR represents an instance of a resource-limited facility location optimisation suitable for the maximal survival problem presented above. Due to specialised skill requirements and the need for post-cannulation in-hospital intensive care, ECPR has been deployed in the most advanced hospitals (for example, in 5 of the 26 hospitals with emergency departments in Sydney, as of 2022). Therefore, these scarce facilities must be utilised efficiently to maximise the benefit of ECPR. There are existing implementations of of pre-hospital ECPR (Lamhaut et al., 2017; Singer et al., 2018; Song et al., 2022) and rendezvous ECPR (Bartos et al., 2020b; Yannopoulos et al., 2020b), so a mobilised team/equipment strategy is both feasible and tested. The concurrent rendezvous strategy described above is an accurate representation of emerging techniques in OHCA care. Building on observed survival rate functions from recent trials, we can compare the performance of rendezvous ECPR with the 12% baseline for OHCA.

The scope of the study is Sydney Greater Metropolitan Area with 4.8 million residents. Mesh blocks (MBs), the smallest census geographies defined in the Australia Statistical Geography Standard, are considered as the basic spatial unit (Australian Bureau of Statistics, 2017). i is the index of the MBs. Each location, including patients, emergency departments, and ambulance stations, is represented by the MB in which it is located. They are equivalent to census tracts in North America and there are 58,000 in Sydney. Patient coverage is based on resident location because 76% of OHCAs occur at a residence (Dyson, 2021). The number of ECPR teams is informed by the status quo in Sydney although other cities may differ.

4.3.1 Emergency Logistics Process

The baseline process for cardiac arrest treatment with ECPR in Sydney is shown in Figure 3.1a. The arrest is assumed to be nearly concurrent with the call to the emergency services dispatcher. The ambulance will travel from an ambulance base or location in the field to the scene of cardiac arrest. Medical personnel will provide on-scene treatment. Finally, the ambulance will transfer the patient to hospital for ECPR. Probability of survival decreases with the interval between arrest and the start of flow for the heart and lung bypass.

Current practice imposes a one hour time threshold on eligibility for ECPR (Randhawa et al., 2022). The one hour includes the ambulance response time, on-scene treatment time, transfer time to move the patient to hospital and cannulation time upon arrival.

Response times are based on historical cardiac arrest based on SA2 level (Statistical Area Level 2 defined by Australian Bureau of Statistics) (Dyson, 2021). The on-scene treatment time is assumed as 20 min for CCPR and 27 min for ECPR. We use the average travel time between each MB pair in November 2019, which are derived from Compass IoT's Connected Vehicle Data (Compass, 2019). Ambulance station and emergency department location information is from Australian Institute of Health and Welfare (Australian Institute of Health and Welfare, 2022). ECMO cannulation time is assumed to be 15 min.

When rendezvous strategy is implemented, the ECPR teams start at the same time as the ambulances and drive directly to the meeting point as shown in Figure 3.1b. The travel time from patient to ECMO is shorter than the baseline because the ECMO team will rendezvous closer to the patient than the in-hospital ECMO facility.

4.3.2 Accessibility Function

In the coverage model, the accessibility to ECMO at each MB is weighted by its population and a MB only contributes to the overall accessibility if it is covered by at least one ECMO team within the time threshold. The patient population n_i for each MB is derived from the corresponding resident counts provided by the Australian Bureau of Statistics (Australian Bureau of Statistics, 2017). For the maximal survival model, time-dependent survival rate will be used as the impedance function. Here we still use the logistic survival rate from Figure 3.2.

Variation in the time to restored flow is determined by differences in response, on-scene and travel times. Notably, the advantage of the rendezvous delivery strategy depends on a shorter travel time to the closest ED than to the closest ECPR treatment hospital. For long travel times, the survival rate drops significantly. In this situation, the CCPR treatment with a short response time may provide a higher chance of survival—therefore CCPR will be practiced for patients awaiting ECPR and in remote locations where ECPR is infeasible.

4.3.3 Solution Method

MCP is an NP-hard problem because the number of possible scenarios will increase drastically when the number of available positions increases. Sydney has approximately 58,000 MBs that each serve as a potential location. Even when we set a fixed number, k, for the allocation of ECMO teams or facilities, the problem still exhibits a high-order polynomial complexity, with a power factor of k. If we only consider EDs as potential ECPR team locations, there are only 26 MBs potential locations to consider. When the ECMO equipment number is 5, there are 26 choose 5 ($C26, 5 = 65, 780$) combinations to test. At this scale, we can enumerate all the cases to find the global optimal result and confirm it with a greedy algorithm. However, if ambulance bases (ABs) are considered as potential facility locations, the number of possible locations increases by 46. The computing time will be proportional to 72 choose 5 ($C72, 5 = 13, 991, 544$) which is 200 times longer. In the most general version of the problem, the ECMO team could be located at any of the 58,000 MBs. This motivates the comparison between the solution to the recast problem using the linear programming solver Gurobi and the results from the greedy and genetic algorithms.

4.4 Result and Analysis

In this section, we will first explore the improvement of ECPR coverage based on the case study with the data introduced in Section 4.3 in order to illustrate the benefit of the rendezvous delivery strategies for predetermined facility locations. Next, we probe the potential of the strategy by searching for the optimal location combinations from 26 EDs. It is expected that the optimal allocation differs under alternative delivery strategies. We incorporate the survival function into the accessibility model creating a gradual coverage or maximal survival problem that reduces the number of solutions with identical and optimal coverage. Finally, 46 extra locations are considered as potential ECMO locations. For this more computationally expensive case, we compare the performance of greedy and genetic algorithms with the benchmark optimum found by solving the recast problem. This will give a performance reference for even larger problems.

4.4.1 Basic Comparison with Rendezvous Strategy

Figure 4.3 shows the non-rendezvous coverage results associated with 5 ECPR hospitals and a 60 min travel time threshold. The current emergency logistics strategy covers 811,091 people within this threshold. The coverage catchment is shown in Figure 4.3a. When the rendezvous strategy is applied, the mobile ECPR team could drive to intermediate EDs to meet an ambulance carrying a cardiac arrest patient. An additional 18 hospitals are involved in ECPR service system under the rendezvous strategy, but no additional ECPR teams or equipment. The coverage area extends noticeably as Figure 4.3b shows, encompassing 2,175,096 people within the threshold.

4.4.2 Allocation Improvement

There are 26 locations that host an emergency department. All of them could be engaged as ECPR team locations. If the number of ECMO equipment does not change, 65,780 combinations of 5 locations are feasible. Some combinations are likely to be better than the current combination.

With the baseline non-rendezvous ECPR strategy, the optimal combination of hospitals shown in Figure 4.4a could cover 929,753 people, which is 118,662 greater than the status quo combination. The optimal coverage with rendezvous strategy grows to 2,247,861. All of the 26 EDs are engaged as an ECPR station or rendezvous

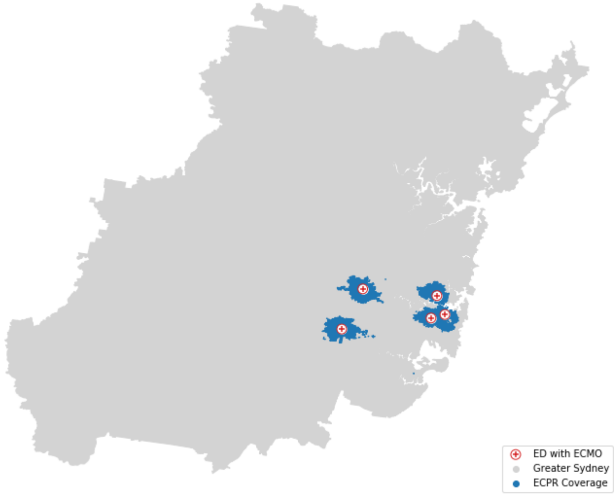

(a) Non-rendezvous baseline

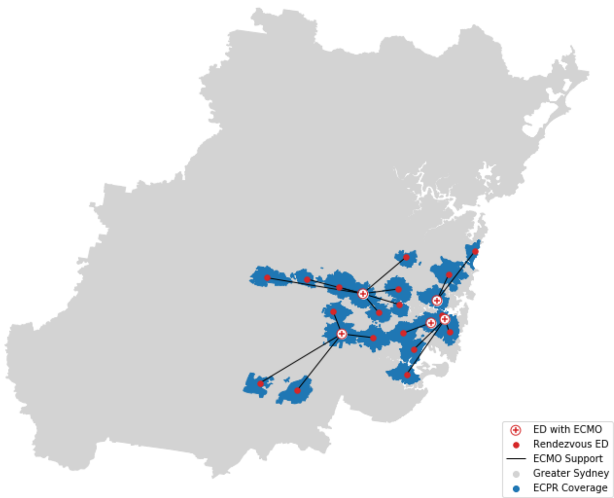

(b) Rendezvous Strategy

Figure 4.3: Coverage provided by 5 ECMO teams when we assume the patient must travel to the ECMO facility (top) and when the patient and equipment can rendezvous at an intermediate location. Without increasing the number of resources, the rendezvous strategy covers a much larger area.

treatment point. However, 1,521 of the 65,780 unique combinations of ECPR locations result in the same coverage because all of them can serve all EDs. Figure 4.4b and 4.4c are two of the 1,521 combinations. This case demonstrates the lack of a unique optimal combination of locations in the maximum coverage problem formulation. The non-unique solutions offer an advantage in this instance because there is flexibility to consider external constraints and secondary objectives.

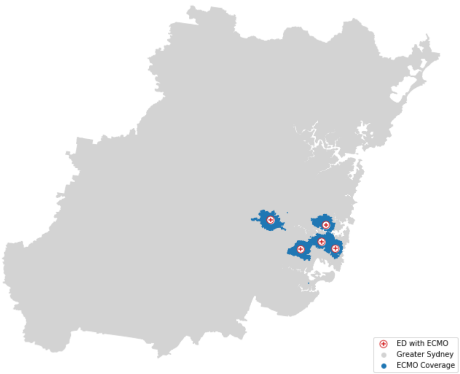

(a) The best five hospitals when rendezvous is not used

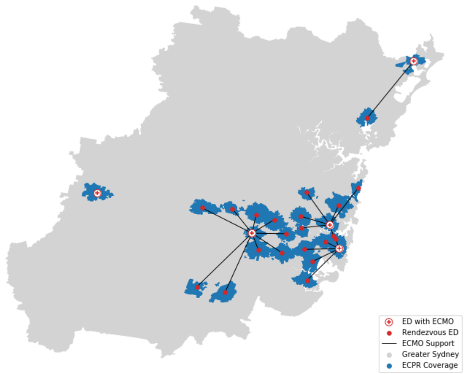

(b) One set of ECPR locations with the optimal rendezvous coverage

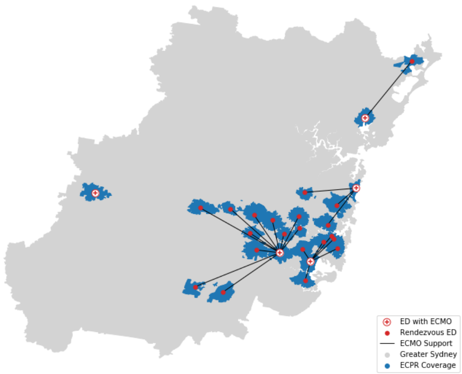

(c) An alternative set of ECPR locations with the same rendezvous coverage as above

Figure 4.4: Optimal ECMO Position Combinations

4.4.3 Optimal Result with Survival Rate Function

The coverage model cannot distinguish between combinations of locations that serve all EDs, but we know that patient outcomes are better when travel times are shorter. Survival rate is introduced as an impedance function so that the quality of the coverage decreases as travel time increases. This improvement finds the best survival rate, which minimises population-weighted travel times across the region.

Figure 4.5 shows the basic survival rate for each location when there is no ECPR at all– survival is higher in high population density areas with many ambulance bases. Between January 2017 and June 2021, there are 914 emergency responses to cardiac arrests that would have been eligible for ECPR in Greater Sydney. Based on the calculated survival rate, the expected survival from non-ECMO CPR is 121 patients in 4.5 years.

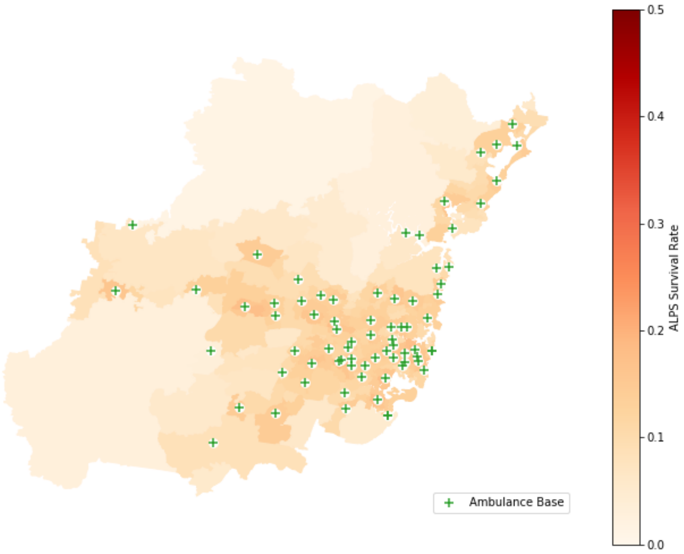

Figure 4.5: Basic survival rate with CCPR. Amiodarone, Lidocaine or Placebo study (ALPS) is used for conventional cardio-pulmonary resuscitation (CCPR) (Kudenchuk et al., 2017)

Building from non-ECMO CPR, Figure 4.6a maps survival resulting from the five existing ECMO hospitals. Without rendezvous, 151 patients are expected to survive by being transported to an ECPR-capable hospital. If the rendezvous strategy is implemented, we observe even better survival (205 expected survivals) for two reasons. First, there are locations that were previously outside ECMO coverage that now have access to ECPR, and second, there are locations that experience higher survival with rendezvous than they did without rendezvous. Conventional CPR and non-rendezvous ECPR are shown in red and blue respectively in Figure 4.6b. Green areas show improved survival when using the rendezvous strategy. The purple regions can be served by rendezvous ECMO treatment but were not covered without rendezvous service.

If 5 ECMO teams are allocated to the best EDs in the rendezvous strategy, the survival rate of many locations increases as shown in Figure 4.7 and 208 patients are expected to survive. There are 5 combinations that yield the same survival rate— 4.7a and 4.7b are 2 of them. Like the maximal coverage model, the maximal survival model does not produce a unique result, but the objective function better represents the true goal of the service and the number of equivalent solutions is small enough to be assessed with respect to other attributes by a health system planner.

4.4.4 Increased number possible locations

In addition to the 26 MBs containing an ED, there are 46 additional locations housing an ambulance station. Although ambulance bases are unsuitable for rendezvous

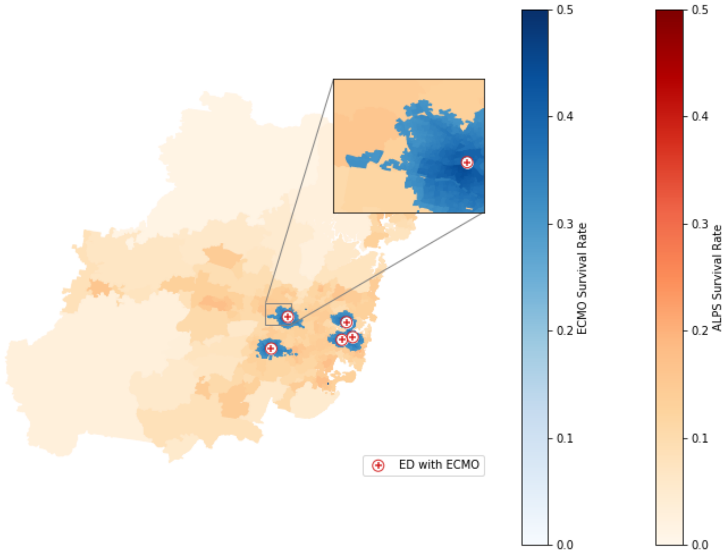

(a) ECPR without Rendezvous Strategy

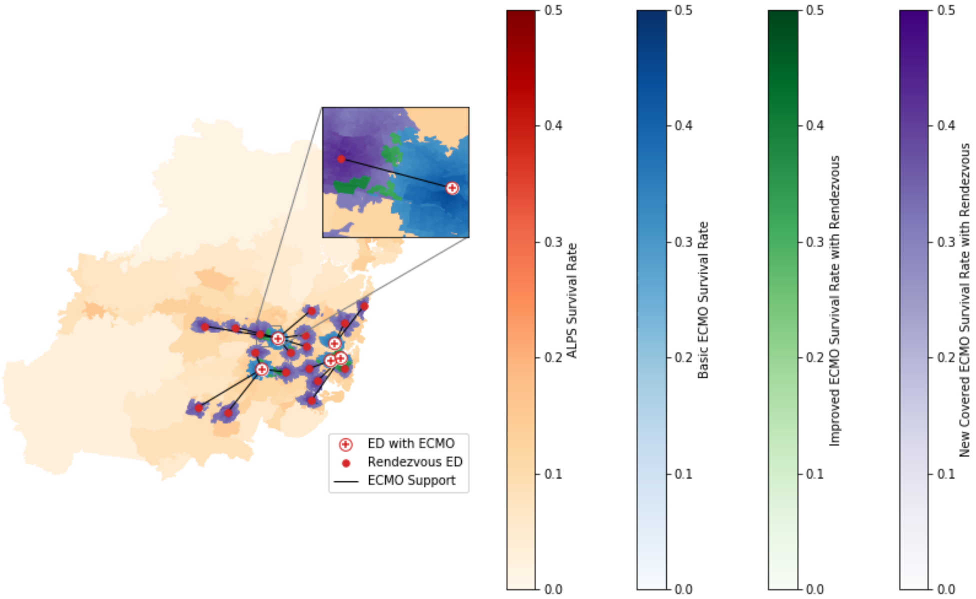

(b) Rendezvous Strategy

Figure 4.6: Basic Combinations with Survival Function

and treatment, they are potential locations to station the ECMO equipment delivery team so as to reduce response time. However, with the number of locations increasing, the number of possible combinations grows to more than ten million. Hence a suitable algorithm may help us to find a tolerable result in a reasonable time.

Enumerating the solutions yields 191 combinations that give the same optimal result. Figure 4.8 is one of them, and we can see that four of the ECMO positions are allocated to ambulance bases rather than emergency departments. This will shorten the travel time of ECPR dispatching in this area. Although the equipment delivery time is reduced, the survival only shows a small increase (0.059%) compared to

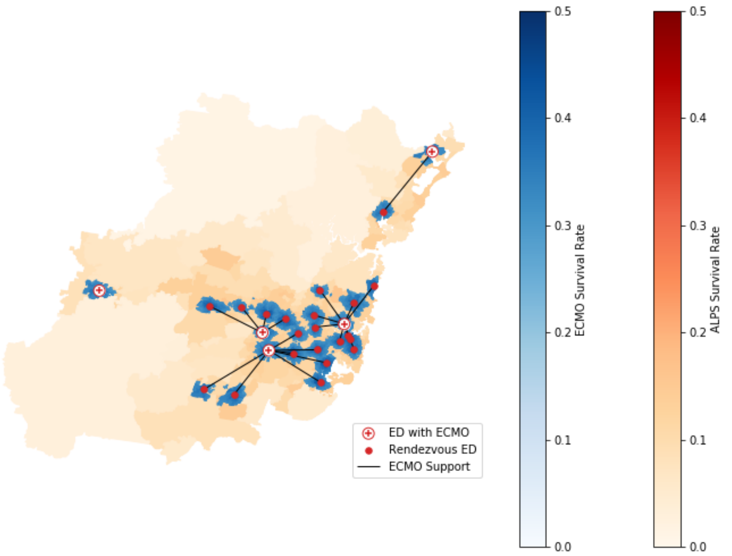

(a) Rendezvous Strategy 1

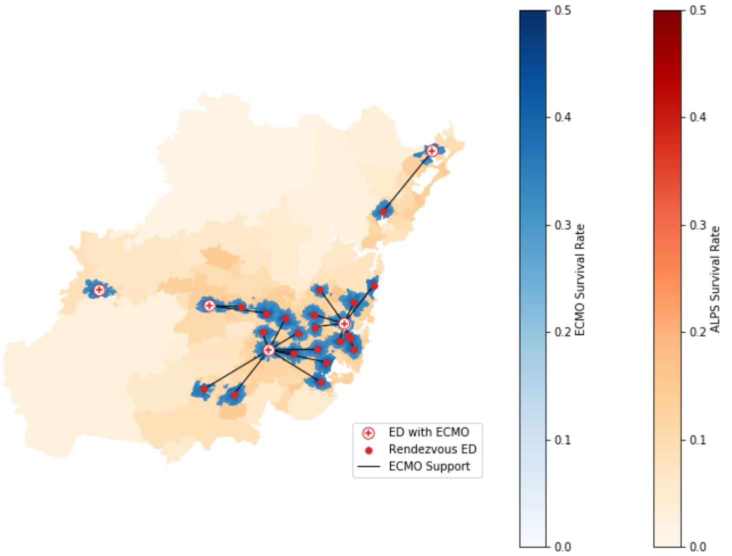

(b) Rendezvous Strategy 2

Figure 4.7: Optimal ECMO Position Combinations with Survival Function

constraining the ECMO teams to wait at emergency departments.

Comparing to the optimal solution from the recast problem as well as the enumerated solutions, we assess the performance of heuristic methods.

The solutions shown in Figure 4.9a are very close to the benchmark for the expected survival, differing by less than 0.001%. However, the ECPR locations are different. Due to the probabilistic aspects of the genetic algorithm, we run the genetic algorithm 5 times with different initialization and random numbers to avoid small probability events. The genetic algorithm could find one of the optimal results in 4 out of 5 tests. The remaining result is also better than the solution we get from the greedy algorithm.

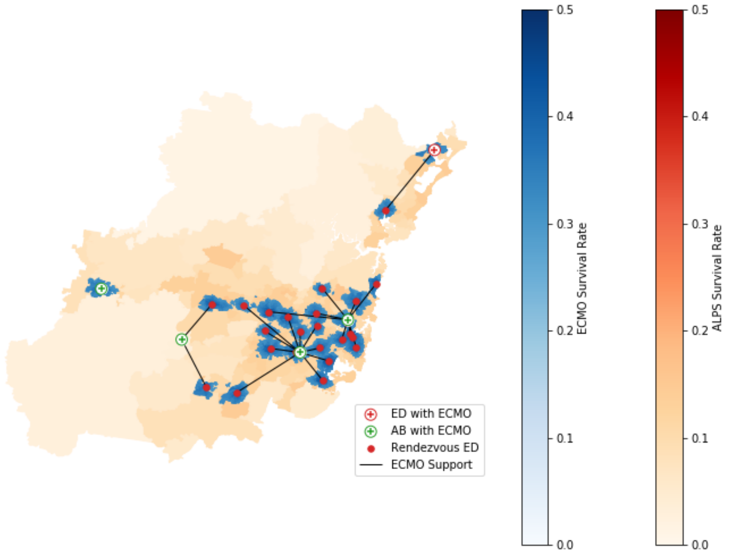

Figure 4.8: One Optimal ECMO Position Combination with Ambulance Bases

4.5 Discussion and Managerial Implications

For both the maximal coverage model and the maximal survival model, the rendezvous strategy can significantly increase the objective, which could be seen from Figure 4.4 and Figure 4.7.

However, the improvement is negligible when considering locating the teams at ambulance stations. The main reason is because rendezvous time is constrained by the patient transfer team rather than the ECPR mobile team. Even when the ECPR teams are allocated at ambulance bases which reduce travel time, the patient transfer ambulance causes the ECPR team to wait for the rendezvous moment. The response time plus on-scene treatment time is generally at least 32 min, which is longer than the ECPR team's response time in most cases. In addition, the ambulance bases cannot provide a suitable treatment environment, so both the patient and ECPR team have to travel to an ED even when the emergency event happens right at the site of an AB.

The survival rates associated with the rendezvous strategy locating ECPR teams only at EDs is close to the optimal results of those with ambulance. However, the best ED-only solution is the 1,515th result in the order of all the 13,991,544 (72C5) combinations when we allow ambulance bases. Both greedy and genetic algorithms can obtain near optimal results. Although genetic algorithm requires more computing time, it always returns results closer to the optimal survival rate than those found by the greedy algorithm and, in most tests, the genetic algorithm solution equals the optimal results. Considering the number of equivalent solutions and trading-off the speed of the heuristics against minor loses in performance, it is

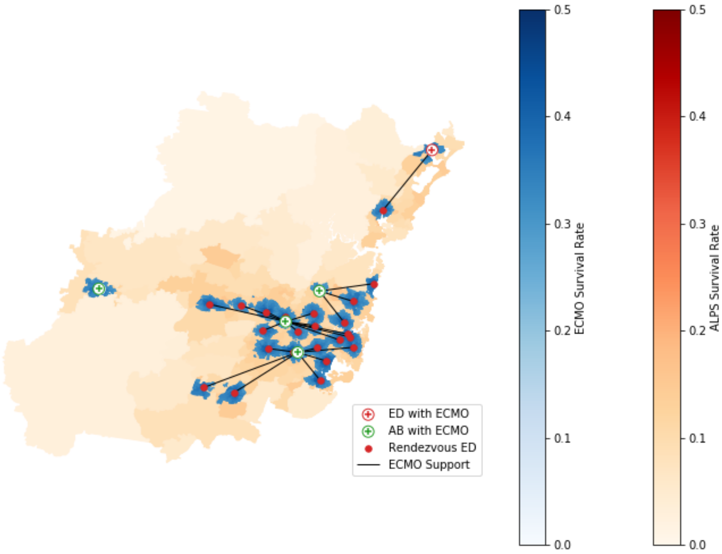

(a) Greedy Algorithm

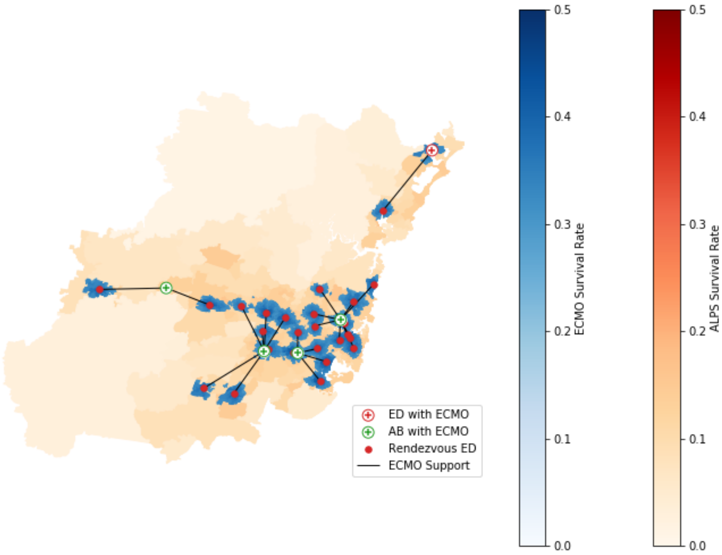

(b) Genetic Algorithm

Figure 4.9: Algorithm Results showing solutions from the heuristic methods. The genetic algorithm outperforms the greedy algorithm and both are within 0.000337% of the optimal result.

clear that decision makers have flexibility to select a (near-)optimal solution that might also address constraints related to budgeting, space or equity as needed.

Chapter 5

Impact of Vehicle Types on EMS Output

5.1 Introduction

The previous chapters' spatial modeling reveals how delivery strategy and facility allocation can heighten planning-level coverage for emergency medical services (EMS). Yet, even with well-placed resources, operational decisions—such as which vehicles respond to which calls—often determine how effectively these resources are used in practice. Specialized ambulances equipped for advanced procedures may boost survival for severe cases, but relying on them too heavily for routine incidents risks diluting their availability when most needed. Conversely, keeping them idle until a perfect high-acuity match may reduce their utilization and leave standard ambulances overwhelmed.

In this chapter, various dispatching policies are developed to enhance operational management. In the context of real-time dispatching, achieving effective outcomes is further complicated by subsequent tasks, as each vehicle can respond to only one incident at a time, and every dispatch decision directly impacts future operations.

For time-varying emergency response systems, agent-based models (ABMs) have emerged as capable tools for simulating and optimising these systems. ABMs have been applied in various domains such as freight transportation, ridesharing, taxi, and delivery (Chen and Chankov, 2017; Di Febbraro et al., 2016; Hörl et al., 2021; Nourinejad and Roorda, 2016). In the context of EMS, ABMs enable the detailed representation of individual agents—such as ambulances, patients, and emergency departments—and their interactions within a dynamic environment (Koch et al., 2020; Laskowski and Mukhi, 2009). This approach is particularly effective in scenarios where environmental factors, such as severe weather or infrastructural disruptions, significantly impact response times. For instance, Wang et al. (2012) explored the use of ABMs in dispatching systems to simulate emergency response under various conditions, highlighting the flexibility and robustness of these models in complex situations.

Agent-based simulation is adopted in this study as it enables the representation of decentralised interactions between vehicles and incidents (Amakama et al., 2025). This approach facilitates the modelling of real time vehicle assignment in multi priority environments and captures the temporal dynamics of participant status (Lam et al., 2024). The characteristics of emergency medical services systems, including sparse demand, high spatial and temporal variability, and the critical requirement for timely response, make mathematical programming methods impractical for real time operational decision making. The stochastic and sparse nature of EMS incidents means that outcomes are often influenced by complex interactions among vehicles, patients, and dispatch decisions (Edjinedja et al., 2024). By simulating individual agents and applying repeated trials, this approach enables the modelling of sparse and noisy events while capturing emergent behaviours that affect survival outcomes.

Dispatching policies determine the rules about allocating emergency vehicles to incidents in the EMS system. Under the ABM framework, both vehicles and incidents are agents defined by a set of relevant attributes. To realistically represent operational conditions, it may be important to consider types of incidents and types of vehicles and the way they interact with each other to produce patient outcomes. In practice, vehicle type represents medical capability, and incident type corresponds to triage categories defining medical need. Incidents of the same type share the same priority, while different types may have the same or different priorities depending on urgency. Vehicle performance also varies across incident types, reflecting the suitability of specific resources for different medical conditions. In this context, the optimal allocation of different types of emergency vehicles can be modelled in an ABM to assess their impact on survival rates. Prior research has explored the use of ABMs to evaluate dispatch policies that address both routine and specialised incidents, offering insights into how EMS systems can be optimised under diverse conditions (Koch et al., 2020; Wang et al., 2012).

Tikani and Setak (2019) explored the challenges associated with multiple patient types and vehicle types, considering the added value of specialised vehicles for specific patients. However, they noted the potential inefficiencies when these specialised vehicles must occasionally respond to general patients due to resource constraints. This dilemma is particularly relevant in cases like the alternative cardiac arrest strategy discussed by Song et al. (2022), where specialised vehicles are critical for treating a small patient population to achieve significant survival benefits.

In this work, we address the research gap in optimizing ambulance dispatch when vehicles are specialized for specific subsets of incidents. This chapter considers two types of pre-hospital vehicles: the first is the conventional cardiopulmonary resuscitation (CCPR) vehicle, which provides standard treatment for all patient types. The second is the extracorporeal membrane oxygenation (ECMO) vehicle, utilized during cardiopulmonary resuscitation (ECPR), offering significant benefits for cardiac arrest incidents. We present an agent-based simulation model representing the emergency medical system and apply it to evaluating vehicle dispatching policies within a mixed fleet of ambulances. The study incorporates incidents of differing types and priorities and illustrates how agent-based modelling can provide insights into the optimal allocation of EMS resources. The methodology is explained in Section 5.2, followed by an experiment with ambulance allocation policies in New South Wales, Australia, presented in Section 5.3. The insights derived from this case study are discussed in Section 5.4.

5.2 Methodology

This section outlines the methodology employed to analyse and optimise emergency medical services (EMS) dispatch strategies through an agent-based simulation model. The following lists the notation that will be used in the chapter.

Symbol	Description
A_i	Agent for incident i, with attributes $(h_i, e_i, o_i, l_i, d_i, z_i)$
A_v	Agent for vehicle v, with attributes $(y_v, x_v, l_v, u_v, z_v)$
$a_{v,t}$	Indicator of vehicle v's availability at time t
$b_{i,t}$	Indicator of incident i's waiting status at time t
$c_{v,i,t}$	Response capacity notation for vehicle v responding to incident i at time t
d_i	Destination for the incident represented by agent A_i within the network; this attribute influences both the vehicle travel time from the location l_i and the final or subsequent starting position of the vehicle for future actions.
e_i	Type of the incident represented by agent A_i
h_i	Time point when the incident represented by agent A_i occurs
I	Set of all incidents, with i representing a specific incident

Symbol	Description
i	A specific incident in the set I
l_i	Location of the incident represented by agent A_i within the network
l_v	Current location of the vehicle represented by agent A_v within the network
o_i	Priority level of the incident represented by agent A_i
P	Set of dispatch policies
$p_{y,e}$	Policy value for vehicle type y and incident type e
$q_t(A_i)$	Position of agent A_i in the queue Q_t at time t
Q_t	Queue of incidents at time t, sorted by priority o_i and time of occurrence h_i
$S_{y,e}$	Survival function for incident type e treated by vehicle type y
T	Timeline, with t representing a specific time point
t	A specific time point in the timeline T
u_v	Route taken by the vehicle represented by agent A_v when responding to an incident
V	Set of all vehicles, with v representing a specific vehicle
v	A specific vehicle in the set V
x_v	Station location of the vehicle represented by agent A_v
y_v	Type of the vehicle represented by agent A_v
z_i	Status of the incident represented by agent A_i during processing
z_v	Status of the vehicle represented by agent A_v during its response to an incident
$\tau_{l_v,l_i,t}$	Travel time from vehicle v to incident i at time t
$\hat{\tau}_{l_v,l_i,t}$	The total expected response time from vehicle v to incident i at time t

5.2.1 Agents for Incidents and Vehicles

The core of the methodology is an agent-based simulation model that represents the interactions between EMS vehicles and incidents. Each agent in the simulation corresponds to an EMS vehicle or an incident, with specific attributes and behaviour rules for incidents and vehicles.

This optimisation identifies the optimal set of vehicle dispatching policies considering multiple types of emergency vehicles responding to multiple types of incidents. The dispatching system will assign vehicles to incidents based on the dynamic simulation information. For the simulation process, we denote T as the whole timeline and $t \in T$ as a time point.

Incident Agents

I is the set of total incidents. Each incident is represented by $i \in I$. We denote A_i as the agent for incident i. The incident agent is defined by an attribute tuple $A_i \equiv (h_i, e_i, o_i, l_i, d_i, z_i)$. e_i is the incident's type, while o_i is the incident's priority level. The time when incident happens is denoted h_i, l_i is the incident's location and d_i is the incident's destination in the network. These attributes are time-invariant for each incident. The incident's status, z_i, will evolve along the timeline.

At the beginning of each time step t, each agent, A_i, will progress according to its status at t. If $t = h_i$, the incident will be added in the system queue ($z_i =$ Waiting) based on its priority o_i. Being in the queue makes the incident eligible for a vehicle allocation. If the incident is allocated a vehicle by the dispatching system, its status z_i matches the responding vehicle's status until the patient is handed over at the destination hospital, d_i. The final status of the patient will be $z_i =$ Finished. Figure 5.1 shows the dynamics of the incident agent.

Vehicle Agent

Let V denote the set of vehicles. Each vehicle $v \in V$ has an agent A_v defined by an attribute tuple $A_v \equiv (y_v, x_v, l_v, u_v, z_v)$. The attributes include y_v for the vehicle's type and x_v for the vehicle's station location. These attributes are time-invariant for each vehicle. Additionally, the vehicle's location, l_v, route, u_v, and status, z_v will evolve along the timeline.

All vehicles will initially sit at their stations with the status Available. When a vehicle is allocated to a patient by the dispatching system, it will change its status to Response, find a route u_v from l_v to l_i and update its location l_v at each time point t. When the vehicle reaches the patient, the status changes to On-Scene. After the on-scene treatment, it will change its status to Transport and find a route u_v from l_i to d_i. When the vehicle arrives at the destination hospital, the status changes to Destination and begins the process of off-loading the patient, cleaning and restocking the vehicle. After finishing the incident, the status updates to Available, and it will find a route u_v to go back to its station locationx_v from the hospital.The vehicle can be allocated to another patient before it reaches its station. Figure 5.2 shows the dynamics of the vehicle agent.

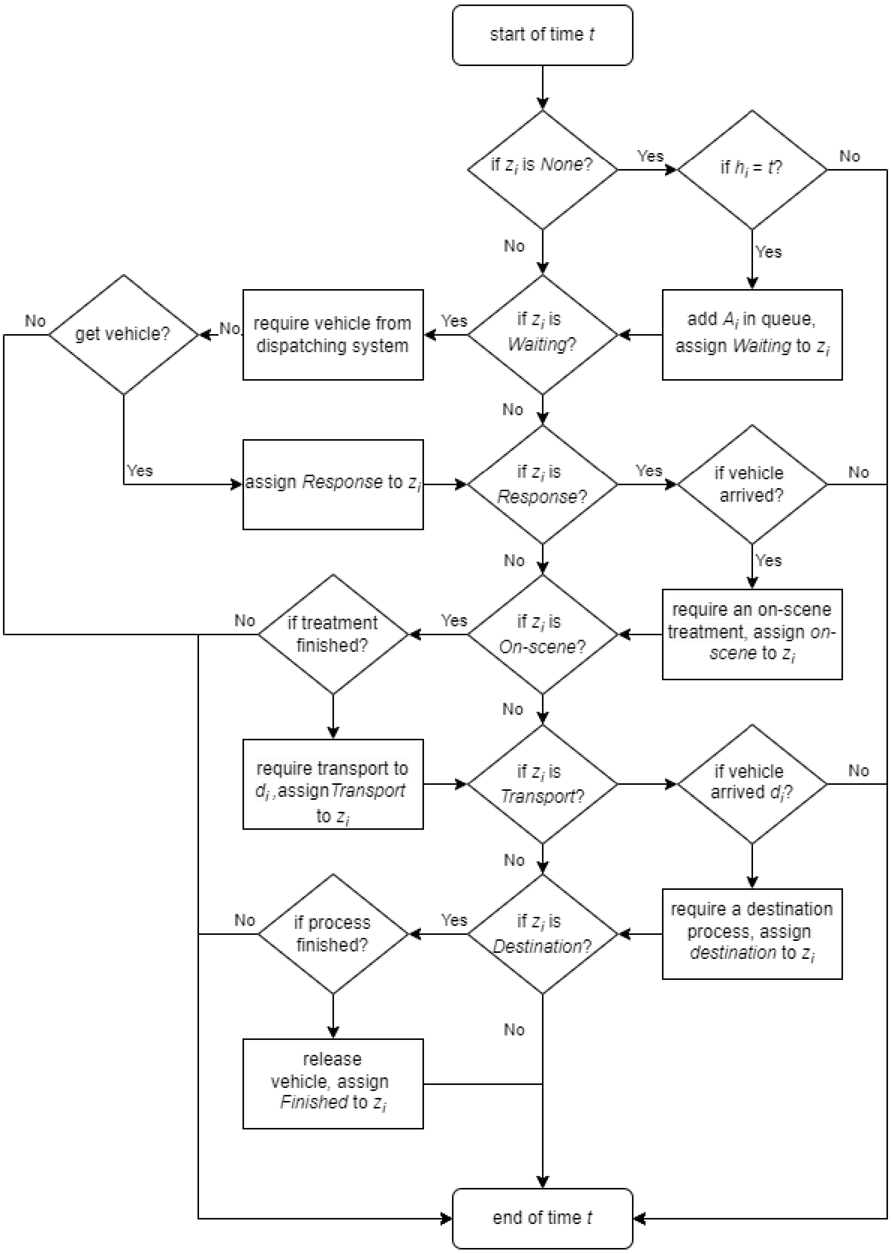

Figure 5.1: Flowchart of Incident Agent (A_i) Dynamics

5.2.2 Dispatching Policy

The interactions between vehicles and incidents are governed by response rules. The dispatching system determines which vehicles are allocated to incidents based on the vehicle and incident attributes and the current conditions.

The dispatching policy, p, selected from the policy set, P defines the eligibility of vehicle type y to respond to incident type e. Each vehicle type-incident time combination has a unique policy where the policy values $p_{y,e}$ follow Eq. 5.1. Y and E are the type sets for vehicles and incidents.

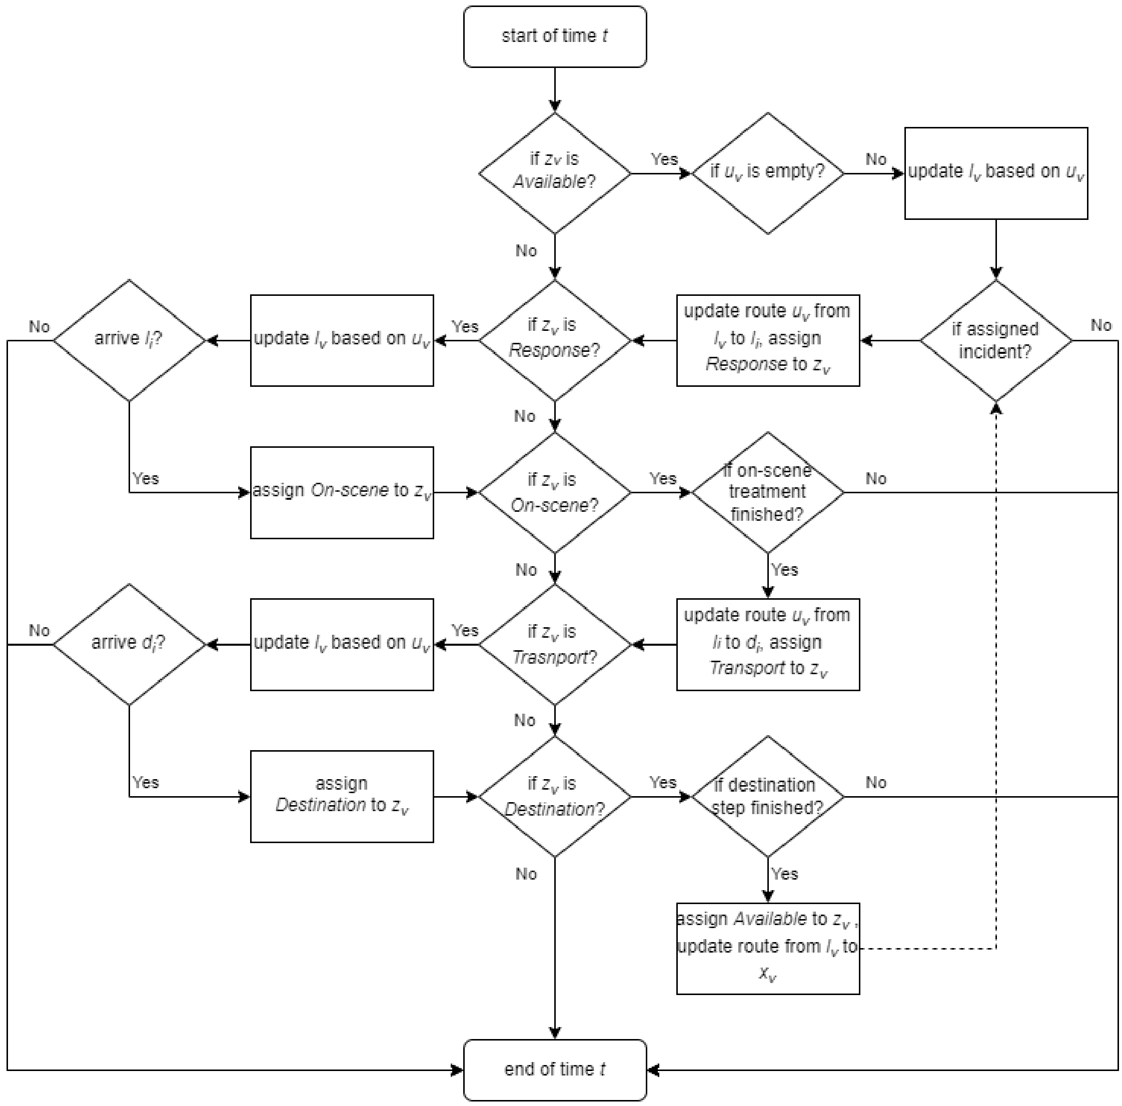

Figure 5.2: Flowchart of Vehicle Agent (A_v) Dynamics

$$p_{y,e} = \begin{cases} 1 & \text{if vehicle type } y \text{ is allowed to respond to incident type } e \\ 0 & \text{else} \end{cases} ; \forall y \in Y, \forall e \in E \tag{5.1}$$

The performance of a vehicle of type, y when responding to an incident of type e is quantified as the survival function, $S_{y,e}$, whereby patient outcomes diminish with response time. The matrix of survival functions has dimensions of all vehicle types $|Y|$ by all incident types $|E|$ as the Table 5.2 shown. Note that under some policies, some elements of the survival matrix are irrelevant because the policy does not allow that type of vehicle to respond to that type of incident.

Vehicle Type y	Incident Type e			
	1	2	3	...
1	$S_{1,1}$	$S_{1,2}$	$S_{1,3}$	...
2	$S_{2,1}$	$S_{2,2}$	$S_{2,3}$	...
3	$S_{3,1}$	$S_{3,2}$	$S_{3,3}$	...
...	...	...	...	...

Table 5.2: Survival Function Table

5.2.3 Simulation Framework

The simulation progresses in discrete time steps, with each step representing a unit of time within the modeled environment. At each time step, the framework will first update the status for each incident and vehicle. Then processes the occurrence of new incidents, allocate available vehicles to waiting incidents according to the queue and update the status for each incident and vehicle. The incident and vehicle numbers are stochastic in the simulation so the number of agents active in each timestep varies.

A queue is used to prioritise incidents that are waiting for an available vehicle. We use $a_{v,t}$ and $b_{i,t}$ to indicate the availability of vehicle v and the waiting status of an incident i at time t respectively. Their values are determined as described in Eq. 5.2 and 5.3.

$$a_{v,t} = \begin{cases} 1 & \text{if vehicle } z_v = \text{`available' at time } t \\ 0 & \text{else} \end{cases} ;\ \forall v \in V,\ \forall t \in T \tag{5.2}$$

$$b_{i,t} = \begin{cases} 1 & \text{if incident } z_i = \text{`waiting' at time } t \\ 0 & \text{else} \end{cases} ;\ \forall i \in I,\ \forall t \in T \tag{5.3}$$

The dispatch order follows a priority queue. To facilitate prioritised dispatch, all incidents join the queue, Q, and Q_t is the queue at time t. The queue sorts the incidents according to the priority o_i and then by the incident occurrence time h_i. Let $q_t(A_i)$ presents the position of A_i in the queue at time t. Hence, the queue could be formulated as Eq. 5.4. The order of the queue is presented as Eq. 5.5 and Eq. 5.6 sorted first on incident priority and then on happening time.

$$Q_t = \{A_i |\ i \in I \text{ and } b_{i,t} \geq 1\} \tag{5.4}$$

$$o_m \leq o_n; \forall A_m, A_n \in Q_t \text{ and } q_t(A_m) < q_t(A_n) \tag{5.5}$$

$$h_m < h_n; \forall A_m, A_n \in Q_t \text{ and } q_t(A_m) \leq q_t(A_n) \text{ and } o_m = o_n \tag{5.6}$$

Figure 5.3 demonstrates the inputs, outputs and the process of the simulation. The simulation generates a range of output metrics that provide insights into the performance of the EMS system under different dispatching policies. The key metric used for assessing the optimal dispatch policy is patient survival rate, as discussed in the next section.

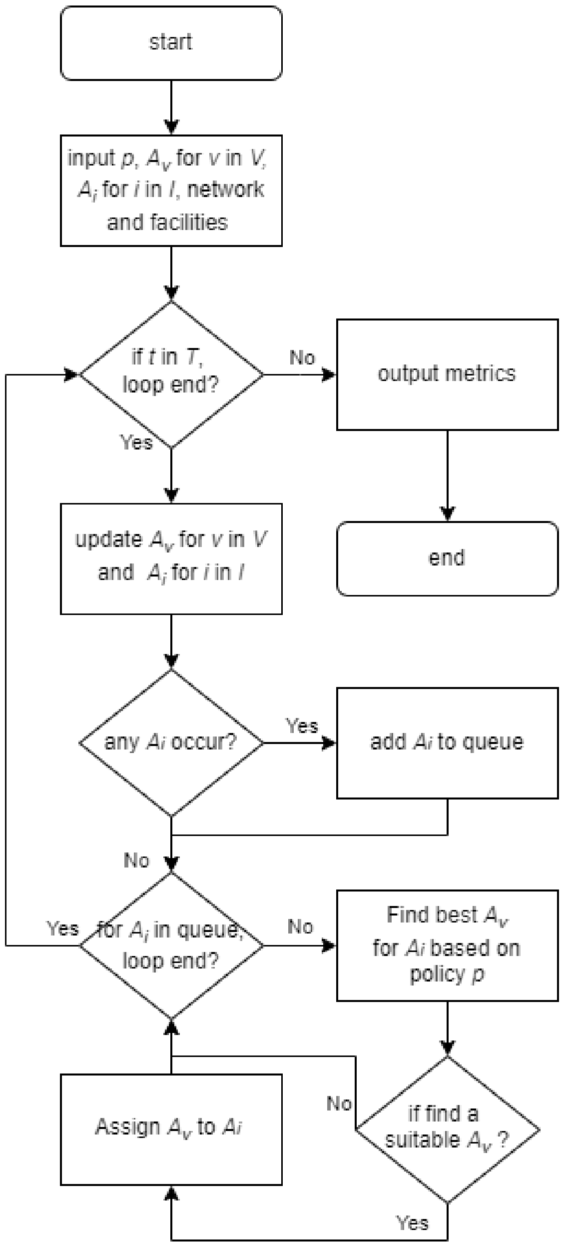

Figure 5.3: Flowchart of Simulation

5.2.4 Optimisation Model

The primary objective of this study is to optimise system performance, particularly patient outcomes, by selecting the best dispatching policies for EMS vehicles. The optimisation problem is built around an agent-based simulation framework that models the dynamic behaviour of incidents and vehicles across a transportation network.

$\tau_{l_v,l_i,t}$ is denoted as the travel time for the vehicle at location l_v to respond to the incident at location l_i at time t. This response time is a critical factor influencing patient survival rates, which are modelled as a function of the response time, vehicle type, and incident type. $\hat{\tau}_{l_v,l_i,t}$ represents the total expected response time to incident A_i.

$$\hat{\tau}_{l_v,l_i,t} = \tau_{l_v,l_i,t} + t - h_i \tag{5.7}$$

The vehicle can be dispatched to at most one incident when their status z_v is Available. Similarly, patients can be allocated at most one ambulance if their status, z_i is Waiting. These limitations are captured in the response capacity notation $c_{v,i,t}$ which says that the maximum number of allocations at time t for each vehicle is less than or equal to that vehicles availability, $a_{v,t}$, and the maximum number of allocations at time t for each incident is less than or equal to the incident's waiting status, $b_{b,t}$. This response capacity notation is used to formulate constraints in Eq. 5.10 and Eq. 5.9.

The optimisation objective is to maximise the overall survival rate, which is formulated as:

$$\max_{p \in P} \sum_{v \in V} \sum_{i \in I} \sum_{t \in T} s_{(y_v,e_i)}(\hat{\tau}_{l_v,l_i,t}) \cdot c_{v,i,t} \cdot p_{y_v,e_i} \tag{5.8}$$

subject to Eq. 5.2, 5.3 and 5.7

$$\sum_{v \in V} c_{v,i,t} \leq b_{i,t}; \forall t \in T \tag{5.9}$$

$$\sum_{i \in I} c_{v,i,t} \leq a_{v,t}; \forall t \in T \tag{5.10}$$

$$A_i = \text{simulation}(A_i, t); \forall i \in I, \ \forall t \in T \tag{5.11}$$

$$A_v = \text{simulation}(A_v, t); \forall v \in V, \ \forall t \in T \tag{5.12}$$

The objective function is the summation over all vehicles, incidents and times of the survival function multiplied by the vehicle capacity and dispatch policy. It is constrained by dynamic vehicle availability, patient occurrence and response time.

5.3 Data and Experiment

Using the agent-based framework described above, the temporal and spatial distribution of emergency incident is simulated subject to a range of vehicle dispatch policies. The progress of the stochastic incidents determines the survival outcome of each simulation. The methodology described above is tested for the case of two types of vehicles serving three types of incidents (two high-priority and one low-priority). This experiment is inspired by a real example of a specialised ECPR vehicle with treatment advantages for certain types of cardiac arrest. While the normal ambulances give basic first aid, the ECPR vehicle can provide advanced treatment for cardiac arrest patients. The scope of the study is the central and eastern parts of Sydney, Australia using historical incident generation rates and travel times.

We assume that the vehicles in each type and the incidents in each type are indistinguishable. If the policy allows, the vehicle will respond to high priority level incidents first. Within each priority level, dispatch follows a first-in-first-out principle.

5.3.1 Traffic Network

The simulation network crops the inner Sydney and Sydney eastern suburbs areas from OpenStreetMap (OpenStreetMap contributors, 2017). Fig. 5.4 shows the geographical scope and the simulation network. The travel time on each link is an average travel time from one week of aggregated connected vehicle data from November 2019 (Compass, 2019). The region outlined by the yellow boundary in Fig. 5.4 (a) is selected as the study area in this chapter. This choice reduces computational complexity compared to the full Greater Sydney area, as the simulation here is more resource-intensive than in previous chapters. The region is bordered by water on three sides, which helps reduce cross-boundary flows and reflects real EMS operations.

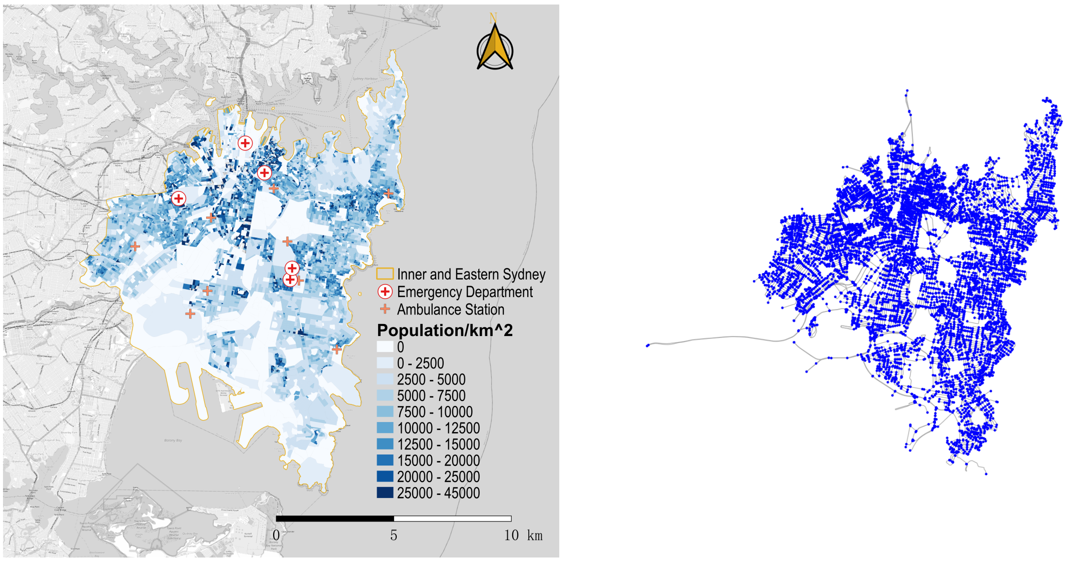

(a) Population density and relevant facilities

(b) Road network

Figure 5.4: Population Density, Facilities and Road Network. The area encloses Sydney's densest areas and is served by 5 emergency departments, including 3 of Sydney's 7 major trauma centres. The road network includes highly congested areas as well as sparse regions around the airport and major parks.

5.3.2 Incident Distribution

Incidents are distributed across on the nodes of the network proportionally to the population density as Figure 5.4 shown. The population is from the Australian census (Australian Bureau of Statistics, 2017). Mesh blocks (MBs), the smallest census geographies defined in the Australia Statistical Geography Standard, are considered as the basic spatial unit. We join the MBs centroids to their nearest network node. The simulated data captures incidents of Priority 1 (P_1) and Priority 2 (P_2) where P_1 cases are more urgent. Based on a one-day sample of anonymised ambulance movements, we take that P_1 and P_2 are half and half (50%) in the total incident number and no consideration is given to lower priority incidents. P_1 has two types of incidents, Incident 1 (I_1) and Incident 2 (I_2). 10% of all incidents are of type I_1, which can get extra benefit from a special type of ambulance Vehicle Type 1 (V_1) offering specialised pre-hospital treatment. 5% and 15% I_1 are used for sensitivity analysis. Vehicle Type 2 (V_2) contributes the same (lower) survival benefit to all types of incidents.

The incidents exhibit a discernible temporal pattern. As illustrated in Figure 5.5a, there is a notable decrease in incident frequency during the hours between 0 AM and 6 AM. An 8-hour simulation captures a full vehicle shift. Figure 5.5b shows the number of incidents that occur over 8-hours starting at each time of day— this represents the total number of incidents to occur in a single simulation. These smoothed

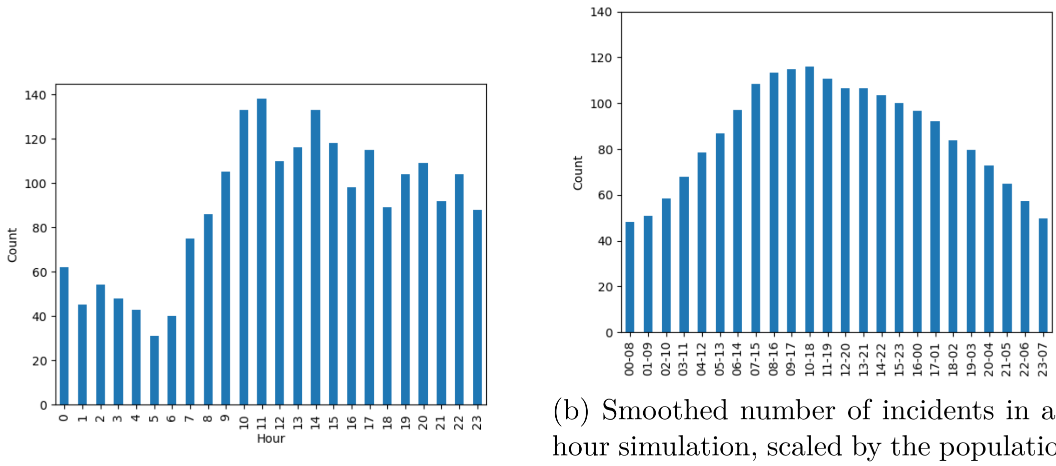

(a) Hourly incidents In NSW on one day

(b) Smoothed number of incidents in an 8-hour simulation, scaled by the population of the study area (12%)

Figure 5.5: Incident Temporal Distribution

Figure 5.6: Incident Temporal Distribution. Hourly distribution of incidents from a single day of anonymised vehicle movements.Incidents are highest during the day and evening. An 8-hour simulation for the study area would cover 50 to 110 incidents depending on the time of day.

values are scaled by 12% to account for the proportion of the state population that falls within the study area. Based on this data, the simulation will randomly generate 80 incidents in the 8-hr base-case simulation, and sensitivity tests will include 50 and 110 incidents in the same interval.

5.3.3 Processing Time Distribution

The duration of the Ambulance Response and Transport steps are determined by the transport network, the locations of the hospitals and the state of the system. On-scene Treatment and Destination Time vary between incidents and are best modeled as a right-tailed distribution with a preponderance of values where no on-scene treatment occurred. In the simulation, On-scene and Destination times are randomly drawn from the zero-inflated gamma distributions in Fig. 5.7 which are fitted to a one-day sample of ambulance data.

5.3.4 Dispatching Policies and Performance Measurement

Three dispatching scenarios are tested:

1. The most restrictive policy is dispatching V_1 to P_1-I_1 only,
2. We can relax the policy by allowing V_1 to be assigned to all incidents in P_1 that includes P_1-I_1 and P_1-I_2,

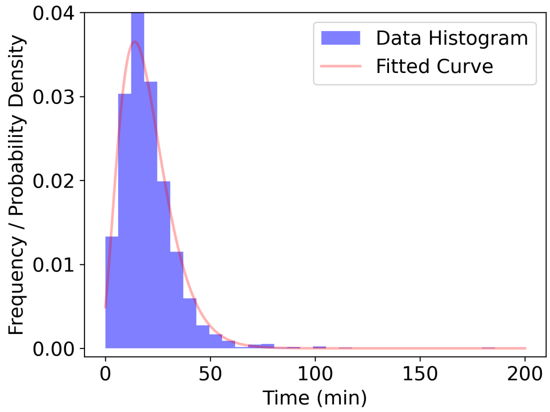

(a) Frequency / Probability Density of On-scene Treatment Times

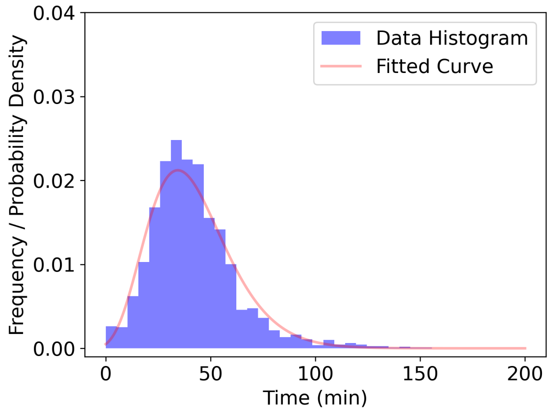

(b) Frequency / Probability Density of Destination Times

Figure 5.7: Treatment and Processing Times Distribution. Zero-inflated gamma distributions fitted to observed On-scene Treatment and Destination times. The Zero-inflation for on-scene time is 0.00626 and destination time is 0.00105.

3. The most flexible policy is when V_1 can respond to all types of incidents including P_2.

$P1$ incidents are the most urgent, life-threatening cases, so we use survival rate in Figure 3.2 of Chapter 3 to measure their response outcomes. Even though the objective function is insensitive to survival for lower priority incidents ($P2$), we consider response time as an indication of the equity of the policy in order to avoid long waiting times.

5.4 Result and Discussion

This experiment examines the relationship between survival rates and vehicle mixture under various dispatch strategies. Here, $V_{total} = V_1 + V_2$ so using 5 V_1 vehicles implies that out of a total of 20, the remaining 15 will be V_2.

We employed the survival rate as the primary performance metric in our study, capturing both the outcome of a single incident type and the overall efficacy of the Emergency Medical Services (EMS) system. An elevated survival rate indicates efficient handling of emergencies and optimised patient outcomes by the service. By contemplating the survival rate in association with the quantity of specialised vehicles and diverse dispatch policies, we identified optimal resource allocation and operational strategies for EMS under varied scenarios.

5.4.1 Diminishing returns to survival from additional special vehicles

Increasing from zero to one special vehicle improves overall patient survival regardless of the dispatching policy or the volume of incidents. This is shown in the initial upwards trajectory of all nine curves in Figure 5.8. The increase is explained by the dramatic improvement in survival of the P1-I1 cases, which improves the average without significantly decreasing survival for any other patients.

For flexible dispatch policies that allow the special vehicles to attend all patients (green in Figure 5.8) or any high priority patients (orange in Figure 5.8), the survival benefit of additional special vehicles can only increase or stay the same. The curves start to level off around six out of 20 vehicles for the base case (80 incidents) and low-demand case (50 incidents)— in these cases, 6 special vehicles are able to serve all of the special incidents and any additional special vehicles will be functioning as a normal vehicle. In the high-demand case, this leveling off is not observed even when more than half of the vehicles are the special type.

For more restrictive dispatch policies, shown with the blue line in Figure 5.8, additional special vehicles diminish the survival on non-P1-I1 patients because they reduce the availability of normal vehicles. This is shown as a peak in survival benefit when there are 4 special vehicles. Notice that for all levels of demand, the peak survival for the restrictive dispatch policy is lower than the maximum survival rate achievable with more permissive policies.

5.4.2 Degradation of response time for non-priority incidents

The objective of the experiment is to improve patient survival, but customer experience is a useful tie-breaker for otherwise similar survival outcomes. Response time is the way we can capture the disbenefit experienced by lower-priority patients when we allocate special vehicles to special incidents.

Figure 5.11 shows how average response time increases under the restrictive policies when the number of special vehicles increases. For each demand level, the red dashed lines benchmark the workload threshold— this is the number of vehicle-hours in the simulation (8 hrs x 20 vehicles), minus the average on-scene, transport and destination times, divided by the number of incidents. If the response time is longer than the workload benchmark, the queue will increase.

Under all demand levels and the most flexible dispatch policy, the total average response time doesn't change. The response time increases strongly with the num-

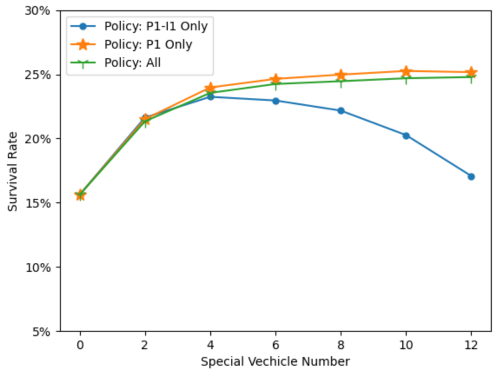

(a) Survival Rate for 50 Incidents

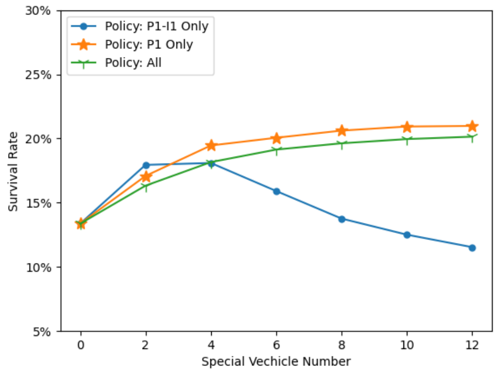

(b) Survival Rate for 80 Incidents

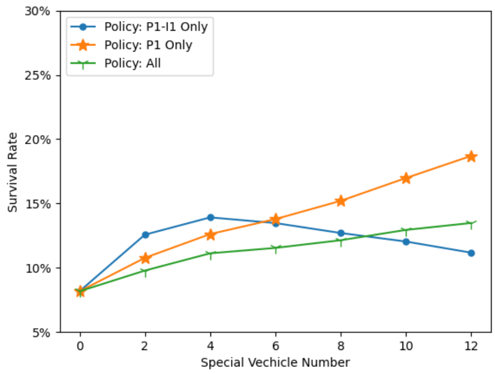

(c) Survival Rate for 110 Incidents

Figure 5.8: Survival Rate in Relation to Special Vehicle Number. All curves depict survival rates for P_1 incidents (P_1-I_1 and P_1-I_2) and 10% of total incidents are special (P_1-I_1). The colours represent different $V1$ dispatch policies scenarios. For low and moderate demand, the survival benefit peaks or plateaus when special vehicles are $\sim 25\%$ of the fleet.

ber of special vehicles under the most restrictive dispatch policy, especially when patient demand is high. This upturn in the blue curve (Figure 5.11) corroborates the finding that having greater than $\sim 25\%$ special vehicles has little benefit and perhaps disbenefits. Combined with Figure 5.8, the response times themselves are generally acceptable for the more permissive policies that perform best on survival rates.

5.4.3 Sensitivity to the ratio of supply and demand

It is expected that dispatching policy becomes more critical when the number of ambulances is small compared to the number of incidents. To reflect this dependency in the experiment, the supply is constant (*i.e.* the number of vehicles is always 20), but the demand varies (*i.e.* 50, 80 or 110 incidents in an 8-hr simulation). In the basic case with no special vehicles, the survival rate decreases as the number of

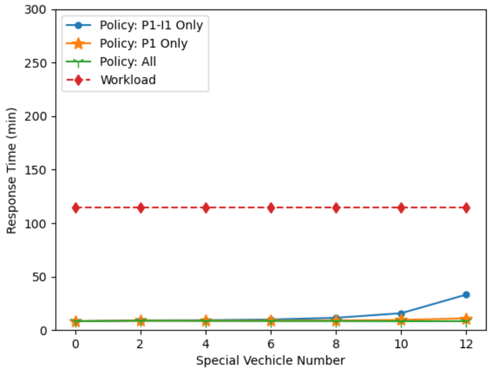

(a) Response Time for 50 Incidents

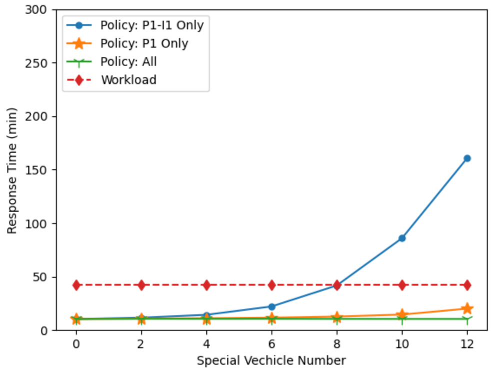

(b) Response Time for 80 Incidents

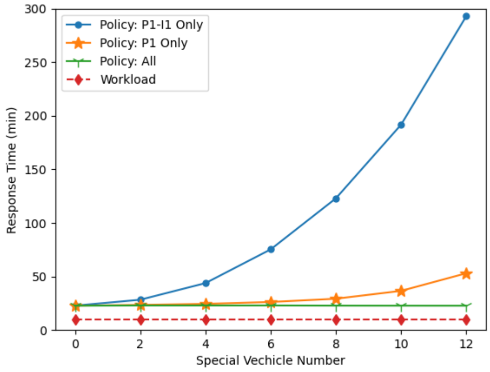

(c) Response Time for 110 Incidents

Figure 5.9: Response Time in Relation to Special Vehicle Number. All curves depict the average response time for all types of incidents, and 10% of incidents are special (P_1-I_1). Only the most restrictive policies (blue) result in increased response time, and when demand is high response time is too large regardless of vehicle number or dispatch policy.

incidents increases (purple line in Figure 5.10)— this is the expected result of more patients competing for fixed ambulance resources.

Across all policies, the survival rates reduce with the incident number increasing. Figure 5.10a shows how the policies are similar when the number of incidents is low because there is oversupply of vehicles. As the demand increases, the policies differentiate themselves.

This pattern differs somewhat in Figure 5.10b and Figure 5.10c when special vehicles represent 30 and 60% of the fleet respectively. When demand is low (50 incidents), the overall survival is slightly higher due to the availability of more special vehicles which dramatically improve survival for P1-I1 patients. However, when the demand is high enough to create queuing, the most restrictive policies (blue lines) have lower survival because P1 incidents are forced to queue.

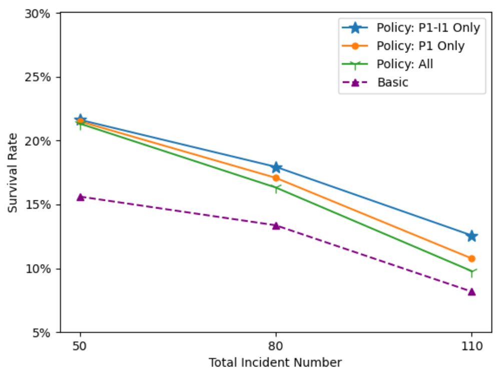

(a) Survival Rate for 2 Special Vehicles

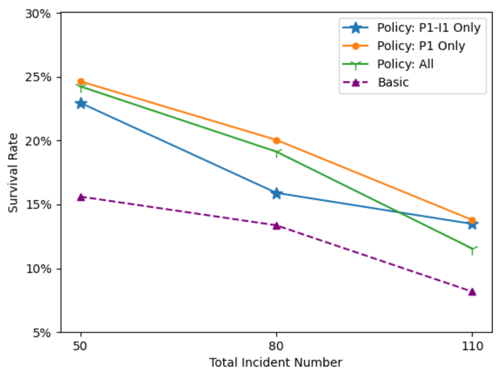

(b) Survival Rate for 6 Special Vehicles

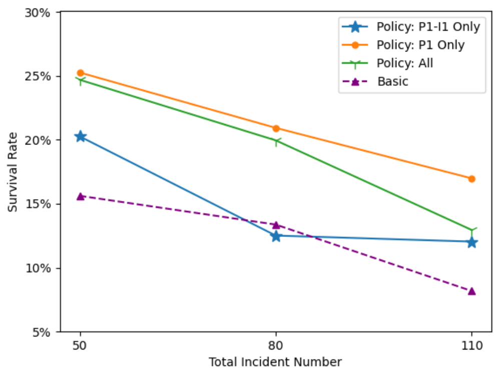

(c) Survival Rate for 10 Special Vehicles

Figure 5.10: Survival Rate in Relation to Special Vehicle Number. All curves depict survival rates for P_1 incidents (P_1-I_1 and P_1-I_2), and 10% of incidents are special (P_1-I_1). The colors represent different policy scenarios, and survival decreases with incident number in all cases. The relationship between policy and survival rate depends on the patient demand and vehicle supply.

Moreover, the overall survival for the most restrictive policy stabilises at the highest demand because the special vehicles grant higher survival to P1-I1 patients. Figure 5.10b shows a nearly perfect trade off between the P1-I1-only and the P1-only policies at the highest demand level— the same overall survival is achieved when the most restricted set of patients gets priority access.

In addition to the survival rate impacts at high demand, restrictive dispatch policies can also increase response times for low-priority incidents, which is a secondary consideration in the experiment. Figure 5.11 shows that the response times increase with total incident numbers and are unacceptably large under the most restrictive policy (blue) and at high demand.

The red lines show the time available for response per patient. Comparison to the policies suggests that there will be queuing under all policies when demand equates to 110 cases per simulation. Even moderate demand (80 cases) will result in queuing

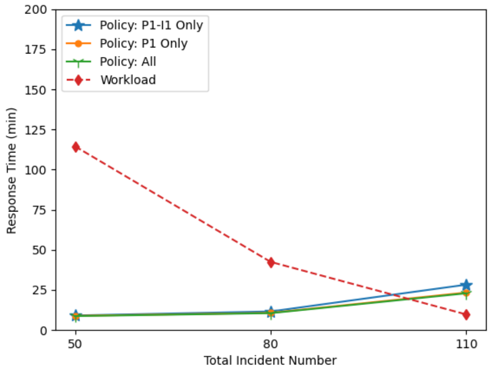

(a) Response Time for 2 Special Vehicles

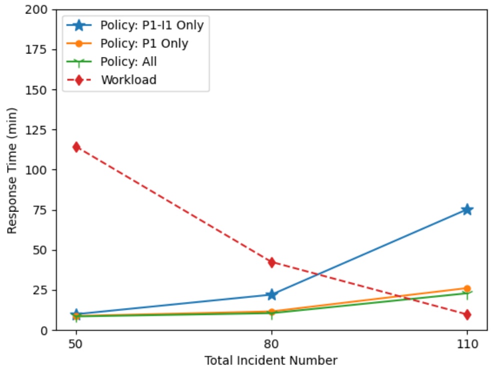

(b) Response Time for 6 Special Vehicles

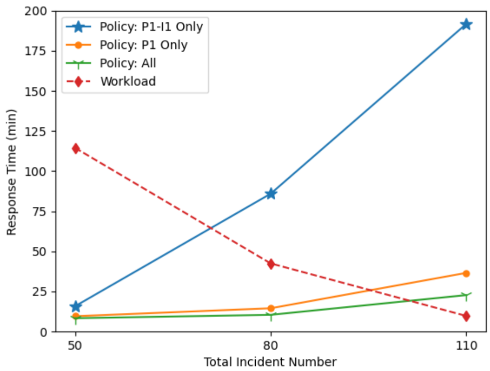

(c) Response Time for 10 Special Vehicles

Figure 5.11: Response Time in Relation to Special Vehicle Number. All curves depict the average response time for all types of incidents, and 10% of incidents are special (P_1-I_1). Under restrictive policies, the response time increases dramatically at high demand. Any point above the red workload line will mean that patients have to queue for a response.

for 50% special vehicles and the most restrictive policy (Figure 5.11c).

5.4.4 Sensitivity to special incident proportion

The results above are based on the special incidents (P1-I1) being 10% of all cases, which aligns with that expected rate in a sample of data. When this percentage is larger, special vehicles are more beneficial as shown in Figure 5.12.

Figure 5.12 illustrates that the survival rate will improve with the proportion of special incidents for all policies. This difference becomes more pronounced as the fraction of special vehicles increases. When 10% of vehicles are special vehicles, as shown in Figure 5.12a, there are only small differences in survival rate between the policies and special incident proportions. But when 50% of the vehicles are special vehicles, as shown in Figure 5.12c, the various dispatch policies show dramatically different effects at high and low rates of special incidents. When only 5% of incidents are P1-I1, the survival rate for the most restrictive policy (blue, $P_1 - I_1$-Only) policy is lower than the basecase with no special vehicles. But for 15% special incidents, the overall survival rate is $\sim$ 10% higher than the base case.

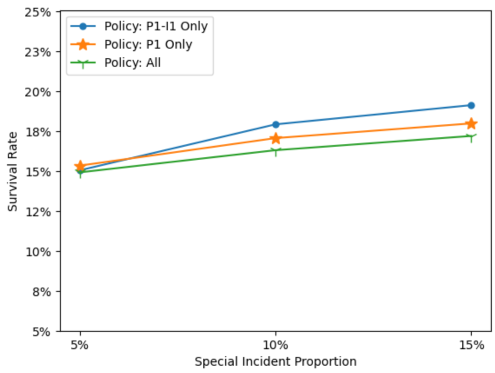

(a) Survival Rate for 2 Special Vehicles

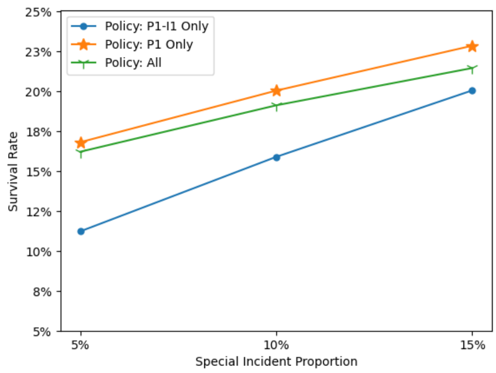

(b) Survival Rate for 6 Special Vehicles

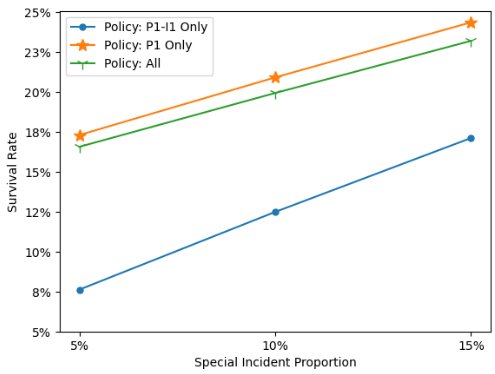

(c) Survival Rate for 10 Special Vehicles

Figure 5.12: Survival Rate in Relation to Special Incident Proportion. All curves depict survival rates for P_1 incidents (P_1-I_1 and P_1-I_2) for simulations with 80 incidents total. The colors depict the dispatch policies, which are highly differentiated when the number of special vehicles is high— when the special incident rate is low, a restrictive dispatch policy makes the survival worse than the baseline of no special vehicles whereas the permissive policies improve it.

The sensitivity of survival rate to special incident rate is not reflected in the response times. As shown in Figure 5.13, the response times are stable across 5, 10 and 15% special incident rates for the more permissive policies (all incidents and all P1 incidents). As expected, for low rates of special incidents, many special vehicles and restrictive policies, the response times are high because a large number of vehicles are allocated to a small number of cases.

Except for the most extreme case shown in Figure 5.13c where 50% of the vehicles are for the exclusive use of 5-15% of cases, the response times are all below the workload line (red). This implies that, when the policy and number of special vehicles are

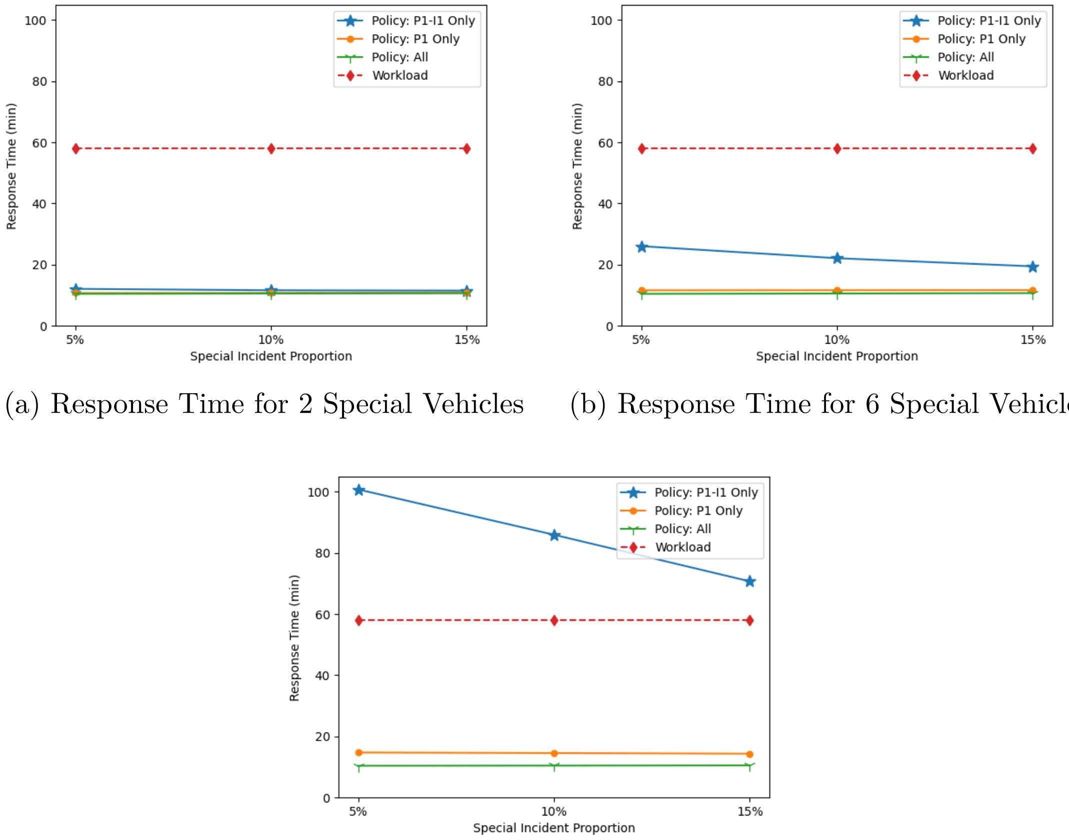

(a) Response Time for 2 Special Vehicles (b) Response Time for 6 Special Vehicles

(c) Response Time for 10 Special Vehicles

Figure 5.13: Response Time in Relation to Special Incident Proportion. All curves depict the average response time for all types of incidents. The legends represent different policy scenarios. $P_1 - I_1$ Only means the special vehicles V_1 only response to special incidents P_1-I_1, while P_1 Only represent dispatch exclusivity for high priority incidents both P_1-I_1 and P_1-I_2, and 'All' signifies that V_1 can be assigned to all incident types. 80 incidents in total.

suitable, queuing can be avoided at high, medium or low rates of special incidents.

5.4.5 Spatial distributions of different policies

To compare the spatial differences among various policies, we choose parameter combinations where all policies have similar outcomes in both survival rate and response time. Both metrics are aggregated to Statistical Area Level 2 (SA2), as defined in the Australian Statistical Geography Standard. SA2s typically have a few tens of thousands of residents living in them and are similar in size to a Local Government Area (LGA).

Figure 5.14 shows the survival rate distributions for different policies. In all three sub-figures, survival rates are generally higher in central areas and lower in the outer regions, particularly in the northeast and southwest. The reason should be that the hospitals are located in the middle so as to that region surrounded by available ambulances. The colors in the center are similar across the policies, but the policies vary more near the edges. This indicates that differences in survival rates are more significant in the boundary areas compared to the inner SA2s. We can see it from the northeast and southeast corners. Among the policies, the most restrictive policy ($P_1 - I_1$ Only) has higher survival rates in the northeast and southwest than the other policies. The P_1 Only policy performs better in the worst cases compared to the $P_1 - I_1$-Only policy but is not as good as the most permissive policy in the best cases. This results in the P_1 Only policy showing smoother changes in survival rates across adjacent regions. It is important to note that variations in the western region should be disregarded due to the geographical limitations of our study area. Although the study excludes ambulances originating from the western boundary, it is conceivable that some vehicles might still be operational in that region, as illustrated in Figure 5.4.

As shown in Figure 5.15, response times are shorter in the inner areas than at the edge of the study area. However, the most permissive policy has shorter response times in most regions compared to the other policies, suggesting that exclusive policies do not fully utilise resources. This pattern is also evident on a smaller scale. For instance, the northeast region has longer response times in Figure 5.15a than in Figure 5.15b and Figure 5.15c, even though the survival rate in this area is higher.

5.5 Discussion

Based on the experiment results of the agent-based simulation, we find that the deployment of advanced special ambulance vehicles can improve the survival rates of emergency incidents, especially for the incidents that correspond to the vehicle's intended function. The effectiveness of the deployment depends on the number of advanced vehicles deployed and the dispatching policy.

More specifically, we find that there is a relationship between the survival rates of emergencies and the number of specialised vehicles used, particularly in the context of incidents that these vehicles are primarily designed to handle. However, the ratio of special vehicles should not exceed the expected ratio of special incidents. If the proportion of special vehicles is higher than the special incidents, the obvious benefits to the special incident is counterbalanced against adverse impacts on the other types of incidents in terms of survival rate and response time.

Furthermore, if dispatched appropriately, these vehicles can substantially improve patient outcomes. With the moderate dispatching policy (special vehicles can respond to all high-priority incidents), the survival rate for high priority incidents can

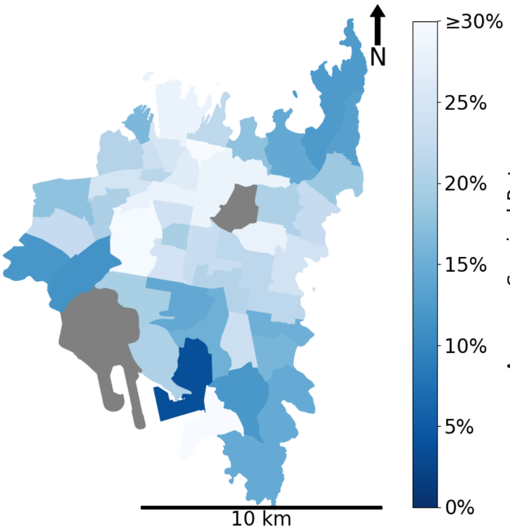

(a) Survival Rate for $P_1 - I_1$ Only Policy

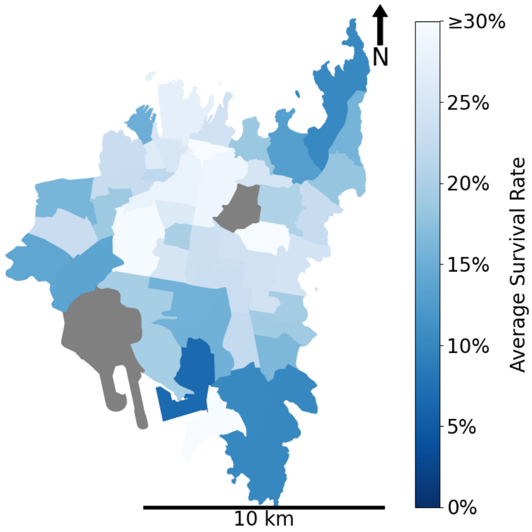

(b) Survival Rate for P_1 Only Policy

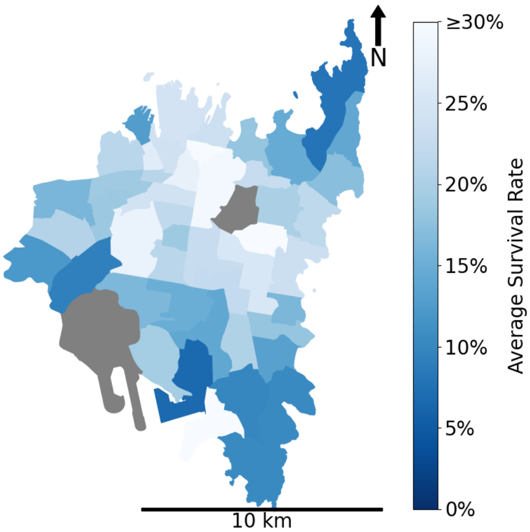

(c) Survival Rate for All Available Policy

Figure 5.14: Spatial Distributions of Survival Rate in Relation to Different Policies. The maps illustrate survival rates from simulations involving 4 special vehicles and 80 incidents, including 10% special incidents. The blue shades depict the average survival rates of $P1$ incidents in each area; lighter colors indicate higher survival rates, while grey indicates areas with no residents.

improve by more than 10% while maintaining acceptable response times for low priority incidents. In contrast, under more restrictive or permissive dispatch policies, survival rate can dip below the baseline of no-special-vehicles or response times can inflate, thus underscoring the importance of effective and informed dispatch policies.

From a geographical perspective, different policies affect subregions differently, even if their overall performances are similar. Exclusive policies perform better for serious incidents, but the most permissive policy has shorter average response times, bene-

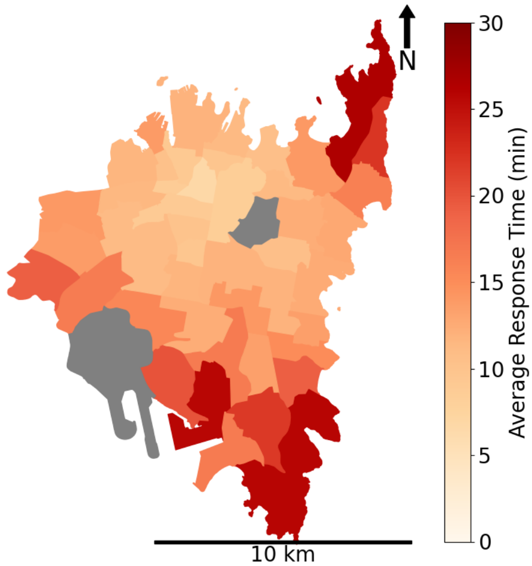

(a) Response time for $P_1 - I_1$ Only Policy

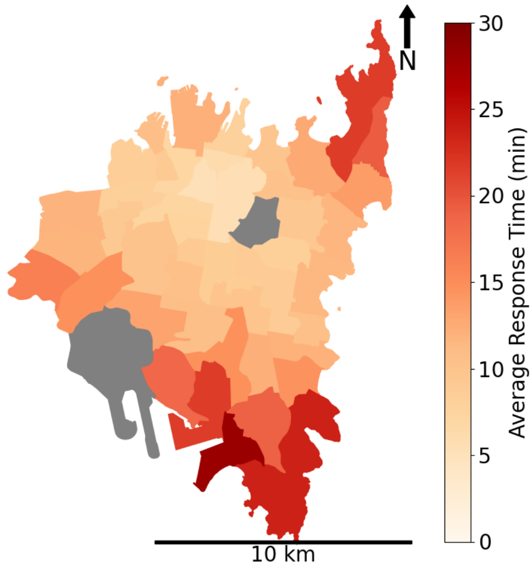

(b) Response time for P_1 Only Policy

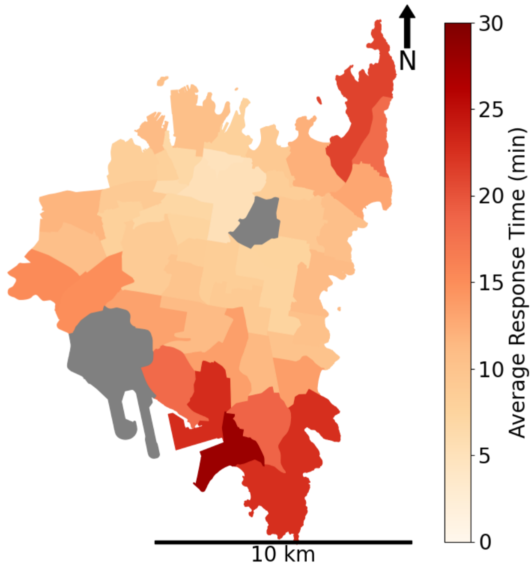

(c) Response time for All Available Policy

Figure 5.15: Spatial Distributions of Response Time in Relation to Different Policies. The maps illustrate response times from simulations involving 4 special vehicles and 80 incidents, including 10% special incidents. The red colors depict the average response time of all incidents in the regions; lighter colors indicate shorter response times, while grey indicates areas with no residents.

fiting less severe incidents. The impact of these policies varies more near the edge of the study area where incidents are less likely to have multiple available vehicles surrounding them. When special vehicles respond only to their designated incidents, outcomes for those incidents improve, especially for long-distance responses, but response times get worse. This effect is also particularly noticeable in edge areas.

Our agent-based simulation framework and methodology used in this study can be adapted to any region to evaluate the potential benefits of deploying advanced emergency vehicles. One of the key strengths of our study lies in the versatility of the agent-based simulation framework. The framework is highly adaptable and can be applied to test sensitivity to dispatch policy, number of special vehicles, number of incidents and special incident rate. By adjusting the parameters to suit the unique requirements of different emergency settings, decision-makers can explore the policy impact on patient survival– an outcome-based (rather than structure- or process-based) performance measure. This can serve as a valuable tool in informing policy decisions and optimising resource allocation.

Chapter 6

Conclusion

6.1 Summary of Research

This thesis explored the integration of transportation engineering principles and logistics strategies to optimize the performance of emergency medical services (EMS). The research aimed to address several critical challenges in EMS operations, including facility location optimization, resource allocation, real-time dispatching, and the evaluation of competing delivery strategies. Through the application of static optimization models, dynamic simulation, and comparative analysis of algorithms, this thesis offers practical insights into how transportation logistics can improve EMS response times, resource use, and patient outcomes.

This chapter revisits the main research questions introduced in Chapter 1 and reflects on how the findings from Chapters 3-5 collectively address the gaps. The thesis set out to understand how emergency medical services can be designed and operated more effectively by improving delivery strategies, optimizing the spatial allocation of resources, and refining dispatching policies. Together, these questions sought to connect strategic planning with real-time operations and to evaluate how logistics decisions ultimately affect accessibility, efficiency, and patient survival.

6.1.1 Logistics Strategies for EMS Delivery

In Chapter 3, the research compared three service delivery strategies: an in-hospital approach, a rendezvous approach, and a pre-hospital approach. In-hospital method is the conventional strategy where patients must be transported from the incident site to a specialized center before advanced interventions begin. Although well established, this model often fails to help patients who live too far from specialized hospitals or who require lengthy on-scene stabilization times. As time to treatment is crucial, in-hospital treatment can leave many potential candidates effectively unserved.

Rendezvous strategy partially alleviates the constraints of in-hospital care by dispatching a mobile team to meet the patient at a middle emergency department. This parallel dispatch can reduce travel time for the specialized team while the patient, transported by a standard ambulance, also heads toward a nearby facility. By splitting responsibilities and minimizing total transfer distances, rendezvous strategies allow further intervention to be initiated sooner, although they demand high coordination and face the logistical challenge of synchronizing movement for both medical staff and patients as well as requiring more EMS resources.

Pre-hospital approach goes a step further by taking the advanced intervention directly to the scene of patient. In this scenario, specialized teams and portable devices are dispatched immediately, for hospital transfer before initiating the procedure. This strategy can notably increase access for populations that would otherwise lie beyond the 60-minute threshold; however, it also involves extensive resource commitments, training requirements, and the complexities of performing advanced procedures in non-local settings. When evaluated through a transport lens, each of these three strategies can be seen as a logistics strategy, where in-hospital one is the simplest single-tier mode, rendezvous parallels the concept of hub-and-spoke, and pre-hospital model is a mobile hub supplying advanced care at the point of demand.

6.1.2 Facility Location Optimization and Resource Allocation

Chapter 4 expanded on the theme of Chapter 3 by applying spatial modeling and facility location optimization to emergency care planning. Using a maximal coverage problem (MCP) framework—augmented by continuous, time-dependent survival curves—the chapter explored how to best position or allocate specialized medical resources in a large metropolitan area like Sydney when implementing rendezvous strategy. Traditional binary coverage measures are limited in that they ignore how each additional minute of delay can diminish the probability of survival. By contrast, using a logistic survival model as an impedance function provided a more nuanced view of coverage.

The analysis indicated that the rendezvous strategy can improve the service coverage significantly. Also, placing advanced equipment capabilities at select nodes or stations, rather than restricting all resources to hospitals in high-population-density regions, can substantially raise the share of patients served within an hour. This approach is particularly potent when paired with parallel or mobile deployments, for example rendezvous strategy. Additionally, the modeling revealed that many areas within a strict coverage boundary can still achieve improved outcomes if their travel times are reduced by even a few minutes. The chapter thus underscored how strategic placement of specialized teams, guided by data on historical incident locations and realistic travel speeds, can optimize population-wide survival benefits.

6.1.3 Dynamic Simulation for Real-Time Dispatching

In Chapter 5, the focus moves from long-range delivery strategy and location planning to operational-level decisions concerning vehicle fleets and dispatch policies. Recognizing that specialized vehicles dedicated exclusively to high-acuity patients risk underutilization, this chapter employs an agent-based simulation to capture dynamic interactions among ambulances, patients, and the dispatch system. Within this simulation, different vehicle types—including standard ambulances and those equipped with specialized, advanced life-support devices—are evaluated to determine how their assigned roles can accelerate or hinder response times.

The results reinforce earlier findings (see Chapter 3), showing that introducing specialized vehicles sharply improves outcomes for the patient population they target. However, committing too many resources to specialized functions can undermine system flexibility if advanced equipment remains idle or demand for it is lower than expected. Dispatch policies that narrowly constrain specialized units to their target incidents minimize misuse but may leave standard ambulances overwhelmed when faced with multiple routine calls. More permissive policies, on the other hand, can spread advanced resources too thin by allowing them to respond to lower-priority incidents, potentially delaying critical care for truly severe emergencies. A moderate approach—one that preferentially assigns specialized units to high-severity cases yet permits these vehicles to assist with general calls when otherwise idle—generally yielded the most balanced and efficient outcomes. This interplay between vehicle composition and dispatch strategy underscores the value of carefully managing specialized and general-purpose ambulances alike, ensuring robust coverage and responsiveness across the EMS system.

6.2 Contributions to EMS Logistics

One of the main contributions of this thesis lies in bridging transportation and medical logistics. By reframing EMS challenges—like resource accessibility or specialized ambulance utilization—into classic logistics concepts (e.g., facility location problems, or dynamic vehicle dispatch), the research demonstrates how established models from transportation engineering can be leveraged to address life-critical problems. This cross-disciplinary approach enriches both fields: medical practitioners gain access to sophisticated analytics tools, while logistics scholars discover new contexts with direct societal impact.

6.2.1 Insights and Advancements from Each Chapter

By contrasting in-hospital, rendezvous, and pre-hospital extracorporeal membrane oxygenation (ECMO) cardiopulmonary resuscitation (ECPR), Chapter 3 elucidates how each approach impacts spatial coverage, survival likelihood, and operational complexity. In-hospital ECPR, while logistically straightforward, is geographically restrictive. Rendezvous and pre-hospital ECPR strategies extend reach by reducing total transport or initiation times. The pre-hospital strategy demonstrates the highest performance but requires large ambulances equipped with ECPR devices and trained paramedics capable of operating outside the hospital environment. Conversely, the rendezvous strategy avoids additional staff training or ambulance purchases but occupies two teams simultaneously. These findings underscore the need for strategic flexibility in EMS logistics to adapt to varying emergency scenarios and enhance patient outcomes.

Building on Chapter 3, Chapter 4 introduces a maximal survival problem that examines how facility locations influence system performance through a tailored delivery strategy. Incorporating real-world constraints, such as time-dependent survival rates, a limited number of specialized teams, and equipment positions, yields a more nuanced measure of effective coverage. This chapter reveals that the rendezvous strategy not only expands coverage but also enhances survival expectations for individuals already covered by the in-hospital strategy. Optimizing facility locations further improves ECPR efficiency, with certain candidate locations significantly boosting overall survival rates. However, an excessive number of facility options diminishes problem-solving efficiency while providing marginal benefits.

In contrast to the previous chapters, Chapter 5 narrows its focus to the operational scale, evaluating logistics strategies for various vehicle types. This chapter highlights the importance of targeted and adaptable dispatch policies. The interaction between vehicle composition and policy rules—such as reserving advanced vehicles for specific incidents or allowing them to respond to any call—has a substantial impact on performance, even when facility locations remain unchanged. A higher ratio of specialized ambulances improves average survival rates if these ambulances also provide standard first aid for typical incidents. However, if they respond only to targeted incidents, their ratio must be carefully managed; an excess of specialized ambulances waiting for rare incidents leads to inefficient resource utilization, negatively impacting other types of incidents. Thus, even with effective delivery strategies and optimized facility locations, dispatch policies remain a critical consideration.

6.2.2 Synthesis of Insights Across Chapters

Taken collectively, the three chapters offer a multi-layered perspective on how logistics strategies shape EMS effectiveness. Chapter 3 provides a macro-level blueprint for alternative delivery pathways, clarifying how each approach modifies access times. Chapter 4 delves deeper into spatial modeling to identify how specialized team locations can optimize survival outcomes. Finally, Chapter 5 focuses on the operational dimension, exploring how dispatch tactics influence the real-time utilization of specialized vehicles.

These layers are mutually reinforcing:

From Chapter 3 to Chapter 4: Understanding that pre-hospital or rendezvous ECPR strategies can extend coverage informs facility location optimization. By identifying the baseline potential of each strategy, planners can determine how many pre-hospital units or rendezvous sites are realistically necessary.

From Chapter 4 to Chapter 5: While facility location optimization highlights where and how many advanced vehicles to position, agent-based modeling (ABM) of dispatch policies captures day-to-day feasibility. An optimal delivery strategy might falter under real-world conditions if multiple calls occur simultaneously, emphasizing the need for nuanced dispatch policies.

From Chapter 5 back to Chapter 3: Observing how specialized vehicles are utilized in real-time further refines decisions about which delivery model to scale. For example, a system heavily reliant on pre-hospital vehicles but lacking supportive dispatch policies may underperform, whereas a rendezvous model could be more cost-effective if advanced vehicles are dynamically reallocated.

Ultimately, integrating insights from all three chapters reveals that EMS logistics strategies cannot be reduced to a single dimension, such as resource placement or policy choice. Success depends on a synchronized approach: selecting an appropriate delivery framework, aligning facility locations and resource levels with that framework, and tailoring dispatch policies to maintain responsiveness under real-world conditions.

6.3 Limitations and Future Work

Designing logistics strategies for emergency medical services confronts inherent limitations arising from shifting technological capabilities, changing patient needs, and system-level complexities. While this thesis has attempted to integrate advanced vehicle types, evolving delivery models, and facility location optimizations under one conceptual framework, several dimensions remain outside its scope and offer opportunities for further investigation.

6.3.1 Evolving Treatment Technologies and Vehicle Configurations

In this thesis, two vehicle types and three incident categories were used for demonstration purposes. However, EMS agencies often manage a broader range of emergencies and may deploy numerous specialized vehicles. The concept of a modular ambulance, where on-board equipment can be rapidly reconfigured to meet specific clinical needs such as installing triage kits or air filtration units, holds considerable promise (Liu et al., 2020; Pena et al., 2025). Future research might explore how modular ambulances could streamline inventory management, reduce duplication of resources, and optimize staff training, thereby creating a more flexible, responsive EMS system.

6.3.2 Integration of Multiple Delivery Modes

Although some chapters discussed delivery strategies (e.g., rendezvous vs. pre-hospital strategies), the actual array of possible "modes" for EMS extends further—helicopters for lengthy transport (Michaels et al., 2019), drones for small medical kit delivery (Zailani et al., 2021), or motorbikes for rapid initial stabilization (Yılmaz et al., 2025). These modes also interact with each other; for instance, Motorbikes allow responders to navigate congested urban areas more swiftly, initiating care ahead of the main ambulance; drones can ferry essential medications or devices to the scene, buying time until a full clinical team arrives. The multimodal EMS framework is further advanced by aerial technologies, including electric vertical take-off and landing (eVTOL) flying ambulances, which facilitate rapid patient transport while circumventing surface-level constraints (Goyal and Cohen, 2022), and drones, which offer reliable and timely delivery of perishable medical supplies (Lakhwani et al., 2025). Notably, patients receiving aerial medical services have demonstrated improved survival outcomes compared to those relying solely on ground-based emergency transport (Andruszkow et al., 2013). Given the diversity of emergent events and geographic contexts, future modeling could delve into multi-modal dispatching algorithms that optimize the synergy among different types of resources.

6.3.3 Dynamic System Inputs and Uncertainty in EMS Modeling

A significant limitation of this thesis lies in its assumption of static conditions for travel time, incident frequency, and resource availability. In practice, EMS operations are affected by dynamic and uncertain inputs—such as traffic congestion, demand surges, and varying patient acuity. Recent studies have applied real-time traffic data and predictive analytics to improve EMS routing and resource allocation (Ji et al., 2019). In parallel, stochastic models have been proposed to capture demand and travel time uncertainty more realistically, enhancing decision robustness under variable conditions (Stratman et al., 2023; Yoon et al., 2021). AI-based decision-support models can help operate EMS systems more efficiently and effectively by learning from historical data and adapting to contextual changes (Tluli et al., 2024).

6.3.4 Additional Resource Requirements

Although alternative strategies—rendezvous strategy—show promise in enhancing timeliness and survival rates, each innovation may demand additional resources that can temper or even undermine its potential benefits. For example, integrating autonomous vehicles into an EMS fleet could significantly reduce driver fatigue, lower operational costs over time, and optimize routing. Yet, strict regulatory controls could limit widespread deployment. Similarly, a broader fleet of drones or motorbikes might improve on-scene arrival times but necessitate extra budgets for pilot training, maintenance, and regulatory compliance. These constraints underscore the importance of conducting cost-benefit analyses and staff readiness assessments before fully committing to new, resource-intensive capabilities.

Bibliography

Abou-Ali, M., Shouman, M., 2004. Effect of dynamic and static dispatching strategies on dynamically planned and unplanned fms. Journal of materials processing technology 148, 132–138.

Aboueljinane, L., Sahin, E., Jemai, Z., 2013. A review on simulation models applied to emergency medical service operations. Computers & Industrial Engineering 66, 734–750.

Ageron, F.X., Debaty, G., Gayet-Ageron, A., Belle, L., Gaillard, A., Monnet, M.F., Bare, S., Richard, J.C., Danel, V., Perfus, J.P., et al., 2016. Impact of an emergency medical dispatch system on survival from out-of-hospital cardiac arrest: a population-based study. Scandinavian journal of trauma, resuscitation and emergency medicine 24, 1–9.

Ahmadi-Javid, A., Seyedi, P., Syam, S.S., 2017. A survey of healthcare facility location. Computers & Operations Research 79, 223–263.

Alanazy, A.R.M., Wark, S., Fraser, J., Nagle, A., 2019. Factors impacting patient outcomes associated with use of emergency medical services operating in urban versus rural areas: a systematic review. International journal of environmental research and public health 16, 1728.

Alotaibi, M.F., Ibrahim, D.M., 2018. Agent-based simulation for coordination emergency response: A review study, in: 2018 1st International Conference on Computer Applications & Information Security (ICCAIS), IEEE. pp. 1–5.

Alshammari, A.D., Alobaid, A.M., Azharuddin, A., 2024. Professionalizing emergency medical service response time. Emergency Health Services Journal , 10–4103.

Altarawneh, L., 2023. Interpretable deep learning and transfer learning-based spatial-temporal modeling for vaccines demand prediction. Ph.D. thesis. State University of New York at Binghamton.

Amakama, N.J., Duserre, G., Cadiere, A., Schuette, R.W., 2025. Predicting the impact of mobile field hospital deployment in reducing patients mean wait times in emergency departments: An agent-based simulation approach .

Amorim, M., Ferreira, S., Couto, A., 2018. Emergency medical service response: Analyzing vehicle dispatching rules. Transportation research record 2672, 10–21.

Andrew, E., Nehme, Z., Cameron, P., Smith, K., 2020. Drivers of increasing emergency ambulance demand. Prehospital Emergency Care 24, 385–393.

Andruszkow, H., Lefering, R., Frink, M., Mommsen, P., Zeckey, C., Rahe, K., Krettek, C., Hildebrand, F., 2013. Survival benefit of helicopter emergency medical services compared to ground emergency medical services in traumatized patients. Critical Care 17, R124.

Aringhieri, R., Bruni, M.E., Khodaparasti, S., van Essen, J.T., 2017. Emergency medical services and beyond: Addressing new challenges through a wide literature review. Computers & Operations Research 78, 349–368.

Atta, S., Sinha Mahapatra, P.R., Mukhopadhyay, A., 2018. Solving maximal covering location problem using genetic algorithm with local refinement. Soft Computing 22, 3891–3906.

Australian Bureau of Statistics, 2017. 2016 census mesh block counts. URL: https://www.abs.gov.au/AUSSTATS/abs@.nsf/DetailsPage/2074.02016?OpenDocument. accessed: July-2021.

Australian Institute of Health and Welfare, 2022. Myhospitals mapping details. URL: https://www.aihw.gov.au/reports-data/myhospitals/content/data-downloads. accessed: May-2022.

Bandara, D., Mayorga, M.E., McLay, L.A., 2014. Priority dispatching strategies for ems systems. Journal of the Operational Research Society 65, 572–587.

Bartos, J.A., Carlson, K., Carlson, C., Raveendran, G., John, R., Aufderheide, T.P., Yannopoulos, D., 2018. Surviving refractory out-of-hospital ventricular fibrillation cardiac arrest: critical care and extracorporeal membrane oxygenation management. Resuscitation 132, 47–55.

Bartos, J.A., Frascone, R., Conterato, M., Wesley, K., Lick, C., Sipprell, K., Vuljaj, N., Burnett, A., Peterson, B.K., Simpson, N., et al., 2020a. The minnesota mobile extracorporeal cardiopulmonary resuscitation consortium for treatment of out-of-hospital refractory ventricular fibrillation: program description, performance, and outcomes. EClinicalMedicine 29.

Bartos, J.A., Grunau, B., Carlson, C., Duval, S., Ripeckyj, A., Kalra, R., Raveendran, G., John, R., Conterato, M., Frascone, R.J., et al., 2020b. Improved survival with extracorporeal cardiopulmonary resuscitation despite progressive metabolic derangement associated with prolonged resuscitation. Circulation 141, 877–886.

Bélanger, V., Lanzarone, E., Nicoletta, V., Ruiz, A., Soriano, P., 2020. A recursive simulation-optimization framework for the ambulance location and dispatching problem. European Journal of Operational Research 286, 713–725.

Belohlavek, J., Smalcova, J., Rob, D., Franek, O., Smid, O., Pokorna, M., Horák, J., Mrazek, V., Kovarnik, T., Zemanek, D., et al., 2022. Effect of intra-arrest transport, extracorporeal cardiopulmonary resuscitation, and immediate invasive assessment and treatment on functional neurologic outcome in refractory out-of-hospital cardiac arrest: a randomized clinical trial. Jama 327, 737–747.

Benjamin, E.J., Muntner, P., Alonso, A., Bittencourt, M.S., Callaway, C.W., Carson, A.P., Chamberlain, A.M., Chang, A.R., Cheng, S., Das, S.R., et al., 2019. Heart disease and stroke statistics—2019 update: a report from the american heart association. Circulation 139, e56–e528.

Berdowski, J., Beekhuis, F., Zwinderman, A.H., Tijssen, J.G., Koster, R.W., 2009. Importance of the first link: description and recognition of an out-of-hospital cardiac arrest in an emergency call. Circulation 119, 2096–2102.

Berman, O., Drezner, Z., Krass, D., 2010. Generalized coverage: New developments in covering location models. Computers & Operations Research 37, 1675–1687.

Berman, O., Krass, D., Drezner, Z., 2003. The gradual covering decay location problem on a network. European Journal of Operational Research 151, 474–480.

Bimpou, K., Ferguson, N.S., 2020. Dynamic accessibility: Incorporating day-to-day travel time reliability into accessibility measurement. Journal of Transport Geography 89, 102892.

Blackwell, T.H., Kaufman, J.S., 2002. Response time effectiveness: comparison of response time and survival in an urban emergency medical services system. Academic Emergency Medicine 9, 288–295.

Borozdukhin, A., Dolinina, O., Pechenkin, V., 2016. Approach to the garbage collection in the "smart clean city" project, in: 2016 4th IEEE International Colloquium on Information Science and Technology (CiSt), IEEE. pp. 918–922.

Boujemaa, R., Jebali, A., Hammami, S., Ruiz, A., 2020. Multi-period stochastic programming models for two-tiered emergency medical service system. Computers & Operations Research 123, 104974.

Boujemaa, R., Jebali, A., Hammami, S., Ruiz, A., Bouchriha, H., 2018. A stochastic approach for designing two-tiered emergency medical service systems. Flexible Services and Manufacturing Journal 30, 123–152.

Bray, J., Nehme, Z., Nguyen, A., Lockey, A., Finn, J., et al., 2020. A systematic review of the impact of emergency medical service practitioner experience and exposure to out of hospital cardiac arrest on patient outcomes. Resuscitation 155, 134–142.

Brotcorne, L., Laporte, G., Semet, F., 2003. Ambulance location and relocation models. European journal of operational research 147, 451–463.

Cabrera, E., Taboada, M., Iglesias, M.L., Epelde, F., Luque, E., 2011. Optimization of healthcare emergency departments by agent-based simulation. Procedia computer science 4, 1880–1889.

Chen, P., Chankov, S.M., 2017. Crowdsourced delivery for last-mile distribution: An agent-based modelling and simulation approach, in: 2017 IEEE international conference on industrial engineering and engineering management (IEEM), IEEE. pp. 1271–1275.

Chen, Y., Lai, Z., 2022. A multi-objective optimization approach for emergency medical service facilities location-allocation in rural areas. Risk Management and Healthcare Policy , 473–490.

Church, R., ReVelle, C., 1974. The maximal covering location problem. Papers of the Regional Science Association 32, 101–118.

Cimellaro, G.P., Arcidiacono, V., Reinhorn, A., Bruneau, M., 2013. Disaster resilience of hospitals considering emergency ambulance services, in: Structures Congress 2013: Bridging Your Passion with Your Profession, pp. 2824–2836.

Comber, A.J., Sasaki, S., Suzuki, H., Brunsdon, C., 2011. A modified grouping genetic algorithm to select ambulance site locations. International Journal of Geographical Information Science 25, 807–823.

Committee on the Future of Emergency Care in the United States Health System, 2007. Emergency Medical Services: At the Crossroads. National Academies Press.

Compass, 2019. Australian Roads Speed Data,Nov. 11 - Nov. 25, 2019. Compass IoT Pty Limited.

Cornuéjols, G., Nemhauser, G., Wolsey, L., 1983. The uncapicitated facility location problem. Technical Report. Cornell University Operations Research and Industrial Engineering.

Coskun, N., Erol, R., 2010. An optimization model for locating and sizing emergency medical service stations. Journal of medical systems 34, 43–49.

Culley, L.L., Clark, J.J., Eisenberg, M.S., Larsen, M.P., 1991. Dispatcher-assisted telephone cpr: common delays and time standards for delivery. Annals of emergency medicine 20, 362–366.

Dami, F., Heymann, E., Pasquier, M., Fuchs, V., Carron, P.N., Hugli, O., 2015. Time to identify cardiac arrest and provide dispatch-assisted cardio-pulmonary resuscitation in a criteria-based dispatch system. Resuscitation 97, 27–33.

Dell'Amico, M., Hadjidimitriou, S., 2012. Innovative logistics model and containers solution for efficient last mile delivery. Procedia-Social and Behavioral Sciences 48, 1505–1514.

Dennis, M., Group, R.S., Group, S.E.R.I., Burns, B., Dyson, S., Austin, D., Arnold, J., Berendsen, S., Coggins, A., Duffin, S., Dutton, N., et al., 2022. Out-of-hospital cardiac arrest outcomes, end-tidal carbon dioxide and extracorporeal cardiopulmonary resuscitation eligibility: New south wales pilot data. Emergency Medicine Australasia 34, 452–455.

Di Febbraro, A., Sacco, N., Saeednia, M., 2016. An agent-based framework for cooperative planning of intermodal freight transport chains. Transportation Research Part C: Emerging Technologies 64, 72–85.

Dick, W.F., 2003. Anglo-american vs. franco-german emergency medical services system. Prehospital and disaster medicine 18, 29–37.

Döyen, A., Aras, N., Barbarosoğlu, G., 2012. A two-echelon stochastic facility location model for humanitarian relief logistics. Optimization letters 6, 1123–1145.

Drabecki, M., Toczyłowski, E., Pieńkosz, K., Honisz, G., Kułak, K., 2023. Multi-criteria assignment problems for optimising the emergency medical services (ems), considering non-homogeneous speciality of the emergency departments and ems crews. Scientific Reports 13, 7496.

Dyson, S., 2021. Cardiac Arrest Registry 2019 Report. Technical Report. NSW Ambulance.

Edjinedja, K., Elfahim, O., Arnaud, É., Barakat, O., Desmettre, T., Robert-Nicoud, S., 2024. Multi-agent modeling and discrete-event simulation for analyzing the responsiveness of pre-hospital emergency services in the event of mass casualties due to rare events, in: 2024 IEEE International Conference on Technology Management, Operations and Decisions (ICTMOD), IEEE. pp. 1–7.

Enayati, S., Mayorga, M.E., Toro-Díaz, H., Albert, L.A., 2019. Identifying trade-offs in equity and efficiency for simultaneously optimizing location and multipriority dispatch of ambulances. International Transactions in Operational Research 26, 415–438.

Erdemir, E.T., Batta, R., Rogerson, P.A., Blatt, A., Flanigan, M., 2010. Joint ground and air emergency medical services coverage models: A greedy heuristic solution approach. European Journal of Operational Research 207, 736–749.

Farcas, A.M., Crowe, R.P., Kennel, J., Little, N., Haamid, A., Camacho, M.A., Pleasant, T., Owusu-Ansah, S., Joiner, A.P., Tripp, R., et al., 2024. Achieving equity in ems care and patient outcomes through quality management systems: A position statement. Prehospital Emergency Care , 1–16.

Farcas, A.M., Joiner, A.P., Rudman, J.S., Ramesh, K., Torres, G., Crowe, R.P., Curtis, T., Tripp, R., Bowers, K., von Isenburg, M., et al., 2023. Disparities in emergency medical services care delivery in the united states: a scoping review. Prehospital Emergency Care 27, 1058–1071.

Gamberini, L., Picoco, C., Del Giudice, D., Zenesini, C., Tartaglione, M., Coniglio, C., Semeraro, F., Bizzi, F., Santini, S., Gordini, G., 2021. Improving the appropriateness of advanced life support teams' dispatch: a before-after study. Prehospital and Disaster Medicine 36, 195–201.

Gendreau, M., Laporte, G., Semet, F., 2001. A dynamic model and parallel tabu search heuristic for real-time ambulance relocation. Parallel computing 27, 1641–1653.

Goyal, R., Cohen, A., 2022. Advanced air mobility: Opportunities and challenges deploying evtols for air ambulance service. Applied Sciences 12, 1183.

Grot, M., 2024. Decision support framework for tactical emergency medical service location planning. Omega 125, 103036.

Guigues, V., Kleywegt, A., Nascimento, V.H., 2022. Operation of an ambulance fleet under uncertainty. arXiv preprint arXiv:2203.16371 .

Gurobi Optimization, LLC, 2023. Gurobi Optimizer Reference Manual. URL: https://www.gurobi.com.

Hadaya, J., Dobaria, V., Aguayo, E., Kwon, O.J., Sanaiha, Y., Hyunh, A., Sareh, S., Benharash, P., 2020. National trends in utilization and outcomes of extracorporeal support for in-and out-of-hospital cardiac arrest. Resuscitation 151, 181–188.

Hansen, W.G., 1959. How accessibility shapes land use. Journal of the American Institute of planners 25, 73–76.

Hawe, G.I., Coates, G., Wilson, D.T., Crouch, R.S., 2015. Agent-based simulation of emergency response to plan the allocation of resources for a hypothetical two-site major incident. Engineering Applications of Artificial Intelligence 46, 336–345.

Hörl, S., Becker, F., Axhausen, K.W., 2021. Simulation of price, customer behaviour and system impact for a cost-covering automated taxi system in zurich. Transportation Research Part C: Emerging Technologies 123, 102974.

Jánošíková, L., Jankovič, P., Kvet, M., Zajacová, F., 2021. Coverage versus response time objectives in ambulance location. International Journal of Health Geographics 20, 1–16.

Ji, S., Zheng, Y., Wang, W., Li, T., 2019. Real-time ambulance redeployment: A data-driven approach. IEEE Transactions on Knowledge and Data Engineering 32, 2213–2226.

Jia, H., Ordóñez, F., Dessouky, M., 2007. A modeling framework for facility location of medical services for large-scale emergencies. IIE transactions 39, 41–55.

Kar, B., Jenamani, M., 2024. Optimal multimodal multi-echelon vaccine distribution network design for low and medium-income countries with manufacturing infrastructure during healthcare emergencies. International Journal of Production Economics , 109282.

Karatas, M., Eriskin, L., 2021. The minimal covering location and sizing problem in the presence of gradual cooperative coverage. European Journal of Operational Research 295, 838–856.

Khoshgehbari, F., Al-e, S.M.J.M., et al., 2023. Ambulance location routing problem considering all sources of uncertainty: Progressive estimating algorithm. Computers & Operations Research 160, 106400.

Koch, Z., Yuan, M., Bristow, E., 2020. Emergency response after disaster strikes: Agent-based simulation of ambulances in new windsor, ny. Journal of Infrastructure Systems 26, 06020001.

Kudenchuk, P.J., Leroux, B.G., Daya, M., Rea, T., Vaillancourt, C., Morrison, L.J., Callaway, C.W., Christenson, J., Ornato, J.P., Dunford, J.V., et al., 2017. Antiarrhythmic drugs for nonshockable-turned-shockable out-of-hospital cardiac arrest: the alps study (amiodarone, lidocaine, or placebo). Circulation 136, 2119–2131.

Lakhwani, T.S., Sinjana, Y., Kapoor, A.P., 2025. Beyond numbers: a novel iot and drone framework for enhancing healthcare logistics. Discover Health Systems 4, 19.

Lam, S., Zhang, Z., Oh, H., Ng, Y., Wah, W., Hock Ong, M., . on behalf of the cardiac arrest resuscitation epidemiology (care) study group (2014) reducing ambulance response times using discrete event simulation. Prehosp Emerg Care 18, 207–216.

Lam, S.S.W., Kumar, A., Yuan, G., Dorosan, M., Ong, M.E.H., Siddiqui, F.J., 2024. Simulation of emergency care systems: A taxonomy and future directions, in: 2024 Winter Simulation Conference (WSC), IEEE. pp. 882–893.

Lam, S.S.W., Zhang, J., Zhang, Z.C., Oh, H.C., Overton, J., Ng, Y.Y., Ong, M.E.H., 2015. Dynamic ambulance reallocation for the reduction of ambulance response times using system status management. The American journal of emergency medicine 33, 159–166.

Lamhaut, L., Hutin, A., Puymirat, E., Jouan, J., Raphalen, J.H., Jouffroy, R., Jaffry, M., Dagron, C., An, K., Dumas, F., et al., 2017. A pre-hospital extracorporeal cardio pulmonary resuscitation (ecpr) strategy for treatment of refractory out hospital cardiac arrest: an observational study and propensity analysis. Resuscitation 117, 109–117.

Laskowski, M., Mukhi, S., 2009. Agent-based simulation of emergency departments with patient diversion, in: Electronic Healthcare: First International Conference, eHealth 2008, London, UK, September 8-9, 2008. Revised Selected Papers 1, Springer. pp. 25–37.

Lewis, M., Stubbs, B.A., Eisenberg, M.S., 2013. Dispatcher-assisted cardiopulmonary resuscitation: time to identify cardiac arrest and deliver chest compression instructions. Circulation 128, 1522–1530.

Li, J., Baker, A.L., D'Ambrosi, G., Monuteaux, M.C., Chung, S., 2023. A statewide assessment of pediatric emergency care surge capabilities. Pediatrics 151, e2022059459.

Lim, C.S., Mamat, R., Braunl, T., 2011. Impact of ambulance dispatch policies on performance of emergency medical services. IEEE Transactions on Intelligent Transportation Systems 12, 624–632.

Liu, Y., Wei, R., Dong, H., Zhu, Y., Zhao, J., 2020. A designation of modular mobile reconfigurable platform system. Journal of Mechanics in Medicine and Biology 20, 2040006.

Lohrasbpoor, M., Arshadi Khamseh, A., Nemati-Lafmejani, R., Naderi, B., 2023.

Four echelons humanitarian network design considering capacitated/lateral transshipment with a destruction radius and abo compatibility: Tehran earthquake. International Journal of Supply and Operations Management 10, 1–22.

Luo, W., Yao, J., Mitchell, R., Zhang, X., Li, W., 2022. Locating emergency medical services to reduce urban-rural inequalities. Socio-Economic Planning Sciences 84, 101416.

MacFarlane, C., Benn, C., 2003. Evaluation of emergency medical services systems: a classification to assist in determination of indicators. Emergency Medicine Journal 20, 188–191.

Malik, A., Sharma, A., Saroha, V., 2018. Greedy algorithm. International Journal of Scientific and Research Publications 3, 2250–3153.

Mansourihanis, O., Maghsoodi Tilaki, M.J., Zaroujtaghi, A., Tayarani, M., Sheikhfarshi, S., 2024. Inequalities in emergency service accessibility: spatial analysis of urban infrastructure. Journal of Place Management and Development 17, 584–614.

Marinaro, J., Guliani, S., Dettmer, T., Pruett, K., Dixon, D., Braude, D., 2020. Out-of-hospital extracorporeal membrane oxygenation cannulation for refractory ventricular fibrillation: A case report. Journal of the American College of Emergency Physicians Open 1, 153–157.

Marini, M., Belleggia, S., Brugiatelli, L., Francioni, M., Battistoni, I., Shkoza, M., Pongetti, G., Angelini, L., Belfioretti, L., Matassini, M.V., 2024. The golden hour in shock management: do a lot, do it quickly. European Heart Journal Supplements 26, i78–i83.

Marla, L., Yue, Y., Ramayya, K., 2017. Data-driven omniscient bounds and greedy policies for ambulance allocation and dynamic redeployment. Available at SSRN 3009043 .

McLay, L.A., 2010. Emergency medical service systems that improve patient survivability. Wiley encyclopedia of operations research and management science .

McLay, L.A., Mayorga, M.E., 2013. A dispatching model for server-to-customer systems that balances efficiency and equity. Manufacturing & Service Operations Management 15, 205–220.

Michaels, D., Pham, H., Puckett, Y., Dissanaike, S., 2019. Helicopter versus ground ambulance: review of national database for outcomes in survival in transferred trauma patients in the usa. Trauma Surgery & Acute Care Open 4, e000211.

Mosadeghrad, A.M., Gebru, A.A., Sari, A.A., Getu, M.A., 2019. Emergency medical services in ethiopia: Drivers, challenges and opportunities. Human antibodies 27, 33–41.

Naji, H.Z., Al-Behadili, M., Kadim, M.S., 2024. A multi-objective facility coverage location problem for emergency medical service decisions in hajj. Baghdad Science Journal 21, 2110–2110.

Nandasena, G., Abeysena, C., 2019. Availability of facilities and equipment to provide emergency care in ambulances, which transport patients to the national hospital of sri lanka. Journal of the College of Community Physicians of Sri Lanka 24.

Nasrollahzadeh, A.A., Khademi, A., Mayorga, M.E., 2018. Real-time ambulance dispatching and relocation. Manufacturing & Service Operations Management 20, 467–480.

National Research Council, 1966. Accidental Death and Disability: The Neglected Disease of Modern Society. National Academies Press.

Newton, R., Kumara, S., Griffin, P., 2021. Emergency equity: Access and emergency medical services in san francisco. arXiv preprint arXiv:2109.03789 .

Nitzschke, R., Doehn, C., Kersten, J.F., Blanz, J., Kalwa, T.J., Scotti, N.A., Kubitz, J.C., 2017. Effect of an interactive cardiopulmonary resuscitation assist device with an automated external defibrillator synchronised with a ventilator on the cpr performance of emergency medical service staff: a randomised simulation study. Scandinavian Journal of Trauma, Resuscitation and Emergency Medicine 25, 1–10.

Nourinejad, M., Roorda, M.J., 2016. Agent based model for dynamic ridesharing. Transportation Research Part C: Emerging Technologies 64, 117–132.

NSW Ambulance, 2019. Nsw ambulance cardiac arrest registry—2017 report. https://www.ambulance.nsw.gov.au/__data/assets/pdf_file/0005/552452/NSW-Ambulance-Cardiac-Arrest-Registry-Report-2017.pdf. Accessed: 2021.

NSW Health, 2021. Cardiac arrest and defibrillators: A guide for consumers. URL: https://www.health.nsw.gov.au/cardiacarrest/Pages/consumers.aspx#:~:text=Over%208%2C500%20people%20experience%20an,Only%2012%25%20will%20survive.

OpenStreetMap contributors, 2017. Planet dump retrieved from https://planet.osm.org . https://www.openstreetmap.org.

Parker, S.A., Kus, T., Bowry, R., Gutierrez, N., Cai, C., Yamal, J.M., Rajan, S., Wang, M., Jacob, A.P., Souders, C., et al., 2020. Enhanced dispatch and rendezvous doubles the catchment area and number of patients treated on a mobile stroke unit. Journal of Stroke and Cerebrovascular Diseases 29, 104894.

Pena, M., Neu, D.T., Feng, H.A., Hammond, D.R., Mead, K.R., Banerjee, R.K., 2025. Use of portable air cleaners within an ambulance workspace. Journal of Occupational and Environmental Hygiene , 1–14.

Petermichl, W., Philipp, A., Hiller, K.A., Foltan, M., Floerchinger, B., Graf, B., Lunz, D., 2021. Reliability of prognostic biomarkers after prehospital extracorporeal cardiopulmonary resuscitation with target temperature management. Scandinavian Journal of Trauma, Resuscitation and Emergency Medicine 29, 1–12.

Peyravi, M., Örtenwal, P., Djalali, A., Khorram-Manesh, A., 2013. An overview of shiraz emergency medical services, dispatch to treatment. Iranian Red Crescent Medical Journal 15, 823.

Pillac, V., Gendreau, M., Guéret, C., Medaglia, A.L., 2013. A review of dynamic vehicle routing problems. European Journal of Operational Research 225, 1–11.

Qi, P., Wang, Y., Lin, ·.X., 2023. Optimization of reverse logistics network for medical waste recycling. Journal of Data, Information and Management 5, 71–88.

Rabbani, M., Oladzad-Abbasabady, N., Akbarian-Saravi, N., 2022. Ambulance routing in disaster response considering variable patient condition: Nsga-ii and mopso algorithms. Journal of Industrial & Management Optimization 18.

Randhawa, A.S., Pariona-Vargas, F., Starkman, S., Sanossian, N., Liebeskind, D.S., Avila, G., Stratton, S., Gornbein, J., Sharma, L., Restrepo-Jimenez, L., et al., 2022. Beyond the golden hour: Treating acute stroke in the platinum 30 minutes. Stroke 53, 2426–2434.

Reuter-Oppermann, M., van den Berg, P.L., Vile, J.L., 2017. Logistics for emergency medical service systems. Health Systems 6, 187–208.

Reynolds, J.C., Frisch, A., Rittenberger, J.C., Callaway, C.W., 2013. Duration of resuscitation efforts and functional outcome after out-of-hospital cardiac arrest: when should we change to novel therapies? Circulation 128, 2488–2494.

Reynolds, J.C., Grunau, B.E., Elmer, J., Rittenberger, J.C., Sawyer, K.N., Kurz, M.C., Singer, B., Proudfoot, A., Callaway, C.W., 2017. Prevalence, natural history, and time-dependent outcomes of a multi-center north american cohort of out-of-hospital cardiac arrest extracorporeal cpr candidates. Resuscitation 117, 24–31.

Ribeiro, D.P., De Oliveira, U.R., da Silva César, A., Aprigliano Fernandes, V., 2021. Evaluation of medicine reverse logistics practices in hospitals. Sustainability 13, 3496.

Richardson, A.S.C., Tonna, J.E., Nanjayya, V., Nixon, P., Abrams, D.C., Raman, L., Bernard, S., Finney, S.J., Grunau, B., Youngquist, S.T., et al., 2021. Extracorporeal cardiopulmonary resuscitation in adults. interim guideline consensus statement from the extracorporeal life support organization. Asaio Journal 67, 221–228.

Ridler, S., Mason, A.J., Raith, A., 2022. A simulation and optimisation package for emergency medical services. European Journal of Operational Research 298, 1101–1113.

Sasaki, S., Comber, A.J., Suzuki, H., Brunsdon, C., 2010. Using genetic algorithms to optimise current and future health planning-the example of ambulance locations. International journal of health geographics 9, 1–10.

Schmid, V., 2012. Solving the dynamic ambulance relocation and dispatching problem using approximate dynamic programming. European journal of operational research 219, 611–621.

Scquizzato, T., Bonaccorso, A., Consonni, M., Scandroglio, A.M., Swol, J., Landoni, G., Zangrillo, A., 2022. Extracorporeal cardiopulmonary resuscitation for out-of-hospital cardiac arrest: a systematic review and meta-analysis of randomized and propensity score-matched studies. Artificial Organs 46, 755–762.

Silva, F., Serra, D., 2008. Locating emergency services with different priorities: the priority queuing covering location problem. Journal of the Operational Research Society 59, 1229–1238.

Singer, B., Reynolds, J.C., Davies, G.E., Wrigley, F., Whitbread, M., Faulkner, M., O'Brien, B., Proudfoot, A.G., Mathur, A., Evens, T., et al., 2020. Sub30: Protocol for the sub30 feasibility study of a pre-hospital extracorporeal membrane oxygenation (ecmo) capable advanced resuscitation team at achieving blood flow within 30 min in patients with refractory out-of-hospital cardiac arrest. Resuscitation Plus 4, 100029.

Singer, B., Reynolds, J.C., Lockey, D.J., O'Brien, B., 2018. Pre-hospital extracorporeal cardiopulmonary resuscitation. Scandinavian Journal of Trauma, Resuscitation and Emergency Medicine 26, 1–8.

Song, C., Dennis, M., Burns, B., Dyson, S., Forrest, P., Ramanan, M., Levinson, D., Moylan, E., 2022. Improving access to extracorporeal membrane oxygenation for out of hospital cardiac arrest: pre-hospital ecpr and alternate delivery strategies. Scandinavian Journal of Trauma, Resuscitation and Emergency Medicine 30, 1–10.

Sporer, K.A., Youngblood, G.M., Rodriguez, R.M., 2007. The ability of emergency medical dispatch codes of medical complaints to predict als prehospital interventions. Prehospital emergency care 11, 192–198.

Stratman, E.G., Boutilier, J.J., Albert, L.A., 2023. Uncertainty in facility location models for emergency medical services. Uncertainty in Facility Location Problems , 213–250.

Su, S., Shih, C.L., 2003. Modeling an emergency medical services system using computer simulation. International journal of medical informatics 72, 57–72.

Sutherland, M., Chakrabortty, R.K., 2023. An optimal ambulance routing model using simulation based on patient medical severity. Healthcare analytics , 100256.

Ter Avest, E., Carenzo, L., Lendrum, R., Christian, M., Lyon, R., Coniglio, C., Rehn, M., Lockey, D., Perkins, Z., 2022. Advanced interventions in the prehospital resuscitation of patients with non-compressible haemorrhage after penetrating injuries. Critical Care 26, 184.

Tikani, H., Setak, M., 2019. Ambulance routing in disaster response scenario considering different types of ambulances and semi soft time windows. Journal of Industrial and Systems Engineering 12, 95–128.

Tluli, R., Badawy, A., Salem, S., Barhamgi, M., Mohamed, A., 2024. A survey of machine learning innovations in ambulance services: Allocation, routing, and demand estimation. IEEE Open Journal of Intelligent Transportation Systems .

Toregas, C., Swain, R., ReVelle, C., Bergman, L., 1971. The location of emergency service facilities. Operations research 19, 1363–1373.

VK, S., Roy, R.B., 2017. Equity-constrained dispatching models for emergency medical services. Team Performance Management: An International Journal 23, 28–45.

Wachs, M., Kumagai, T.G., 1973. Physical accessibility as a social indicator. Socio-Economic Planning Sciences 7, 437–456.

Wajid, S., Nezamuddin, N., 2022. A robust survival model for emergency medical services in delhi, india. Socio-Economic Planning Sciences 83, 101342.

Wang, W., Wu, S., Wang, S., Zhen, L., Qu, X., 2021. Emergency facility location problems in logistics: Status and perspectives. Transportation Research Part E: Logistics and Transportation Review 154, 102465.

Wang, X., Zhao, J., Cheng, C., Qi, M., 2023. A multi-objective fuzzy facility location problem with congestion and priority for drone-based emergency deliveries. Computers & Industrial Engineering 179, 109167.

Wang, Y., Luangkesorn, K.L., Shuman, L., 2012. Modeling emergency medical response to a mass casualty incident using agent based simulation. Socio-Economic planning sciences 46, 281–290.

Wengenmayer, T., Rombach, S., Ramshorn, F., Biever, P., Bode, C., Duerschmied, D., Staudacher, D.L., 2017. Influence of low-flow time on survival after extracorporeal cardiopulmonary resuscitation (ecpr). Critical care 21, 1–6.

World Health Organization, n.d.a. Emergency and critical care. https://www.who.int/teams/integrated-health-services/clinical-services-and-systems/emergency-and-critical-care. Accessed: 2024-09-20.

World Health Organization, n.d.b. Emergency medical teams (emt). https://www.who.int/westernpacific/wpro-emergencies/response/emergency-medical-teams. Accessed: 2024-09-20.

Wu, H., Levinson, D., 2020. Unifying access. Transportation Research Part D: Transport and Environment 83, 102355.

Yang, X.S., 2020. Nature-inspired optimization algorithms. Academic Press.

Yannopoulos, D., Bartos, J., Raveendran, G., Walser, E., Connett, J., Murray, T.A., Collins, G., Zhang, L., Kalra, R., Kosmopoulos, M., et al., 2020a. Advanced reperfusion strategies for patients with out-of-hospital cardiac arrest and refractory ventricular fibrillation (arrest): a phase 2, single centre, open-label, randomised controlled trial. The lancet 396, 1807–1816.

Yannopoulos, D., Bartos, J., Raveendran, G., Walser, E., Connett, J., Murray, T.A., Collins, G., Zhang, L., Kalra, R., Kosmopoulos, M., et al., 2020b. Advanced reperfusion strategies for patients with out-of-hospital cardiac arrest and refractory ventricular fibrillation (arrest): a phase 2, single centre, open-label, randomised controlled trial. The lancet 396, 1807–1816.

Yılmaz, C., Livatyalı, H., Gür, E., Arslan, M.S., 2025. Design of a motorcycle ambulance with an integrated passenger seat mechanism for emergency response. Proceedings of the Institution of Mechanical Engineers, Part D: Journal of Automobile Engineering , 09544070241312959.

Yoon, S., Albert, L.A., White, V.M., 2021. A stochastic programming approach for locating and dispatching two types of ambulances. Transportation Science 55, 275–296.

Zailani, M., Azma, R., Aniza, I., Rahana, A., Ismail, M., Shahnaz, I., Chan, K., Jamaludin, M., Mahdy, Z., 2021. Drone versus ambulance for blood products transportation: an economic evaluation study. BMC health services research 21, 1–10.

Zhu, L., Ding, T., Liu, Z., 2024. Reverse logistics network design for medical waste disposal under the scenario of uncertain proposal demand. Sustainability 16, 2996.