UNIVERSITY OF CALIFORNIA,
IRVINE

Leveraging Geospatial Systems for Disaster Resilience in Healthcare

DISSERTATION

submitted in partial satisfaction of the requirements
for the degree of

DOCTOR OF PHILOSOPHY

in Networked Systems

by

Modeste Mefenya Kenne

Dissertation Committee:
Professor Nalini Venkatasubramanian, Chair
Professor Ardalan Amiri Sani
Professor Sharad Mehrotra
Professor Marco Levorato
Doctor Julie Rousseau

2026

DEDICATION

To my wife and family, whose love, sacrifice, and support carried me through this journey.

Contents

List of Figures

List of Tables

LIST OF ALGORITHMS

ACKNOWLEDGMENTS

First and foremost, I thank Almighty God for His grace, guidance, and protection throughout this journey. This dissertation would not have been possible without His blessings and the many people He placed in my life along the way.

My deepest gratitude goes to my advisor, Professor Nalini Venkatasubramanian, for her exceptional mentorship, trust, encouragement, and support throughout my Ph.D. journey. Her guidance provided me with the opportunity to grow as a researcher and to pursue work at the intersection of computer science, healthcare, emergency management, and community resilience. Through CareDEX and related projects, she created opportunities that truly shaped the direction of this dissertation. I am deeply grateful to the entire CareDEX team; their collective efforts, insights, and dedication were vital to the success of this research.

I am also grateful to my committee members, Professor Sharad Mehrotra, Professor Marco Levorato, Professor Ardalan Amiri Sani, and Dr. Julie Rousseau, for their time and valuable feedback. I would like to thank Dr. Prasanna Date, my research mentor at Oak Ridge National Laboratory, whose guidance and technical insights were instrumental in the development of my first research paper. I am also thankful to Professor Shangping Ren for her guidance and feedback.

This work benefited immensely from the collaboration and support of many partners and domain experts. I am thankful to our collaborators at ImageCat, UCI Health, the Orange County Health Care Agency (OCHCA), and our senior health facility partners for the invaluable domain knowledge they provided. I would especially like to thank Chad Cossey for his guidance on emergency preparedness and ZhengHui Hu for his continuous support.

I gratefully acknowledge the financial support that made this research possible, including the National Science Foundation under Grant Nos. 2044107, 2133391, 2008993, 2245372, 2420846, and 1952247, as well as the DARPA TIPPERS program under Contract No. FA8750-16-2-0021. Most notably, I extend my profound gratitude to the Eugene Cota-Robles Fellowship for providing the primary financial support that sustained me throughout my doctoral journey.

To Nailah Saleh Alhassoun, Fangqi Liu, Guoxi Wang, Mowhebat Bazargani, and Elahe Khatibi, thank you for supporting me during the early stages of my Ph.D. journey. I also extend my gratitude to the members of the Distributed Systems Middleware group, including Rummana Rahman, Rahul Bhope, Jared Macshane, Tung Chun Chan, Ryan Hildebrant, Yuqiao Li, Andrew Chio, and Fernanda Ventorim. Their collaboration, technical discussions, and encouragement made this work stronger. A special thank you goes to Andrew for the many late nights in the lab working together on papers, and to Fernanda for her immense help with experiments. I would also like to thank Surabhi Honnavalli and Mahika Chimata for their help in implementing an early prototype of BlackOps.

My path to UC Irvine was shaped by mentors who believed in me long before graduate school. To Ashley Garrin, Thelma Harding, Lynn Lundy, Jack Polifka, Erin Todey, Arnold Woods, and the rest of the McNair team, thank you for being part of the support system that propelled me toward graduate school. I would also like to thank Nancy Woods from Des Moines Area Community College for her encouragement during a pivotal stage of my academic path.

To my children, Brial, Brielle, and Brianna: as you grow and one day read this, I want you to know that you are my constant source of joy and motivation. Even during the most demanding moments, you reminded me why this journey mattered.

To my family, especially my parents, Mefenya and Ngangmo, as well as Nembot and Lagne, thank you for your love, prayers, sacrifices, and unwavering belief in me. Your endless support paved the way for my education, and I could not have crossed the finish line without you. This milestone truly belongs to all of us, and I hope it brings you the immense pride and joy you so richly deserve.

I am also deeply grateful to Nina Rasmusson for her care, encouragement, and steady support over the years. You have watched me grow into the person I am today, and I am thankful for the many ways you have supported me.

To my dearest wife, Arolle Metsefang, your love, patience, sacrifice, and constant encouragement carried me through the long days, late nights, uncertainty, and many challenges of this Ph.D. journey. *Chérie Coco*, you believed in me when the path was difficult and gave me the strength to keep going. This accomplishment is not mine alone; it is a milestone we reached together.

Finally, to my friends and to everyone who supported me directly or indirectly along the way, thank you!

VITA

Modeste Mefenya Kenne

EDUCATION

Doctor of Philosophy in Networked Systems **2026**
University of California, Irvine *Irvine, California*

Bachelor of Science in Computer Engineering **2020**
Iowa State University *Ames, Iowa*

Associate of Arts in Marketing **2012**
University of Dschang *Dschang, Cameroon*

RESEARCH EXPERIENCE

Graduate Research Assistant **2020–2026**
University of California, Irvine *Irvine, California*

TEACHING EXPERIENCE

Teaching Assistant **2021–2026**
University of California, Irvine *Irvine, California*

PROFESSIONAL EXPERIENCE

Research Intern **Summer 2022**
Oak Ridge National Laboratory *Oak Ridge, Tennessee*

REFEREED CONFERENCE PUBLICATIONS

BlackOps: GIS-Informed Planning for Healthcare Resilience Amid Power Grid Outages — **Jun 2026**
IEEE International Conference on Smart Computing (SMARTCOMP)

RADAR: Resource Allocation for Disaster Resilience in Senior Health Care — **Dec 2025**
ACM International Conference on Advances in Geographic Information Systems (SIGSPATIAL)

iFair: Achieving Fairness in the Allocation of Scarce Resources for Senior Health Care — **Jun 2024**
IEEE International Conference on Smart Computing

REFEREED JOURNAL PUBLICATIONS

Testing CareDEX: SAGAT Methodology Deployed for a Disaster Drill at a Senior Living Community — **Oct 2023**
Issues in Information Systems (IACIS)

ABSTRACT OF THE DISSERTATION

Leveraging Geospatial Systems for Disaster Resilience in Healthcare

By

Modeste Mefenya Kenne

Doctor of Philosophy in Networked Systems

University of California, Irvine, 2026

Professor Nalini Venkatasubramanian, Chair

Healthcare systems face significant resilience challenges during disasters, as hazards can simultaneously disrupt facilities, transportation networks, utilities, staffing, and resource availability. These challenges are especially acute in senior healthcare, where older adults have heterogeneous medical profiles, mobility limitations, equipment dependencies, and continuity-of-care requirements. This dissertation frames healthcare disaster resilience as a multi-scale geospatial decision-support problem, in which computational models integrate heterogeneous data about people, facilities, hazards, infrastructure, policies, and scarce resources to support efficient, fair, and operationally actionable decisions.

This perspective is explored through three connected systems. First, at the facility scale, we develop *iFair*, a fairness-aware resource allocation framework for senior healthcare facilities. In this work, residents, caregivers, equipment, care tasks, facility layouts, and dynamic emergency conditions are modeled through a digital-twin-inspired representation. Allocation is formulated as a constrained scheduling problem that balances task efficiency, resident criticality, individualized needs, and preference satisfaction during urgent events such as evacuation. The results show that accounting for resident needs and preferences can improve both fairness and operational efficiency.

Second, at the regional scale, we develop *RADAR*, a geospatial decision-support framework

for disaster response. The framework integrates multisource GIS feeds, facility status, transportation conditions, resource availability, and policy-driven ranking functions to support evacuation, relocation, and resource-sharing decisions. It formulates regional coordination as a coupled matching-and-routing problem, combining stable matching with hazard-aware routing to generate interpretable recommendations for Emergency Operations Centers and healthcare stakeholders. This enables regional plans that account for facility compatibility, capacity, route safety, hazard evolution, and operational policy constraints.

Third, at the infrastructure and community scale, we develop *BlackOps*, a GIS-informed preparedness framework for healthcare resilience amid prolonged power outages. The framework models facility autonomy, backup-power shortfall, infrastructure dependencies, road-network accessibility, hazard exposure, social vulnerability, and regulatory requirements to identify facilities most at risk of losing critical services during multi-day outages. It also uses bounded retrieval-augmented language-model workflows to extract facility-specific autonomy requirements from unstructured regulatory documents, while keeping risk computation grounded in geospatial analysis, infrastructure modeling, and interpretable preparedness indices.

Together, these contributions demonstrate how geospatial systems can move beyond visualization to serve as computational foundations for healthcare disaster resilience. Across facility-level allocation, regional matching and routing, and infrastructure-aware preparedness planning, this dissertation shows that spatially grounded and human-centered models can help decision-makers reason more effectively about vulnerable populations, scarce resources, cascading disruptions, and time-sensitive operational constraints.

Chapter 1

Introduction

Disasters expose a central challenge for healthcare systems: sustaining care and making critical decisions when infrastructure is disrupted, information is incomplete, and resources are scarce. During wildfires, earthquakes, floods, heat waves, and prolonged utility failures, healthcare stakeholders must determine which facilities can continue operating safely, which patients or residents should receive limited attention first, where medically fragile people can be safely relocated, and how scarce resources should move through disrupted transportation and service networks. These decisions are difficult not only because they are urgent, but also because they are computationally complex. They require reasoning over people, facilities, hazards, infrastructure, and policies whose states evolve and whose interactions are strongly spatial. Recent federal hazard accounting highlights the growing operational pressure that extreme events place on critical services. NOAA's National Centers for Environmental Information reports 403 U.S. weather and climate disasters with losses exceeding $1 billion each from 1980–2024, with an annual average of 23 such events during 2020–2024 alone [124, 125]. Large-scale disruptions are therefore not rare edge cases in modern operations, but recurring stressors that increasingly test the resilience of health and emergency systems.

This dissertation studies healthcare disaster resilience through decision support grounded in geospatial systems. The central premise is that resilience in healthcare is not only a planning problem, a logistics problem, or a visualization problem. It is a computational problem that requires spatial-temporal data management, human-centered modeling, optimization, routing, and interpretable decision support to work together. Computational methods are not intended to replace decision-makers. Instead, they are designed to help stakeholders reason more effectively under uncertainty, time pressure, and incomplete information.

The stakes of this problem are visible in real disasters. In the 2018 Camp Fire in California, older adults accounted for a large share of fatalities [116]. After Hurricane Katrina, a mortality analysis of Louisiana victims found that 49% were age 75 or older, making that age group the most affected population cohort [21]. During Washington State's 2021 heat wave, 106 of 157 reported heat-related deaths were among people age 65 and older [157]. Internationally, epidemiological analysis of the 2003 France heat wave found that 11,731 of 14,729 excess deaths occurred among people age 75 and older [69]. These events show that healthcare disruption is not merely an infrastructure problem. It is also a human problem concentrated among older adults and other medically fragile populations whose safety depends on timely, compatible, and spatially informed decisions.

Across these settings, the central decision-support question is both spatial and temporal: where are the people, hazards, resources, facilities, and infrastructure dependencies, and how quickly must decisions be made before conditions change? The answer may shape which resident receives assistance first, which destination can safely receive evacuees, which route remains usable, or which healthcare facility is most likely to lose critical services during a prolonged outage. This spatial-temporal view motivates the approach developed in this dissertation.

We focus particularly on senior health care because it is both societally important and technically demanding. Older adults often have complex medical needs, reduced mobility, cognitive impairments, specialized equipment dependencies, and continuity-of-care requirements that make generic disaster-response strategies inadequate. In these settings, poor decisions can produce cascading harm through delayed evacuation, inappropriate relocation, inaccessible destinations, insufficient backup services, and prolonged care disruption. As a result, healthcare disaster resilience for older adults requires decision-support methods that are personalized, spatially aware, adaptive, and interpretable.

1.1 Healthcare Disaster Resilience as a Computational Problem

Healthcare disaster resilience involves maintaining or restoring the ability of care systems to protect life and sustain essential services under hazardous and degraded conditions. In practice, this requires decisions that connect care delivery, transportation, infrastructure, and resource availability. Staff, equipment, and time must be assigned to people whose clinical needs, mobility constraints, and resource requirements differ substantially and may change during an event. Healthcare facilities may need to coordinate relocation, mutual aid, and resource sharing across sites with different capabilities, capacities, accessibility conditions, and hazard exposure. Agencies must also identify facilities that are likely to lose critical services during prolonged disruptions and determine where preparedness investments or pre-staged support can reduce risk most effectively.

These problems are technically challenging because the decision environment is heterogeneous, dynamic, resource-constrained, and spatial. People differ in acuity, mobility, preferences, and resource requirements. Facilities differ in service levels, operational status, available capacity, and infrastructure dependencies. Hazards differ in footprint, severity, and temporal evolution. At the same time, viable actions may change as roads close, smoke spreads, utility services fail, or facility status updates arrive. Decision-makers rarely have the luxury of optimizing with perfect knowledge; they must act using partial situational awareness while balancing efficiency, safety, fairness, and operational feasibility. Because hazard zones, facilities, roads, power infrastructure, and vulnerable communities are geographically distributed, response quality depends critically on the ability to integrate and reason over spatial context.

A central argument of this dissertation is that healthcare disaster resilience should be approached as a geospatial decision-support problem. This view treats maps, routes, infrastructure layers, facility states, and hazard footprints not as presentation aids, but as core computational objects. Under this framing, spatial data become part of the logic of decision-making itself: they determine which facilities are reachable, which reallocations are feasible, which routes are safe, and which sites are structurally or socially vulnerable under different scenarios. This computational perspective motivates the use of GIS-backed data systems, digital-twin-inspired facility models, and optimization methods that explicitly incorporate evolving spatial constraints.

1.2 Why Senior Health Care Requires Human-Centric Models

Senior health care presents a particularly important case for disaster resilience because the population being served is highly heterogeneous and often medically fragile. Older adults cannot be modeled adequately as interchangeable demand units. Their needs depend on clinical condition, functional limitations, equipment dependence, medication requirements, and social or behavioral factors. Two individuals may face the same hazard, yet differ substantially in the level of assistance needed, the suitability of possible destinations, or the consequences of delayed care. For this reason, one-size-fits-all disaster logistics are poorly aligned with the realities of senior care.

Demographic trends make this problem larger and geographically distributed. The U.S. Census Bureau reports that the population age 65 and older reached 61.2 million in 2024, and that older adults outnumber children in 11 states and nearly half of U.S. counties [148]. At the same time, long-term care is delivered through a large and distributed facility ecosystem. In 2022, the United States had 14,700 nursing homes serving about 1.2 million residents, and 32,200 residential care communities serving about 988,800 residents [120, 121]. Figure 1.1 illustrates this broader demographic shift, showing the projected growth of the older adult population relative to children in the United States. This combination of demographic aging and distributed care infrastructure amplifies the consequences of mismatches between individual needs, facility capabilities, and disaster-time constraints.

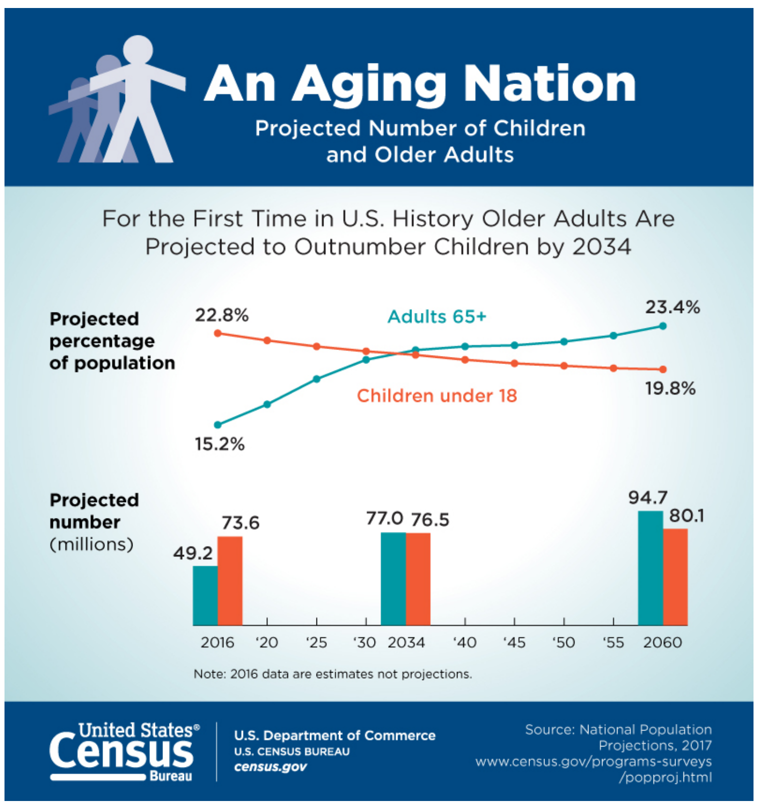

Figure 1.1: Projected demographic shift in the United States, highlighting the growth of the population age 65 and older relative to children.

These conditions motivate human-centric computational models. In this context, human-centricity does not simply mean that humans remain in the loop. The computational representation itself must encode clinically meaningful distinctions across individuals and care settings. Allocation decisions should account for differences in urgency and vulnerability rather than assuming homogeneous demand. Relocation decisions should respect care compatibility, residency type, and continuity needs rather than relying only on distance or bed counts. Risk assessments should reflect the operational consequences of infrastructure failure for medically dependent populations rather than measuring exposure alone.

A human-centric view also complicates traditional optimization objectives. In many classical settings, efficiency is measured through throughput, utilization, completion time, or travel time. In senior health care disaster settings, these objectives remain important, but they are not sufficient. A system that minimizes average relocation time while repeatedly assigning incompatible destinations or systematically disadvantaging higher-need residents is operationally unacceptable. Accordingly, fairness, compatibility, and explainability are treated as first-class design requirements rather than afterthoughts.

1.3 Why Geospatial Systems Matter

The decisions studied in this dissertation are inherently geospatial. Facilities are located in hazard fields and communities with unequal vulnerability. Residents and resources must move along transportation networks that may become congested, degraded, or impassable. Power outages propagate through infrastructure dependencies that are regionally distributed rather than confined to a single building. For these reasons, healthcare disaster resilience requires more than static facility records or isolated optimization modules. It requires systems that can ingest, organize, query, and reason over spatially distributed information in a form usable by operational stakeholders.

This view is reflected in federal operational tooling. FEMA's National Risk Index provides baseline measurements of expected annual loss, social vulnerability, and community resilience at county and Census-tract scales, helping identify communities most at risk from natural hazards [58, 59]. These data can also be accessed and interpreted through FEMA's Resilience Analysis and Planning Tool (RAPT), which supports GIS-backed planning by combining risk, infrastructure, population, and resilience-related layers. Figure 1.2 illustrates this type of GIS-backed hazard decision support, where National Risk Index information shown through RAPT can be queried, visualized, and interpreted across regions. Such systems demonstrate the value of spatial decision support, but they also highlight an important gap for healthcare operations: broad regional risk layers do not by themselves determine how to allocate scarce resources within facilities, where medically fragile people should be relocated, or which facilities may lose critical services under infrastructure failure. These questions require tighter coupling between spatial analytics and operational decision models.

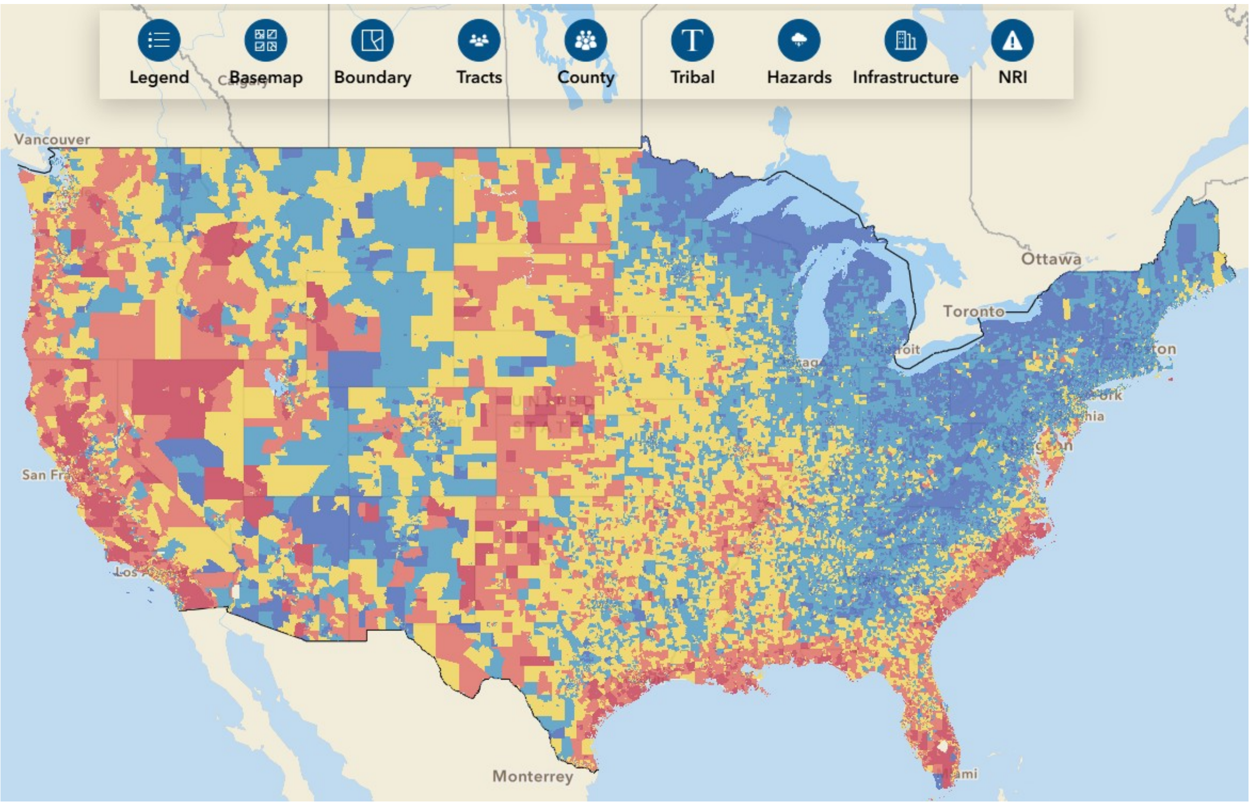

Figure 1.2: FEMA National Risk Index information shown through the Resilience Analysis and Planning Tool (RAPT), illustrating GIS-backed hazard decision support for regional risk, vulnerability, and resilience planning [58, 59].

Spatial data systems are especially valuable because they integrate information that is often fragmented across agencies, tools, and formats. Hazard products, facility data, transportation conditions, utility dependencies, environmental context, and community-level indicators can all be represented as interoperable spatial layers. Once unified, these layers can support computations such as nearest-feasible facility search, route filtering under disruption, dynamic matching under shifting hazard boundaries, and vulnerability assessment across many scenarios. In this dissertation, geospatial systems are therefore treated as both data infrastructure and decision substrate: they provide the representational and computational foundation needed to move from raw inputs to actionable recommendations.

However, spatial reasoning alone is not enough. A map can show where hazards and facilities are located, but it does not by itself determine how scarce resources should be allocated, which pairings satisfy operational policies, or how facility risk should be quantified and prioritized. To become useful in disaster operations, geospatial systems must be combined with formal models that can express fairness, compatibility, capacity, routing feasibility, infrastructure dependency, and uncertainty.

1.4 Core Research Challenges

This dissertation addresses a set of connected research challenges that arise when healthcare systems must make decisions under scarcity, uncertainty, and disruption. The common difficulty is that these decisions are not purely logistical. They involve vulnerable people, constrained resources, changing hazards, degraded infrastructure, and operational policies that must be interpreted under time pressure.

One challenge is fair and efficient allocation during urgent care operations. In senior care settings, staff, equipment, and time must be assigned while accounting for residents with different medical conditions, mobility limitations, preferences, task requirements, and levels of urgency. The objective is not only to complete tasks quickly, but also to ensure that allocation decisions reflect resident needs and do not treat medically fragile individuals as interchangeable demand units. This requires models that reason jointly about people, care tasks, resources, and the physical layout of the care environment.

Another challenge is coordination across a disrupted healthcare network. During disasters, relocation and resource-sharing decisions cannot be separated from routing decisions. A destination may have available capacity, but still be inappropriate if it lacks the needed services, is operationally degraded, or cannot be reached safely. Similarly, a short route may not be useful if it leads to a facility that cannot support the people being relocated. This creates a coupled matching-and-routing problem in which hazard evolution, road conditions, facility status, resource availability, and operational policies must be considered together.

A third challenge is identifying facilities that are likely to lose critical services during prolonged infrastructure failures. In multi-day outages, exposure to a hazard is only one part of the problem. Planners also need to understand which facilities may exhaust backup capacity, experience delayed resupply, depend on vulnerable grid components, or serve communities with limited recovery capacity. Addressing this challenge requires combining facility autonomy, regulatory requirements, power-grid dependencies, transportation-based resupply difficulty, backup shortfall, and community vulnerability into interpretable risk characterizations.

Together, these challenges motivate a geospatial decision-support approach that connects urgent care delivery, cross-facility coordination, and risk assessment for prolonged disruptions.

1.5 Research Scope and Contributions

This dissertation develops a human-centric geospatial approach for healthcare disaster resilience. The work focuses on three connected decision contexts: assigning scarce resources to older adults during urgent in-facility events, coordinating relocation and resource sharing across healthcare facilities during disasters, and identifying facilities at risk of losing critical services during prolonged power disruptions. Across these contexts, the central claim is that healthcare disaster resilience can be improved by combining geospatial grounding, human-centered modeling, and interpretable algorithmic decision support. The main contributions are as follows:

iFair: human-centric allocation in senior health facilities. We develop a fine-grained resource allocation framework for senior health facilities. The model represents residents, caregivers, equipment, care tasks, facility layouts, and changing emergency conditions within a digital-twin-inspired structure. It introduces fairness-aware allocation methods that balance task efficiency with resident criticality, individualized needs, and preference satisfaction during time-sensitive events such as evacuation.

RADAR: policy-aware matching and routing for disaster response. We extend allocation from individual facilities to cross-facility coordination. This geospatial decision-support framework integrates hazard conditions, facility status, transportation constraints, resource availability, and policy-driven ranking. By coupling weighted policy scoring with stable matching and hazard-aware routing, it supports evacuation, relocation, and resource-sharing decisions across heterogeneous healthcare facilities.

BlackOps: GIS-informed risk assessment for prolonged power disruptions. We develop a risk assessment framework for identifying healthcare facilities that are most likely to lose critical services during multi-day power outages. The model combines facility-level autonomy requirements, electric-grid dependencies, hazard exposure, transportation-based resupply difficulty, backup shortfall, and social vulnerability into interpretable risk indices and facility scorecards. It also explores the bounded use of retrieval-augmented generation to extract information from unstructured regulatory documents.

Together, these contributions show how geospatial systems can serve as a computational foundation for healthcare disaster resilience.

1.6 Dissertation Organization

The remainder of this dissertation is organized as follows. Chapter 2 reviews prior work on healthcare disaster resilience, including resource allocation, fairness, geospatial decision support, regional matching and routing, infrastructure resilience, and bounded AI-assisted decision support. Chapter 3 presents the unified problem statement and spatial-temporal approach that connects the three frameworks developed in this dissertation. Chapter 4 focuses on fair allocation of scarce resources within senior health facilities, where resident needs, preferences, care tasks, and facility layouts shape urgent response decisions. Chapter 5 extends the problem to cross-facility coordination by modeling policy-driven matching and hazard-aware routing during disasters. Chapter 6 shifts to prolonged power disruptions and develops a GIS-informed risk assessment framework for identifying healthcare facilities that may lose critical services. Chapter 7 concludes with key takeaways and future directions.

Chapter 2

Related Work

This chapter reviews prior work through the lens of healthcare disaster resilience across spatial and temporal scales. The dissertation draws from several research traditions that are often studied separately, including healthcare scheduling, fairness under scarcity, geospatial decision support, evacuation and resource coordination, infrastructure resilience, and AI-assisted knowledge extraction. Rather than organizing the literature around the three systems developed later, the discussion is organized around the kinds of decisions healthcare stakeholders must make before, during, and after disasters.

Within facilities, decision-makers must allocate limited staff, equipment, and time to people whose clinical needs, mobility constraints, preferences, and urgency differ. Across healthcare networks, stakeholders must coordinate relocation, resource sharing, and routing while accounting for facility capabilities, road conditions, hazards, and operational policies. During prolonged disruptions, planners must identify which facilities are likely to lose critical services and determine where resilience investments or pre-staged support can reduce risk. Existing work provides strong foundations for many parts of this problem, including resource allocation, fairness metrics, GIS-based situational awareness, evacuation planning, infrastructure vulnerability assessment, and AI-assisted decision support. However, healthcare disaster resilience for older adults requires these elements to be considered together because care delivery, regional coordination, and infrastructure risk are tightly connected during real disasters.

2.1 Healthcare Resilience Across Spatial and Temporal Scales

Healthcare disaster resilience involves maintaining or restoring essential care when hazards disrupt facilities, infrastructure, transportation networks, and resource availability. Rather than being a single optimization problem, it consists of connected decisions that unfold at different operational levels and time horizons. A senior healthcare facility may need to determine which residents should receive assistance first during an evacuation. An Emergency Operations Center may need to identify which facilities can receive evacuees, where scarce resources should be routed, and which roads remain usable. Regional planners may also need to identify facilities that are likely to lose critical services during prolonged infrastructure failure and decide where resilience investments or pre-staged support can reduce risk.

These decisions also unfold over different time horizons. In-facility response may occur within minutes to hours, especially when staff must assist residents with mobility limitations, cognitive impairments, or equipment dependencies. Regional coordination often unfolds over hours to days as hazards evolve, facility status changes, transportation networks degrade, and demand for resources shifts. Longer-horizon resilience assessment relies on scenario analysis, hazard models, and vulnerability indicators to identify weaknesses before an event occurs. Thus, healthcare resilience involves a family of related decision problems that differ in geography, urgency, data availability, and operational constraints.

Disaster medicine and emergency preparedness frameworks recognize that normal healthcare delivery may need to adapt when demand overwhelms available capacity. Crisis Standards of Care (CSC) provide system-level guidance for allocating care when usual operations cannot be maintained, emphasizing consistency, transparency, proportionality, and accountability [88]. In acute shortage settings, ethical guidance similarly emphasizes the need to balance benefit maximization with protection of vulnerable populations [51]. These principles are important, but they do not by themselves specify how decision-support systems should integrate spatial constraints, heterogeneous facility capabilities, route feasibility, infrastructure dependencies, or individual care needs.

Older adults and medically fragile populations make this problem especially difficult. During disasters, evacuation, relocation, sheltering, and service disruption can all create additional harm if they are not matched to individual needs. Empirical studies of nursing home residents exposed to major hurricanes report increases in mortality and hospitalizations, suggesting that evacuation decisions can worsen outcomes for frail residents when poorly planned or poorly matched to care requirements [47, 48, 158]. Related work on assisted living residents during Hurricane Irma shows that evacuation was associated with increased short-term emergency department use and nursing home placement, further illustrating that relocation itself can become a clinical risk [87]. Preparedness guidance therefore emphasizes that evacuation and sheltering decisions must account for functional needs, transportation resources, assistance requirements, and continuity of care [149].

Prior work has made substantial progress on individual dimensions of healthcare disaster resilience, but the literature remains fragmented. Healthcare scheduling emphasizes allocation and efficiency within care settings. Evacuation and logistics research focuses on regional movement and transportation. Infrastructure resilience work examines exposure, cascading failures, and service continuity. Geospatial decision support provides methods for integrating spatial data, but often stops short of embedding those data into allocation, routing, and risk-assessment models. AI-assisted methods can extract or summarize complex information, but they require grounding and validation in high-stakes settings. The following sections review these areas in more detail and identify the gaps that motivate the geospatial decision-support approach developed in this dissertation.

2.2 Facility-Scale Allocation, Scheduling, and Fairness

Within healthcare facilities, resilience depends on allocating limited resources to people and tasks under time pressure. A substantial literature in operations research investigates the optimization of personnel, facility space, and medical assets within healthcare systems. Core problem classes include staff scheduling and rostering, admission and bed management, operating room scheduling, and outpatient appointment scheduling. Surveys emphasize that these problems are commonly modeled using mixed-integer programming (MIP), constraint programming (CP), and metaheuristics, with objectives that balance service quality, regulatory compliance, staffing constraints, and operational cost [6, 22, 53, 76]. More recent reviews highlight the growing role of data-driven methods, including predictive models for demand and length of stay, as inputs to optimization and simulation pipelines [143]. Hospital bed management illustrates this coupling of analytics and optimization through metrics such as bed occupancy rate (BOR) and length of stay (LOS), which are often optimized jointly with throughput and waiting-time objectives [115].

Although these methods provide a strong foundation, senior healthcare settings complicate the use of classical schedulers. Tasks are heterogeneous, resource compatibility is clinically meaningful, and feasibility depends on human factors such as mobility impairments, cognitive limitations, caregiver familiarity, and continuity of care. Surveys of home health care routing and scheduling highlight the breadth of objectives that arise in this setting, including travel cost, timeliness, skill compatibility, workload balance, and continuity [64]. Continuity has itself been formalized as an optimization objective, reflecting the importance of maintaining stable caregiver-resident relationships over time horizons that matter clinically and socially [80, 100]. Dynamic formulations further recognize that requests and constraints evolve over time and must often be handled online rather than through a fixed offline schedule [98]. Related literature also uses Markov decision processes (MDPs) and hierarchical planning models to reason about sequential adaptation under uncertainty, especially when decisions must balance present service with future flexibility [99, 130].

Uncertainty and congestion effects are frequently modeled through queueing theory and discrete-event simulation. Canonical results and modeling assumptions appear in standard texts [139], while surveys summarize healthcare applications spanning waiting-time analysis, capacity planning, emergency department flow, and appointment systems under stochastic arrivals and service times [67]. Queueing and simulation studies have been used to evaluate intensive care and emergency department capacity decisions under stochastic demand and priority structures [15, 73, 81]. At a more abstract level, baseline policies such as First-Come-First-Served (FCFS), Shortest-Job-First (SJF), Round Robin (RR), and priority scheduling remain useful reference points for reasoning about trade-offs among simplicity, responsiveness, and fairness. However, these policies are too coarse to capture resident preferences, inter-task dependencies, spatial traversal costs, and changing care conditions in senior health settings [7, 90, 92, 136].

Many scheduling models assume that jobs are well specified, processing requirements are stable, and resource compatibility can be represented through fixed constraints. Senior healthcare challenges these assumptions. A resident's condition may change during an event. A task may require specific caregiver skills, equipment, or physical assistance. A resident's location within a facility may affect how long assistance takes. A non-preferred resource may be technically feasible but operationally slower or less appropriate. In this setting, allocation is not simply temporal scheduling. It is a spatial-temporal decision problem over people, tasks, resources, facility layouts, and changing resident conditions.

Fairness adds another source of complexity. In resource allocation, fairness is defined in different ways depending on the operational context and the values being prioritized. In the "carpool problem," fairness was introduced as a measure of cumulative proportionality, where a scheduling algorithm is considered fair if the service a participant has provided remains close to the service they were equitably expected to provide over time [55]. Foundational research in social choice expands this view by defining fairness through principles such as maximin welfare, which prioritizes the least advantaged individuals [132], and envy-freeness, which requires that no participant prefer the allocation received by another [151]. These ideas are important because fairness becomes central whenever people, rather than abstract jobs, are the recipients of limited resources.

Fairness is not a single mathematical objective, but a family of normative commitments that may conflict with efficiency goals. Work on the fairness-efficiency trade-off formalizes the potential loss in system utility induced by fairness constraints, often described as the price of fairness [16]. Within operations research and networking, families of fairness objectives generalize proportional fairness and connect smoothly to max-min behavior through tunable parameterizations such as α-fairness [118]. Classical fair-division literature also emphasizes properties such as Pareto efficiency, which remain useful conceptual baselines but are difficult to enforce simultaneously when resources are heterogeneous, indivisible, or time-critical [55, 104]. In multi-resource settings, fairness definitions such as Dominant Resource Fairness (DRF) and related frameworks illustrate how heterogeneous demands across multiple resource types complicate classical single-resource fairness notions [75, 91]. These definitions are useful, but they operationalize different notions of what it means for an allocation to be fair. Some emphasize fairness of *shares*, others fairness of *outcomes*, and others fairness of *procedure*. This distinction matters in healthcare because the ethically relevant quantity is rarely identical across individuals. Equal allocation can be inequitable when the underlying need, acuity, or vulnerability is unequal. Conversely, unequal allocation may be justified when it reflects clinically meaningful differences in risk or benefit. Thus, fairness in healthcare is not reducible to equal division alone.

In healthcare, fairness is inseparable from ethics because allocation decisions encode judgments about need, urgency, risk, benefit, and moral obligation. Influential bioethics work argues that allocation principles often combine treating people equally, favoring the worst-off, maximizing total benefits, and recognizing instrumental value, and that no single principle suffices across all scarcity contexts [128]. Related triage and prioritization work further shows that decision-makers often combine severity, prognosis, and waiting time when allocating scarce resources under emergency conditions [161]. For senior healthcare, fairness must also account for functional limitations, medical needs, mobility constraints, resource dependencies, and the resident-specific consequences of delay or displacement.

Fairness also provides an evaluative lens. Studies commonly report efficiency metrics, such as delay, completion time, utilization, or throughput, alongside distributional measures that quantify dispersion or inequality of outcomes. Measures such as Jain's fairness index, the Gini coefficient, standard deviation, and entropy offer generic proxies for allocation evenness [32, 89, 103]. These metrics are informative, but they must be interpreted carefully in healthcare because unequal allocations can be ethically justified when they reflect clinically meaningful differences in urgency, vulnerability, or continuity needs.

Together, this literature shows that healthcare allocation during disasters cannot be reduced to a purely logistical scheduling problem. It must incorporate heterogeneous task requirements, clinically meaningful compatibility constraints, resident preferences, dynamic state updates, indoor spatial layout, and time-critical adaptation. The remaining challenge is to design allocation models that are both operationally efficient and human-centered, while preserving enough interpretability for decision-makers to understand why certain residents, tasks, or resources are prioritized.

2.3 Regional-Scale Coordination, Matching, and Routing

When disasters affect a broader healthcare network, resilience depends on coordinating facilities, people, transportation networks, resources, and hazards across space. Evacuation and emergency logistics problems frequently integrate allocation decisions with transportation constraints. Extensive literature examines emergency facility location, shelter siting, and medical logistics under stochastic demand and disrupted networks [112]. Many formulations are inherently two-stage: planners optimize facility selection or pre-stage resources before an event, then dynamically assign individuals, supplies, or vehicles once disruptions occur. This logic is foundational to risk-averse shelter-location models and robust optimization formulations that explicitly account for link failures, degraded capacity, and evolving accessibility [108, 160].

Beyond optimization, Multi-Criteria Decision Analysis (MCDA) methods are widely used to rank shelters, staging areas, or candidate destinations when accessibility, capacity, hazard exposure, and service availability must be considered jointly. Techniques such as the Analytic Hierarchy Process (AHP) and the Technique for Order Preference by Similarity to Ideal Solution (TOPSIS) are especially common because they are interpretable and facilitate communication with practitioners [8]. Their primary limitation is that they are often static and ranking-oriented. They are less effective when travel times, hazard footprints, or facility status updates must be ingested continuously during live operations.

Routing and trip planning research has long recognized that evacuation is inherently multi-objective because minimizing travel time can conflict with congestion avoidance, hazard mitigation, safety, and fairness across populations. Work on scalable evacuation routing under dynamic congestion highlights the computational challenges of generating advice that remains useful as traffic conditions evolve [138]. Systematic reviews of transit-based evacuation planning similarly emphasize that optimization and simulation methods must account for realistic constraints on transit capacity, information availability, and compliance [97]. Equity considerations are increasingly prominent in this literature as well, including formulations that incorporate proportional fairness into trip planning and access to evacuation services [5].

For senior healthcare evacuations, the interdependency between matching and routing is especially strong. Route feasibility depends on transportation resources, road status, congestion, hazard evolution, and the ability to safely move medically fragile people. Destination suitability depends on medical needs, functional status, facility type, available capacity, service compatibility, and continuity of care. Allocation can therefore be framed as a constrained matching problem between relocating facilities, evacuees, resources, and receiving facilities. Stable matching theory provides a principled way to model two-sided preferences and avoid assignments that would be undermined by mutually preferred alternatives [71]. However, recent work shows that stability and fairness can conflict, meaning that naively composing classical stable matching with fairness constraints may produce significant welfare loss [94]. This tension is especially relevant in disaster settings.

A further technical complication is that destination suitability and route feasibility are not independent. Hazard evolution can invalidate a previously desirable route. Facility degradation can invalidate a previously desirable match. Traffic dynamics can reshape both travel time and effective accessibility. A robust coordination model therefore requires iterative coupling between matching and routing, rather than a one-time assignment followed by independent trip planning.

Geospatial information is central to this coordination problem because hazards, infrastructure damage, resource availability, facilities, and vulnerable populations are spatially distributed. Geographic Information Systems (GIS) have long been used for hazard mapping, evacuation planning, accessibility analysis, and damage assessment. Early work noted both the promise and practical challenges of real-time, model-driven GIS for emergency response, including data heterogeneity, rapid change, and the persistent gap between research prototypes and operational adoption [101, 140, 162]. Work on geospatial data fusion further emphasized that multi-agency response requires shared data models and mechanisms that reduce duplication and improve data quality under time pressure [155, 156].

Spatial Decision Support Systems (SDSS) extend GIS by coupling spatial data management with analytical decision models and simulation. Prior research emphasizes that SDSS architectures must balance model sophistication against usability and latency, as disaster-response decisions are frequently executed under severe time pressure and informational uncertainty [44, 50]. Modern SDSS increasingly integrate GIS with sensor feeds, simulation, and Internet of Things (IoT) infrastructure to support real-time "what-if" reasoning and adaptive response [35]. This integration is especially important for healthcare coordination because planners must move beyond static maps toward continuously updated situational awareness.

Simulation remains essential for evaluating policy alternatives and stress-testing evacuation plans against realistic behavior. Agent-based and microsimulation models capture heterogeneous mobility, staff assistance, and facility layouts. Prior studies have characterized fall risk and mobility limitations during evacuation procedures for older adults [49]. More recent research explicitly models disabled older adults and nursing staff behavior to evaluate evacuation strategies in elderly care facilities under varying caregiver-to-resident ratios and staffing policies [107, 109, 113]. While these simulation approaches add behavioral realism, they are often disconnected from real-time, multi-source GIS feeds and from optimization models that jointly address equity, clinical suitability, and infrastructure disruption.

Together, these research trajectories show that disaster logistics for senior healthcare is neither only a routing problem nor only a matching problem. It is an integrated matching-and-routing challenge under uncertainty, shaped by human, institutional, spatial, and policy constraints. Existing work provides important foundations for facility location, routing, evacuation simulation, stable matching, and SDSS, but healthcare resilience requires these ideas to be connected in a way that accounts for facility compatibility, route safety, changing hazards, capacity constraints, and policy-driven prioritization.

2.4 Infrastructure Resilience, Community Vulnerability, and Risk Assessment

Healthcare resilience also depends on whether facilities can sustain critical services when the infrastructure around them fails. The goal is not only to identify which facilities are exposed to a hazard, but also to determine which facilities are likely to lose essential functions, why they are vulnerable, and what interventions could reduce risk. This is especially important for prolonged power outages, where healthcare facilities may remain physically intact but still become unable to support clinical operations, communications, refrigeration, water systems, and other essential services.

Regional resilience efforts often combine geospatial data with facility-level vulnerability assessments. At a planning level, frameworks such as Threat and Hazard Identification and Risk Assessment (THIRA) and Hazus support scenario-based hazard analysis and loss estimation [4, 60]. At a facility level, organizations frequently rely on Hazard Vulnerability Analysis (HVA) tools and instruments such as the Risk Identification and Site Criticality (RISC) Toolkit to structure emergency planning and assess reliance on external services [84, 93]. These approaches are valuable because they standardize assessment protocols and make planning more auditable. Their primary limitation is that outputs are often ordinal, checklist-oriented, or spreadsheet-based, which complicates their integration with fine-grained spatial exposure, transportation degradation, and facility-specific autonomy requirements during prolonged outages.

Disaster impacts on healthcare are further shaped by critical infrastructure interdependencies. Foundational research demonstrates that disruptions can cascade across tightly coupled physical, cyber, and organizational systems [134]. In healthcare, power system reliability is especially consequential because clinical operations, medical devices, cold-chain storage, water systems, and information systems are electricity-dependent. Empirical accounts of major storms and outages document how backup generator failures, fuel constraints, and prolonged utility disruptions can trigger rapid evacuations and strain surrounding facilities [9, 43, 147]. These cases illustrate that power outages are not only utility failures; they can become regional healthcare disruptions when critical services cannot be sustained.

Recent grid-resilience literature also emphasizes the role of mobile power sources, including electric vehicles and mobile energy storage, as flexible resources that can be dispatched to critical loads such as hospitals and emergency shelters [127]. These mobile units bridge the transportation and energy sectors, but their effectiveness depends on road access, travel time, and the condition of the transportation network during a disaster. A central challenge during prolonged outages is therefore that resilience is not dictated solely by the presence of backup hardware. It depends on how long critical services can be sustained, how quickly consumables and fuel can be replenished, how exposed upstream infrastructure is to hazards, and how fragile the surrounding transportation network becomes under cascading disruption. Existing facility-level tools often treat these factors in isolated workflows. As a result, planners are left with a difficult question: which facilities are most likely to exhaust their autonomous capacity first, and which will require the most intensive support?

Guidance for utility failures emphasizes preventive maintenance, backup-power testing, fuel availability, and planning for long-duration outages [40]. Beyond infrastructure readiness, spatial analyses indicate that protecting medically vulnerable communities requires analyzing specific physical locations and barriers rather than relying only on administrative tabulation [142]. Emerging work on electric vehicles as a resilience solution demonstrates their potential to reduce health hazards for medically fragile residents by providing vehicle-to-home (V2H) or vehicle-to-building (V2B) power during extreme weather, potentially sustaining life-critical equipment for several days [111]. Agent-based modeling approaches have also shown that a coordinated fleet of electric vehicles can enhance outage resilience by providing both energy and emergency supply delivery in a decentralized manner [36]. However, the equitable distribution of these mobile resources remains a concern, as disparities in electric vehicle ownership and charging infrastructure can lead to uneven resilience outcomes for vulnerable communities [114].

Facility risk is also shaped by the surrounding community. Indices such as the Social Vulnerability Index (SVI) use demographic and socioeconomic indicators to identify populations that may have lower capacity to prepare for, respond to, and recover from hazards [10, 39, 65]. These indices are valuable for screening and prioritization, but they are often coarse-grained and relatively static compared with the pace of an unfolding event. SVI is most useful when combined with facility-specific and infrastructure-specific indicators, such as backup autonomy, outage duration, road accessibility, and service dependency. Social vulnerability can identify where consequences may be more severe, but it does not by itself determine whether a facility can maintain critical services.

These limitations are especially consequential for healthcare facilities serving older adults and medically fragile populations, where even brief interruptions can trigger cascading clinical complications. The literature therefore motivates risk assessment models that link HVA-style planning concepts to GIS-based analysis.

2.5 Geospatial Data Systems and Bounded AI Support

Healthcare resilience depends on the ability to turn heterogeneous disaster data into usable computational representations. Modern disaster operations increasingly rely on multi-layer spatial representations that combine vector infrastructure, raster hazard products, facility attributes, transportation networks, social vulnerability layers, and infrastructure dependencies. The technical challenge is not only representing these layers, but also supporting frequent updates, incremental recomputation, and interactive querying under operational latency constraints. These requirements make back-end data architecture as important as front-end visualization.

Operational SDSS often rely on spatial databases to ingest, index, and query high-velocity, multi-source geospatial feeds. Spatial indexing structures, notably R-trees, enable efficient search across multidimensional spaces [79]. Modern systems such as PostGIS, Oracle Spatial, and SQL Server build upon these foundations to support index-assisted k-nearest neighbor (k-NN) queries using structures such as GiST (Generalized Search Tree) or R-trees to rapidly retrieve proximate objects [38]. These capabilities are essential for the iterative identification of candidate facilities, resources, suppliers, or routes as disaster conditions evolve. They also shift foundational geospatial computations into scalable back-end architectures, enabling real-time dashboards and multi-user operational workflows rather than static analyses.

Digital twin concepts provide a complementary framing by emphasizing an evolving digital representation of a physical system that supports monitoring, simulation, and intervention planning. Prior work proposes "disaster-city" digital twins that integrate artificial and human intelligence for situational assessment and coordination [56]. Recent studies further emphasize the value of geospatial digital twins for decision-making and post-disaster risk management, especially when multi-source sensing, simulation, and human oversight must be coordinated in a common representational framework [74, 96, 102]. In healthcare settings, digital-twin-inspired representations are particularly useful because they can connect facility layouts, resident and resource locations, hazard dynamics, transportation networks, and infrastructure status.

Decisions also depend on unstructured information, including facility guidance, status reports, regulatory requirements, mutual-aid documentation, and after-action reports. Transforming such text into computable representations predates recent generative AI and includes curated logic engines, semantic parsing, and relation extraction pipelines [62, 70]. These earlier approaches established the value of structured knowledge extraction, but they often struggled to bridge free-form regulatory language with the operational and geospatial representations required in emergency workflows.

Recent research demonstrates that AI foundation models, including Large Language Model (LLM) families such as GPT, Gemini, and Llama, can support schema-driven extraction of structured data from technical text, powering downstream databases, dashboards, and decision-support analytics [41]. To improve grounding and factual reliability, Retrieval-Augmented Generation (RAG) supplements models with external knowledge and has emerged as a common design pattern for domain-sensitive applications where the data source must be verifiable and traceable [72, 106]. These capabilities are especially relevant for healthcare resilience because policy documents and regulatory standards are often lengthy, heterogeneous, and difficult to translate into operational variables.

At the same time, the technical value of these models depends heavily on how narrowly their role is scoped. In high-stakes operational settings, foundation models are most defensible when used for bounded tasks such as schema extraction, knowledge-grounded summarization, metadata normalization, query assistance, or explanation generation. They are less reliable when used as unrestricted decision-makers over rapidly changing, safety-critical environments. This distinction is especially important when model outputs must interface with formal optimization, GIS layers, or operational policies.

Factuality remains a central concern. LLMs can hallucinate unsupported content, conflate sources, or produce outputs that are syntactically well-formed but operationally misleading. These limitations are especially problematic in healthcare and emergency management, where unsupported recommendations can compromise planning and coordination [154, 159].

For this reason, AI-assisted decision support in disaster settings should be designed conservatively. Rather than treating LLMs as autonomous decision-makers, the literature increasingly motivates human-in-the-loop workflows in which extraction, summarization, and recommendation are paired with validation steps, provenance tracking, and bounded operational roles.

From the perspective of healthcare resilience, geospatial systems and AI play complementary roles. Geospatial systems provide the computational grounding for location, exposure, routing, accessibility, infrastructure dependency, and regional context. AI-assisted methods can help translate unstructured policy and preparedness information into structured representations that can be queried, validated, and connected to those geospatial models. The open challenge is to ensure that AI supports accountable decision-making without replacing the formal spatial, optimization, and risk models that high-stakes planning requires.

2.6 Summary of Gaps

The literature reviewed in this chapter shows that healthcare disaster resilience requires methods that connect care delivery, spatial coordination, infrastructure dependencies, and human vulnerability. Prior work on healthcare scheduling and resource allocation provides useful optimization tools, but often assumes relatively stable jobs, fixed constraints, and generic efficiency objectives. Senior healthcare disaster response requires more personalized models that account for resident needs, urgency, preferences, resource compatibility, and changing conditions.

Research on evacuation planning, facility location, routing, stable matching, MCDA, and SDSS provides important foundations for coordination under disruption. However, healthcare evacuations require tighter integration between matching and routing because destination suitability depends on care compatibility, capacity, facility status, and policy constraints, while route feasibility depends on hazards, road conditions, traffic, and changing accessibility.

Existing approaches often treat these components separately or lack the ability to update decisions as conditions evolve.

Work on HVA, THIRA, Hazus, RISC, critical infrastructure interdependencies, mobile energy resources, and social vulnerability provides important tools for assessing disaster risk and resilience. However, these methods often remain fragmented across checklists, spatial layers, regulatory documents, and infrastructure models. Healthcare planners need integrated methods that can assess facility-specific risk under prolonged outages, account for backup autonomy and resupply feasibility, and explain why particular facilities are more vulnerable.

Geospatial data systems and AI-assisted decision support introduce both opportunities and risks. Spatial databases, GIS layers, digital twins, and SDSS enable the integration of facilities, hazards, routes, resources, and infrastructure into computational workflows. AI and RAG can support bounded extraction of regulatory and preparedness information, but these tools must remain grounded, auditable, and subordinate to formal geospatial, optimization, and risk models. These gaps motivate the overall approach presented in the next chapter, which frames healthcare disaster resilience as a geospatial decision-support problem connecting in-facility allocation, regional matching and routing, and facility risk assessment.

Chapter 3

The Approach: Geospatial Systems for Healthcare Resilience

The preceding chapters motivate healthcare disaster resilience as a human-centered, geospatial, and computational problem. Chapter 2 reviewed related work on allocation, fairness, regional coordination, infrastructure resilience, geospatial decision support, and bounded AI-assisted operations. Building on that discussion, this chapter presents the unified problem statement and approach that guide the three technical frameworks developed in this dissertation.

The central argument is that healthcare disaster resilience cannot be addressed at only one operational level or at only one moment in time. A disaster may require urgent decisions inside a senior health facility, coordination across a broader healthcare network, and risk assessment for infrastructure disruptions that unfold over several days. These decisions differ in scope and urgency, but they share a common need: decision-makers must reason about where people, facilities, hazards, resources, routes, and infrastructure dependencies are located, and how these conditions change over time.

This perspective frames healthcare disaster resilience as a geospatial decision-support problem. Geospatial systems serve as more than visualization tools; they provide the computational substrate for representing spatial relationships, evaluating feasibility, routing people and resources, modeling infrastructure dependencies, and producing interpretable recommendations. Figure 3.1 summarizes the overall structure of the dissertation. The first framework focuses on personalized allocation under scarce resources inside senior health facilities. The second extends allocation to cross-facility coordination through policy-driven matching and hazard-aware routing. The third shifts toward facility risk assessment under prolonged power outages. Together, these frameworks show how spatial data and computational models can support allocation, routing, risk modeling, and resilience planning in healthcare disaster settings.

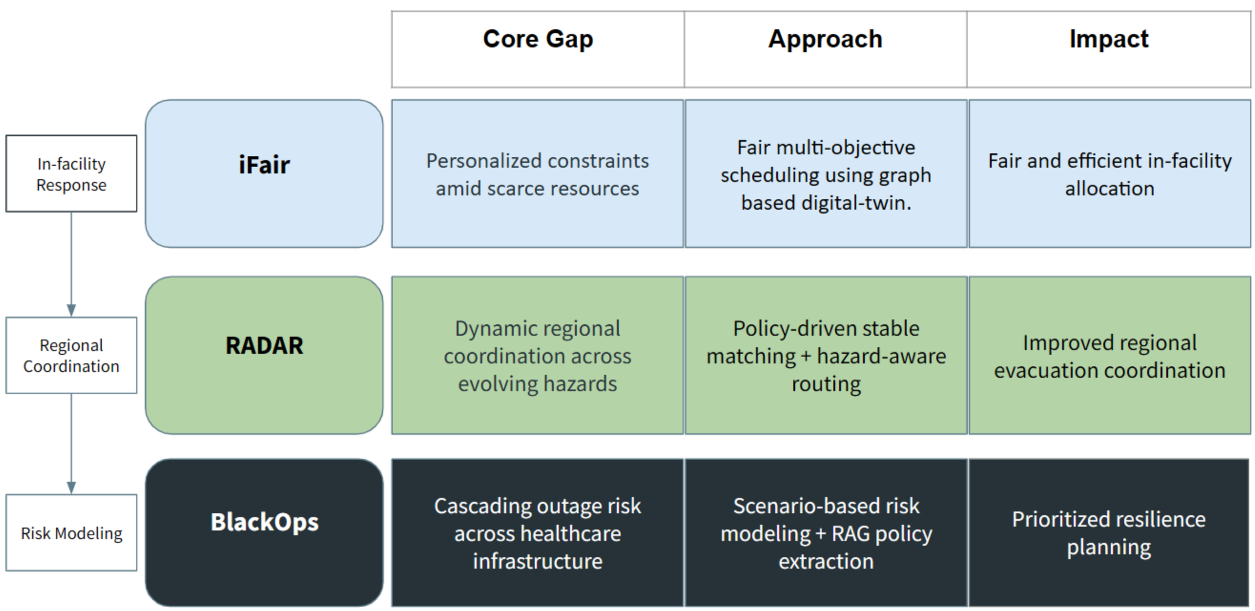

Figure 3.1: Dissertation roadmap illustrating core gaps, methodological approaches, and impacts across distinct operational scales.

3.1 Multi-Scale Problem Statement

The overarching problem addressed in this dissertation is how to support healthcare disaster resilience through computational models that are spatially grounded, human-centered, and operationally interpretable. In this context, resilience is not treated as a single outcome or a single optimization objective. Instead, it is treated as the ability to support decisions that preserve care continuity under scarcity, uncertainty, and disruption.

One part of this problem concerns urgent allocation inside senior health facilities. During events such as fires or evacuations, limited staff, equipment, and time must be assigned to residents whose needs vary significantly. Decision-makers must determine which residents require assistance first, which resources are appropriate, and how the physical layout of the facility affects task completion time. The technical objective is to balance efficiency with resident-centered fairness, where fairness means accounting for heterogeneous needs, criticality, preferences, and care constraints rather than treating residents as interchangeable demand units.

Another part concerns coordination across multiple healthcare facilities and transportation networks. A facility may need to relocate residents, request resources, or share available capacity. Emergency managers must decide which facilities should be matched, which routes should be used, and how assignments should change when hazards expand, roads become unavailable, or facility status changes. Here, matching and routing must be modeled together because a destination is useful only if it is compatible, has capacity, and can be reached safely.

A third part concerns identifying facilities that may lose critical services during prolonged disruptions. A facility may be outside the direct hazard footprint but still become vulnerable if the electric grid fails, backup autonomy is insufficient, resupply is delayed, or the surrounding community has limited recovery capacity. This requires risk assessment models that combine facility autonomy, infrastructure dependency, road-network accessibility, hazard exposure, social vulnerability, and regulatory requirements into interpretable outputs for planners.

These decision contexts are connected. Facility evacuation may require relocation to other sites. Relocation depends on transportation networks, available receiving facilities, and infrastructure status. Infrastructure failures may determine whether facilities can continue operating, receive evacuees, or support resource sharing. Rather than treating these as isolated applications, this dissertation frames healthcare disaster resilience as a connected set of allocation, coordination, and facility risk assessment problems.

3.2 Spatial-Temporal Framing of Healthcare Resilience

The problem described above is both spatial and temporal. Spatially, healthcare resilience depends on where residents, caregivers, equipment, facilities, roads, hazards, suppliers, communities, and infrastructure assets are located. Temporally, it depends on how quickly decisions must be made and how rapidly conditions change.

Inside a senior health facility, the physical environment is relatively small, but the timing is urgent. Residents, staff, equipment, rooms, corridors, exits, and hazards all shape what actions are possible and how long they may take. During an evacuation, decisions may need to be made in minutes. The challenge is not only to move people quickly, but also to account for differences in mobility, medical conditions, care needs, and preferences. A purely distance-based or first-come-first-served strategy may appear efficient, but it can fail to prioritize residents whose conditions make delay more dangerous or whose needs require specific resources.

Across a healthcare network, the decision environment expands to many facilities, roads, resources, and jurisdictions. A facility may need to relocate residents, a hospital may experience bed surges, or emergency managers may need to move resources such as oxygen, staff, water, or transportation assets. These decisions often unfold over hours to days. The problem is no longer only which resident should be helped first within a facility, but which facilities should be matched, which routes should be used, which destinations have capacity, and how assignments should change when hazards, roads, and facility conditions evolve.

For prolonged infrastructure disruptions, the time horizon shifts further. Healthcare facilities depend on electricity, transportation, communications, suppliers, and regional partners. A facility may be physically outside a hazard footprint but still lose critical services if the power grid fails, backup autonomy is insufficient, or resupply routes become inaccessible. Here, the decision problem is not only how to respond after failure occurs, but how to identify vulnerable facilities in advance, estimate the consequences of infrastructure disruption, and prioritize interventions that reduce risk.

This spatial-temporal framing motivates the structure of the dissertation. The first framework addresses urgent in-facility response, the second addresses cross-facility coordination under evolving hazards, and the third addresses facility risk assessment under prolonged power outages. These are different views of the same broader challenge: supporting healthcare disaster resilience through spatially aware and time-sensitive decision-making.

3.3 Design Principles

The three frameworks developed in this dissertation differ in technical formulation, but they share four design principles.

Human-centered representation: Healthcare resilience decisions affect people whose needs are heterogeneous. Older adults may differ in mobility, diagnoses, equipment dependence, cognitive status, medication needs, care preferences, and vulnerability to delay. The models developed in this dissertation therefore avoid treating people as interchangeable units whenever individual or facility-level information is available. For in-facility allocation, this means modeling resident-specific criticality, preferences, and task requirements. For cross-facility coordination, it means incorporating facility compatibility, residency type, triage priority, and care needs. For outage risk assessment, it means accounting for the social and clinical consequences of losing critical services.

Geospatial grounding: Decisions depend on location, distance, routes, hazard footprints, facility layouts, and infrastructure networks. The approach therefore treats geospatial information as part of the decision logic rather than as a final visualization layer. Facility layouts, road networks, hazard zones, supplier locations, power-grid assets, and community vulnerability layers are used to determine feasibility, priority, and risk.

Adaptation across time: Disasters are dynamic. Resident condition may change, hazards may expand, roads may close, facility status may degrade, and outages may last longer than expected. The models therefore distinguish between relatively static information, such as facility attributes or resident profiles, and dynamic information, such as current hazard proximity, operational status, route availability, and scenario-dependent outage duration. This allows recommendations to be updated as conditions evolve.

Operational interpretability: Decision support must be understandable to emergency managers, healthcare providers, and facility stakeholders. The goal is not only to compute an assignment, route, or risk score, but also to make the reasoning behind it visible. Policy scores, criticality values, route feasibility, autonomy gaps, shortfall indices, and facility scorecards help stakeholders understand why a recommendation was produced and how it relates to operational constraints.

These principles distinguish the approach in this dissertation from models that optimize a single efficiency objective, rely only on static GIS overlays, or use AI outputs without structured validation. They also provide the bridge between the related work reviewed in Chapter 2 and the technical chapters that follow.

3.4 Geospatial Systems as the Unifying Substrate

A central idea in this dissertation is that geospatial systems should be treated as more than visualization tools. Maps are essential for situational awareness, but the underlying spatial data can also drive computation. Facility layouts can be represented as graphs. Road networks can be queried for travel time and accessibility. Hazard footprints can be overlaid with facility locations. Social vulnerability can be joined to healthcare facilities based on geography. Power-grid assets can be connected to facilities to model infrastructure dependency. These spatial relationships provide the basis for allocation, matching, routing, and risk analysis.

Inside a senior health facility, geospatial reasoning appears through facility layouts. Rooms, hallways, exits, and other spaces can be modeled as nodes, while traversal paths between them become weighted edges. This representation allows the model to reason about resident locations, resource locations, hazard proximity, evacuation paths, and task completion time. Across a healthcare network, GIS layers help determine which facilities are impacted, which candidate destinations are safe, which routes are feasible, and how travel time changes with traffic or road conditions. Under prolonged infrastructure disruption, spatial layers connect healthcare facilities to power-grid assets, road networks, hazard scenarios, supplier locations, and social vulnerability indicators. This makes it possible to evaluate not only direct facility exposure, but also indirect vulnerability caused by outage duration, backup-energy shortfall, resupply difficulty, and community context.

Across these settings, geospatial data are combined with optimization, graph algorithms, policy scoring, stable matching, simulation, bounded AI-assisted extraction, and interpretable risk modeling. The result is a decision-support approach in which spatial awareness is directly connected to operational decisions.

3.5 iFair: Facility-Scale Response

The first framework focuses on urgent allocation inside senior health facilities. The motivating question is how decision-makers can use personalized knowledge about older adults to allocate scarce resources fairly and efficiently during emergencies. This is a human-centered allocation problem because residents differ in mobility, diagnoses, equipment dependencies, cognitive status, care requirements, and preferences. During an evacuation, these differences directly affect how long tasks take, which resources are appropriate, and which residents should be prioritized.

iFair models a senior health facility as a graph-based digital representation. The graph captures the physical structure of the facility, while resident profiles capture medical conditions, locations, and preferences. Resources include human resources such as nurses, caregivers, and first responders, as well as equipment resources such as wheelchairs, oxygen tanks, and medication. Tasks include movement, basic care, critical care, and evacuation-related actions. This representation allows the system to reason jointly about people, resources, tasks, and space.

The fairness formulation is designed as a practical part of allocation rather than as a separate ethical statement. In this work, fairness does not mean giving every resident the same resource. It means accounting for what each resident needs in order to receive appropriate care. Resident preferences are used to compute satisfaction, and assigning non-preferred resources can introduce additional time penalties. This reflects an operational reality: ignoring a resident's needs or preferences can make care less efficient, not merely less fair.

The model combines offline and online criticality. Offline criticality is computed from resident medical profiles and risk-adjustment scores. Online criticality updates priority based on the current event, such as a resident's proximity to a fire. The allocation process then prioritizes residents, identifies required tasks and resources, considers preferences, assigns available resources, computes task completion time, and records the allocation. The full model, algorithms, and experimental evaluation are presented in Chapter 4.

This framework shows that fairness and efficiency are not always competing goals. When a model captures the real factors that shape response time, including resident criticality, preferences, and facility layout, fairness-aware allocation can also improve operational efficiency.

3.6 RADAR: Regional Coordination

The next part of the dissertation extends allocation from a single facility to a broader healthcare network. This transition was motivated by CareDEX drills and workshops involving wildfire evacuation, earthquake response, medical surge, and water-contamination scenarios. These exercises showed that facility-level readiness is not enough during major disasters. Decision-makers also need to determine where people should be relocated, which facilities can receive them, which resources can be shared, and which routes remain usable.

RADAR addresses this coordination problem by combining GIS hazard feeds, facility data, resource availability, policy-driven scoring, stable matching, and routing. The system ingests and processes data about facilities, hazards, transportation networks, and operational status. A policy engine then converts decision rules into score functions that capture factors such as same residency type, proximity, safe routes, capacity, operational status, triage priority, and social vulnerability. These scores are used to rank candidate facilities and generate preference lists.

The optimization problem asks how evacuees or resources can be matched to facilities while maximizing policy scores and satisfying operational constraints. These constraints ensure that each facility has at most one primary match, receiving facilities do not exceed capacity, assigned pairs are connected by safe routes, and travel distances remain reasonable. Stable matching is then used to produce allocation pairs that are feasible, interpretable, and aligned with the policy preferences of relocating and receiving facilities.

After matching, routes are computed between paired facilities. Candidate destinations are filtered using hazard and facility-status information, and shortest paths are computed over the road network using Dijkstra's algorithm and GIS network analysis. The system also supports dynamic updates. When roads become blocked, hazards expand, or a facility changes status, valid assignments are preserved when possible, newly impacted facilities are reassigned, and the candidate pool is expanded only when needed. This allows the allocation plan to adapt without unnecessary disruption.

This work shows how policy-aware matching and hazard-aware routing can be combined to support evacuation planning, bed surge management, and resource movement across a healthcare network. The full formulation, policies, matching and routing algorithms, and experiments are presented in Chapter 5.

3.7 BlackOps: Facility Risk Assessment Under Power Outages

The final framework shifts from response coordination to risk assessment for prolonged power outages. This shift is important because healthcare facilities depend heavily on electricity for life-saving equipment, refrigeration, water systems, communications, transportation coordination, and basic operations. A disaster can therefore affect healthcare delivery not only by damaging a facility directly, but also by disrupting the infrastructure systems that allow the facility to function. *BlackOps* addresses the following question: given a prolonged power outage in a region, which healthcare facilities are most at risk of losing critical services, and how should planners reduce that risk?

The framework combines regulatory requirements, facility backup-power characteristics, hazard scenarios, power-grid dependencies, road-network accessibility, supplier availability, and community vulnerability. One challenge is that facilities must comply with requirements distributed across heterogeneous documents, including NFPA standards, FEMA guidance, CMS requirements, state guidance, and local plans. Because these documents are difficult to translate into facility-specific operational requirements, the framework uses retrieval-augmented generation in a bounded and grounded way. Facility attributes are embedded, relevant regulatory chunks are retrieved using cosine similarity, and a language model extracts structured autonomy requirements and metadata into a validated format. The extracted requirement is then used to compute an autonomy gap index, which compares required autonomy with the facility's estimated critical autonomy.

Regulatory compliance alone does not guarantee resilience. A facility may have some backup capacity, but still lose critical services if the outage lasts longer than expected or if resupply cannot arrive on time. To model this, the road network is represented as a weighted graph, where intersections are nodes, road segments are edges, and weights represent travel time based on historical congestion. This supports a resupply difficulty index that combines response delay with partner availability and route redundancy. A shortfall index then measures the fraction of the outage not covered by on-site backup and feasible external support.

The risk model also incorporates social vulnerability and physical hazard exposure. Social vulnerability captures the demographic and socioeconomic context of the community in which a facility is located, while exposure captures how hazards affect both facilities and the infrastructure they depend on. For earthquake scenarios, the framework uses stochastic earthquake catalogs, ground-motion simulation, damage estimation, and a physics-based power-flow model to estimate cascading effects across the electric grid. These components are aggregated into a facility-scenario risk score, mapped to interpretable risk tiers, and translated into scorecards and recommendations.

This work shows how geospatial systems can support facility risk assessment under cascading infrastructure disruptions. The output is not only a risk score, but an interpretable planning artifact that explains why a facility is vulnerable, which factors contribute to its risk, which suppliers may support it, and how much backup energy should be pre-staged. The full risk formulation, preparedness indices, scorecard design, and evaluation are presented in Chapter 6.

3.8 Connection Across the Frameworks

The three frameworks form a progression from urgent care delivery to cross-facility coordination and facility risk assessment. The first begins inside a senior health facility, where the focus is on residents, caregivers, equipment, tasks, and indoor movement. The second expands to the healthcare network, where the focus is on matching facilities, routing people and resources, and adapting to evolving hazard conditions. The third examines prolonged infrastructure disruptions, where the focus is on identifying vulnerable facilities before outages interrupt critical services.

This progression also reflects a temporal shift. In-facility response decisions must often be made within minutes. Cross-facility coordination unfolds over hours to days as hazards, transportation conditions, and facility capacities change. Risk assessment for prolonged outages requires scenario-based modeling to identify vulnerabilities and prioritize interventions before or during an event.

The unifying theme is that healthcare disaster resilience requires decision support that is spatially aware, temporally adaptive, human-centered, and interpretable. Spatial awareness helps decision-makers understand where people, hazards, routes, resources, and infrastructure dependencies are located. Temporal awareness captures how urgency, travel time, outage duration, and recovery windows affect decisions. Human-centered modeling ensures that resident needs, preferences, medical conditions, and social vulnerability are not ignored.

Interpretability helps stakeholders understand and act on system recommendations.

Together, the frameworks developed in this dissertation connect indoor facility graphs, regional road networks, hazard layers, facility status feeds, social vulnerability data, regulatory requirements, and infrastructure models into computational workflows for healthcare resilience. The result is a dissertation that moves from fair allocation, to regional coordination, to facility risk assessment under prolonged infrastructure disruption.

3.9 Chapter Summary

This chapter presented the overall approach used in the dissertation. The work is grounded in a spatial-temporal view of healthcare disaster resilience, where decisions depend both on where critical entities are located and on how quickly conditions change. The approach supports fair and efficient allocation of scarce resources to older adults, policy-driven matching and hazard-aware routing across healthcare facilities, and risk assessment for identifying facilities likely to lose critical services during prolonged power outages. The next chapter develops the in-facility allocation model in detail.

Chapter 4

Achieving Fairness in the Allocation of Scarce Resources

Fair resource allocation is essential in senior care settings, where limited staff, equipment, and time must be assigned to older adults with heterogeneous medical needs, mobility limitations, preferences, and levels of urgency. During emergencies, this problem becomes more difficult because decisions must be made quickly, often with incomplete information, scarce resources, and changing conditions. We develop *iFair*, a fairness-aware resource allocation framework created in collaboration with domain experts and Senior Health Facility (SHF) partners in Orange County, California. Prototyped via CareDEX [1], the framework combines resident profiles, preferences, facility layouts, resource availability, and dynamic event conditions to support allocation during urgent events. Using fire evacuation as the motivating scenario, we model intra-facility allocation as a human-centered scheduling problem and show how accounting for resident criticality and preferences can improve resident satisfaction while also supporting operational efficiency.

4.1 Chapter Overview

This chapter focuses on the allocation problem that arises inside senior health facilities. In these settings, limited staff, equipment, and time must be aligned with residents whose needs vary significantly. Scheduling and allocation methods are useful for managing limited resources in healthcare and emergency services [99, 130], but senior care introduces additional human-centered constraints that are difficult to capture with generic allocation policies. Residents may differ in mobility, diagnoses, cognitive status, equipment dependence, care preferences, and urgency. These differences affect both what resources are needed and how long tasks may take. For example, during a fire evacuation, assigning residents randomly or relying only on shortest completion time may appear simple, but it can overlook residents who require specific assistance, specialized equipment, or higher priority because of their medical condition or proximity to danger. Prior knowledge of resident needs and preferences can therefore make allocation more informed, more equitable, and more efficient.

The main challenge is to balance fairness and efficiency under urgent and uncertain conditions. Fairness is not treated as equal allocation of identical resources. Instead, it is modeled as a resident-centered objective that accounts for criticality, preferences, and satisfaction while still considering task completion time. This formulation reflects the practical reality that assigning a technically available but poorly matched resource can increase care time, reduce satisfaction, and delay response.

To address this problem, the framework models a senior health facility as a graph-based digital representation. The model captures the facility layout, residents, resources, tasks, preferences, and event dynamics. It combines offline criticality, computed from resident medical profiles, with online criticality, updated using event-specific information such as proximity to danger. The allocation process then prioritizes residents, identifies required tasks and resources, incorporates resident preferences, and computes task completion time and satisfaction.

Although the focus here is intra-facility allocation, the problem connects directly to the broader healthcare resilience argument developed in this dissertation. Internal allocation helps determine which residents need assistance first, what resources are required, and when a facility may no longer be able to resolve all needs on its own. These conditions motivate the regional coordination problem addressed in Chapter 5. The contributions are as follows:

- A mathematical model that captures a graph-based digital representation of a senior health facility, including residents, resources, tasks, facility layout, and dynamic event conditions - §4.2.1.
- A fairness formulation for human-centered resource allocation that balances resident criticality, preference satisfaction, and task efficiency - §4.2.2.
- A task-based assessment and allocation framework that integrates expert insights and senior care facility experience to score, prioritize, and assign resources to older adults during urgent events - §4.3.
- An emulation-based evaluation using data adapted from a real senior health facility to examine the trade-off between fairness, resident satisfaction, and task completion time - §4.4.

4.2 Problem Formulation

4.2.1 Modeling Key Components

We model the static and dynamic elements of a senior care facility, including residents, resources, tasks, events, and allocation decisions. The goal is to capture the information needed to reason about fair and efficient resource assignment during urgent conditions.

Facility: We model a facility as a graph $\mathcal{G} = (V, E)$ where nodes $v_i \in V$ represent different areas (e.g., rooms, nurse's office) and weighted edges $e_{ij} \in E$ denote a pathway between

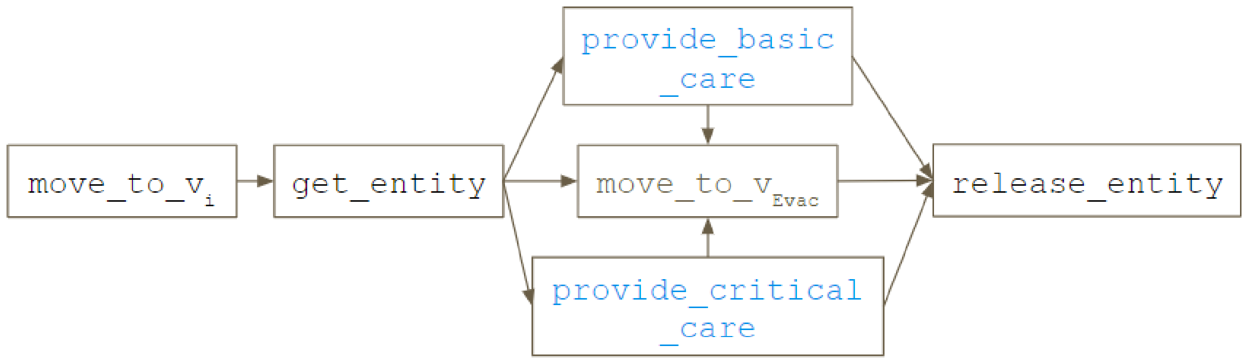

Figure 4.1: Task Dependency Graph.

v_i and v_j with travel time ω_{ij}. We denote two key areas in the facility: exits v_{Exit} and an evacuation area v_{Evac}. The dynamic state of a node v_i is captured by $state(v_i, t)$ representing the condition of v_i at time t, which can be one of $\{Impacted, Clear, Unknown\}$. This is used to pinpoint emergency sites and their proximity to residents. We also model the traversal delays using $access(v_i, t)$, which may occur during an emergency due to smoke, narrow exits, or other access constraints.

Resources: We consider two types of resources: *humans* such as nurses and firefighters, and *equipment* such as medication and wheelchairs. Human resources $h_i \in \mathcal{H}$ are characterized by a headcount $\mathcal{C}_i$ and skill-set $\mathcal{S}_i$. We model $\mathcal{S}_i$ as a tuple of personal attributes (e.g., gender) and skills defined in the O*NET resource center skill ontology [3], which includes certifications (e.g., basic/critical care) and abilities (e.g., mobility aid). Equipment resources, $e_i \in \mathcal{E}$ are characterized by their quantity q_i, reusability r_i and type γ_i. In our context, we define reusability to distinguish between single-use items (e.g., syringes, medication) and multi-use items (e.g., portable tanks). The type of equipment resource defines its purpose and capabilities, e.g., a wheelchair is used for mobility and can be manual or electric. For our allocation problem, we also consider the real-time location $loc(\cdot, t)$ and status $avail(\cdot, t)$ of human and equipment resources.

Tasks: We support a set of tasks $\tau_i \in \mathcal{T}$ with dependencies denoted by $\mathcal{T}_D$, such as in Fig. 4.1. Tasks range from relocating resident or resource entities (which we denote abstractly as $\mathcal{X}$) to specific facility locations $move\_to(v_i)$, to managing entities via $get\_entity()$ and $release\_entity()$. The core tasks, color-coded in blue in Fig. 4.1, involve providing basic and/or critical care. Denoted as $provide\_basic\_care()$ and $provide\_critical\_care()$, they range from ADLs such as dressing and toileting to more advanced medical needs. Each task is defined by its name $\mathcal{N}$, entities involved $\mathcal{X}$, and the resources needed $Req(\mathcal{E}, \mathcal{H})$. We model the changing conditions of an emergency through task updates $task\_update(\tau_i, \mathcal{X}, t)$. We define the overall task completion time in Eqn. 4.1, which encompasses the path traversal times ω_e, delays for space access $access(\cdot)$, and additional time incurred for allocating a resident's non-preferred resources λ_p (described later). To optimize task execution, Dijkstra's shortest path algorithm is employed.

$$TTC() = \sum_e \omega_e + \sum_v access(v_i, t) + \lambda_p, \forall_{e,v} \in \mathcal{G} \tag{4.1}$$

Residents: Let $P_i \in \mathcal{P}$ denote residents in the senior care facility. Each resident is characterized by a 3-tuple $\{\mathcal{B}, \mathcal{D}, \mathcal{F}\}$, denoting basic attributes (e.g., name, gender, age, DOB, room), their diagnoses (e.g., dementia, arthritis, obesity, and related diagnoses), and preferences for specific resources. In our model, preferences are quantified on a numerical scale ranging from -5 (strong dislike) to $+5$ (strong preference). Table 4.1 indicates these preferences for basic care tasks. This scaling allows us to measure the impact on task efficiency and residents' satisfaction. Positive values indicate resources that enhance operational performance, while negative values increase task completion time, directly reflecting the operational costs of utilizing less favored resources.

A resident's satisfaction level $\mathcal{LS}()$ is the sum of their preferences $\mathcal{F}_i$, for all resources used in a task τ_i, i.e., $\mathcal{LS}() = \sum_{i=0}^{n_{P_i}} \mathcal{F}_i$, where n_{P_i} is the number of resources allocated to resident P_i. The extra time λ_p, for using less preferred resources is computed from summing negative $\mathcal{F}_i$ values, i.e., $\lambda_p = -\sum \alpha_i \cdot \min(0, \mathcal{F}_i)$. α_i quantifies how much a preference impacts task

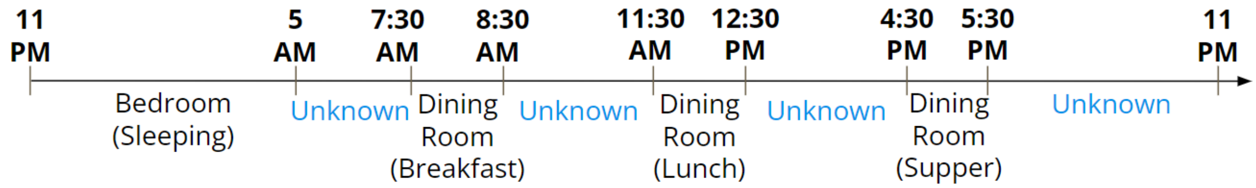

Figure 4.2: Resident Expected Location.

efficiency. For an emergency scenario, let the criticality score C_{score}, assess the urgency of residents' needs based on a set of important attributes $\mathcal{I}$ with corresponding weights ω_a. We track changes in resident conditions (e.g., head injury from falling) using a health report form $\mathcal{HRF}$ and adjust attribute weights and necessary resources, as described later in §4.3.1. A resident's location is captured using $loc(P_i, t)$; we assume that residents have expected locations at specific times, as depicted in Fig. 4.2.

Consider a resident, Joe, who requires two-person assistance and equipment e_2 as shown in Table 4.1. If he was assigned resources h_2, h_3, and e_2 under impact rates α_{h_2}=3, α_{h_3}=2,

Table 4.1: Resident Preferences for Basic Care.

	Joe	Mary	Bob	Carol	Don
h_1	+5	-5	0	-4	0
h_2	+2	-3	0	-1	0
h_3	-5	+3	0	+4	0
e_1	0	0	0	+5	0
e_2	0	-1	0	-2	-5

and α_{e_2}=4 (in minutes per $\mathcal{F}_i$), we find that λ_{Joe}=10 extra minutes, and $\mathcal{LS}(Joe)$=-3, suggesting slight dissatisfaction with his allocated resources.

Event and Allocation: Our analysis prioritizes urgent events like fires that require immediate action. An *event* is defined by its starting point, with v_o=$loc(event, t$=$0)$, and its propagation, $prop(v_i, v_j)$ within the facility, which increases edge weights ω near the affected nodes. Allocations $\mathcal{A}$ are denoted by $\{P_i, h_i \,||\, e_i, t_s, t_e\}$, indicating that a resident P_i is allocated a resource h_i or e_i during time period (t_s, t_e). Notably, resources assigned to the same task might start simultaneously but end at different times, such as a nurse and a consumable (medication). We also note that the human resource count in senior facilities is much lower than that of residents, i.e., $\sum_{i=1}^{|\mathcal{H}|} \mathcal{C}_i << \sum_{j=1}^{|\mathcal{P}|} P_j$.

4.2.2 Problem Statement

In our model, fairness is defined as the balance between prioritizing urgent resident needs, reflected by criticality scores C_{score}, and achieving equitable satisfaction levels $\mathcal{LS}()$ across residents. This dual objective ensures resources are initially allocated to those in immediate need based on their conditions and needs while striving to meet the satisfaction of all residents uniformly. Yet, emergencies introduce a balance challenge between fairness and efficiency; for instance, during a fire, swift evacuation might override individual preferences for everyone's safety. Hence, fairness is prioritized without compromising urgent safety measures, leading to the formulation of the fair resource allocation problem as shown in Eqn. 4.2.

$$\min \quad \sum (\omega_1 \cdot \lambda_p - \omega_2 \cdot \mathcal{LS}()) \tag{4.2a}$$

$$\text{subject to} \quad \lambda_p = \sum |\mathcal{F}_i| \times \alpha_i, \forall_{\mathcal{F}_i} < 0 \tag{4.2b}$$

$$\mathcal{LS}() = \sum_{i=0}^{n_{P_i}} \mathcal{F}_i \tag{4.2c}$$

$$\omega_1 + \omega_2 = 1 \ , \ 0 \leq \omega_1, \omega_2 \leq 1 \tag{4.2d}$$

This corresponds to a multi-objective optimization problem with two objectives: improving efficiency by minimizing the extra time λ_p incurred when assigning non-preferred resources to residents, and achieving fairness by maximizing residents' satisfaction $\mathcal{LS}()$. Thus, ω_1 and ω_2 are the efficiency and fairness weights, respectively. For instance, setting $\omega_1 = 0.4$ and $\omega_2 = 0.6$ indicates a greater emphasis on fairness over efficiency in the allocation.

Theorem: *Our fair resource allocation problem is NP-hard.*

Proof: We demonstrate its NP-hardness by reducing from the makespan optimization problem, a classic NP-hard problem in scheduling theory. Consider an instance of the makespan problem with n jobs and m machines, where each job i has a processing time p_{ij} on machine j. We construct an analogous instance of our resource allocation problem as follows: each job corresponds to a resident, and each machine corresponds to a resource, either human or equipment. Let $p_{ij} = TTC()$ represent the processing time of resident i when assigned to resource j. In the simple case of identical machine scheduling problems where the processing time of each job is the same on each machine, minimizing the weighted average completion time is NP-hard by reduction from the Knapsack problem [86]. Even if there are only two machines, the problem remains NP-hard by reduction from the partition problem [135]. In our resource allocation problem, the processing times are unrelated. For instance, machine i could have $p_{ij} > p_{ij'}$, while machine i' could have $p_{i'j} < p_{i'j'}$. This corresponds to an instance of the unrelated parallel machines scheduling problem where the objective is to minimize the makespan (Eqn. 4.3).

$$\min \quad \sum_{j=1}^{n} p_{ij} x_{ij} \leq C_{max} \tag{4.3a}$$

$$\text{subject to} \quad \sum_{i=1} x_{ij} = 1, \forall_{ij} \in \{0, 1\} \tag{4.3b}$$

Our reduction from the well-established NP-hard makespan problem shows that our resource allocation problem is at least as complex as the makespan problem. Consequently, our problem inherits the NP-hard classification. In the subsequent section, we outline our methodology for achieving fairness.

4.3 The iFair Approach

4.3.1 Overview of the iFair Approach

We prototyped this framework within a data exchange platform that aims to enhance the resilience of older adults in aging communities (CareDEX). We assume the availability of static healthcare and dynamic smart space information about residents; sensors embedded in the space are used to collect data about residents' location and movement [34, 110, 117]. Information from health records and caregiver logs can provide additional data about health and ADLs to create a comprehensive picture of each individual's needs. This information is then shared with other stakeholders including first responders, family, and caregivers [1]. Fig. 4.3 provides an overview of key components in the iFair framework.

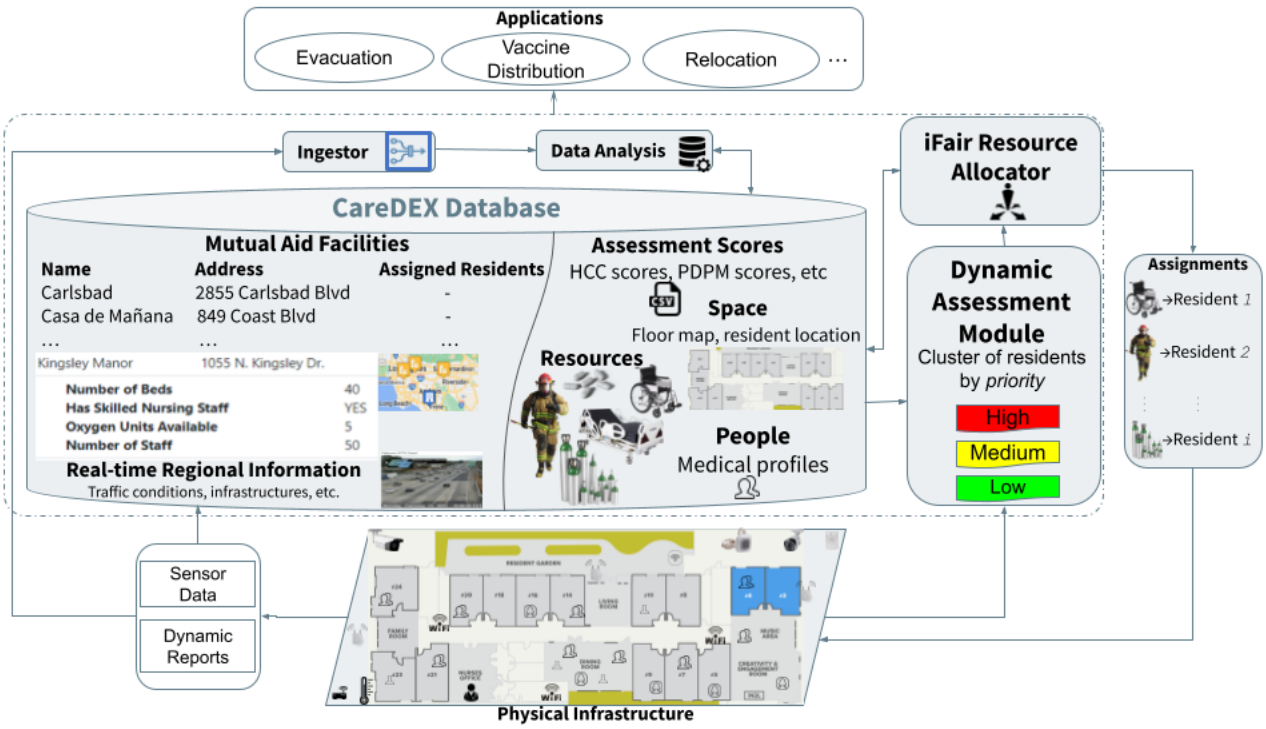

Figure 4.3: Overview of the iFair Framework.

(a) Physical Infrastructure and Occupants: First, we collect data about the residents and resources in the facility using a variety of records and sensors for localization and tracking, such as motion sensors, emergency call buttons, and Wi-Fi sensing. Indoor localization is not the main focus of this work; however, through the CareDEX project, we conducted multiple drills that allowed us to approximate participant locations. Thus, we have some real-world experience in this area. All data collection complies with the Health Insurance Portability and Accountability Act of 1996 (HIPAA) regulations, utilizing encryption and data anonymization to ensure the privacy and security of resident information [117]. Detailed

privacy-preserving sensing mechanisms are outside the scope of this dissertation.

(b) Resident/Resource Database: Next, we store and analyze this data to gain insights about the residents. We also use the database for storing the residents' medical profiles and ICD-10-CM (International Classification of Diseases, Tenth Revision, Clinical Modification) [28] coding data, including real-time regional information about mutual aid facilities (e.g., number of beds or oxygen units available).

(c) Dynamic Assessment Module: We use this module to assess the residents and update their criticality scores, a measure of their wellness and needs. Furthermore, we cluster the residents into three groups based on their criticality. This module also handles real-time changes that affect the priority of residents. For example, if a resident classified as low-priority based on criticality is injured during an evacuation.

(d) Resource Allocator: Based on all known information, this module assigns the resources (see Alg. 3). Overall, this framework can serve a range of purposes such as facilitating the evacuation of residents, improving the distribution of vaccines during an outbreak, and enabling swift relocation of residents in the face of a natural disaster.

4.3.2 Task-Based Assessment and Scoring

We collaborated with geriatric experts at the University of California, Irvine (UCI) to explore leading assessment tools for evaluating older adults. This exercise led to the adoption of the ICD-10-CM system for resident assessment and scoring. The ICD-10-CM, developed by the Centers for Disease Control and Prevention (CDC) and based on the World Health Organization's (WHO) alphanumeric disease classification codes, is widely used in the U.S. healthcare system for disease classification. In line with HIPAA regulations, healthcare professionals use these codes for patient assessments, assigning risk scores for various purposes such as determining care levels, classifying healthcare services and pricing, and assessing cognitive and mobility functions among older adults. Assessments are performed when a resident is admitted to a facility, quarterly, and annually. In addition, residents are assessed whenever they experience a significant change in status, and whenever the facility identifies a significant error in a prior assessment.

a) Offline Criticality Score - (computed a priori based on medical profile): For evaluating residents' physical conditions, we analyze attributes from their medical profiles, employing a risk-adjustment model to calculate their Risk Adjustment Factor (RAF) scores. This study adopts the Hierarchical Condition Category (HCC) model [12] for this purpose.

Consider a resident with a severe head injury coded as T31.11, equating to an HCC value of 0.486. Adding dementia increases their score by 0.346. A BMI between 40.0-44.9 adds another HCC value of 0.25. Summing these values gives $C_{score} = 0.486 + 0.346 + 0.25 = 1.082$ (refer to lines 4-5 of Alg. 1). This method is applied to all residents, noting that those with a score over 1.0 are regarded as having severe health issues. During an urgent evacuation, attributes like the weight of a resident and their ability to move independently are crucial. To address this, we define important attributes $\mathcal{I}$ and desired values v_{α}, e.g., $\mathcal{I} = \{BMI, Ambulatory\}$ with $v_{BMI} > 25$ for overweight, and $v_{Ambulatory} = False$ for mobility issues. A set of weights is then applied to emphasize these attributes' importance. Consequently, residents matching these criteria have their HCC values adjusted, thereby updating their criticality score as outlined in line 6 of Alg. 1. This adjustment yields a weighted C_{score}, our measure for offline criticality.

b) Online Criticality Score - (computed in real-time based on event urgency): During emergencies, criticality scores are updated to reflect the urgency of care needed by residents, as outlined in Alg. 2.

In a fire scenario, for instance, a resident's criticality score is modified based on their closeness to the fire, enabling a reevaluation of their priority. This could lead to assigning a higher

Algorithm 1: Offline Criticality Score Assessment

Input: Residents $\mathcal{P}$, Important attributes $\mathcal{I}$, Desired attribute value v_α, Weight of attribute ω_α

Output: List of Criticality Scores

$List\_C_{score} \leftarrow \varnothing$

for $resident\ in\ \mathcal{P}$ **do**

  $C_{score} \leftarrow 0$ `// to update using ICD-10-CM`

  for $\mathcal{D}_i \leftarrow resident.diagnoses$ **do**

    $C_{score} \leftarrow C_{score} + hcc\_val(\mathcal{D}_i)$

  $C_{score} \leftarrow \sum_{\alpha \in \mathcal{I}} \omega_\alpha \cdot HCC\_value_\alpha, \forall_{\alpha = v_\alpha}$

  Add C_{score} to $List\_C_{score}$

return SortDescending($List\_C_{score}$)

Algorithm 2: Online Criticality Score

Input: $\mathcal{P}$, $\mathcal{G}$, $List\_C_{score}$

Output: $List\_C_{score}$

$v_o \leftarrow \mathcal{G}.loc(event, t = 0)$

for $P_i \leftarrow \mathcal{P}$ **do**

  $d_{P_i} \leftarrow ShortestDistance(loc(P_i, t), v_o)$

  $U_{P_i} \leftarrow \frac{1}{d_{P_i}}$ `// compute urgency`

  $U\_List.add(U_{P_i})$

for $U_{P_i} \leftarrow U\_List$ **do**

  Normalized $U'_{P_i} = \frac{U_{P_i} - min(U\_List)}{max(U\_List) - min(U\_List)}$

  Update $List\_C_{score}[P_i] \leftarrow List\_C_{score}[P_i] * U'_{P_i}$

return SortDescending($List\_C_{score}$)

priority to those in immediate danger, thus allowing for a dynamic response to the unfolding situation. By combining the CDC's established medical scoring framework with real-time data from sensors across the facility, we dynamically gauge the new criticality of each resident.

Next, we allocate resources to residents as shown in Alg. 3, cycling through sorted criticality scores to match residents with tasks, which involves setting task names (e.g., $provide\_basic\_care()$), identifying dependencies, and determining necessary resources with preferences (lines 3-6). We assign preferred resources within d_{max} and, to avoid starvation, assign any available human resource when no preferred resource is available (lines 7-12). We then compute task completion times and satisfaction (lines 13-14), update resource states, and store these values in the allocation record (lines 15-21). Additionally, we maintain a health report form ($\mathcal{HRF}$) where we include residents experiencing changes in health conditions and use it to update their criticality scores.

Algorithm 3: Fair Resource Allocation

Input: Online criticality scores $List\_C_{score}$, Distance threshold d_{max}, Health report form $\mathcal{HRF}$

Output: Allocation records $\mathcal{A}_R$

$\mathcal{A} \leftarrow \emptyset; \mathcal{A}_R \leftarrow \emptyset$
$t_s \leftarrow$ GetCurrentTime()
while $\neg$ *isEmpty($List\_C_{score}$)* **do**
  $C_{score} \leftarrow \max(List\_C_{score})$
  $P_i \leftarrow$ GetCorrespondingResident(C_{score})
  $R_{eq}, \mathcal{F}_i \leftarrow$ GetTaskDetailsAndPrefs(P_i)
  for $rr \in R_{eq}$ **do**
    $r \leftarrow$ GetResource($\mathcal{F}_i, max\_dist = d_{max}$)
    if $r ==$ **null** // Assign human
    then $r \leftarrow$ NextAvailableHuman() ;
    $\mathcal{A} \leftarrow \mathcal{A} \cup \{(P_i, r, t_s)\}$
    SetResourceState(r, "in-use")
  // Get Task completion time, satisfaction
  $TTC() \leftarrow \sum \omega_e + \sum access(v_i, t_s) + \lambda_p$
  $\mathcal{LS}() \leftarrow \sum_{i=0}^{n_{P_i}} \mathcal{F}_i$
  for r *in* $\mathcal{A}$ **do**
    SetResourceState(r, "unavailable")
    if *reusable(r)* **then**
      $t_s, t_e \leftarrow t_e, t_e + TTC()$
      SetResourceState(r, "available")
      $\mathcal{A} \leftarrow \mathcal{A} \cup \{(P_i, r, t_e)\}$
  $\mathcal{A}_R \leftarrow \mathcal{A}_R \cup \{(TTC(), \mathcal{LS}(), \mathcal{A})\}$
  if $\neg$*isEmpty*($\mathcal{HRF}$) **then**
    for $i, sc \leftarrow$ *GetUpdatedCriticalityScores*($\mathcal{HRF}$) **do** $List\_C_{score}[i] \leftarrow sc$;
return $\mathcal{A}_R$

4.4 Experimental Evaluation

Experimental Setup: We evaluated our framework through emulation, utilizing anonymized data modeled after a real-world senior care facility in Orange County, USA, including 30 elderly residents' profiles (age, gender, medical information). With input from experts in the Department of Family Medicine, Division of Geriatric Medicine and Gerontology at UCI Health, we generated synthetic data for 400 resident profiles to facilitate larger population experiments. The offline criticality score was evaluated using the CASAS dataset [37], featuring real-world data on Activities of Daily Living (ADL), activity scores, and diagnoses for 400 residents from Washington State University. We mapped ICD-10-CM diagnoses to compute scores using HCC coding, based on residents' medical information. A simulator (written in Python) was developed to emulate resource allocation within a facility, incorporating actual floor plans (Fig. 4.4) from a partner facility.

Figure 4.4: Floor Plan of Small Facility.

Experimental Plan: When a disaster strikes, older adults are disproportionately affected [11]. During this type of event, the human resources (i.e., medical staff and first responders) needed to evacuate this frail population from their SHF to a safe location are minimal. Hence, we crafted a use case where rapid evacuation is imperative. We map this scenario to a parallel machine scheduling problem, translating staff to machines and residents to jobs. Using the formulation in Section 4.2, our goal is to reduce evacuation times (i.e., optimize the makespan C_{max} [52]) as well as the residents' satisfaction. Accordingly, the key metrics we utilize in our evaluation are time-to-completion (an efficiency/latency metric) and level of satisfaction (a fairness metric).

We next discuss strategies that are used as baselines and comparison points for iFair. Senior care facilities have emergency plans in place that correspond to the following algorithms: Shortest Job First (SJF), which aims to evacuate residents who can be moved with the least effort and in the shortest amount of time. This includes those who are fully ambulatory, require minimal assistance, or are closest to the exits. Longest Job First (LJF), on the other hand, prioritizes the most critical residents, requiring the most time and resources to evacuate. This could include those with mobility issues, those requiring medical equipment, or residents located in the most challenging parts of the facility.

To ensure a fair distribution of tasks among staff members, we also developed a Clustered Round Robin (C-RR) strategy. Here, each staff is assigned a fixed number/group of residents. For example, staff and equipment may be assigned to specific units/floors for a specific period to accomplish the task before moving on to the subsequent floor. We also developed a Criticality-aware Clustered Round Robin (CC-RR) strategy which is an improved version of the C-RR technique. CC-RR consists of assigning resources first to high-priority residents, then when there are no more residents in that cluster, it proceeds as C-RR alternating between medium-priority and low-priority residents for a given amount of time. Let us note that strict adherence to any single plan could lead to liability issues based on negligence [85]. LJF policies may overlook urgent needs for less critical, resource-intensive cases. Conversely, SJF risks neglecting high-need residents for faster evacuations, while RR's fixed order may delay aid to urgent cases.

Table 4.2: Comparative Outcome when Evacuating Residents with Preferences (WP) vs. without Preferences (NP).

Priority	Residents	C_{score}	$\mathcal{LS}$-WP	TTC-WP	$\mathcal{LS}$-NP	TTC-NP	λ_p
HIGH	Joe	2.741	5	10	4	10	0
	Mary	2.634	5	12	2	27	15
	Bob	2.527	6	6	-6	6	0
	...	...	...	...	...	...	...
MED	Don	1.813	1	10	-6	18	8
	Georgia	1.806	1	12	-6	22	10
	Laura	1.799	1	14	-6	18	4
	...	...	...	...	...	...	...
LOW	Mona	0.708	-1	17	2	18	1
	Gennifer	0.601	0	18	2	20	2
	Karole	0.0	5	11	-4	15	4
	...	...	...	...	...	...	...

We use these algorithms to assess the effectiveness and efficiency of our framework across various facility sizes. By implementing these in both smaller settings, where resource constraints may differ, and larger institutions with more complex logistical challenges, we comprehensively evaluate our approach. This enables us to adapt and refine our strategies, optimizing evacuation procedures and ensuring equitable care delivery in diverse operational environments.

Experimental Results: Although the evacuation time for some residents, such as Joe, remains the same whether preferences are used or not, most evacuation times differ substantially between scenarios where resident preferences are considered (TTC-WP) and scenarios where they are not included (TTC-NP), as shown in Table 4.2. Notably, Mary experiences a 15-minute increase in evacuation time without preferences, rising from 12 to 27 minutes. This indicates that ignoring preferences can lead to inefficient routing or resource allocation, significantly increasing evacuation times.

Moreover, computing λ_p across different resident profiles further illustrates this disparity. This measure quantifies the additional time or delay incurred when resident-specific preferences and needs are overlooked. For instance, Mary's high λ_p value of 15 suggests considerable

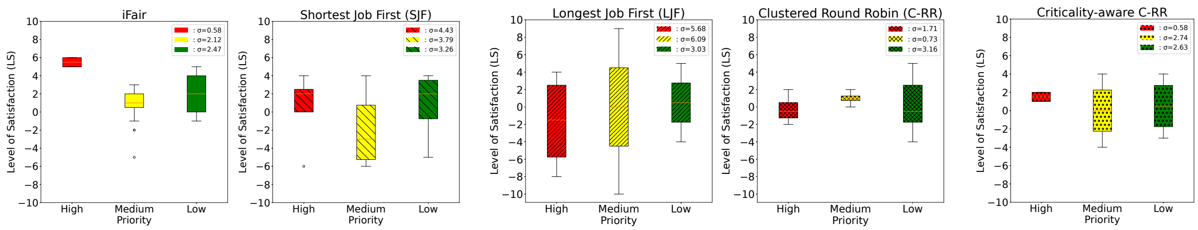

Figure 4.5: Residents' Satisfaction Across Different Approaches for a Small Facility

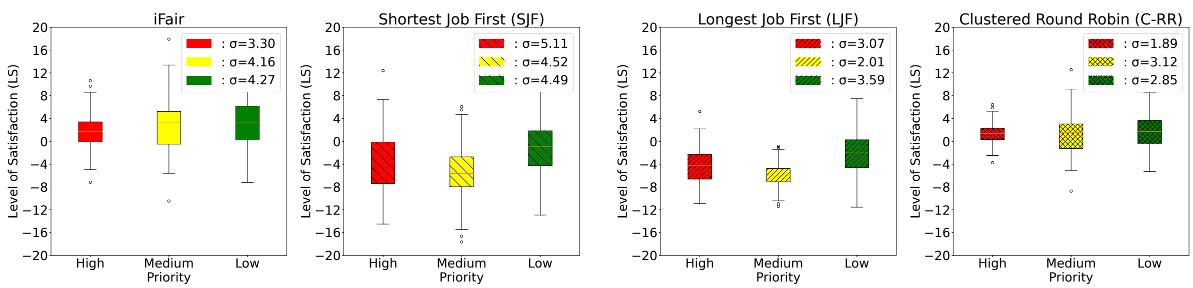

Figure 4.6: Residents' Satisfaction Across Different Approaches for a Large Facility

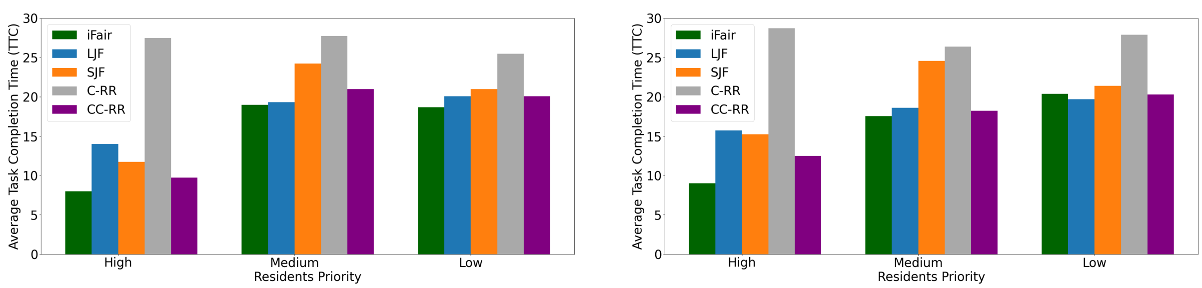

(a)

(b)

Figure 4.7: Comparing Average Task Completion Times: (a) Small Facility vs. (b) Large Facility.

neglect of her specific needs or location, which is adequately addressed when preferences are considered, resulting in a more optimized and faster evacuation process under the TTC-WP scenario. iFair, particularly in a small facility as shown in Fig. 4.5, achieves the lowest variability in satisfaction, with standard deviations of 0.58 for high, 2.12 for medium, and 2.47 for low priority levels, indicating more equitable satisfaction among residents. It surpasses the other methods in terms of satisfaction spread. The Clustered Round Robin method shows moderate variability, suggesting the importance of a more personalized approach.

In larger facilities, as illustrated in Fig. 4.6, iFair maintains lower variability in satisfaction, highlighting its effectiveness and scalability in evacuations that consider resident preferences. The figure shows that iFair outperforms other methods by achieving higher satisfaction levels across high, medium, and low-priority residents. The satisfaction levels depicted in Fig. 4.6 for iFair are higher compared to SJF and LJF, especially for residents of medium and low priority, where the need to balance resources can often leave these groups less attended in other models. Even when compared to Clustered Round Robin (C-RR), iFair demonstrates a more favorable satisfaction profile.

In Fig. 4.7, iFair consistently provides the shortest average task completion time (TTC) for high-priority residents in both facilities, demonstrating efficient prioritization. LJF and SJF show moderate TTCs, indicating less effective prioritization, while C-RR records the longest TTC, showing a slower response. Likewise, the efficiency gap between iFair and the other methods, especially for high-priority tasks in large facilities, highlights iFair's scalability and effective prioritization. CC-RR, despite high TTCs, suggests potential scalability close to iFair's performance.

The Gantt charts in Fig. 4.8 indicate that iFair achieves a 68-minute makespan, setting a benchmark. Compared to other methods, it improves the makespan by 20% over LJF, 35.85% over SJF, 22.47% over CC-RR, and 37.04% over C-RR. iFair's efficiency stems from prioritizing urgency and preferences, whereas the comparison methods lead to longer completion times. Overall, our findings demonstrate that our framework is equitable, as supported by the satisfaction of the residents, and it is also an efficient system.

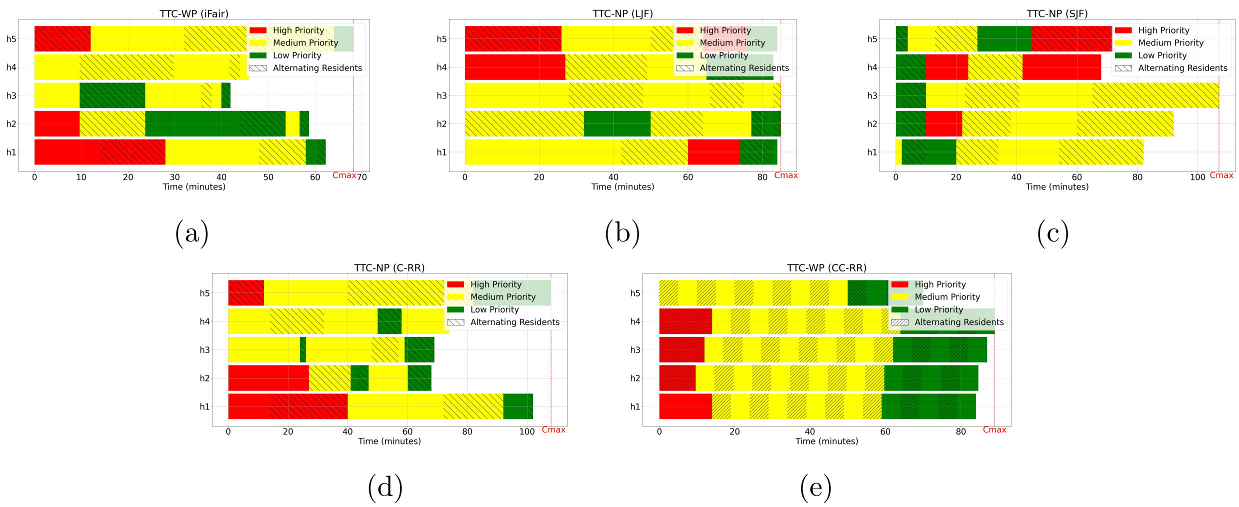

Figure 4.8: Gantt Charts of Human Resource Allocation, using (a) iFair; (b) Longest Job First; (c) Shortest Job First; (d) Clustered Round Robin; and (e) Criticality-aware Clustered Round Robin for a Small Facility.

4.5 Extensions for Regional Resource Allocation and Decision Support

Although *iFair* focuses on intra-facility allocation, our work with senior health facilities through *CareDEX* [1] also revealed the need for regional coordination during larger disasters. A facility may not always be able to resolve all needs internally. Residents may need to be relocated, resources may need to be shared across mutual-aid facilities, and decision-makers may need to account for hazards, road conditions, facility status, and available capacity across a broader healthcare network. *CareDEX* provided an early regional-awareness capability for viewing facility impacts, resource availability, evacuation needs, and potential relocation options. This included GIS-based visualization of hazard conditions, facility locations, mutual-aid resources, and disaster-specific information such as proximity and severity. These capabilities helped stakeholders reason about regional impacts and highlighted the importance of connecting facility-level resident needs with cross-facility resource coordination. However, this early regional tool primarily supported situational awareness and relocation exploration. It helped decision-makers see impacted facilities, available resources, and possible relocation options, but it did not fully answer which facilities should be paired, which routes should be used, or how recommendations should change as hazard and facility conditions evolve. These limitations motivate the next chapter, which presents a framework for regional matching and routing during disasters.

4.6 Chapter Summary and Discussion

We presented *iFair*, a fairness-aware framework for allocating scarce resources within senior health facilities during urgent events. The central motivation was that older adults should not be modeled as interchangeable units in an allocation problem. Residents differ in medical conditions, mobility, preferences, resource needs, and urgency, and these differences directly affect both fairness and response efficiency. To capture this setting, the framework models the facility as a graph-based representation, incorporates resident profiles and preferences, and updates resident criticality using both offline medical information and real-time event conditions.

The results show that incorporating resident preferences and criticality into allocation decisions can improve resident satisfaction while also supporting operational efficiency. This finding is important for the broader dissertation because it shows that fairness does not necessarily have to come at the cost of response time. In senior care settings, ignoring individualized needs can introduce delays, reduce satisfaction, and produce less effective care. By accounting for these needs directly, allocation decisions can become more equitable and better aligned with the practical realities of emergency care delivery.

Although fire evacuation was used as the motivating scenario, the modeling approach can apply to other facility-level disruptions. Equipment shortages, staff shortages, infectious disease outbreaks, localized infrastructure failures, or internal emergencies may also require administrators to prioritize residents and assign scarce resources under uncertainty. In each case, the allocation process must account for facility layout, resource availability, task requirements, resident urgency, and the human factors that shape care delivery.

At the same time, in-facility allocation is only one part of healthcare disaster resilience. During larger disasters, a senior health facility may not be able to resolve all needs internally. Residents may need to be relocated, hospitals may face surge demand, and resources may need to be shared across a broader healthcare network. This limitation motivates the next chapter, which extends the allocation problem from a single facility to regional matching and routing across multiple healthcare facilities.

Chapter 5

Optimizing Regional Matching and Routing for Disaster Response

The previous chapter focused on fair allocation within a single senior health facility. During larger disasters, however, internal allocation may not be enough. A facility may need to evacuate residents, hospitals may experience surge demand, roads may become congested or unsafe, and scarce resources may need to be shared across a broader healthcare network. In these situations, decision-makers must determine not only who needs support, but also where people and resources should go and how they can safely get there. We develop *RADAR*, a geospatial decision-support framework for matching and routing during disasters. It integrates multisource GIS feeds, facility status, transportation data, resource availability, and operational policies to support time-sensitive coordination across healthcare facilities. Rather than treating relocation or resource sharing as a nearest-facility assignment, the framework uses policy-driven ranking, stable matching, and hazard-aware routing to identify feasible and operationally appropriate assignments.

The framework is evaluated across several disaster settings, including an emulation of the

Palisades wildfire and tabletop drill scenarios involving earthquake-related hospital surge and county-wide water contamination. These scenarios demonstrate how geospatial data, policy-driven matching, and routing can support evacuation planning, resource delivery, and healthcare coordination during regional disruptions.

5.1 Chapter Overview

This chapter focuses on the coordination problem introduced in Chapter 3. While Chapter 4 addressed allocation within a single senior health facility, the problem now expands to a healthcare network in which multiple facilities, hazards, resources, and transportation constraints interact. Decision-makers must determine which facilities should be matched, which resources or evacuees should be moved, and which routes remain feasible as conditions evolve.

Through the CareDEX disaster resilience effort [1], we participated in drills and tabletop exercises with senior care facilities, healthcare partners, public health stakeholders, and emergency response personnel. These exercises revealed a recurring operational gap: stakeholders need a systematic way to match impacted facilities with appropriate receiving facilities and to route people or resources through disrupted transportation networks. Manual coordination may be feasible in small scenarios, but it becomes difficult as hazards evolve, roads close, facility status changes, and resource requests arrive from multiple locations.

To address this gap, *RADAR* integrates hazard feeds, facility data, resource availability, and operational policies into a GIS-centric decision-support workflow. Rather than assigning evacuees or resources to the nearest facility by distance alone, the system uses policy-driven scoring to generate preference lists for candidate matches. These policies capture operational considerations such as capacity, same residency type, proximity, safe routes, triage priority, facility status, and social vulnerability. Stable matching is then used to produce assignments that are feasible, interpretable, and aligned with policy preferences.

The framework supports both aggregate and fine-grained allocation. The aggregate model is designed for rapid coordination, where impacted facilities may need to relocate groups of patients or residents to receiving facilities with available capacity. This is useful in short-term care, hospital surge, and emergency operations settings where decisions must be made quickly. The fine-grained extension incorporates resident-level medical profiles, preferences, care continuity, and social constraints, making it more appropriate for long-term senior care settings where personalized placement matters.

An important feature of the framework is its ability to adapt as conditions change. Facility status, transportation conditions, and hazard boundaries may evolve during a disaster. Instead of recomputing all assignments from scratch, the framework preserves valid matches when possible and updates only the assignments affected by new conditions. Routing is then computed over the transportation network using GIS-based network analysis and shortest-path methods, allowing the system to evaluate feasible routes between matched facilities.

The contributions are as follows:

- A GIS-centric decision-support framework that couples hazard data, facility status, transportation information, resource availability, and policy-driven matching.
- A policy-based scoring model for generating facility preference lists under operational constraints such as capacity, safety, proximity, resource availability, operational status, and social vulnerability.
- A stable matching approach for assigning impacted facilities to appropriate receiving facilities in a way that is interpretable and compatible with emergency operations workflows.
- A dynamic update strategy that preserves valid assignments while recomputing affected matches as hazards, roads, or facility conditions change.

- Scenario-based evaluation across wildfire, earthquake, and water-contamination settings, demonstrating the framework's ability to support evacuation planning, hospital surge coordination, and resource-routing decisions.

The remainder of this chapter is organized as follows. Section 5.2 presents the overall approach and system architecture. Section 5.3 formulates the regional matching and routing problem. Section 5.4 describes the policy functions, matching model, and dynamic update strategy. Section 5.5 presents the experimental evaluation across disaster scenarios. Finally, Section 5.7 summarizes the chapter and discusses its role in the broader dissertation.

5.2 The RADAR Approach

Building on the regional coordination gap identified in Chapter 2, there is a clear need for a framework that integrates assignment, routing, and human oversight under uncertainty. With RADAR, we ingest live facility and hazard feeds, quantify multidimensional policy objectives, and provide actionable recommendations to Emergency Operations Centers (EOCs). This combines real-time data integration, dynamic matching, and adaptive routing within a unified, modular architecture. By coupling spatio-temporal GIS streams with predictive hazard forecasts, the system can proactively recompute allocations as conditions evolve, while incremental matching updates preserve allocation stability and bound computational overhead. Human oversight is embedded at key decision points through an interactive dashboard that presents policy rationales, supports "what-if" scenario analysis, and enables operators to adjust weights or override plans based on new events. Moreover, RADAR's multi-objective policy engine explicitly balances efficiency, equity, and resilience, allowing EOC teams to tailor trade-offs in real time.

Our framework comprises two main components: (i) an aggregate model that allocates resources using coarse policies and (ii) a fine-grained extension that incorporates resident-

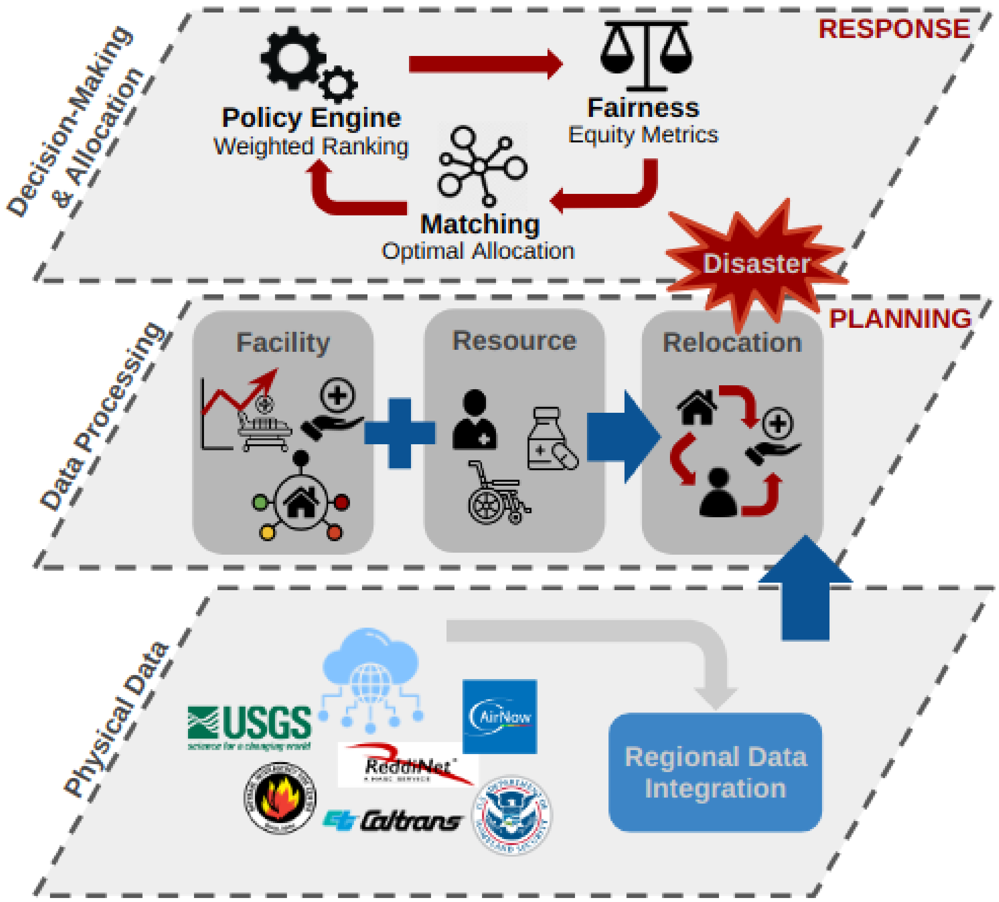

Figure 5.1: RADAR Architecture

specific information (e.g., medical needs, personal preferences) into the allocation. In both cases, facilities rank candidate partners based on weighted policy scores, and through stable matching, we pair those evacuating with the receiving facilities. The fine-grained extension further adjusts assignments using individual preferences, medical requirements, and historical relocation data. The architecture consists of several modules (see Fig. 5.1).

At the physical layer, a regional module collects real-time data from external sources, including seismic activity, damage assessments, wildfire perimeters, air quality, hospital capacity, bed availability, and traffic conditions. Static and dynamic facility data are maintained in the facility information module, while the resource management module monitors demand and processes resource requests in real time. The relocation module ranks potential receiving facilities based on compatibility metrics such as medical needs, residency type, and safety criteria. A policy engine applies weighted ranking rules to generate comprehensive policy scores. These scores feed into a stable matching module, which produces allocations that are stable and envy-free. Finally, a fairness evaluator computes penalties to capture deviations from preferred allocations, aggregating these into concise equity metrics that EOC staff can monitor and balance against efficiency goals.

Thus, by combining dynamic data-driven matching with hazard-aware, human-in-the-loop routing, RADAR converts static allocation paradigms into an adaptive decision-support framework for disaster resilience. In the next section, we will discuss the mathematical formulation of this matching and routing problem.

5.3 Problem Formulation

We formulate the allocation as a policy-driven optimization that integrates real-time facility conditions, mobility, and social vulnerability. The following section outlines the key components, decision variables, and constraints of this problem.

5.3.1 Modeling Key Components

Modeling Facilities: Each facility F_i in the set $\mathcal{F} = \{F_1, F_2, \ldots, F_n\}$ has an identifier ID_i, name N_i, and coordinates L_i (latitude/longitude). Each facility is classified by a type T_i (Senior Healthcare Facility, Hospital, and Non-Healthcare Facility) and has a maximum capacity C_i. The availability of resource R_k at facility F_i is denoted by A_{ki}, and inter-facility distances and travel times are measured by the function $D(F_i, F_j)$. Their operational status O_i is one of: G (Green), normal operations; Y (Yellow), under control; O (Orange), modified services requiring some assistance; R (Red), limited services with significant assistance; and B (Black), no services. To capture preparedness, we use the Social Vulnerability Index $SVI_i \in [0, 1]$, which quantifies a community's susceptibility to a disaster based on 16 socioeconomic factors defined by the CDC [68].

Modeling People: A person P_j in the set $\mathcal{P} = \{P_1, P_2, \ldots, P_m\}$ is identified by ID_j with a location L_j. We use R_j to specify the facility type in which they reside. Every individual has specific resource needs $R_n(P_j)$ and is assigned a triage category $T_j \in \{R, Y, G, B\}$, where R denotes critical, Y urgent, G minor, and B deceased or non-responsive as defined by the START triage model [13].

Modeling Resources: We represent a resource R_k as a component of the set $\mathcal{R} = \{R_1, R_2, \ldots, R_p\}$. For each facility, the available quantity of resource R_k is denoted by Q_{ki}. A binary parameter M_{ki} indicates resource mobility, where $M_{ki} = 1$ if the resource is relocatable and 0 otherwise. Similarly, resource reusability U_{ki} is a binary variable whose value is 1 when the resource is reusable and 0 otherwise. We consider three categories of resources: *Human Resources* (H_R), such as doctors, nurses, caregivers, and EMTs, essential for operating equipment and delivering critical services; *Healthcare Equipment Resources* (E_R), e.g. beds or oxygen tanks that depend on human resources for effective use; and *Service Resources* (S_R), which includes essential services like rehabilitation therapy or postmortem care that rely both on human and equipment resources.

Modeling Hazards: We model a hazard H_h as an event in $\mathcal{H} = \{H_1, H_2, \ldots, H_l\}$. Each hazard is identified by ID_h and characterized by its type T_h (e.g., earthquake, wildfire) and location L_h. The hazard's impact is delineated into two regions: the *Impacted Region* Z_{imp}, corresponding to the red zone directly affected, and the *Warning Region* Z_{warn}, the broader at-risk (yellow) zone. To model hazards in real time, we ingest road-condition feeds (lane/full closures, highway incidents) and label segments as open or closed, inferring short interior egress stretches so evacuees inside Z_{imp} can reach the boundary on open, outward paths. For example, Caltrans QuickMap [27] publishes road feeds with most layers refreshing about every minute and some every 5–10 minutes. Cascading effects, denoted by S_h, such as power outages, gas leaks, or structural collapses, may exacerbate the overall disaster impact.

Modeling Tasks: We define a task $T_x \in \mathcal{T} = \{T_1, T_2, \ldots, T_q\}$ as an action driving the allocation process. Tasks are categorized into: (1) *Receiving People*, (2) *Relocating People*, (3) *Requesting Resources*, and (4) *Sending Resources*. A facility receiving evacuees indicates its capacity C_i and available resources A_{ki}. Similarly, a facility requesting resources specifies a request time t_r, a request type R_t (initial, update, or final), and its resource needs $R_n(F_i)$ (i.e., human, equipment, or service resources). For relocation tasks, a timestamp t_m is recorded, and the assignment depends on the triage priority.

Modeling Policies: Our final key component is the set of policy-driven ranking functions. Each policy p is formalized as a score function $S_p(i, j) \in [0, 1]$ (e.g., triage priority, operational status, proximity, safe-route availability, and social vulnerability). These scores define the preference lists used in our stable matching procedures. We discuss them in depth in §5.4. Therefore, we address this problem by first using stable matching techniques to maximize the overall policy score; then, routing using Google's Distance Matrix API and ArcGIS Pro historical traffic data [77, 141] to capture real-world travel conditions. Hence, we integrate them into a unified approach that optimizes both matching and routing efficiency.

5.3.2 RADAR Optimization Problem

We formulate this problem around two central research questions: (1) How can evacuees be **matched** to facilities to maximize a composite policy-driven objective, integrating safety thresholds, capacity utilization, and response timeliness under dynamic hazard and operational constraints? (2) How can **routing** be computed to minimize travel time and distance for relocation, while considering evolving facility statuses and network accessibility constraints?

Let $S_p(F_i, F_j)$ denote the policy score for matching a facility F_i relocating-people with a facility F_j receiving-people, and let w_p be the corresponding weight of the policy. We introduce a binary decision variable $x_{ij} \in \{0, 1\}$ such that $x_{ij} = 1$ if facility F_i (with $i \in S_{rel}$) is matched to facility F_j (with $j \in S_{rec}$), and 0 otherwise.

$$
\begin{aligned}
\max_{x_{ij}} \quad & \sum_{i\in S_{rel}} \sum_{j\in S_{rec}} w_p\, S_p(F_i,F_j)\, x_{ij} && \text{(5.1a)}\\
\text{s.t.} \quad & \sum_{j\in S_{rec}} x_{ij} \le 1, \quad \forall\, i \in S_{rel} && \text{(5.1b)}\\
& \sum_{i\in S_{rel}} x_{ij} \le 1, \quad \forall\, j \in S_{rec} && \text{(5.1c)}\\
& \sum_{i\in S_{rel}} R_i\, x_{ij} \le C_j, \quad \forall\, j \in S_{rec} && \text{(5.1d)}\\
& x_{ij} \le S_{\Delta}(F_i,F_j), \quad \forall\, i \in S_{rel}, \forall\, j \in S_{rec} && \text{(5.1e)}\\
& D(F_i,F_j) \le D^{\max}, \quad \forall\, i \in S_{rel}, \forall\, j \in S_{rec} && \text{(5.1f)}\\
& x_{ij} \in \{0,1\}, \quad \forall\, i \in S_{rel}, \forall\, j \in S_{rec}. && \text{(5.1g)}
\end{aligned}
$$

Furthermore, R_i denotes the number of evacuees (or resource demand) at facility F_i. In addition, we use C_j to indicate the capacity available at facility F_j, and $S_{\Delta}(F_i,F_j)$ is a binary parameter that equals 1 if a safe route exists between F_i and F_j and 0 otherwise. The distance between facilities is given by $D(F_i,F_j)$, with a maximum acceptable threshold $D^{\max}$. Hence, the overall optimization problem is expressed in Eqn. (5.1).

Constraints (5.1b) and (5.1c) ensure a one-to-one matching between facilities relocating people and facilities receiving people. Constraint (5.1d) guarantees that the allocated people do not exceed the resource capacity at the receiving facilities. Constraint (5.1e) limits assignments to facility pairs connected by safe routes, while Constraint (5.1f) ensures that travel distances remain within a predefined threshold, as detailed in Section 5.4. Together, this forms an integrated framework, simultaneously considering dynamic and static attributes of *facilities*, *people*, *resources*, *hazards*, and *tasks* in a disaster response.

5.4 Decision-making in the RADAR Model

In this section, we present the set of policy-driven ranking rules used for evacuee relocation (§5.4.1) and introduce our model (§5.4.2). In addition, we illustrate two extensions of our approach: the application of policy-driven rules for requesting and sending resources (§5.4.3) and a personalized extension that incorporates individual-level evacuee preferences (§5.6.1).

5.4.1 Defining Policies for Receiving and Relocating People

A receiving facility F_{rec} prioritizes evacuees from relocating facility F_{rel} based on the following rules:

Triage status: Individuals are prioritized by triage category [13], ensuring that limited resources are allocated efficiently. Those in critical condition receive the highest priority for survival:

$$S_{\text{tri}}(F_{\text{rel}}, P_i) = \begin{cases} 1, & \text{if } T_i = \text{Red (Immediate)}, \\ 0.75, & \text{if } T_i = \text{Yellow (Delayed)}, \\ 0.5, & \text{if } T_i = \text{Green (Minor)}, \\ 0, & \text{if } T_i = \text{Black (Deceased)}. \end{cases} \tag{5.2}$$

Same residency type: They also prioritize evacuees from facilities of the same type (e.g., SHF to SHF).

$$S_{\text{res}}(F_{\text{rel}}, F_{\text{rec}}) = \begin{cases} 1, & \text{if } T_{\text{rel}} = T_{\text{rec}}, \\ 0, & \text{otherwise.} \end{cases} \tag{5.3}$$

Proximity for faster relocation: To minimize delays, facilities closer to the receiving facility are prioritized. We normalize this score over the candidate set so that it remains within [0,1]. Hence, this policy is defined as follows:

$$S_D(F_{rel}, F_{rec}) = 1 - \frac{D(F_{rel}, F_{rec})}{\max\limits_{F_x \in N_{rel}} D(F_{rel}, F_x)}, \tag{5.4}$$

where $D(F_{rel}, F_{rec})$ is the distance between F_{rel} and F_{rec}, and N_{rel} denotes the set of n nearest candidate facilities. Under this formulation, closer facilities receive higher scores, while the farthest candidate in the set receives a score near zero. To compute distances and travel times between multiple origins and destinations, we used Google's Distance Matrix API [77], which incorporates real-time traffic and network conditions. The API processes concurrent HTTPS requests and returns JSON responses, enabling rapid distance assessments. This capability allows us to integrate real-world routing data into the decision-making process.

Operational status: Relocating-facilities with lower operational readiness are given higher priority. That is, facilities with limited services (Red) or no services (Black) are prioritized.

$$S_O(F_{rel}, F_{rec}) = \begin{cases} 1, & \text{if } O_{rel} = \text{Red or Black}, \\ 0.75, & \text{if } O_{rel} = \text{Orange}, \\ 0.5, & \text{if } O_{rel} = \text{Yellow}, \\ 0, & \text{if } O_{rel} = \text{Green} \end{cases} \tag{5.5}$$

When relocating evacuees, relocating facilities F_{rel} apply similar criteria. They prefer receiving facilities of the same residency type (Eqn. 5.3) and prioritize those not nearing full capacity (e.g., bed limits). Only receiving facilities with operational status $O_{rec} = \textit{Green}$ are eligible to receive evacuees.

Proximity to impacted region: Facilities can receive evacuees if they are located at a safe distance from the impacted region (as shown in Fig. 5.2). Let $d = D(F_{rec}, C)$ be the distance from a facility to the impacted zone, and r the radius of the impacted zone. We define x as

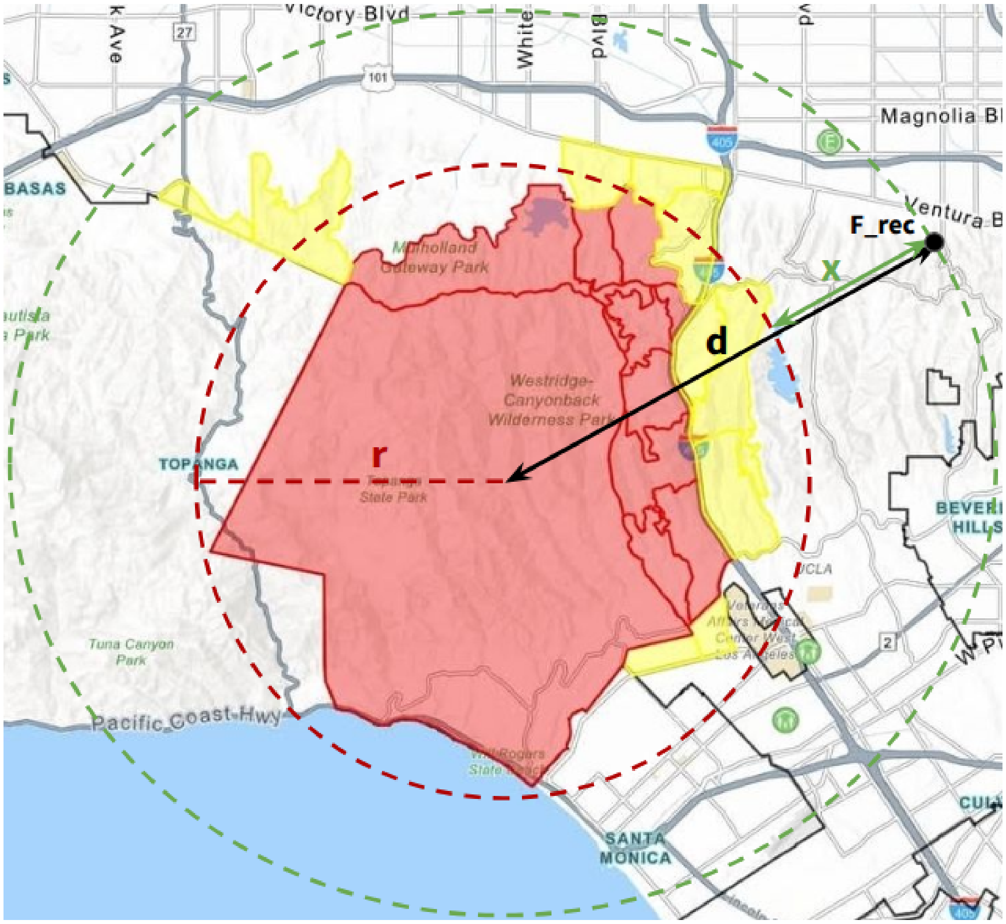

Figure 5.2: Impacted Region (red) and Mandatory Evacuation Warning Zone (yellow) (Palisades Fire 01/10/25, 7 pm PST)

the additional safety margin such that a facility is considered safe only if $d > r + x$. Hence, facilities prioritize those that are nearest to that threshold. We define this as follows:

$$S_{\text{imp}}(F_{\text{rec}}) = \begin{cases} 0, & \text{if } d \le r + x, \\ \\ \frac{r+x}{d}, & \text{if } d > r + x, \end{cases} \tag{5.6}$$

Safe routes: The availability of safe transportation paths is critical. Hence, facilities that can be reached via safe routes are prioritized. We denote this as follows:

$$S_{\Delta}(F_{rec}, F_{rel}) = \begin{cases} 1, & \text{if safe routes exist,} \\ \\ 0, & \text{otherwise.} \end{cases} \tag{5.7}$$

Social Vulnerability Index (SVI): Facilities in areas with high SVI [68] are also prioritized as shown in Eqn. 5.8. The higher the index, the more vulnerable the region. Hence, the maximum SVI is taken over a set N_{rec} of candidate receiving facilities (Eqn. 5.9). The range is as follows:

$$SVI = \begin{cases} \textit{Very Low}, & \text{if } 0 < SVI \leq 0.25, \\ \\ \textit{Low}, & \text{if } 0.25 < SVI \leq 0.5, \\ \\ \textit{High}, & \text{if } 0.5 < SVI \leq 0.75, \\ \\ \textit{Very High}, & \text{if } 0.75 < SVI \leq 1 \end{cases} \tag{5.8}$$

$$S_{SVI}(F_{rec}, F_{rel}) = \frac{SVI_{rel}}{\max\limits_{F_x \in N_{rec}} SVI_x}. \tag{5.9}$$

Prior to the matching process, we compute the final score for relocating people to facilities as the sum of all policy scores. Although many of these scores are linear (i.e., proportional and constant-slope forms), they are grounded in real-world operational data feeds.

5.4.2 The RADAR Model and Algorithms

This is our baseline model, which minimizes computational complexity by allocating resources in clusters rather than on an individual basis. It uses aggregate information to provide EOCs with rapid situational awareness and effective resource distribution in time-critical disaster scenarios. Inspired by the Gale-Shapley stable matching (SM) problem [71], we use predefined policies to score and rank facilities' preferences. On one side, facilities that are relocating, and on the other side, receiving facilities (potential candidates) that can accommodate the evacuees. Once ranked, we aim to achieve a two-sided optimal pairing between the facilities. This involves optimally pairing two disjoint sets while ensuring no unmatched pairs prefer each other over their assigned matches.

Given two sets: $A = \{a_1, a_2, \ldots, a_n\}$ and $B = \{b_1, b_2, \ldots, b_m\}$, each agent a_i in A has a preference ranking over agents in B, and vice versa. Each agent maintains a preference ordering, where $P_{a_i} = \{b_j \succ b_k \succ \cdots\}$ indicates the preference list of a_i and $P_{b_j} = \{a_k \succ a_l \succ \cdots\}$ that of b_j, with $\succ$ denoting a preference relation. A matching function $M : A \cup B \to A \cup B$ maps agents to each other, such that $M(a_i) = b_j$ and $M(b_j) = a_i$, ensuring that each agent is either matched or remains unassigned. A matching is stable if no blocking pair exists. Hence, a blocking pair (a_i, b_j) exists if a_i prefers b_j over its current match $M(a_i)$ and b_j prefers a_i over its current match $M(b_j)$; that is, both would benefit from pairing together instead with their current partners. Therefore, the objective is to find a matching M such that no blocking pairs remain. Initially, each agent a_i proposes to its highest-ranked, unproposed b_j. Each b_j tentatively accepts the best-ranked offer while rejecting lower-ranked agents. If a rejected agent finds a better match, it updates its proposal. The process repeats until no further improvements occur and terminates in $O(n^2)$ time complexity. Hence, the objective function maximizes the sum of ranking scores:

$$\max \sum\nolimits_{i=1}^{n} \sum\nolimits_{j=1}^{m} P_{a_i}(b_j) X_{ij}, \tag{5.10}$$

where $X_{ij} = 1$, if a_i is matched to b_j, and 0 otherwise.

Exploring Stable Matching in RADAR: Let $\mathcal{F}_{rel} = \{F_{rel}^1,\ F_{rel}^2,\ \ldots,\ F_{rel}^n\}$ be the set of relocating-facilities and $\mathcal{F}_{rec} = \{F_{rec}^1, F_{rec}^2, \ldots, F_{rec}^m\}$ the set of receiving-facilities. Each facility $F_{rel}^i \in \mathcal{F}_{rel}$ ranks $F_{rec}^j \in \mathcal{F}_{rec}$ based on their policies. The objective is to maximize preference scores of both F_{rel}^i and F_{rec}^j. Hence, a stable match means no facility in $\mathcal{F}_{rel}$ and facility in $\mathcal{F}_{rec}$ mutually prefer each other over their assigned partners. To illustrate this, let's assume we have 10 facilities among which 5 are relocating and the other 5 are potential receiving-facilities (As shown in

Table 5.1: Relocating-Facilities (F_{rel}^{i})

ID	Type	$\mathbf{O_i}$	$\mathbf{SVI_i}$	Needs
F_{rel}^{1}	H	Y	0.30	Medication
F_{rel}^{2}	S	R	0.65	Oxygen
F_{rel}^{3}	N	O	0.50	Food
F_{rel}^{4}	S	G	0.40	Nurse
F_{rel}^{5}	H	O	0.75	Ventilator

Table 5.2: Receiving-Facilities (F_{rec}^{j})

ID	Type	$\mathbf{O_i}$	$\mathbf{SVI_i}$	Resource
F_{rec}^{1}	N	G	0.20	Nurses(5)
F_{rec}^{2}	S	Y	0.50	Oxygen(10)
F_{rec}^{3}	S	R	0.80	Med(50),Bed(10)
F_{rec}^{4}	H	G	0.35	Ventilator(15)
F_{rec}^{5}	N	O	0.60	Oxygen(5)

Table 5.3: Distances $D(F_{rel}^{i}, F_{rec}^{j})$

	F_{s1}	F_{s2}	F_{s3}	F_{s4}	F_{s5}
F_{r1}	5	10	12	15	9
F_{r2}	12	8	14	11	10
F_{r3}	7	9	20	5	15
F_{r4}	10	16	9	8	14
F_{r5}	3	13	25	6	19

Tables 5.1 and 5.2). Using the policies for relocating (e.g., proximity, safe routes) and for receiving people (e.g., triage, operational status), each facility computes the score of each potential candidate and ranks them as shown in Alg. 4. For example, if we consider F_{rel}^{2}, which needs to relocate people to facilities with available oxygen resources, it prefers F_{rec}^{2} and F_{rec}^{5} since only these have oxygen. In addition Table 5.3 shows that F_{rec}^{2} is 8 miles away vs. F_{rec}^{5} which is 10 miles. Assuming there is a safe route to get to those two facilities, F_{rel}^{2} ranking is $\{F_{rec}^{2} \succ F_{rec}^{5} \succ \dots\}$. Following similar process, all facilities in the set $\mathcal{F}_{rel}$ rank those in $\mathcal{F}_{rec}$ and vice versa. Next, we perform matching between the two sets as shown in Alg. 5.

We apply the same logic to the set of facilities that relocate people (i.e., nursing homes or hospitals that need to evacuate residents) and the set of facilities receiving the evacuees (i.e., shelters or other healthcare facilities with available capacity). Given N facilities and P policies, the matching is done in $O(N^2)$ time complexity. To reduce the search space complexity, we only consider the top k closest facilities that are safe, instead of all possible matches. This restriction reduces proposals

Algorithm 4: Policy Score Computation & Preference List Generation

Input: Set of relocating-facilities $\mathcal{F}_{rel}$, Set of receiving-facilities $\mathcal{F}_{rec}$, Policy functions $\{S_p(f,c)\}_{p=1}^{P}$ and weights $\{w_p\}_{p=1}^{P}$

Output: Preference list $L(f) = \langle c_1, c_2, \ldots, c_{|L(f)|} \rangle \; \forall c \in \mathcal{F}_{rec}$

foreach $f \in \mathcal{F}_{rel}$ **do**
    Initialize List $\leftarrow \varnothing$;
    foreach $c \in \mathcal{F}_{rec}$ **do**
        Compute policy score $S(f,c) \leftarrow \sum_{p=1}^{P} w_p \, S_p(f,c)$;
        Append the pair $(c, S(f,c))$ to List;
    Sort List in descending order by $S(f,c)$;
    $L(f) \leftarrow$ sequence of candidates c from the sorted List;
return $\{L(f) \mid f \in \mathcal{F}_{rel}\}$;

Algorithm 5: Resource Matching

Input: Facilities $\mathcal{F}_{rel}$, $\mathcal{F}_{rec}$; Preference list $L(f) = \langle c_1, c_2, \ldots, c_{|L(f)|} \rangle \subseteq \mathcal{F}_{rec}$, $\forall f \in \mathcal{F}_{rel}$

Output: Matching $M \subseteq \mathcal{F}_{rel} \times \mathcal{F}_{rec}$

$M \leftarrow \varnothing$ `// Initialize matching`
for $f \in \mathcal{F}_{rel}$ **do** Mark f as free ;
while *there exists a free* $f \in \mathcal{F}_{rel}$ *with* $L(f) \neq \varnothing$ **do**
    Let $c \leftarrow$ first candidate in $L(f)$;
    Remove c from $L(f)$;
    if c *is free* **then**
        Add (f,c) to M;
    else
        Let f' be the current match of c in M;
        if c *prefers* f *over* f' **then**
            Remove (f',c) from M and add (f,c) to M;
            Mark f' as free;
return M;

from $O(N^2)$ to $O(Nk)$ and bounds routing queries. When approaching nearby capacity, we expand the candidate set in batches of Δk until a feasible pair is found.

To handle evolving disaster conditions, we periodically (e.g., every 15–30 min) or upon detection of a major hazard update (such as a newly impassable road segment or change in facility status), recompute the matching (Alg. 6). Rather than recomputing all assignments, we restrict this procedure only to the subset of newly impacted relocating facilities $\mathcal{I} \subseteq \mathcal{F}_{rel}$ whose initial match has been invalidated by the new hazard state H. As shown in lines 7–11, for each $f \in \mathcal{I}$ we first reuses its precomputed top-k preference list $L_k(f)$ and apply both capacity checks ($\texttt{demand}(f)$) and $\texttt{safeRoute}(f, c, H)$ test to find an immediate reassignment. If none of those candidates passes (lines 13–15), we dynamically expand the candidate set in increments of Δk, recompute the combined policy scores $S(f, c)$, and repeat until a feasible match is found or all options are exhausted. By operating only on affected facilities and growing the search space in controlled batches, this approach preserves valid assignments, bounds computational complexity, and provides rapid, localized adaptation to evolving hazards.

While this model provides an efficient aggregate-based allocation mechanism, it does not incorporate individual-level evacuee preferences (i.e., preferences specified by the individuals themselves). In addition, individuals are grouped in clusters with the same relocation needs, which may not always hold in real-world settings. To address these limitations, we first extend our framework with policies for requesting and sending resources to manage resource exchanges (§5.4.3), and then introduce a fine-grained improvement of our model that integrates resident-specific preferences, including their medical information, into the allocation process (§5.6.1).

5.4.3 RADAR Extensions

a) Policies for Requesting and Sending Resources

Several of the spatio-temporal policies introduced for relocating people, i.e., proximity (Eqn. 5.4), safe-route viability (Eqn. 5.7), operational readiness (Eqn. 5.5), and social vulnerability (Eqn. 5.9) also apply in the resource requesting and sending context, since both problems rely on accurate

Algorithm 6: Dynamic Rematching with Evolving Hazard

Input : M, $\mathcal{I} \subseteq \mathcal{F}_{rel}$, $L_k(f)$, $\mathcal{F}_{rec}$, capacities $C[\cdot]$, hazard H, policies $\{(S_p, w_p)\}$, batch Δk

Output: Updated matching M'

$M' \leftarrow M$;
foreach $f \in \mathcal{I}$ **do**
  assigned $\leftarrow$ **false**; considered $\leftarrow \emptyset$; batch $\leftarrow L_k(f)$;
  while *not assigned* ***and*** *considered* $\neq \mathcal{F}_{rec}$ **do**
    foreach $c \in$ batch **do**
      if $C[c] \geq$ demand(f) ***and*** *safeRoute*(f, c, H) **then**
        $M'[f] \leftarrow c$;
        $C[c] \leftarrow C[c] -$ demand(f);
        assigned $\leftarrow$ **true**;
        break;
      considered $\leftarrow$ considered $\cup \{c\}$;
    if *assigned* **then break**;
    remaining $\leftarrow \{ c \in \mathcal{F}_{rec} \mid c \notin$ considered$\}$;
    if *remaining*$= \emptyset$ **then break**;
    foreach $c \in$ remaining **do**
      $S(f, c) \leftarrow \sum_p w_p \, S_p(f, c)$;
    batch$\leftarrow$top-Δk from remaining by $S(f, c)$;
  if *not assigned* **then** log_warning (“No feasible match for f”);
return M';

distance measurements, hazard overlays, and facility status feeds. In addition, resource exchanges require policies to capture inventory levels, request criticality, delivery consolidation, and sender flexibility. Hence, we use a **slack-ratio** (a principle of inventory management that quantifies a sender's buffer capacity relative to its total capacity) [23] to prevent over-exhaustion of any single source by favoring senders with more available inventory.

$$S_{\text{slack}}(F_{\text{send}}) = \frac{\text{available\_stock}(F_{\text{send}})}{\text{capacity}(F_{\text{send}})} \tag{5.11}$$

The available stock is the count of uncommitted units, such as beds, water bottles, and related supplies, while capacity represents the maximum stock. Next, the **request urgency** captures the relative criticality of each demand. Drawing on priority-queueing theory [16, 18], we normalize a facility's current demand against the peak regional demand.

$$S_{\text{urg}}(F_{\text{req}}) = \frac{\text{demand}(F_{\text{req}})}{\max_{f' \in \mathcal{F}_{\text{req}}} \text{demand}(f')} \tag{5.12}$$

Here, $\text{demand}(F_{\text{req}})$ aggregates all requested units (e.g., ICU beds, oxygen tanks), so that higher-need requests automatically receive higher scores within the $[0, 1]$ range. To improve logistical efficiency, we introduce a **consolidation potential** metric based on classical vehicle-routing heuristics [42]. We identify clusters of co-located requests $\mathcal{R}_{\text{cluster}}(F_{\text{req}})$ and compare their size to the total local request set $\mathcal{R}_{\text{all}}(F_{\text{req}})$.

$$S_{\text{grp}}(F_{\text{req}}) = \frac{|\mathcal{R}_{\text{cluster}}(F_{\text{req}})|}{|\mathcal{R}_{\text{all}}(F_{\text{req}})|} \tag{5.13}$$

A higher value indicates a stronger potential for multi-stop deliveries, reducing overall travel distance and vehicle utilization. Finally, the **commitment-ratio** [131] policy ensures senders retain flexibility for future surges by penalizing near-exhaustion of capacity.

$$S_{\text{commit}}(F_{\text{send}}, F_{\text{req}}) = 1 - \frac{\text{allocated}(F_{\text{send}})}{\text{capacity}(F_{\text{send}})} \tag{5.14}$$

In Eqn. 5.14, allocated(F_{send}) indicates the current reserved stock. Values near 1 indicate ample remaining capacity, while values closer to 0 discourage further assignments. All request-side and send-side policy scores S_p are then combined to obtain the final policy score as shown in Eqn. 5.15, where weight w_p reflects the EOC priorities (e.g. speed versus equity trade-offs). These aggregated scores produce preference lists that drive the matching and routing of resources.

$$S(F_{req}, F_{send}) = \sum_{p} w_p \, S_p(F_{req}, F_{send}) \tag{5.15}$$

b) Integrating Medical Profiles and Preferences

We developed a fine-grained extension of RADAR that integrates resident-level medical profiles, care continuity constraints, and social preferences into the allocation process. This enhancement enables the system to take into account factors such as specialized clinical service requirements, existing provider affiliations, and personal preferences. Incorporating these preferences increases the complexity of the matching and routing process, but yields more person-centered relocation plans. This extension is particularly relevant in long-term scenarios or for populations with chronic care needs, where inappropriate placements may exacerbate health risks. A complete description of the RADAR fine-grained model (including formal definitions, scoring metrics, and computational challenges) is provided in Section 5.6.1.

5.5 Experimental Evaluation

We evaluated RADAR across three realistic disaster scenarios: (1) A **wildfire** in the Palisades region with rapid reallocations amid shifting fire perimeters and facility statuses (5.5.1); (2) A 7.8-magnitude **earthquake** drill in Orange County, CA, showcasing hospital bed surge management and safe routing amid infrastructure damage (5.5.2); and (3) A countywide **water contamination** drill exploring relocation and water resupply for those sheltering-in-place to sustain operations (5.5.3).

Setup: We obtained facility data (including facility type, licensed bed capacity, location, and related attributes) from California's Health and Human Services Agency [25] and HCAI [83] open datasets (10409 facilities). We stored them in a PostgreSQL/PostGIS database to enable efficient geospatial queries. We integrated OpenStreetMap road network data and real-time/historical traffic information with Google's Distance Matrix API [77] and ArcGIS Pro to model distances and travel times between facilities. Facility operational statuses and resource availability were updated dynamically based on ReddiNet situational reports. Hazard footprints and progression data were obtained from authoritative sources: CalFire for wildfire perimeters [24], USGS ShakeMap for earthquake impact [150], and OCHCA and ReddiNet reports for water-related disruptions [126, 133]. We acknowledge that in disaster situations, these feeds may not fully capture real-time conditions accurately. Therefore, supplementing or replacing them may be necessary to improve accuracy. For evaluating allocations, we generated synthetic resident profiles guided by domain experts, including individualized medical needs and preferences. Across all scenarios, we used stable matching to assign evacuees or resources while comparing RADAR's performance against baseline and state-of-the-art allocation strategies.

5.5.1 Wildfire Use Case: Evacuation Based on the Palisades Fire

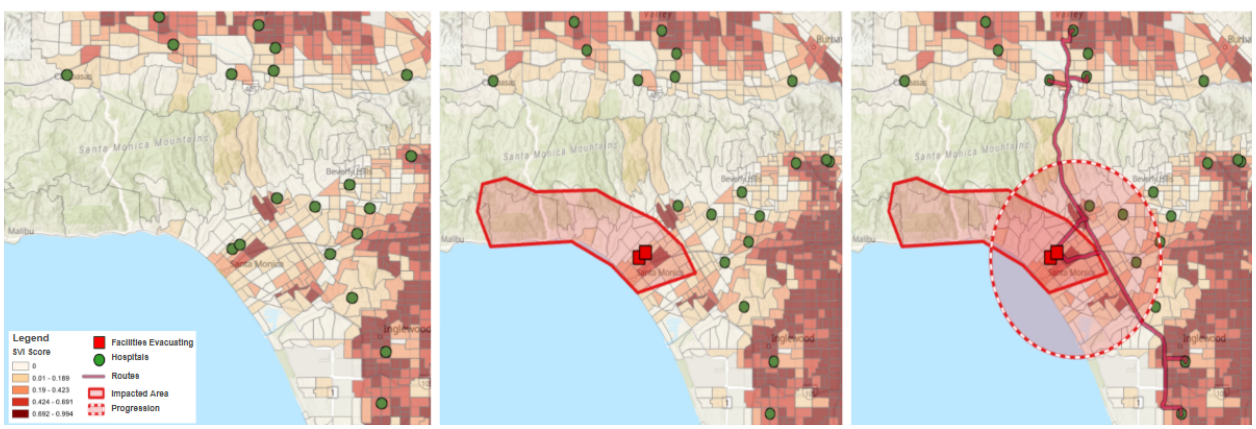

Figure 5.3: Assignment and Routing for Wildfire Scenario

We simulated evacuations based on real conditions during the Palisades fire (Fig. 5.3). Three facilities in the projected fire zone were selected as evacuees. We compare RADAR with several other allocation methods. First, we consider three variants of a greedy method: assigning evacuees to (1) the nearest facility by distance (Greedy-D); (2) minimizing estimated travel time (Greedy-

TT), and (3) based on available capacity (Greedy-C). We also include a (4) Random allocation method, assigning evacuees arbitrarily to facilities of the same type with available resources. This method serves as a conservative lower bound, where we select only from a pool of eligible facilities. We also compare RADAR against two state-of-the-art optimization methods: (5) Non-dominated Sorting Genetic Algorithm II (NSGA-II) [45] and (6) Ant Colony Optimization (ACO) [153].

Key metrics include: average travel time for evacuee relocation; preferences via a penalty-based satisfaction score, where lower penalties indicate better alignment with resident preferences; and efficiency loss (”price of fairness”[16]), which assesses the trade-off between fairness and optimal travel time. These metrics collectively provide a comprehensive evaluation of efficiency and equity.

Each impacted facility receives its top five candidate receiving facilities, which are ranked based on their policy scores (§5.4). These correspond to the closest and safest preferred facilities with the available resources, as determined by the policies. Similarly, each evacuating facility generates a ranked list of preferred destinations. We then apply stable matching to pair evacuating facilities (or individual evacuees, in the fine-grained model) with receiving facilities, subject to capacity constraints. Matching optimizes policy scores and ensures stability, preventing any facility from preferring another feasible match. The matching process repeats as the fire expands, updating allocations as new facilities become impacted.

We considered three types of impacted facilities: Emergency Room (ER), Intensive Care Unit (ICU), and General Acute Care Hospital (GACH), totaling 182 bed assignments and ensuring no single facility becomes overwhelmed (Fig. 5.4a and 5.4e). Outliers with higher travel times indicate cases where only distant facilities had sufficient capacity. We also measured efficiency loss by comparing makespans against the optimal baseline of 37 minutes (the shortest travel time to the nearest safe facility). Fig. 5.4c shows RADAR incurs an efficiency loss of only 8 minutes, which is modest compared to other approaches. This slight trade-off supports RADAR’s policy-driven improvements in stability and fairness.

Fig. 5.4b and 5.4f illustrate the average makespan (average travel time, in minutes, across all bed categories). During daytime, RADAR achieved a makespan of 45 min., compared to 50 min.

(Greedy-D), 48 min. (Greedy-TT), 55 min. (Greedy-C), and 78 minutes (Random). Although NSGA-II and ACO showed slightly better makespans (43 and 42 min., respectively), they lack policy integration and stable matching guarantees offered by RADAR. We also assessed preference using a penalty-based satisfaction metric, reflecting deviations from evacuees' preferred allocations. Fig. 5.4d and Fig. 5.4h show RADAR has a significantly lower standard deviation, indicating more consistent satisfaction across evacuees compared to Greedy and Random methods. NSGA-II and ACO, despite relatively low variance, showed reduced overall resident satisfaction due to a lack of policy integration.

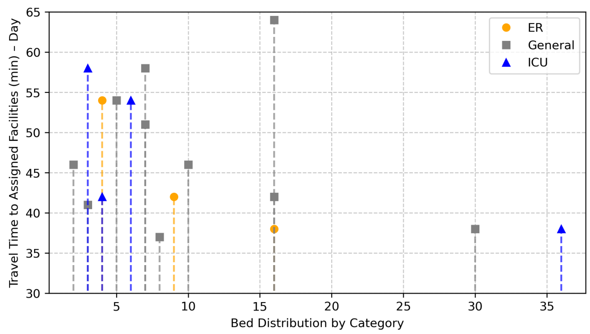

(a) Bed Coverage (Daytime)

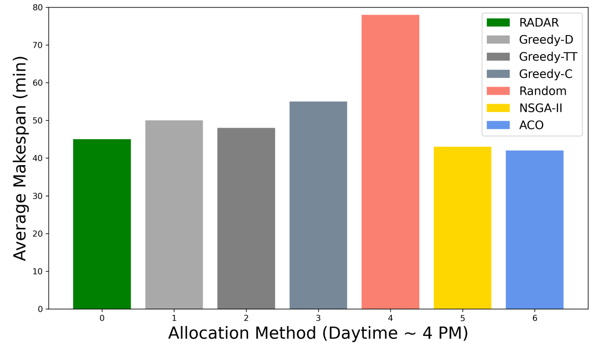

(b) Average Makespan

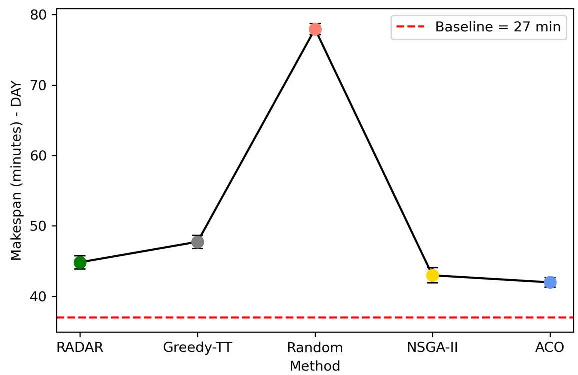

(c) Efficiency Loss

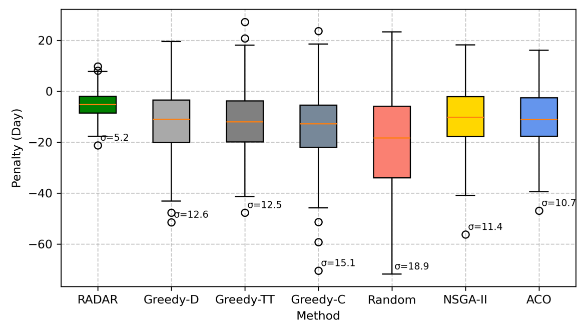

(d) Preference Distribution

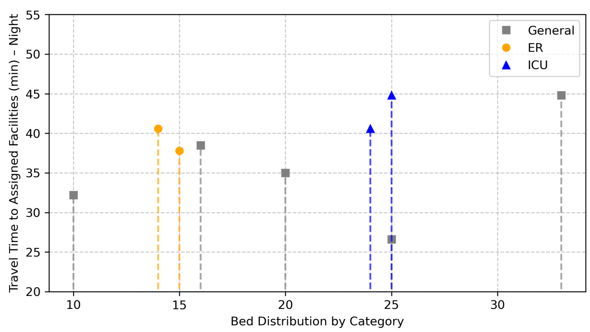

(e) Bed Coverage (Nighttime)

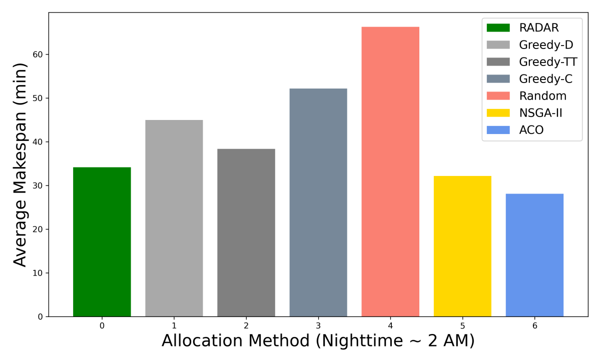

(f) Average Makespan

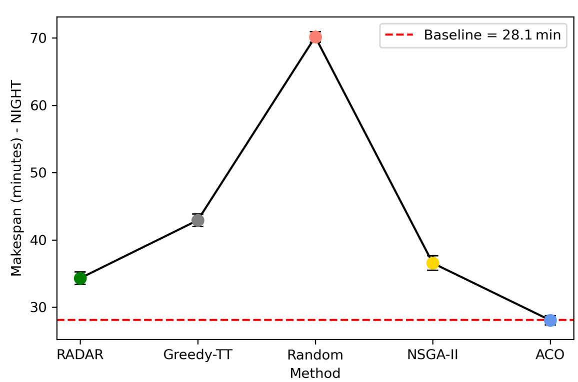

(g) Efficiency Loss

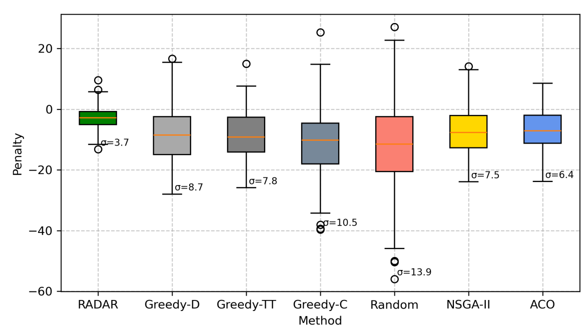

(h) Preference Distribution

Figure 5.4: Comparison of Daytime (a–d) vs. Nighttime (e–h) allocation: bed-category coverage (a,e), average makespan (b,f), efficiency loss (c,g), preference distribution (d,h).

5.5.2 Earthquake Use Case: Managing Hospital Bed Surge and Safe Routing

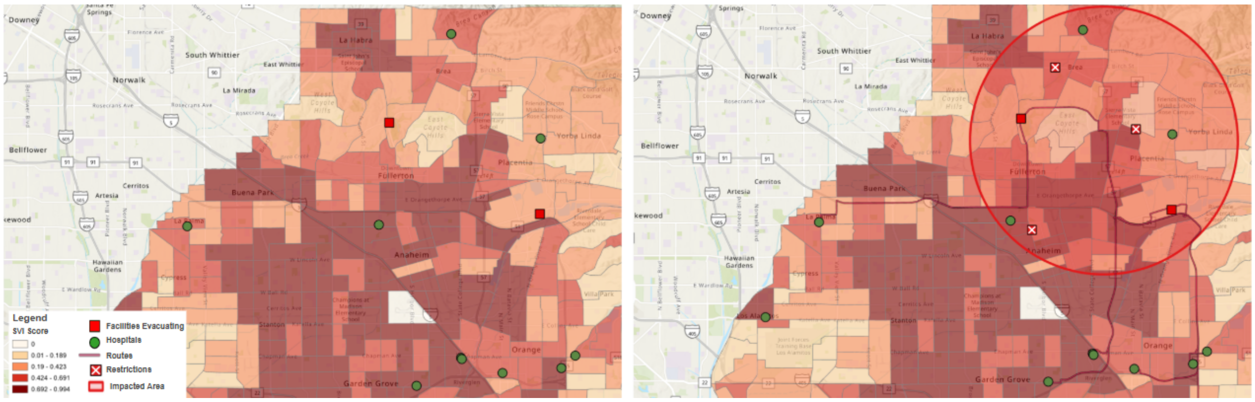

Figure 5.5: Assignment and Routing for Earthquake Scenario

As part of a county-wide emergency preparedness drill, we simulated the relocation of two hospitals (Kaiser Foundation Hospital and St. Jude Medical Center) under disaster conditions (Fig. 5.5). Together, the hospitals had a combined licensed capacity of 582 beds. However, ReddiNet reports indicated 274 available beds in nearby facilities, but only 109 were in facilities operating under "Normal Operations" status, resulting in a surge scenario where 308 patients required rapid relocation under constrained capacity. Using RADAR, we generated recommended destination facilities based on real-time constraints such as capacity, proximity, and infrastructure status. The resulting dashboard (Fig. 5.10) presents a summary of the recommended allocations, including destination options and relevant operational metrics to support decision-making.

Fig. 5.6 compares estimated travel distances and durations across 20 origin-destination facility pairs (routes) at four times of day: morning (9am), afternoon (3pm), night (9pm), and dawn (3am). As demonstrated in Fig. 5.5, given the affected area, the 3 closest facilities were unable to receive people. As expected, the travel distances remain relatively stable across time periods since those physical routes do not change. However, small variations in distance (especially for routes R5 through R14) suggest that routing algorithms may adjust paths slightly in response to traffic conditions, selecting alternate roads to optimize travel time.

(a) Distance by Time of Day

(b) Distance by Day of the Week

(c) Travel Time by Time of Day

(d) Travel Time by Day of Week

Figure 5.6: Earthquake results for Distance (km) and Travel Time Distribution (min) by Time of Day (a, c), and by Time of Week (b, d)

In contrast, travel time varies more significantly by time of day. Afternoon (3pm) consistently shows the longest travel durations, likely due to increased traffic congestion during mid-day. Morning (9am) times are also elevated, reflecting typical commute-hour delays. Night (9pm) and dawn (3am) generally offer the shortest travel times, corresponding with lower traffic volumes. For example, R19 shows a difference of over six minutes between the afternoon and dawn. These findings highlight how traffic dynamics drive temporal variation in travel time, which is important for planning time-sensitive facility relocations or emergency response strategies.

5.5.3 Water Contamination Use Case: Relocation Based on Water Constraints

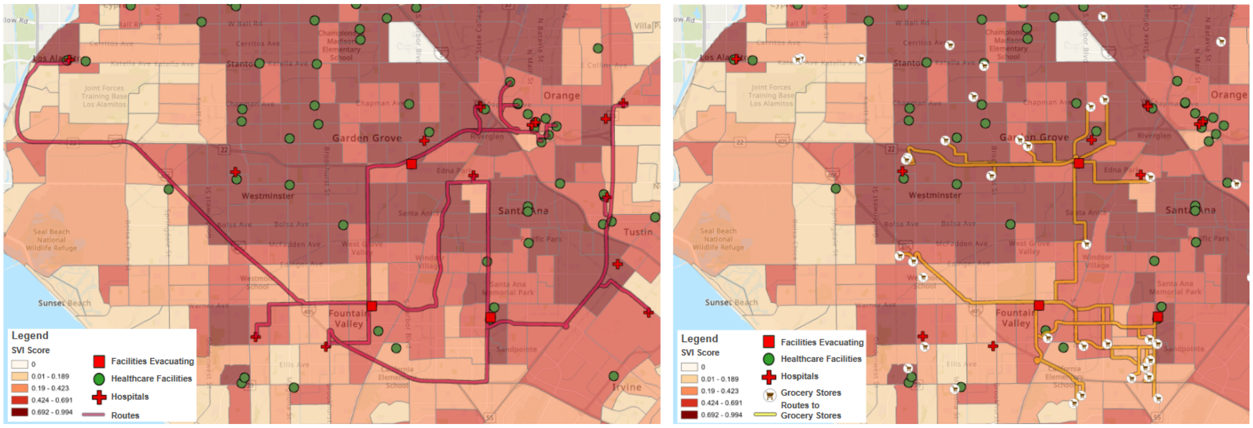

Figure 5.7: Assignment and Routing for Water Scenario

For this scenario (Fig. 5.7), each facility was initialized with water reserves according to California's General Acute Care Hospital (GACH) emergency standards: 150 gallons per bed for 72 hours, plus a 5,000-gallon refillable tank [30, 152]. We updated facility-level water availability using real-time ReddiNet [133] reports and enabled congestion-aware delivery routing via CalTrans [26] feeds and the Google Distance Matrix API [77]. For re-supply, we identified each facility's ten nearest potable-water vendors and incorporated five official Points-of-Distribution (PODs) from OCHCA [126]. Facilities still under contamination were excluded as relocation targets, and candidate sites were filtered to ensure sufficient water reserves for incoming residents.

Our results illustrate how RADAR provides resident relocation as the scale of water-contamination impacts grows. Fig. 5.8 shows three paired histograms and scatter-and-fit plots for scenarios in which 3, 10, and 20 facilities require evacuation. Distance distributions shift steadily outward: when only 3 facilities relocate, the mean move is 9.10 km (σ=4.33 km), expanding to 10.10 km (σ=5.52 km) at 10 facilities and 11.74 km (σ=7.17 km) at 20. These results exhibit right skew, particularly in the 20-facility case, where most relocations fall between 5 km and 15 km, but a long tail extends beyond 30 km. An overlaid normal curve in that plot captures the central mass but underestimates both the very short moves (1.5–5km) and the extended relocations, highlighting occasional assignments to distant, uncontaminated sites once nearby capacity is exhausted. Travel-

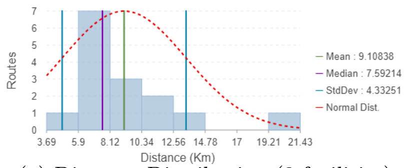

(a) Distance Distribution (3 facilities)

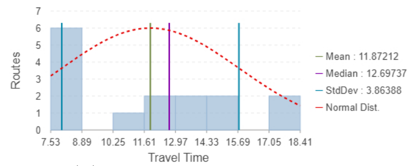

(b) Travel-time Distribution

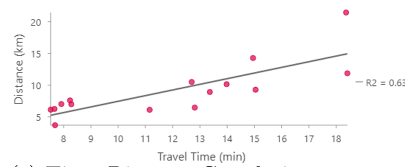

(c) Time–Distance Correlation ($R^2 = 0.63$)

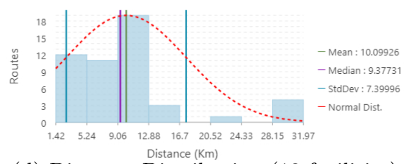

(d) Distance Distribution (10 facilities)

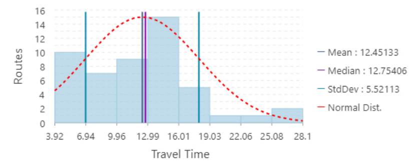

(e) Travel-time Distribution

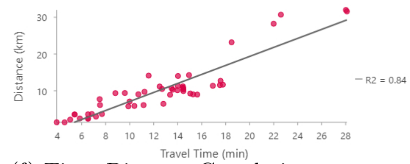

(f) Time–Distance Correlation ($R^2 = 0.84$)

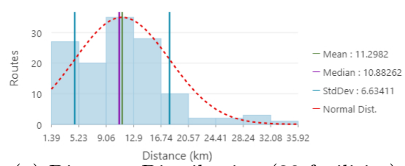

(g) Distance Distribution (20 facilities)

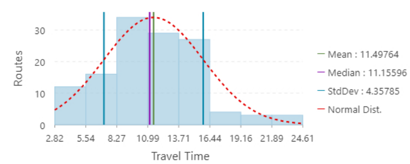

(h) Travel-time Distribution

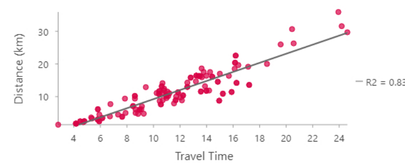

(i) Time–Distance Correlation ($R^2 = 0.83$)

Figure 5.8: Water-contamination results for 3 (Top row), 10 (Middle row), and 20 (Bottom row) impacted facilities

time distributions reflect a progression where the mean trip increases from 11.87 min (σ=3.86 min) through 12.45 min (σ=5.52 min) to 14.01 min (σ=5.32 min). In addition, all three cases retain heavy-tailed shapes, where most relocations complete in roughly 9–20 min, but some demand up to 30 min on the road. In practical terms, even under large-scale disruptions, most residents and patients reach a safe facility within manageable time windows for first responders.

Likewise, the scatter plots quantify the strength of the distance–time relationship under our safe route constraints. The variance proportion R^2 is 63% when relocating 3 facilities; rising to 84% with 10 facilities and settling at 83% for 20. This means that as more facilities relocate, transit time becomes increasingly predictable from a straight-line distance alone as congestion and detours play a smaller relative role when the candidate pool increases. A high proportion indicates that the decision-makers can reliably estimate evacuation loads, which simplifies resource staging and route planning under time constraints. In addition, Fig. 5.12 illustrates vendor-access distances and water-delivery times across all 144 re-supply assignments, revealing a similarly compact delivery radius.

These metrics demonstrate RADAR's ability to balance capacity, safety, and timeliness across realistic emergency scenarios. Mean distances and times increase roughly linearly with the number of impacted facilities, while variances remain bounded. Even in the most demanding scenario, the majority of relocations conclude in under 20 minutes over distances of 12 km, indicating that this framework can support fast, hazard-aware evacuations for vulnerable populations during widespread water-contamination events.

In addition to quantitative metrics, we received valuable feedback from stakeholders during drills and tabletop exercises. EOC staff reported that RADAR's automated, data-driven, policy-aware recommendations substantially reduced manual matching effort and served as a computational check to rapidly compare and refine their own allocations (see the dashboard in Fig. 5.11). Senior-care facility administrators emphasized that some of the dashboards improved shared situational awareness and made it faster to communicate actionable options to first responders.

5.6 Supplementary RADAR Extensions

5.6.1 Integrating Medical Profiles and Preferences

Our model assumes that all individuals within a facility have homogeneous resource needs and relocation preferences, making it suitable for rapid, bulk evacuations of short-term facilities (hospitals). However, people have heterogeneous needs and exhibit diverse medical conditions, personal preferences, and social constraints. Leveraging domain experts' knowledge and based on our experience working with Senior Health Facilities (SHFs) [1, 95], we capture fine-grained information about older adults in long-term care settings (e.g., Assisted Living Facilities, Skilled Nursing Facilities). Therefore, we extend our model to incorporate resident-level information into the decision-making process (Fig. 5.9).

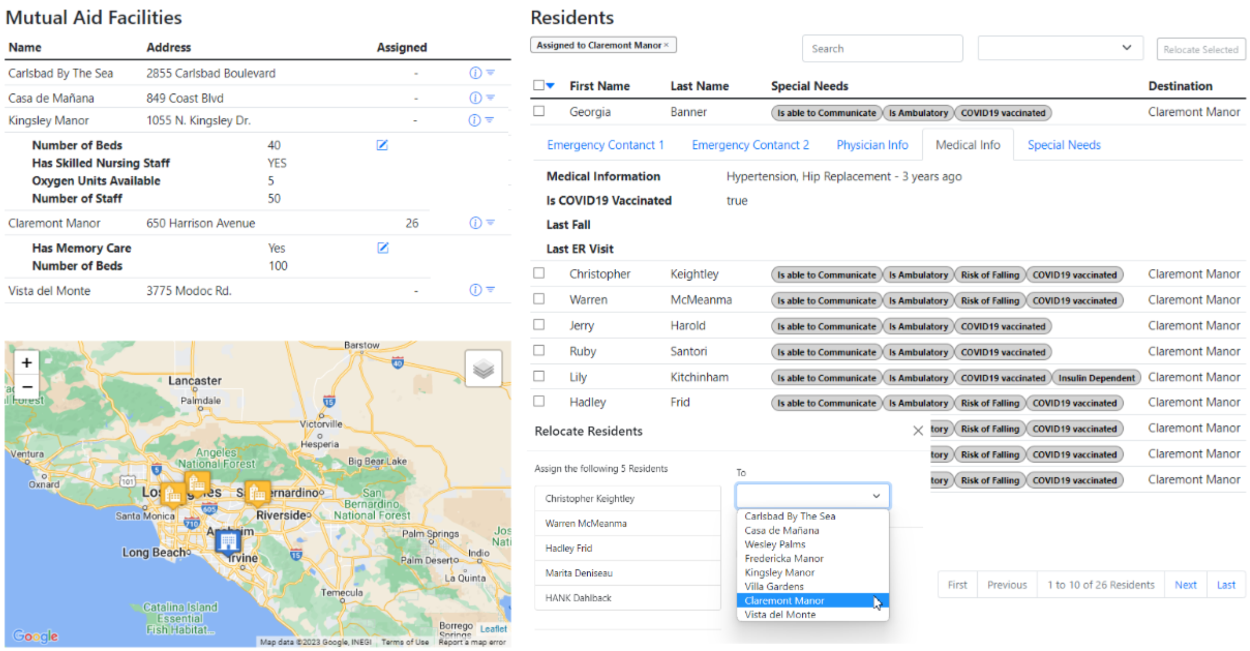

Figure 5.9: Dashboard with Integrated Medical Profiles

We characterize each resident r not only by static attributes (age, gender, room assignment) but also by detailed medical profiles, dynamic care requirements, and individualized preferences. Hence, the key personalized enhancements are as follows:

Medical compatibility: We adapt the concept of recall from information retrieval [137] to quantify how well a candidate receiving facility F_{rec} can satisfy the set of specialized services $\mathcal{S}_{\text{req}}$ required by relocating facility F_{rel}. Formally:

$$S_{\text{med}}(F_{\text{rel}}, F_{\text{rec}}) = \frac{\left|\mathcal{S}_{\text{req}} \ \cap \ \mathcal{S}_{\text{avail}}(F_{\text{rec}})\right|}{|\mathcal{S}_{\text{req}}|}, \tag{5.16}$$

where $\mathcal{S}_{\text{avail}}(F_{\text{rec}})$ is the set of specialized services (e.g. dialysis, ventilator support) that F_{rec} can provide. By normalizing to $[0, 1]$, we prioritize facilities that cover the largest fraction of required clinical services. For example, if F_{rel} requires $\mathcal{S}_{\text{req}} = \{\text{dialysis, ventilator, wound care}\}$ and F_{rec} offers $\mathcal{S}_{\text{avail}} = \{\text{dialysis, wound care}\}$, then $S_{\text{med}} = \frac{|\{\text{ dialysis, wound care}\}|}{3} = \frac{2}{3} \approx 0.67$, indicating that F_{rec} can fulfill two out of three required services.

Continuity of care: Established resident–provider relationships are preserved through a binary affiliation akin to linkage analysis in health networks [82]. For resident r and candidate facility c:

$$S_{\text{cont}}(r, c) = \begin{cases} 1, & \text{if network}(r) = \text{network}(c), \\ 0, & \text{otherwise,} \end{cases} \tag{5.17}$$

where network$(\cdot)$ returns the provider of the entity. For example, if resident r is associated with the "SunriseCare" network and c_1 belongs to "SunriseCare," then $S_{\text{cont}}(r, c_1) = 1$. If c_2 is part of a different group, say "ElderHealth Partners," then $S_{\text{cont}}(r, c_2) = 0$. This policy ensures that allocations maintain continuity of care by favoring facilities within the same organizational partnership.

Social preference: We capture each social-preference attribute a (e.g., pet-friendly facility, spouse accommodation, facility where staff speak resident's native language (language barrier)) as a binary indicator $S_a(r, c) \in \{0, 1\}$, and aggregate them into a single score:

$$S_{\text{soc}}(r, c) = \sum_{a \in \mathcal{A}} w_a \, S_a(r, c), \quad \sum_a w_a = 1, \tag{5.18}$$

where $\mathcal{A}$ is the set of relevant social attributes and w_a their weights. For example, $S_{\text{spouse}}(r, c) = 1$ if the spouse of resident r can also be accommodated at c (0 otherwise), and $S_{\text{pet}}(c) = 1$ if c allows pets (0 otherwise). Other attributes, such as facilities where staff speak the resident's native language, can be similarly defined and weighted. For example, let $\mathcal{A} = \{\text{spouse}, \text{pet}\}$ with $w_{\text{spouse}} = 0.6$ and $w_{\text{pet}} = 0.4$. Consider a candidate facility c_1 that can accommodate their spouse, but it's not pet-friendly. Hence, $S_{\text{spouse}} = 1$ and $S_{\text{pet}} = 0 \Rightarrow S_{\text{soc}}(r, c_1) = 0.6 \cdot 1 + 0.4 \cdot 0 = 0.6$. On the other hand, c_2 allows pets but cannot accommodate their spouse. Thus $S_{\text{spouse}} = 0$; $S_{\text{pet}} = 1$ $\Rightarrow S_{\text{soc}}(r, c_2) = 0.6 \cdot 0 + 0.4 \cdot 1 = 0.4$. Therefore, c_1 is preferred over c_2 according to the resident's social needs.

Finally, we incorporate penalties to evaluate non-preferred allocations. Suppose a resident has the preference list $L(r) = \{f_2, f_3, f_1, f_5, f_4\}$. If r is assigned to f_2, the penalty incurred is 0; if assigned to f_3, it is -1; if assigned to f_1, it is -2; and so forth. These penalties are used as feedback to trigger reassignment when cumulative dissatisfaction exceeds a threshold, thereby steering the matching process toward solutions that balance efficiency and fairness [14, 16, 95, 161].

From a computational standpoint, incorporating resident-specific constraints significantly increases the solution space. Our reduction from the NP-hard makespan scheduling problem shows that even a simplified instance with personalized constraints remains NP-hard [95]. Although traditional assignment and routing can be solved in polynomial time, the combinatorial complexity from resident preferences requires heuristics and approximation algorithms to achieve feasible solutions. Despite the complexity, the profound benefits, including significantly improved resident outcomes, enhanced continuity of care, and a demonstrably fairer allocation of resources, justify the adoption of this approach. Furthermore, older adults are particularly vulnerable to a range of negative consequences following disasters, such as heightened anxiety, depression, exacerbation of chronic health conditions, social isolation, and even post-traumatic stress disorder [123]. Therefore, by integrating resident-specific preferences and medical profiles, we proactively address these vulnerabilities and mitigate the associated risks, leading to a more person-centered and effective disaster response that promotes resident well-being and resilience.

As shown in §5.5, we use synthetic resident profiles in the wildfire scenario (developed with domain experts) to capture older adults' preferences, which we evaluate using penalties. By contrast, the earthquake and water contamination case studies use aggregate, facility-level representations (e.g., ER/ICU/GACH bed classes) to enable scalable analysis since per-resident data were unavailable.

5.6.2 Additional Experiments and Visualizations

The following additional results illustrate RADAR across different disaster contexts, including spatiotemporal analyses of vendor access for water re-supply and dashboard snapshots that demonstrate situational awareness and coordination support.

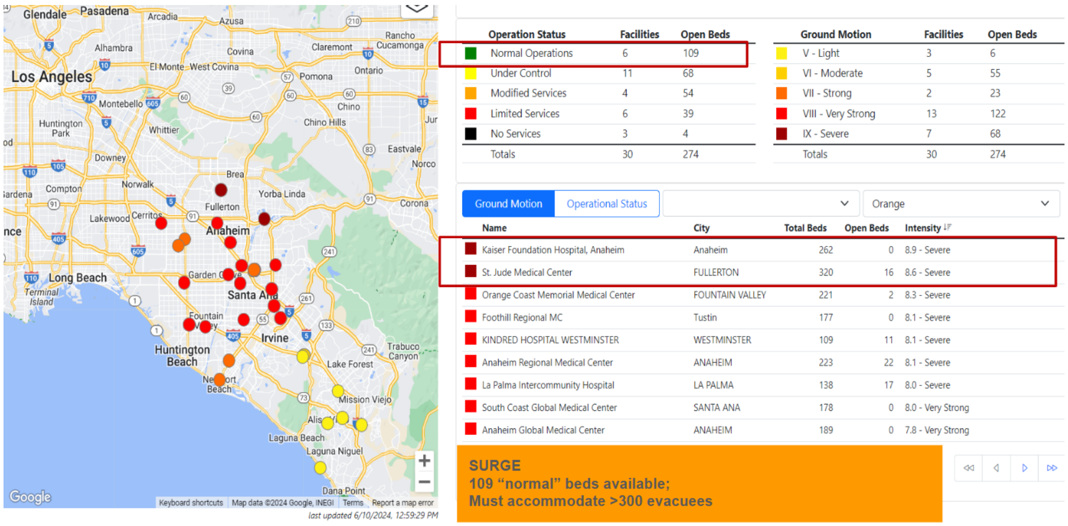

Figure 5.10: Dashboard of the Earthquake Drill Showing Surge Capacity.

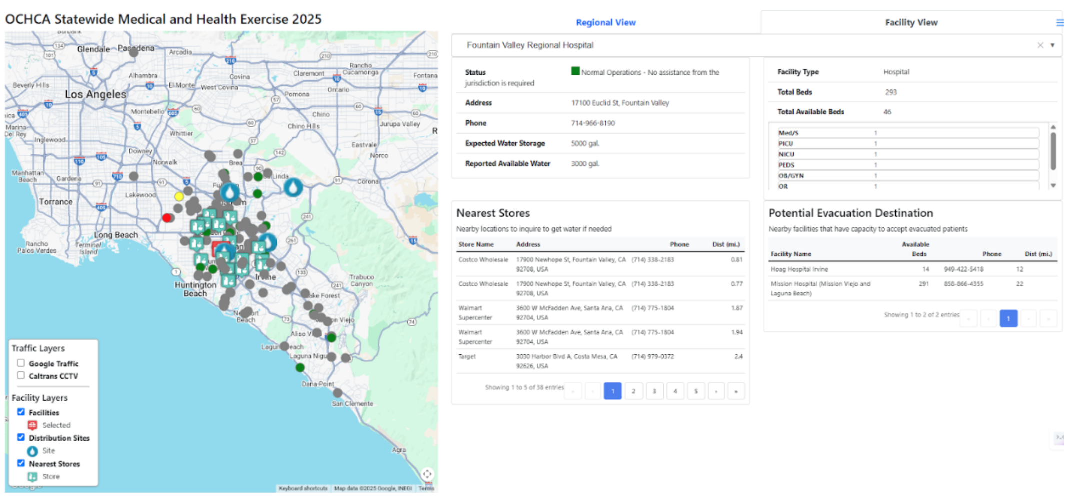

Figure 5.11: Dashboard of Water Contamination Drill.

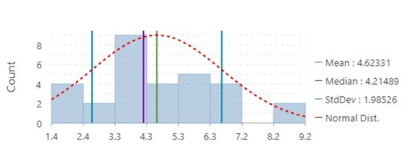

(a) 3-Facilities Impacted: Distance to Stores

(b) 10-Facilities Impacted: Distance to Stores

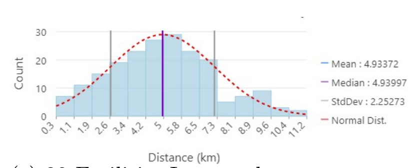

(c) 20-Facilities Impacted: Distance to Stores

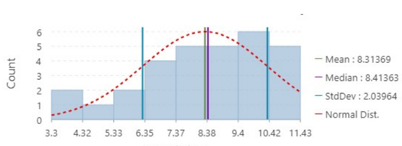

(d) 3-Fac. Impacted: Travel Time to Stores

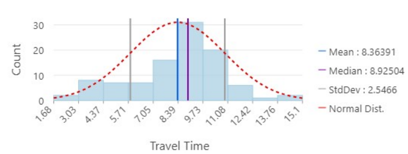

(e) 10-Fac. Impacted: Travel Time to Stores

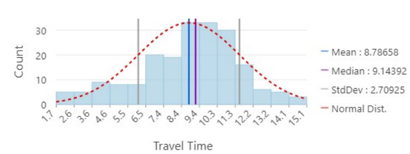

(f) 20-Fac. Impacted: Travel Time to Stores

Figure 5.12: Distributions of vendor-access distances (top row) and travel times (bottom row) for water re-supply.

5.7 Chapter Summary and Discussion

We presented *RADAR*, a GIS-centric decision-support framework for matching and routing across healthcare facilities during disasters. The central challenge was to determine where people and resources should go, which facilities can receive them, and which routes remain feasible as hazards evolve. To address this problem, the framework integrates hazard feeds, facility status, transportation information, resource availability, and policy-driven preferences into a unified matching and routing workflow. A central contribution is the use of policy-driven stable matching to support interpretable allocation across facilities. Rather than assigning evacuees or resources only to the nearest destination, the model incorporates operational policies such as capacity, residency type, proximity, safe-route availability, facility status, resource compatibility, and social vulnerability. These policies are translated into preference rankings that guide the matching process and help produce assignments that are feasible, explainable, and aligned with emergency response priorities.

The results also show that matching and routing must be treated together. A destination may be appropriate from a capacity or policy perspective, but unusable if routes are unsafe or travel times are prohibitive. Similarly, a route may be short, but the destination may lack the resources or operational status needed to support evacuees. By combining policy-aware matching with hazard-aware routing, *RADAR* supports more practical relocation and resource-sharing decisions under changing conditions.

The experimental scenarios demonstrated the usefulness of this approach across different disaster contexts. In the wildfire scenario, *RADAR* supported relocation decisions under an evolving hazard footprint while maintaining a modest efficiency trade-off relative to optimization baselines. In the earthquake drill, it supported hospital surge coordination under constrained bed availability. In the water-contamination scenario, it showed how routing and resource access can support both relocation and resupply decisions. Together, these evaluations illustrate how geospatial data, policy constraints, and adaptive routing can improve healthcare coordination during disasters.

These results also reveal a broader resilience challenge. Matching and routing assume that facilities, routes, and resources remain available enough to support relocation or resupply. However, disasters can also disrupt the infrastructure systems that healthcare facilities depend on, especially the electric grid. A facility may be a poor candidate for receiving evacuees if it is likely to lose power, exhaust backup resources, or become unreachable during a prolonged outage. This motivates the next chapter, which shifts from response coordination to facility risk assessment under prolonged power-grid outages.

Chapter 6

GIS-Informed Planning Amid Power Grid Outages

The previous chapter focused on matching and routing during disaster response. Those decisions assume that healthcare facilities, transportation routes, and critical services remain available enough to support relocation or resupply. During prolonged power outages, this assumption may no longer hold. A facility may be outside the most visible hazard footprint, but still lose critical services if backup resources are exhausted, the electric grid fails, or external support cannot reach the facility in time. Here, we shift from response coordination to facility risk assessment under prolonged power disruptions. The central research question is the following: *Given a prolonged power outage in a region, which healthcare facilities are most at risk of losing critical services, and how should planners reduce that risk?* Answering this question requires connecting facility-level backup capacity, regulatory requirements, hazard exposure, power-grid dependencies, transportation accessibility, regional suppliers, and community vulnerability into a unified risk assessment workflow.

To address this problem, we develop *BlackOps*, a GIS-informed decision-support framework created with input from healthcare, emergency management, and hazard analytics experts. The framework integrates geospatial data, infrastructure models, regulatory requirements, and interpretable preparedness indices to identify vulnerable healthcare facilities and translate risk into actionable scorecards and recommendations.

6.1 Chapter Overview

This chapter focuses on the facility risk assessment problem introduced in Chapter 3. The goal is to identify healthcare facilities that may lose critical services when regional infrastructure, especially the electric grid, fails for multiple days. In this setting, the main concern is not only whether a facility has backup hardware, but whether it can sustain essential operations long enough, obtain external support in time, and remain connected to the infrastructure and transportation systems needed for continuity of care. Healthcare facilities are especially vulnerable to prolonged outages because they rely on continuous electricity for life-sustaining equipment, refrigeration of medications and blood supplies, water systems, communications, electronic health records, and real-time patient monitoring. Empirical studies of major disasters have shown that power loss can contribute to generator failures, forced evacuations, and increased morbidity and mortality among medically fragile populations [9, 17]. The evacuation of hundreds of patients from a Los Angeles hospital after generator failures during Tropical Storm Hilary further illustrates how a local outage can quickly become a regional healthcare emergency [146].

Our prior disaster resilience work through CareDEX [1, 96] involved wildfire, earthquake, water-contamination, and surge-planning exercises with senior care facilities, healthcare partners, and first responders. These engagements revealed a persistent operational gap. When outages last for multiple days, resilience depends on a facility's functional autonomy, which we define as its ability to sustain essential services using on-site resources and coordinate resupply under degraded transportation conditions and uncertain outage durations. Current planning tools, including Hazard Vulnerability Analyses (HVAs), FEMA Hazus, the Social Vulnerability Index (SVI), and the ASPR RISC Toolkit, provide important pieces of this assessment [4, 29, 60, 84]. However, these pieces often remain separated across GIS dashboards, spreadsheets, regulatory documents, facility records, and local planning processes.

This fragmentation creates an operational challenge for emergency operations centers and regional planners. A GIS dashboard may show where facilities and hazards are located, but it may not indicate whether a specific facility can meet emergency-power requirements. A facility checklist may document generator capacity, but it may not account for road closures, supplier availability, social vulnerability, or cascading grid damage. Regulatory documents may define autonomy requirements, but those requirements are often embedded in unstructured text and are difficult to connect to facility-specific risk calculations. As a result, planners may struggle to determine which facilities are most likely to exhaust their autonomy, which facilities require immediate support, and where backup resources should be pre-staged.

To address this gap, we develop *BlackOps*, a GIS-informed decision-support framework for assessing healthcare facility risk during prolonged power outages. The framework integrates facility data, hazard scenarios, power-grid infrastructure, road networks, supplier locations, regulatory requirements, and social vulnerability indicators into a unified risk assessment workflow. It computes interpretable indices for autonomy gaps, backup shortfall, resupply difficulty, social vulnerability, and hazard exposure. These indices are aggregated into facility risk scores and translated into scorecards that explain why a facility is vulnerable and what actions may reduce that risk.

An essential component of *BlackOps* is the extraction of facility-specific required autonomy from emergency-power and healthcare regulatory documents. These requirements are often distributed across heterogeneous sources and embedded in unstructured text. To bridge this gap, the system uses a bounded retrieval-augmented generation workflow to retrieve relevant regulatory content and extract structured autonomy requirements and metadata. These outputs are then validated and used as inputs to the autonomy gap calculation, while the final risk assessment remains grounded in geospatial, infrastructure, and preparedness models. The contributions are as follows:

- A GIS-informed risk formulation for assessing healthcare facility vulnerability under prolonged power outages, linking facility autonomy, infrastructure exposure, resupply feasibility, and community vulnerability.
- A set of interpretable preparedness indices that capture autonomy gaps, backup shortfall,

resupply difficulty, social vulnerability, and physical hazard exposure.

- A bounded retrieval-augmented generation workflow that extracts facility-specific required autonomy and metadata from unstructured emergency-power and healthcare regulatory documents.
- A physics-based infrastructure modeling component that estimates cascading power-grid impacts and facility loss of function under stochastic earthquake scenarios.
- Facility scorecards and recommendations that translate multi-dimensional risk into priority tiers, supplier options, and actionable backup-energy planning guidance.
- Scenario-based experiments that identify vulnerable healthcare facilities and spatial clusters of elevated risk across thousands of simulated outage conditions.

The remainder of this chapter is organized as follows. Section 6.2 provides background on regional preparedness and presents the framework overview. Section 6.3 formulates the preparedness risk problem. Section 6.4 defines the static and scenario-dependent indices used to characterize vulnerability. Section 6.5 describes the scorecard and recommendation workflow. Section 6.6 presents the experimental evaluation. Finally, Section 6.7 summarizes the chapter and discusses its connection to broader healthcare resilience planning.

6.2 BlackOps: Background and Overview

Regional-scale preparedness typically combines spatial data with facility-level assessments using frameworks such as FEMA's Threat and Hazard Identification and Risk Assessment (THIRA) [60]. Tools such as Hazus model hazard impacts [4], while indices like the CDC SVI [29] and other resilience frameworks [20, 95, 96] incorporate social context. From an infrastructure perspective, grid resilience studies use topology-aware metrics to identify vulnerabilities and estimate disruption [57, 105], remote sensing to quantify exposure [119, 144], and network modeling to characterize hospital access during disruptions [46]. However, during prolonged failures, it is difficult to synthesize these methods into the site-specific indicators needed to model backup autonomy. Along with GIS analytics, healthcare facilities use readiness instruments to structure emergency planning. While HVA tools assign ordinal scores to the probability and impact for a range of hazard events [93], the ASPR Risk Identification and Site Criticality (RISC) Toolkit uses customizable scoring and guided assessments to evaluate facility reliance on external utilities [84]. This typically produces spreadsheet-oriented outputs that lack the spatial context needed during disasters. Transforming regulatory text into computable representations has also evolved from curated logic engines [70] to semantic parsing and relation extraction [62]. However, these techniques fail to connect linguistic extraction with geospatial modeling: regulatory requirements are not mapped to the spatial attributes and operational timelines needed for evaluating resilience. Automated extraction requires rigorous validation before integration into high-stakes emergency management.

Key limitations: Current practices yield fragmented views of outage preparedness. Qualitative instruments create an objective gap where ordinal scores are not easily comparable across facilities, and disconnects between regulatory text extraction and geospatial layers prevent the integration of policy requirements with physical infrastructure models. These shortcomings motivate a GIS-informed risk framework that synthesizes disparate data into actionable insights.

To address these limitations, BlackOps integrates three primary data sources: (i) physical infrastructure data describing healthcare facilities, electric-grid components, and hazards; (ii) emergency-power regulatory documents processed using retrieval-augmented generation (RAG) to extract facility-relevant autonomy requirements and context-aware metadata; and (iii) GIS layers that provide published spatial services such as historical traffic patterns, typical peak-hour congestion data, and the Social Vulnerability Index (SVI) to model the socio-economic context. Scenarios are instantiated from the USGS National Seismic Hazard Model Project [129] and simulated through the NGA-West2 ground motion simulator [19].

As shown in Fig. 6.1, BlackOps computes a set of interpretable indices, including static measures such as autonomy gaps, and scenario-dependent measures such as routing-based resupply difficulty, backup shortfall, and exposure, as defined in Section 6.4. These indices are then combined by an aggregator into a unified risk score and translated into per-facility scorecards and recommen-

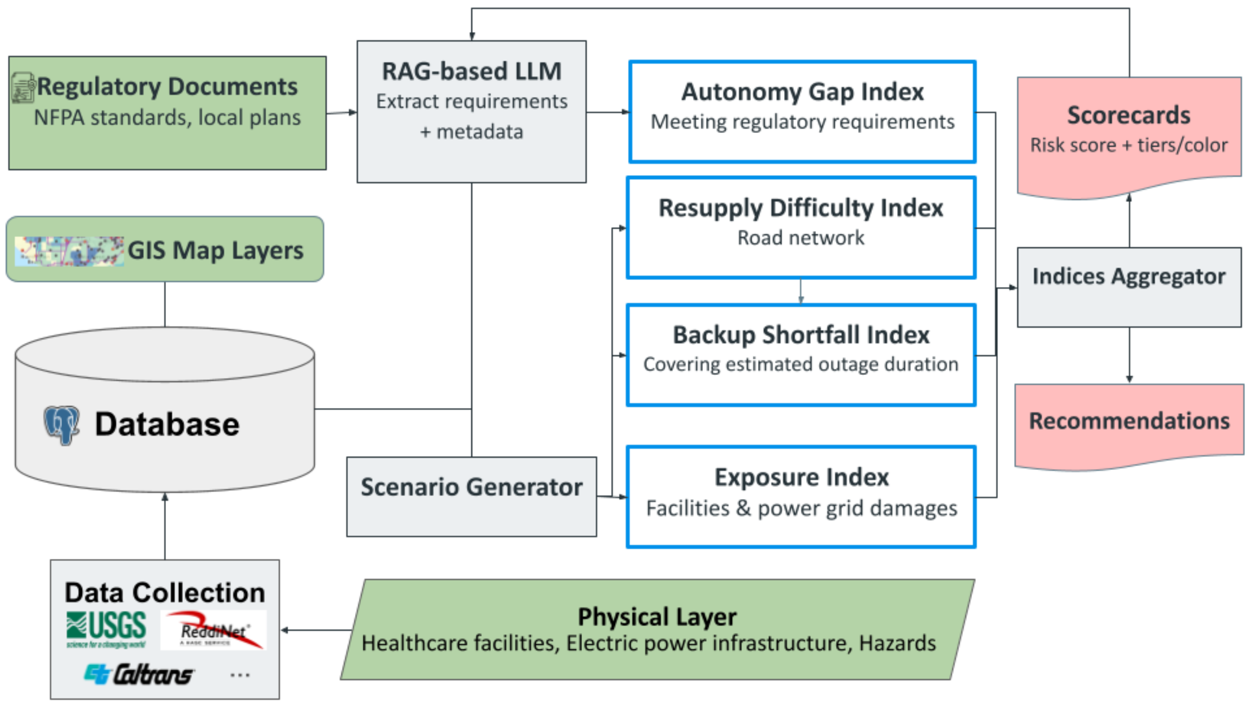

Figure 6.1: The BlackOps Overview

dations. Scorecards summarize a facility's overall vulnerability, while recommendations provide derived, actionable planning guidelines for decision-makers and stakeholders. A key design decision is to use GIS layers as operational data sources rather than static map illustrations. Figure 6.2 illustrates these spatial layers used in the framework, including healthcare facilities, substations, power distribution lines, roads, and additional regional infrastructure layers. These layers enable the system to connect facility-level preparedness with the surrounding infrastructure systems that influence outage exposure, resupply feasibility, and regional response planning.

By overlaying these layers, BlackOps can evaluate not only whether a facility is located in a hazard-prone area, but also whether the infrastructure supporting that facility, such as nearby substations, distribution lines, and transportation routes, may become a source of cascading operational risk.

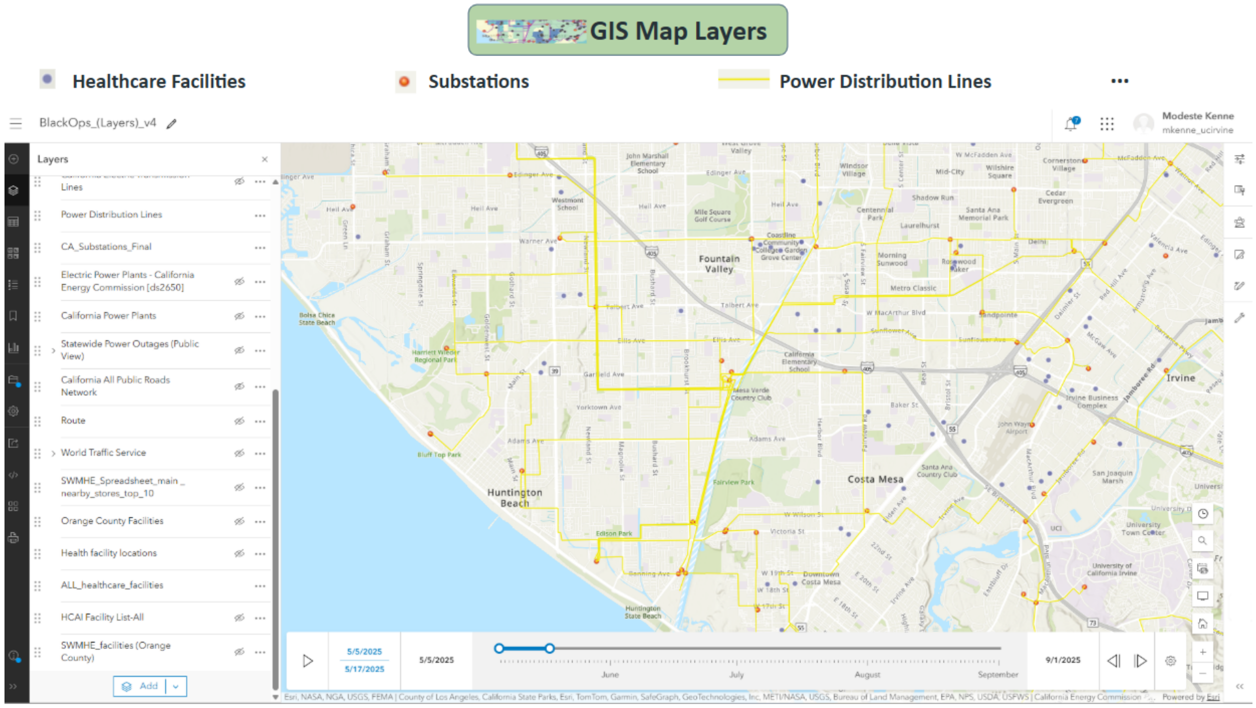

Figure 6.2: Geospatial layers illustrating the intersection of healthcare facilities and electrical power infrastructure for risk modeling

6.3 Problem Formulation

We begin by specifying the inputs to the preparedness model: hazard and grid-outage scenarios, a catalog of healthcare facilities with backup-power characteristics, transportation-network information for resupply, and requirements from emergency-power regulatory standards. For each facility and scenario, we then model a set of interpretable indices that quantify exposure, backup shortfall, resupply difficulty, autonomy gaps, and social vulnerability. While social vulnerability v_f and the autonomy gap g_f are static components, the remaining indices are scenario-dependent and are described in Section 6.4. Formally, let $\mathcal{F}$ be the set of healthcare facilities, indexed by f, with attribute vector a_f capturing facility type, location, estimated critical electrical demand, generator capacity, usable fuel or battery storage, operational status, and related operational attributes. Let $\mathcal{S}$ be the set of outage scenarios, where each $s \in \mathcal{S}$ represents a stochastic earthquake event and its induced damage state, including expected grid-outage duration $T^{\text{out}}_{f,s} \geq 0$ hours, with $T^{\text{out}}_{f,s}$=0 when grid power remains available at facility f. Let $\mathcal{B}$ denote the set of electric substations. We use $z_{b,s} \in 0, 1$ to indicate whether substation b remains operational in scenario s. In addition, $\pi(f) \in \mathcal{B}$ maps facility f to its serving substation. We define the road network as a graph $\mathcal{G} = (V, E)$. For each road segment $e \in E$, $\kappa_{e,s} \in 0, 1$ indicates whether e is traversable in scenario s.

During multi-day outages, autonomy can be extended if resupply begins early enough and routes remain usable. Therefore, for each scenario s, we estimate the earliest feasible resupply start time $T^{\text{start}}_{f,s} \geq 0$ (hours) by combining shortest-path travel times on the road network with scenario-specific closures (blocked roads) and historical traffic information [54]. We compute route redundancy $\rho_{f,s} \in [0, 1]$ from the road network graph as a normalized count of K-shortest paths from supply depots to f under scenario s, such that $\rho_{f,s} \approx 0$ indicates a bottleneck route, while larger values indicate multiple alternatives. To quantify the operational challenge of restoring resources to a facility, we introduce a ***resupply difficulty index***, $d_{f,s} \in [0, 1]$. This captures the logistical strain caused by both the timing of the response and the vulnerability of the local transport network, as shown in Eq. (6.1):

$$d_{f,s} = \begin{cases} 0 & T^{\text{out}}_{f,s}=0, \\ \text{clip}_{[0,1]}\Big(\alpha_d \frac{T^{\text{start}}_{f,s}}{T^{\text{out}}_{f,s}} + (1-\alpha_d)(1-\frac{\rho_{f,s}}{K})\Big) & T^{\text{out}}_{f,s}>0 \end{cases} \tag{6.1}$$

Here, $\alpha_d \in [0,1]$ weights the relative importance of operational bottlenecks: $\alpha_d \approx 1$ prioritizes timeliness (i.e., how late resupply begins relative to the outage), while $\alpha_d \approx 0$ prioritizes the availability of unblocked K-shortest paths from supply origins to facility f in scenario s. This formulation makes $d_{f,s}$ interpretable and tunable for jurisdictions that view response speed versus network robustness as a key constraint. In addition, we express how much of the outage horizon can be covered by on-site energy plus feasible deliveries as the effective autonomy. We define this in Eq. (6.2) as follows:

$$T^{\text{eff}}_{f,s} = T^{\text{crit}}_{f} + E^{\text{del}}_{f,s}/P^{\text{crit}}_{f} \tag{6.2}$$

where $E^{\text{del}}_{f,s} \geq 0$ is the delivered usable energy (kWh) that can reach f under scenario s within the outage window, given access and resource availability assumptions. Therefore, the ***shortfall index***, as depicted in Eq. (6.3) is the fraction of outage duration *not covered* by effective autonomy:

$$s_{f,s} = \begin{cases} 0, & T^{\text{out}}_{f,s} = 0, \\ \text{clip}_{[0,1]}\left(\frac{T^{\text{out}}_{f,s} - T^{\text{eff}}_{f,s}}{T^{\text{out}}_{f,s}}\right), & T^{\text{out}}_{f,s} > 0, \end{cases} \tag{6.3}$$

with $s_{f,s} = 0$ indicating full coverage of the outage horizon, and larger values indicate expected loss of critical services. Furthermore, ***exposure index*** $e_{f,s} \in [0,1]$ captures how strongly scenario s is expected to impact facility f (see §6.4.2). This is derived from established GIS-based hazard models, including shaking intensity, earthquake fragility curves, and power-grid feeder vulnerability [33]. Hence, for each facility-scenario pair (f,s), we aggregate indices into an overall preparedness risk score as follows:

$$r_{f,s} = \sum_{i=1}^{n} w_i \, \phi_i(f,s), \quad s.t. \sum_{i=1}^{n} w_i = 1, \quad w_i \geq 0, \tag{6.4}$$

where components are $\phi_i = \{e_{f,s}, s_{f,s}, d_{f,s}, g_f, v_f\}$ and w_i are policy weights set by decision-makers. The weighted sum is used as an interpretable planning score rather than a fitted predictor. Hence, all components are normalized to $[0,1]$, so each term contributes a comparable marginal risk. Equal weights provide a neutral default when local priorities are unavailable. During deployment, w_i can be tuned through expert elicitation or sensitivity analysis. For example, regions with severe transit constraints may emphasize $d_{f,s}$, while local SVI availability may motivate reweighting. Under this framework, the most vulnerable facility is identified as:

$$R_s^{\max} = \max_{f \in \mathcal{F}} r_{f,s} \tag{6.5}$$

Finally, to measure regional operational efficiency, we cluster facilities into C_s risk tiers as shown in Eq.(6.6). Here, $\mu_{c,s}$ is the centroid risk score for tier c in scenario s and $c_{f,s}$ is the tier assignment.

$$\min_{\{c_{f,s}\},\{\mu_{c,s}\}} \sum_{f \in \mathcal{F}} \left(r_{f,s} - \mu_{c_{f,s}}\right)^2, \quad c_{f,s} \in \{1, \ldots, C_s\}. \tag{6.6}$$

6.4 Modeling Preparedness Indices

In this section, we illustrate how we model the indices used for assessing the vulnerability of healthcare facilities.

6.4.1 Modeling Static Indices for Communities

a) Social Vulnerability Index (SVI)

Risk is typically modeled as a function of hazard, vulnerability, and available resources [66]. While hazard is an external factor, vulnerability can be modeled as a pre-existing community condition that can amplify disaster impacts. The CDC uses a percentile ranking $[0, 1]$ derived from 16 U.S. Census indicators aggregated into four themes: socioeconomic status, household composition/disability, minority status/language, and housing/transportation [29]. Its computation follows a hierarchical ranking pipeline where raw indicators are converted to percentiles, aggregated, and re-ranked into themes. This is then combined into a composite regional score [66]. We map healthcare facilities to these scores by overlaying their coordinates onto the SVI layer; this mapping is presented later in our experiments in Fig. 6.5. Hence, smaller facilities can exhibit higher vulnerability than major centers. However, SVI alone does not capture a facility's ability to maintain critical services during grid failure. To characterize vulnerability to prolonged power interruptions, we combine this with operational metrics.

b) Autonomy Gap Index

Prolonged outages threaten continuity of care when on-site energy reserves are exhausted before utility restoration. We quantify this as a temporal deficit when a facility's estimated **critical autonomy** T_f^{crit} is less than the **required autonomy** T_f^{req} prescribed by regulatory and accrediting guidance (e.g., NFPA, FEMA, or CMS).

Characterizing critical autonomy (T_f^{crit}): This is derived from facility maintenance records, generator logs, and Emergency Operations Plans (EOPs) (reviewed and validated through drills and tabletop exercises). If $P_f^{\text{crit}}>0$ represents the critical electrical power demand (kW) required to sustain life-safety and essential branches during an outage, and $E_f^{\text{store}} \geq 0$ denotes usable on-site energy (kWh) available from backup resources, then the critical autonomy obtained as follows:

$$T_f^{\text{crit}} = E_f^{\text{store}} / P_f^{\text{crit}} \tag{6.7}$$

```
RAG Prompt
Task:Extract mandated backup hours using ONLY the excerpts
Excerpts:<relevant vector DB passages>
Return:{"hours":<num>,"condition":<text>,"doc_id":<id>}|"unknown"
Rule:If hours are not found in the excerpts,return "unknown"(no guessing)
```

Thus a facility with 100,800 kWh of fuel energy (E_f^{store}) and a critical demand load (P_f^{crit}) of 1,400 kW, its critical autonomy would be T_f^{crit}=72 hours. Per NFPA safety standards [122], demand estimates incorporate a 10% to 25% safety margin to account for unexpected surges that occur during disasters.

Extracting required autonomy (T_f^{req}): This is challenging because requirements are scattered across heterogeneous federal, state, and local documents. We address this using Large Language Models (LLMs) integrated with Retrieval-Augmented Generation (RAG) to mitigate hallucinations and ensure regulatory grounding (Alg. 7). The extraction follows a five-step pipeline. First, facility attributes a_f are stringified and embedded into a high-dimensional vector space (Step 1). We perform a similarity search against a vector database $\mathcal{V}$ of regulatory codes, using cosine similarity to identify the k most relevant chunks (Steps 2–3). These chunks are appended to a structured instruction prompt ρ, which directs the LLM to perform inference and return results in a JSON format (Step 4). Finally, the system validates the output; if a valid numeric value is found, it is returned as T_f^{req} along with its citation metadata M_f. Otherwise, it is flagged for manual review (Step 5). This yields the autonomy gap index g_f:

$$g_f = \text{clip}_{[0,1]}\left(\frac{T_f^{\text{req}} - T_f^{\text{crit}}}{T_f^{\text{req}}}\right) \tag{6.8}$$

Hence, a facility with T_f^{req}=96 hours but maintaining only T_f^{crit}=72 hours, has g_f=0.25, which indicates a 25% deficit.

Algorithm 7: RAG-based Extraction of Requirements

Input: Vector database $\mathcal{V}$, Facility attributes a_f, Parameter k

Output: Required autonomy T_f^{req}, Metadata M_f

`// STEP 1: Convert attributes into strings and embed`

$s_q, v_q \leftarrow \text{StringifyAndEmbed}(a_f)$

`// STEP 2: Search database for the best matches`

$Scores \leftarrow [\,]$

for $c \in \mathcal{V}$ **do**

  $sim \leftarrow \text{CosineSimilarity}(v_q, c.\text{emb})$

  $Scores.\text{append}((sim, c))$

`// STEP 3: Retrieve top-`k` (most relevant) chunks`

Sort $Scores$ descending by sim

$Context \leftarrow \emptyset$

for $i \leftarrow 1$ **to** k **do**

  $(sim_i, c_{top}) \leftarrow Scores[i]$

  $Context \leftarrow Context \oplus c_{top}.\text{text}$

`// STEP 4: Construct prompt and run inference`

$\rho \leftarrow \text{Instruction}() \oplus Context \oplus s_q$

$J \leftarrow \text{LLM}(\rho, \text{mode} = \text{JSON})$

`// STEP 5: Parse and validate output`

if *isValidNumber(J.hours)* **then**

  $T_f^{\text{req}}, M_f \leftarrow J.\text{hours}, \{J.\text{condition}, J.\text{doc\_id}\}$

else $T_f^{\text{req}}, M_f \leftarrow \text{FlagManualReview}()$

return T_f^{req}, M_f

6.4.2 Modeling Scenario-dependent Indices

a) Assessing the Resupply Difficulty Index

A facility's preparedness also depends on its ability to replenish exhausted resources from external partners. While on-site autonomy is finite, it can be extended if external resupply begins early and transport routes remain usable. We quantify this logistical challenge through the resupply difficulty index $d_{f,s}$ (Eq. (6.1)), which balances temporal delays against structural network fragility. This metric shifts the focus from static building attributes to the dynamic strain caused by response timing and transport network vulnerability. To model the infrastructure, we use the ArcGIS network analyst extension [54] to represent the road system as a weighted graph where edges incorporate historical traffic data to account for typical congestion. Health care facilities typically maintain partnerships with other sister facilities (within the same network) or contracted vendors to secure resources during emergencies [96]. Consequently, we evaluate the earliest feasible resupply start time $T^{\text{start}}_{f,s}$ using Dijkstra's algorithm to identify a set of K potential shortest paths between a facility f and its suppliers.

We define the route redundancy $\rho_{f,s}$ as the specific count of these K paths that remain accessible after removing edges to reflect scenario-specific road closures (detailed in §6.4.2). As defined in Eq. (6.1), the difficulty index captures two primary operational bottlenecks. *Temporal constraints*: higher resupply start time $T^{\text{start}}_{f,s}$ relative to the outage duration increases $d_{f,s}$; *Structural constraints*: difficulty is inversely related to route availability. That is, as the number of accessible partners drops toward zero, the term $(1-\rho_{f,s}/K)$ approaches one, signaling a near single point of failure. We use α_d as a tunable knob to prioritize the timeliness of delivery and the availability of access routes. Ultimately, this index provides a comprehensive measure of the logistical constraints faced by facilities during a disaster. This index is critical because it dictates the feasibility and timing of resource arrivals, which in turn determines the total energy supply available to mitigate the electric power demand to cover the estimated duration of the outage.

b) Quantifying the Backup Shortfall Index

Defined as $s_{f,s} \in [0,1]$, this index measures the fraction of an outage duration not covered by a facility's total accessible power. A value of $s_{f,s}=0$ indicates full coverage, while higher values represent periods where critical services are projected to fail due to energy exhaustion. A facility's effective autonomy is thus the duration it can sustain operations by combining its static on-site capacity with dynamic replenishment resources. Per Eq. (6.2), it is computed as $T^{\text{eff}}_{f,s} = T^{\text{crit}}_{f} + (E^{\text{del}}_{f,s}/P^{\text{crit}}_{f})$, indicating the combined critical autonomy (minimum to sustain critical equipment and services) and the autonomy produced by the electrical backup equipment delivered by the partnered facilities. The expected outage duration $T^{\text{out}}_{f,s}$ is derived using FEMA Hazus-MH's methodology, which estimates restoration timelines by applying fragility curves to utility infrastructure (substations and distribution nodes) under scenario-specific hazard intensities (see §6.4.2). Consider a facility with P^{crit}_{f}=90 kW during a 72-hour outage ($T^{\text{out}}_{f,s}$=72). With only 24 hours of critical autonomy (T^{crit}_{f}=24), it requires 48 hours of supplemental power, necessitating a delivery of $E^{\text{del}}_{f,s} = 4,320$ kWh. If logistical difficulty ($d_{f,s}$) limits the feasible delivery to only 2,160 kWh (providing 24 hours of power), the effective autonomy reaches only 48 hours ($24 + 24$). Therefore, we quantify this facility's shortfall as $s_{f,s}$=$(72 - 48)/72 = 0.33$.

c) Physics-Based Hazard Exposure and Damage Estimation

In Alg. 8, we detail the process of capturing the total damage and loss of function of facilities caused by the hazard. This requires modeling the initial damage (lines 2-7) and its resulting downstream impacts (lines 8-12). To keep this work focused, we consider earthquake hazard events.

We start by leveraging historical catalogs of hazard events to instantiate and run N stochastic simulations. The likelihood of a hazard scenario depends on probability distributions that describe where such events are likely to occur, e.g., proximity to fault lines. For earthquake events, this data can be obtained from USGS [129], and simulated using NGA-West2 [63]; this estimates aspects such as the scale and intensity of ground motion, site effects, and epistemic uncertainty. We then analyze each hazard event for its potential to cause physical damage and loss of function to both facilities and electric grid components using FEMA Hazus [61]. This tool uses engineering-based fragility and functionality models to estimate structural damage, service disruption, and downtime. This yields a *damage ratio* and *loss of function* value for each facility f and grid substation x, as the extent of damage and subsequent reduction in operational capacity, respectively.

Following substation failures, the grid reroutes power subject to power flow balance and capacity constraints, but may need to shed load to maintain safe operation. This leads to an increased loss of function in nearby facilities. We extend our prior physics-based grid model [33] to compute post-hazard flows in Eq. (6.9). When a substation's damage ratio exceeds a threshold τ, we assume that it is no longer safe to operate. Such failures propagate through the network, starting with the power lines connected to the substation. To capture the cascading failure, we relate several power grid concepts[1]: generation p^g, demand p^d, load shed p^s, flows p, maximum generation $\bar{p}^g$,

[1]Subscripts are omitted for ease of presentation

Algorithm 8: Computing Physics-based Damage and Loss of Function

Input: Hazard Event Catalog $catalog$, Community cm, Electric Grid $\mathcal{G}$, Facilities $\mathcal{F}$, int N

Output: Damage States $\mathcal{D}$, Estimated Loss of Function $\mathcal{L}$

$\mathcal{D} \leftarrow \{\}$ $\mathcal{L} \leftarrow \{\}$ `// Damage, Loss of function repository`

for $i \leftarrow 1$ **to** N **do** `// Phase 1: Model Initial Damage`

  $s \leftarrow InitializeHazard(catalog, cm)$

  $res \leftarrow SimulateHazard(s)$

  $damage, loss \leftarrow EstimateDamageAndLoF(res, \mathcal{F})$

  $\mathcal{D}.addDamageRatioScenario(s, damage)$

  $\mathcal{L}.addLossOfFunctionScenario(s, loss)$

for $s \leftarrow \mathcal{S}$ **do** `// Phase 2: Model Cascading Impact`

  $X \leftarrow GetHighlyDamagedComponents(s)$

  $\mathcal{G}^* \leftarrow ExtractReducedGrid(\mathcal{G}, X)$

  $\{p_i^s\} \leftarrow$ Solve for power flow in (6.9)

  $\mathcal{L}.updateLossOfFunction(s, \{p_i^s\})$

return $\mathcal{D}, \mathcal{L}$

line capacity t, reactance X, and phase angles θ to the failures x.

$$\max \quad \sum\nolimits_{i \in \mathcal{N}} (p_i^d - p_i^s) \tag{6.9a}$$

$$\text{s.t.} \quad p_i^g - p_i^d + p_i^s = \sum_{(i,j) \in \mathcal{E}^f(i)} p_{i,j} - \sum_{(j,i) \in \mathcal{E}^t(i)} p_{j,i} \qquad \forall i \in \mathcal{N} \tag{6.9b}$$

$$p_{i,j} = X_{ij}^{-1}(\theta_i - \theta_j)(1 - x_{i,j}) \qquad \forall (i,j) \in \mathcal{E} \tag{6.9c}$$

$$0 \leq p_i^g \leq (1 - x_i)\bar{p}_i^g \qquad \forall i \in \mathcal{N} \tag{6.9d}$$

$$-t_{i,j}(1 - x_{i,j}) \leq p_{i,j} \leq t_{i,j}(1 - x_{i,j}) \forall (i,j) \in \mathcal{E} \tag{6.9e}$$

The objective in (6.9a) maximizes served power in the network, subject to power flow physics (6.9b, 6.9c) and grid capacity limits (6.9d, 6.9e). Solving this optimization problem can then be accomplished by using linear program solvers [78] offline.

NB: Subscripts are omitted for ease of presentation.

6.5 Towards Facility Preparedness Scorecards and Recommendations

Overall, we synthesize the individual indices into a unified preparedness risk score, denoted as $r_{f,s} \in [0,1]$. This composite score is computed by aggregating all the indices via a weighted sum as defined in Eq. 6.4 (lines 2–12 in Alg. 9). Based on Eq. 6.5, $R_s^{\max}$ represents the most vulnerable facility in each scenario s. However, for operational efficiency, we partition facilities into C_s discrete risk tiers. We implement this using k-means clustering by grouping facilities into risk tiers based on the proximity of their risk scores to computed tier centroids (line 13). We record the average risk score $\frac{1}{|\mathcal{S}|}\sum_{s\in\mathcal{S}} r_{f,s}$ (line 15) and the worst-case risk $\max_{s\in\mathcal{S}} r_{f,s}$ and map the average score to a tier/color as shown in Table 6.1. Tier labeling and interpretations are inspired by the Simple Triage And Rapid Treatment (START) convention [13], where tiers use fixed risk-score ranges for consistent interpretation across scenarios (line 16). Because we evaluate the risk score through many outage scenarios ($\sim$1000), the *scorecard* is reported per facility by summarizing risk over the scenario ensemble.

BlackOps also supports the following recommendations: #1) The optimal supplier(s) selected by minimizing the resupply difficulty ($d_{f,s}$) (based on top K suppliers). #2) The amount of power needed to cover the estimated deficit. Here, the deficit-hours x correspond to the unmet autonomy duration $max(0, T_f^{\text{req}} - T_f^{\text{crit}})$ and P_f^{crit} indicates the capacity of the recommended power supply (generator). #3) The LLM snapshot summarizes the relevant excerpts discussed in §6.4.1. Note that the snapshot does not affect the risk score, but rather serves as a context-aware summary of the facility.

Algorithm 9: Scorecards & Recommendations

Input: Facilities $\mathcal{F}$, Scenarios $\mathcal{S}$, GIS layers $\mathcal{L}$, Requirements $\{T_f^{\text{req}}, M_f\}$, Number of tiers C_s

Output: Risk scores $\{r_{f,s}\}$, Tiers $\{c_{f,s}\}$, Scorecards SC, Recommendations REC

$\mathsf{SC} \leftarrow \{\}$; $\mathsf{REC} \leftarrow \{\}$ `// scorecard, recommendation`

for $f \in \mathcal{F}$ **do**

  $T_f^{\text{crit}} \leftarrow \textsc{CriticalAutonomy}(P_f^{\text{crit}}, E_f^{\text{store}})$

  $g_f \leftarrow \textsc{AutonomyGap}(T_f^{\text{req}}, T_f^{\text{crit}})$

  `//` $\mathcal{L}(\cdot)$ `are queryable GIS layers (feature class (point, line, polygon)/raster/network dataset)`

  $v_f \leftarrow \textsc{SVIJoin}((\text{lat}_f, \text{lon}_f), \mathcal{L}.\text{SVI})$

for $s \in \mathcal{S}$ **do**

  $(H_s, \{T_{f,s}^{\text{out}}\}) \leftarrow \textsc{ScenarioGen}(\mathcal{L}.\text{Hazards}, s)$

  for $f \in \mathcal{F}$ **do**

    $e_{f,s} \leftarrow \textsc{Exposure}((\text{lat}_f, \text{lon}_f), H_s, \mathcal{L}.\text{Grid})$

    $d_{f,s} \leftarrow \textsc{ResupplyDiff}((\text{lat}_f, \text{lon}_f), H_s, \mathcal{L}.\text{Roads})$

    $s_{f,s} \leftarrow \textsc{Shortfall}(T_f^{\text{crit}}, P_f^{\text{crit}}, T_{f,s}^{\text{out}}, d_{f,s})$

    $r_{f,s} \leftarrow \textsc{RiskScore}(e_{f,s}, s_{f,s}, d_{f,s}, g_f, v_f, \{w_i\})$

  $\{c_{f,s}\}_{f\in\mathcal{F}} \leftarrow \textsc{KMeansCluster}(\{r_{f,s}\}_{f\in\mathcal{F}}, C_s)$

for $f \in \mathcal{F}$ **do**

  $\bar{r}_f \leftarrow \textsc{Mean}(\{r_{f,s}\}_{s\in\mathcal{S}})$

  $(\text{Tier}_f, \text{Color}_f, \text{Interp}_f) \leftarrow \textsc{TierMap}(\bar{r}_f)$

  $\Delta E_f \leftarrow \textsc{EnergyDeficit}(T_f^{\text{req}}, T_f^{\text{crit}}, P_f^{\text{crit}})$

  `// Closest accessible supplier(s)`

  $\text{Sup}_f \leftarrow \textsc{OptSupplier}(f, \{d_{f,s}\}_{s\in\mathcal{S}}, \mathcal{L}.\text{Suppliers})$

  $\text{Snap}_f \leftarrow \text{LLM}(M_f)$ `// LLM snapshot from metadata`

  $\mathsf{SC}[f] \leftarrow (\bar{r}_f,\ \text{Tier}_f,\ \text{Color}_f,\ \text{Interp}_f)$

  $\mathsf{REC}[f] \leftarrow (\text{Sup}_f,\ \Delta E_f,\ \text{Snap}_f)$

return $\{r_{f,s}\}, \{c_{f,s}\}, \mathsf{SC}, \mathsf{REC}$

Table 6.1: Facility Risk Tiers

$r_{f,s}$	**Tier**	**Description**
$[0, 0.25)$	Low (Green)	Stable: minimal external action expected.
$[0.25, 0.50)$	Moderate (Yellow)	Monitor: secondary actions as needed.
$[0.50, 0.75)$	High (Orange)	Urgent: resources likely needed mid-outage.
$[0.75, 1]$	Extreme (Red)	Critical: immediate resource/pre-staging needed.

Scorecard

Facility: *[name, type, location]*

Risk score (avg): $\frac{1}{|\mathcal{S}|}\sum_{s\in\mathcal{S}} r_{f,s}$

Risk score (max): $\max_{s\in\mathcal{S}} r_{f,s}$

Tier/Color: *[Tier, Color from Table 6.1]*

Interpretation: *[row meaning from Table 6.1]*

Recommendations: #1) Optimal supplier(s): *[facility/vendor]*

#2) Based on x hrs autonomy deficit $\Rightarrow$ provision $\approx x \times P_f^{crit}$ kWh usable backup energy.

#3) LLM snapshot: *["context-aware summary of facility f"]*

6.6 Experiments

Setup: We evaluated BlackOps using stochastic earthquake events from the USGS 2018 NSHMP catalog [129], and estimated ground motion using NGA-West2 [19]. Damage analysis then applied the FEMA Hazus Earthquake Model [61]. Facility data (974 sites) were obtained from the California Department of Public Health (CDPH) [31], and power grid infrastructure data was sourced from the California Test System (CATS) [145]. We stored these in a PostgreSQL/PostGIS database to enable efficient geospatial queries. Road network data and historical traffic information were integrated via ArcGIS Pro to model distances and travel times between facilities. Requirements and metadata were extracted via a distributed client-server architecture, offloading inference to a workstation (AMD Ryzen7 5800X, NVIDIA RTX3080 GPU). Using the Ollama API [2], we deployed GPT-OSS-20b and Gemma 3 12B models with 32k context windows.

To validate the LLM-based requirement extraction, we ran a small sanity-check with curated ex-

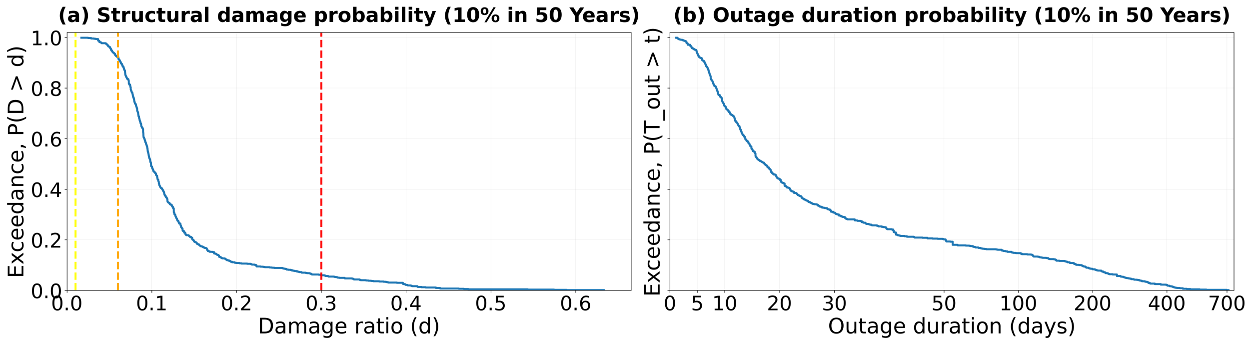

Figure 6.3: Exceedance Curves (Long-term Earthquake Risk)

cerpts (snapshots of NFPA 99 and 110 standards). We evaluated (i) extraction accuracy by querying for requirements with known T_f^{req} values present in the excerpts (e.g., 96 hours) and (ii) hallucination resistance by issuing negative-control queries whose requirements were intentionally absent from the provided text. The model was required to return "unknown" and trigger manual review rather than invent values. This check ensures outputs are grounded in retrieved documents and that missing requirements do not produce fabricated autonomy values. The next set of experiments ranged from long-term seismic risk quantification and spatial damage characterization to resupply routing optimization, end-to-end risk integration, and comparative benchmarking against an existing federal assessment tool (ASPR RISC 2.0).

We quantified long-horizon seismic risk using the "10% in 50 years" exceedance baseline. Fig. 6.3(a) shows the structural damage probability $P(D>d)$, where the curve identifies the fraction of facilities exceeding critical thresholds. We then estimate the probability $P(T_{\text{out}}>t)$ that facilities will exceed specific outage durations (Fig. 6.3(b)). This provides a predictive stress-test for response timelines by highlighting a "prolonged failure tail" where estimated recovery durations (t) significantly exceed the standard 96-hour backup autonomy.

In Fig. 6.4, we illustrate spatial damage results for Orange County facilities. Damage severity is classified based on the damage ratio d: None (<1%), Slight (1-6%), Moderate (6-30%), Extensive (30-75%), and Complete (>75%). The high-quantile damage (Q90) captures high-impact outcomes, indicating the level of exposure a facility experiences in approximately 1 out of 10 scenarios. In contrast, the mean damage summarizes the typical (central-tendency) impact across the scenario

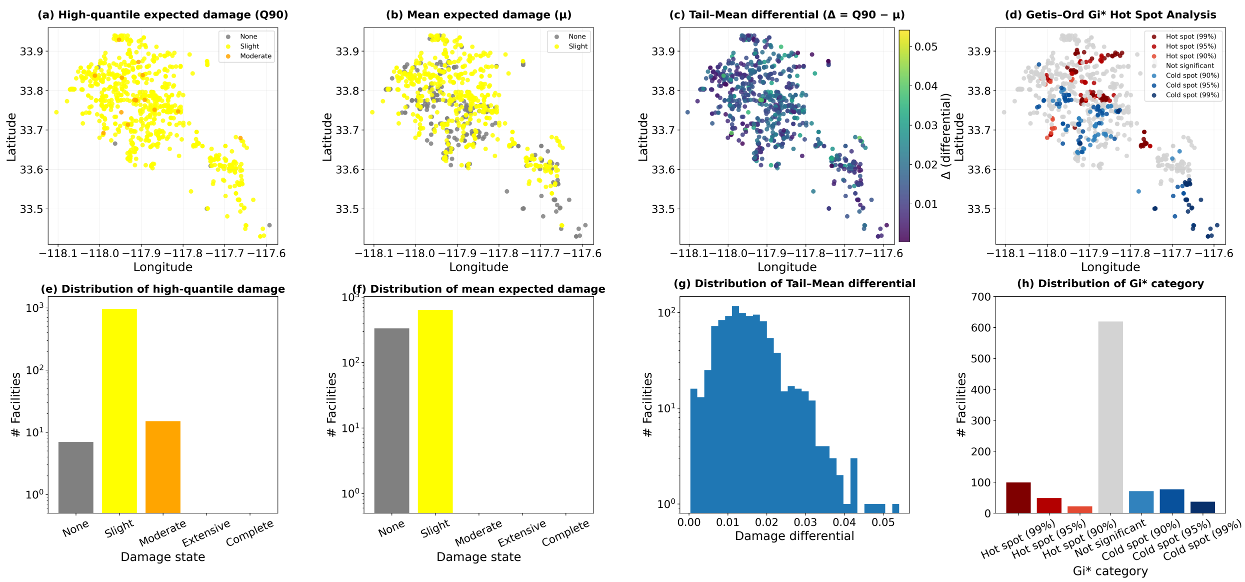

Figure 6.4: High-quantile Damage (Q90), Mean Damage (μ), Tail–mean Damage Differential (Δ), Getis–Ord G_i^* Hot/cold Spots with Distributions

ensemble. With these marginal metrics, we compute the tail–mean damage differential (Δ), which quantifies cross-scenario dispersion and highlights facilities whose tail-risk is disproportionately larger than their typical outcome. Finally, the Getis–Ord G_i^* identifies statistically significant spatial correlation where high (or low) damage forms non-random hot/cold clusters (computed via z-scores and confidence levels). The G_i^* statistic distinguishes isolated outliers from statistically significant clusters by evaluating each facility in the context of its neighboring values.

We also use SVI to prioritize resource acquisition from low-vulnerability areas, thereby minimizing the burden on highly vulnerable communities. Travel times between origins and top-ranked partners (Fig. 6.5) are computed using Dijkstra's shortest-path algorithm, which incorporates historical traffic data to account for congestion (at different times of the day). To simulate infrastructure degradation, road segments within a defined buffer radius of facilities and substations reaching 'Extensive' or 'Complete' damage states are assumed to be impassable. Furthermore, to optimize the routing search space, we prioritize higher-order road networks. Hence, the resupply difficulty index d_f is obtained by mapping each facility's top-K partner set to the earliest feasible start time T_f^{start} and the availability count ρ_f capturing how many partners are reachable. For one of the illustrative facilities in Fig. 6.5 with K=5 candidate suppliers, we obtain d_f=0.240 with

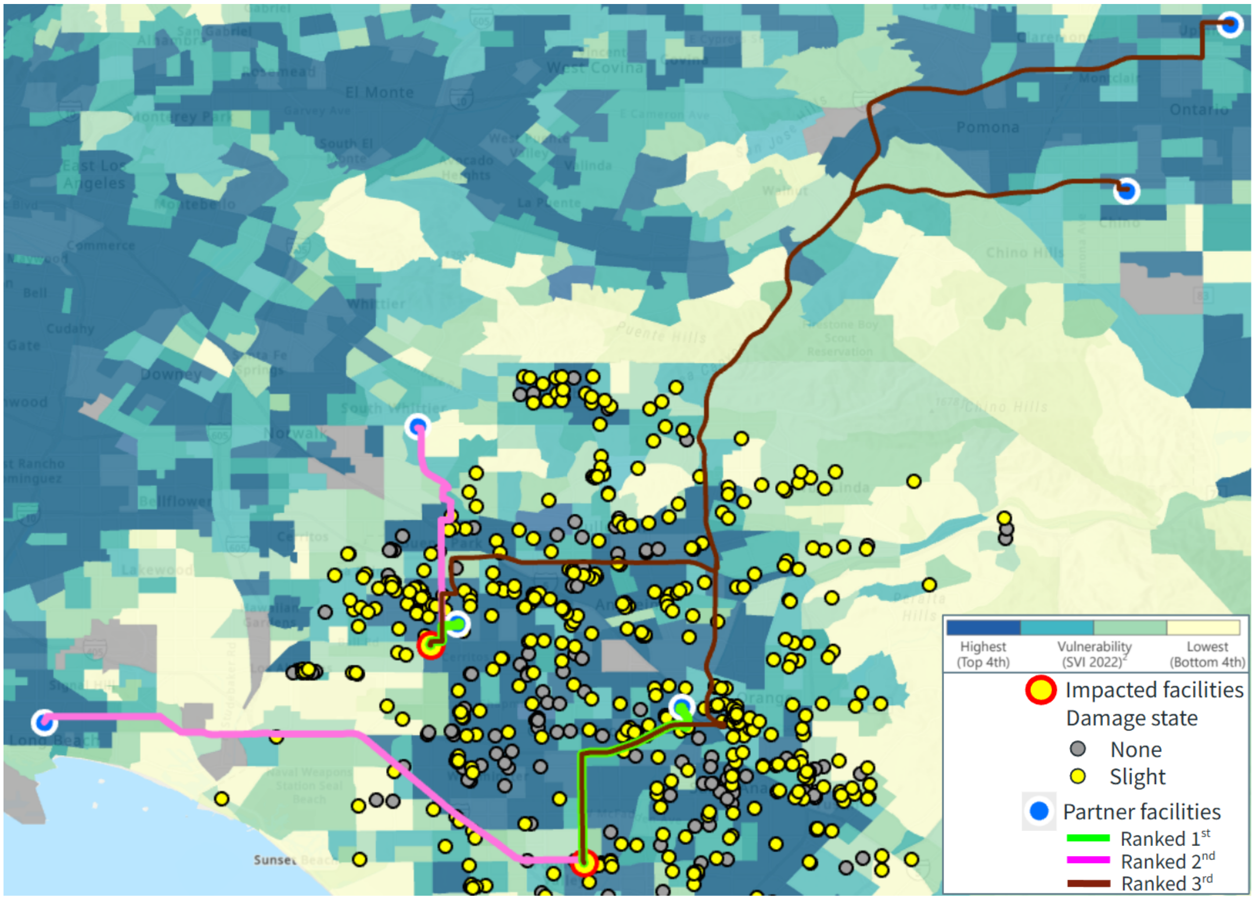

Figure 6.5: Distribution of Suppliers for Selected Facilities Relative to Social Vulnerability Index (SVI)

T_f^{start}=3.42 h, and ρ_f=3 corresponding to 9.03/45.46/84.00 km.

Fig. 6.6 provides end-to-end results of the overall risk score after combining all indices. As shown in Fig. 6.6(a), the overall risk distribution reveals a non-linear gradient with a steep, small subset of high-priority facilities (left), whereas the long, flatter tail indicates that the majority of the network clusters within moderate-to-lower risk levels. This is further quantified by the risk exceedance (Fig. 6.6(b)) and the resulting vulnerability classification (Fig. 6.6(c)), which together provide a visual breakdown of how many facilities are really high risk versus those that are moderate. However, the most critical insights emerge from the indices comparison (Fig. 6.6(d)), where it is evident that the top-100 high-risk facilities exhibit significantly higher mean intensities across all dimensions. Most notably, in the shortfall, SVI, and exposure indices.

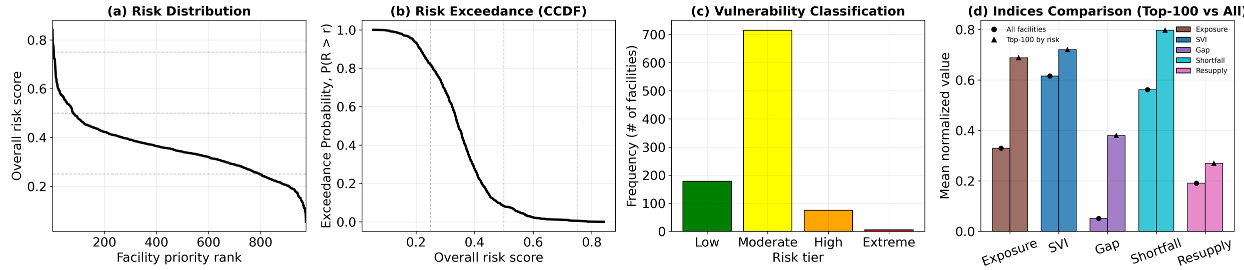

Figure 6.6: End-to-End Risk Evaluation

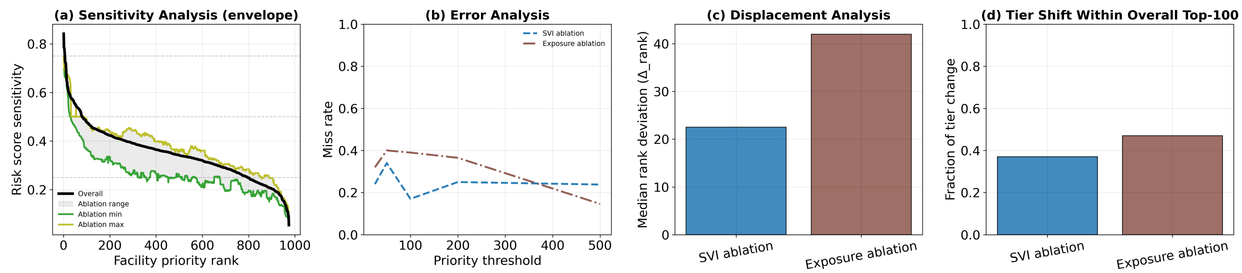

Figure 6.7: Risk Model Ablation Study

To validate the integration of all indices, we conducted an ablation study as shown in Fig. 6.7. The sensitivity analysis defines the operational bounds of the model by mapping the minimum and maximum possible risk scores across all indices (Fig. 6.7(a)). The shaded grey area (the envelope) represents the sensitivity of our model. Hence, wide shading indicates where the facility's rank is highly dependent on a specific index. For example, a facility could have very high social vulnerability but low structural damage (exposure). Therefore, by dropping the SVI, its risk score will considerably decrease. Similarly, a narrow shading indicates a "unanimous" risk profile where the indices converge (at the very high and very low ends of the rank). Thus, by removing the SVI and exposure indices, we highlight the importance of a multi-dimensional approach, as shown in Fig. 6.7(b), which reveals significant miss rates of 20% to 40% when neglecting SVI or exposure. Furthermore, the displacement analysis (Fig. 6.7(c)) and the high degree of tier disagreement (Fig. 6.7(d)) confirm that without this integration, nearly 40% of the most critical facilities would be misclassified by decision-makers.

In addition, to benchmark against an established federal preparedness framework, we utilize the ASPR RISC Toolkit 2.0 [84] for a representative facility (UCI Health Fountain Valley) and compare its risk ratings to the corresponding BlackOps estimates. RISC, which computes risk as the product of Threat/Hazard, Vulnerability, and Consequence modules (THAM×RIST-V×RIST-C), yielded an earthquake exposure score of 0.96, whereas BlackOps produced an exposure value of 0.67, corresponding to a 30.2% reduction relative to the national-default assessment. In contrast, RISC reported a vulnerability score of 0.02 (national value) for the same facility, while the fine-grained census-tract and facility-specific inputs used in BlackOps resulted in an SVI of 0.76. Because RISC documentation explicitly notes that THAM hazard ratings are derived from national-level data and may be overwritten when local data or subject-matter expertise indicates different likelihoods, we interpret these differences as methodological rather than contradictory, with BlackOps as a systematic local refinement layer that can complement standardized RISC assessments for regional outage planning.

6.7 Chapter Summary and Discussion

We presented *BlackOps*, a GIS-informed decision-support framework for assessing healthcare facility risk amid prolonged power outages. While the previous chapter focused on matching and routing during disaster response, this work shifts attention to the question of whether healthcare facilities can sustain critical services when power-grid failures, backup-resource limitations, road-network constraints, and community vulnerability interact. The central concern is not only whether a facility is exposed to a hazard, but whether it can continue operating when the infrastructure systems it depends on are degraded.

The framework integrates hazard footprints, healthcare facility attributes, electric-grid dependencies, road networks, supplier locations, regulatory requirements, and social vulnerability indicators into a unified risk assessment workflow. Its interpretable indices capture different dimensions of outage risk, including autonomy gaps, backup shortfall, resupply difficulty, social vulnerability, and physical hazard exposure. By aggregating these indices into a facility-scenario risk score, the framework supports the identification of high-priority facilities and translates risk estimates into scorecards and recommendations for decision-makers.

The results illustrate the value of connecting geospatial analysis with infrastructure modeling. Physics-based simulations over the Orange County healthcare network showed how facility damage, power-grid disruption, and cascading infrastructure impacts can affect healthcare continuity. The results also showed that partial or coarse assessments can miss important vulnerabilities. In particular, the ablation analysis demonstrated that removing dimensions such as social vulnerability or exposure can change which facilities are identified as high priority. This supports the broader dissertation argument that healthcare disaster resilience requires multi-dimensional and spatially grounded decision support.

The work also demonstrates the bounded use of AI-assisted extraction for facility risk assessment. Regulatory and emergency-power requirements are often embedded in heterogeneous documents and are difficult to translate into facility-specific planning variables. By using retrieval-augmented generation to extract required autonomy and metadata, the framework helps bridge the gap between broad regulatory guidance and computable assessment. Importantly, the language model is used as a constrained extraction component, while the final risk assessment remains grounded in GIS layers, infrastructure models, and interpretable indices.

Overall, this work completes the dissertation's progression from in-facility allocation, to cross-facility coordination, to facility risk assessment under prolonged infrastructure disruption. Together with the previous chapters, it shows how geospatial systems can support healthcare resilience by helping facilities allocate resources fairly, helping regions coordinate relocation and resource sharing, and helping planners identify where infrastructure failures may threaten continuity of care. Future work can further calibrate the model with ground-truth outage, facility, and utility data; expand the hazard scope to include real-time fire, flood, and compound-event scenarios; and evaluate scorecard recommendations with emergency managers and healthcare facility operators. These directions are discussed more broadly in the final chapter.

Chapter 7

Conclusions and Future Directions

7.1 Overview

This dissertation examined healthcare disaster resilience as a spatial-temporal decision-support problem. The central motivation was that healthcare systems must make difficult decisions when hazards evolve, infrastructure is disrupted, information is incomplete, and resources are scarce. These decisions are not only logistical. They are also human-centered, spatially constrained, time-sensitive, and shaped by infrastructure and organizational dependencies. Across this work, geospatial systems were used to model where people, facilities, resources, hazards, and supporting infrastructure are located, and how those relationships affect preparedness and response.

The dissertation began with urgent allocation inside senior health facilities, where staff and equipment must be assigned to residents with heterogeneous medical needs, mobility limitations, preferences, and levels of urgency. It then expanded to coordination across healthcare networks, where facilities, roads, resources, hazards, and operational policies must be considered together during disasters. Finally, it examined facility risk under prolonged power outages, where healthcare facilities may lose critical services even when they are not directly damaged by the hazard itself. Together, these perspectives provide a connected view of healthcare resilience: supporting immediate care delivery, coordinating relocation and resource sharing, and identifying facilities vulnerable to cascading infrastructure disruptions.

A central lesson from this work is that healthcare disaster resilience depends on both spatial awareness and temporal awareness. Spatial awareness helps decision-makers understand where residents, hazards, resources, roads, vulnerable communities, and infrastructure assets are located. Temporal awareness helps them understand how quickly decisions must be made, how urgency changes over time, how long routes may take, how fast hazards evolve, and how long facilities can sustain critical services during outages. These two dimensions are inseparable. A facility may be close to an impacted site but unreachable because of road closures. A resident may be physically near an exit but require specialized assistance. A hospital may have available beds but be a poor destination if its operating status is degraded. Likewise, a facility may appear safe from direct hazard exposure but still become vulnerable if its backup autonomy is insufficient. These examples show why healthcare disaster resilience requires decision support that integrates people, places, infrastructure, and time.

7.2 Summary of Contributions

The first contribution addressed resource allocation within senior health facilities. This work showed that older adults should not be modeled as interchangeable units in a scheduling problem. Residents differ in clinical condition, mobility, equipment needs, preferences, and sensitivity to delay. By modeling facility layouts as graphs, computing resident criticality, and incorporating preferences into task completion and satisfaction, the framework demonstrated that fair allocation can also improve operational efficiency. This is important because fairness is often treated as a competing objective that slows down response. In senior care settings, however, ignoring individualized needs can create delays and reduce the effectiveness of care. When those needs are modeled directly, allocation can become both more equitable and more efficient.

The second contribution expanded the problem from one facility to a healthcare network. During disasters, facilities may need to relocate residents, receive evacuees, share resources, or coordinate with Emergency Operations Centers. These decisions cannot be made using distance alone.

They require reasoning about capacity, operational status, residency type, resource availability, safe routes, hazard exposure, and social vulnerability. This work showed how policy-driven stable matching and hazard-aware routing can be combined to support regional coordination. It also showed that matching and routing should not be treated as separate decisions. A destination may be appropriate from a policy perspective but unreachable, while a route may be short but lead to a facility that lacks the required capacity or services. Integrating these decisions provides a more practical foundation for disaster response.

The third contribution shifted from response coordination to facility risk assessment amid prolonged infrastructure disruption. Power outages pose a serious risk to healthcare facilities because critical services depend on electricity, transportation, suppliers, and regional partners. This work connected regulatory requirements, backup autonomy, road-network accessibility, supplier availability, social vulnerability, hazard exposure, and power-grid dependencies into a multi-dimensional risk assessment. The resulting scorecards and recommendations were designed to help planners understand not only which facilities are vulnerable, but also why they are vulnerable and what actions may reduce that risk.

Together, these contributions show that geospatial systems can serve as a computational foundation for healthcare disaster resilience. Spatial information supports indoor movement, resident location, resource assignment, and hazard proximity within facilities. It supports matching, routing, and coordination across healthcare networks. It also helps identify facilities and communities whose service continuity may be threatened by cascading disruptions. The broader contribution is therefore not only the development of three separate systems, but the demonstration that spatial data and computational models can be connected to support more human-centered, adaptive, and interpretable disaster resilience decisions.

7.3 Key Takeaways

One of the clearest takeaways is that fairness and efficiency should not be treated as opposing goals by default. In senior healthcare, efficiency depends on understanding the people being served. A resident's mobility limitation, medical condition, preference, or need for specialized assistance can directly affect how long a task takes and how successful the response will be. When these factors are ignored, an allocation may appear simple or efficient in theory, but become slower or less effective in practice. By modeling resident criticality, preferences, and facility layout together, this dissertation demonstrated that fairness-aware decisions can also improve response efficiency.

The work also reinforced the importance of grounding computational models in real operational contexts. The technical direction of this research was shaped by engagement with senior living partners, healthcare experts, emergency managers, public health stakeholders, and hazard analytics collaborators. These collaborations revealed practical gaps that are difficult to identify from algorithmic formulations alone, including resident tracking, family reunification, relocation planning, hospital surge coordination, resource sharing, backup-power planning, regulatory compliance, and the need for interpretable dashboards. The value of the models developed here is therefore not only that they optimize formal objectives, but that those objectives were shaped by real preparedness and response needs.

Another takeaway is that geospatial systems are most powerful when they are connected to decision models rather than used only for visualization. Maps help decision-makers see where hazards, facilities, roads, resources, and infrastructure assets are located, but spatial awareness alone does not determine what action should be taken. The frameworks in this dissertation show how spatial data can become part of the decision logic itself. Facility layouts support in-building allocation, road networks support matching and routing, and infrastructure layers support facility risk assessment under cascading outages. This turns GIS from a passive display layer into an active computational substrate for healthcare resilience.

The role of AI in this work also points to a broader design principle for high-stakes systems. In the facility risk assessment framework, large language models were not used as autonomous decision-makers. Instead, they were used for a bounded extraction task: identifying facility-specific autonomy requirements from heterogeneous regulatory documents. The extracted information was then connected to GIS layers, infrastructure models, and interpretable risk indices. This approach preserves the benefits of AI, such as semantic retrieval and structured extraction, while limiting the risks of unsupported recommendations. For healthcare disaster resilience, AI is most useful when it improves traceability, interpretation, and access to complex information, while final decisions remain grounded in validated models and human oversight.

Finally, the dissertation shows that resilience requires a multi-dimensional view of risk. A facility's vulnerability cannot be explained by physical damage, social vulnerability, backup autonomy, road access, or travel time alone. High-risk facilities may emerge from different combinations of exposure, shortfall, resupply difficulty, autonomy gaps, and community vulnerability. This matters because partial assessments can miss facilities that require support or misclassify their level of need. A more complete approach must integrate physical, operational, social, and infrastructure-aware indicators.

Overall, these takeaways suggest that healthcare disaster resilience is fundamentally a systems problem. Decisions made inside a facility can create regional demand. A relocation plan can fail if infrastructure dependencies are ignored. A risk assessment can be incomplete if it does not account for the needs of the people who depend on the facility. The value of a geospatial decision-support approach is that it can connect these layers into a more coherent operational picture.

7.4 Future Directions

A natural next step is to connect the three frameworks into a more unified disaster resilience platform. In the current work, the in-facility, cross-facility, and facility-risk models address related but distinct decision contexts. Future systems should allow information to flow across these contexts. For example, a prolonged outage risk identified during preparedness could trigger relocation analysis. Regional assignments could then incorporate resident-level constraints when relocating older adults from long-term care facilities. Similarly, facility status updates observed during response could update coordination and risk models. Such integration would move the work toward a continuous decision-support pipeline spanning preparedness, response, and recovery.

Another important direction is to strengthen validation with operational data. The evaluations in this dissertation used anonymized facility data, synthetic resident profiles, public GIS layers, hazard models, drill data, and scenario-based simulations. These were appropriate for developing and evaluating the frameworks under controlled conditions. Broader deployment, however, would require deeper validation with after-action reports, facility maintenance records, emergency operations data, utility outage histories, transportation disruptions, supplier agreements, and real resource-sharing workflows. These data would support better calibration of model weights, travel-time assumptions, outage-duration estimates, autonomy requirements, and scorecard recommendations.

The hazard scope can also be expanded. This dissertation considered fire evacuation, earthquake response, water contamination, and prolonged power outages. The same spatial-temporal modeling approach could be extended to floods, extreme heat, public safety power shutoffs, cyber-physical attacks, communication outages, and compound disasters. This is especially important because healthcare disruptions often emerge from interacting systems. Hazards can trigger evacuations, road closures, air-quality threats, power shutoffs, staffing shortages, and communication failures at the same time. Modeling these interactions more deeply would improve the realism of future decision-support systems.

Privacy and governance remain central open challenges. The data needed to support effective decisions can be sensitive. Resident location, medical needs, mobility status, care preferences, facility vulnerabilities, and backup-power limitations all raise privacy, security, and governance concerns. Future systems should investigate role-based access control, policy-aware data sharing, anonymization, secure computation, and minimum-necessary disclosure. The goal should be to ensure that emergency responders receive the information they need, when they need it, without exposing more personal or operational data than necessary.

Finally, future work should study how decision-support systems can balance automation with human judgment. The frameworks developed here are intended to support decision-makers, not replace them. This distinction is important in healthcare disaster response, where accountability, trust, local knowledge, and ethical judgment are essential. Future studies should examine how emergency managers interact with algorithmic recommendations, when they override them, what explanations they need, and how systems can learn from expert feedback while maintaining transparency and safety.

7.5 Closing Remarks

Healthcare disaster resilience requires coordinated decision-making across diverse stakeholders. This dissertation addresses this complexity by developing geospatial decision-support frameworks for equitable in-facility allocation, regional matching and routing, and risk assessment modeling. While each framework targets a distinct decision context, collectively they demonstrate how computational models can enhance operational decision-making. The broader contribution of this work lies in the integration of geospatial systems for healthcare resilience. By grounding computational models in spatial context, human needs, policy constraints, and operational realities, this dissertation advances toward decision-support systems that genuinely empower the individuals responsible for protecting vulnerable populations during disasters. However, the challenge moving forward is to design systems that improve safety and resilience without compromising privacy. Achieving this will require continued collaboration with the stakeholders responsible for planning, delivering, and receiving care during crises. Ultimately, the goal is to build human-centric decision-support systems that fundamentally strengthen healthcare disaster resilience.

Bibliography

[1] CareDEX. `https://sites.uci.edu/caredex/`.

[2] Ollama API. `https://ollama.com/`.

[3] O*NET resource center. `https://www.onetcenter.org//`.

[4] Hazus: FEMA's methodology for estimating potential losses from disasters. `https://www.fema.gov/flood-maps/products-tools/hazus`, 2024.

[5] S. Aalami and L. Kattan. Fair transit trip planning in emergency evacuations: A combinatorial approach. *Transportation Research Part C: Emerging Technologies*, 122:102760, 2021.

[6] A. Abdalkareem. Healthcare scheduling in optimization context: A review. *Health and Technology*, 2021.

[7] A. M. Abdelhamid and A. A. Zamel. Implementing and measuring the performance of pb, rr and pbrr scheduling algorithms on atmega32a using freertos. In *2023 5th Novel Intelligent and Leading Emerging Sciences Conference (NILES)*, pages 18–22. IEEE, 2023.

[8] M. F. Abdullah, S. Siraj, and R. E. Hodgett. An overview of multi-criteria decision analysis (mcda) application in managing water-related disaster events: analyzing 20 years of literature for flood and drought events. *Water*, 13(10):1358, 2021.

[9] A. A. Adalja, M. Watson, N. Bouri, K. Minton, R. C. Morhard, and E. S. Toner. Absorbing citywide patient surge during hurricane sandy: a case study in accommodating multiple hospital evacuations. *Annals of emergency medicine*, 64(1):66–73, 2014.

[10] Agency for Toxic Substances and Disease Registry (ATSDR). CDC/ATSDR social vulnerability index (svi): Documentation and data, 2026.

[11] N. Aldrich and W. F. Benson. Disaster preparedness and the chronic disease needs of vulnerable older adults. *PCD*, 2008.

[12] American Academy of Family Physicians. Hierarchical condition category coding, 2026. Accessed: April 3, 2026.

[13] J. Bazyar, M. Farrokhi, and H. Khankeh. Triage systems in mass casualty incidents and disasters: a review study with a worldwide approach. *Open access Macedonian journal of medical sciences*, 7(3):482, 2019.

[14] L. Bedoya-Valencia and E. Kirac. Evaluating alternative resource allocation in an emergency department using discrete event simulation. *Simulation*, 92(12), 2016.

[15] S. Belciug and F. Gorunescu. Improving hospital bed occupancy and resource utilization through queuing modeling and evolutionary computation. *Journal of Biomedical Informatics*, 53:261–269, 2015.

[16] D. Bertsimas, V. F. Farias, and N. Trichakis. The price of fairness. *Operations research*, 59(1):17–31, 2011.

[17] M. F. Blum, Y. Feng, G. B. Anderson, D. L. Segev, M. McAdams-DeMarco, and M. E. Grams. Hurricanes and mortality among patients receiving dialysis. *Journal of the American Society of Nephrology*, 33(9):1757–1766, 2022.

[18] T. Bonald, L. Massoulié, A. Proutiere, and J. Virtamo. A queueing analysis of max-min fairness, proportional fairness and balanced fairness. *Queueing systems*, 53(1):65–84, 2006.

[19] Y. Bozorgnia, N. A. Abrahamson, L. A. Atik, T. D. Ancheta, G. M. Atkinson, J. W. Baker, A. Baltay, D. M. Boore, K. W. Campbell, B. S.-J. Chiou, et al. Nga-west2 research project. *Earthquake Spectra*, 30(3):973–987, 2014.

[20] M. Bruneau, S. E. Chang, R. T. Eguchi, G. C. Lee, T. D. O'Rourke, A. M. Reinhorn, M. Shinozuka, K. Tierney, W. A. Wallace, and D. Von Winterfeldt. A framework to quantitatively assess and enhance the seismic resilience of communities. *Earthquake spectra*, 19(4):733–752, 2003.

[21] J. Brunkard, G. Namulanda, and R. Ratard. Hurricane Katrina Deaths, Louisiana, 2005. *Disaster Medicine and Public Health Preparedness*, 2(4):215–223, 2008.

[22] E. K. Burke, P. De Causmaecker, G. Vanden Berghe, and H. Van Landeghem. The state of the art of nurse rostering. *Journal of Scheduling*, 7(6):441–499, 2004.

[23] G. P. Cachon and M. Fisher. Supply chain inventory management and the value of shared information. *Management science*, 46(8):1032–1048, 2000.

[24] CALFIRE. Palisades fire, n.d. Accessed Feb 2025.

[25] CalHHS. California Health and Human Services Agency, n.d. Accessed Apr 2024.

[26] California Department of Transportation. California department of transportation (caltrans) official website. `https://dot.ca.gov/`. Accessed: April 3, 2026.

[27] California Department of Transportation. Quickmap: Real-time traffic and road conditions. `https://quickmap.dot.ca.gov/`, 2026. Accessed: April 3, 2026.

[28] CDC. ICD-10-CM Tool. `https://icd10cmtool.cdc.gov/`.

[29] CDC/ATSDR. Social vulnerability index (SVI). `https://www.atsdr.cdc.gov/place-health/php/svi/index.html`, 2024.

[30] CDPH. AFL Emergency Water Requirements, 2023. Accessed Apr 2025.

[31] CDPH, 2026.

[32] L. Ceriani and P. Verme. The origins of the Gini index: Extracts from Variabilità e Mutabilità (1912) by corrado Gini. *The Journal of Economic Inequality*, 10(3):421–443, 2012.

[33] A. Chio, R. Bent, K. Sundar, and N. Venkatasubramanian. Sequin: A network science and physics-based approach to identify sequential nk attacks in electric power grids. In *Proceedings of the ACM/IEEE 16th International Conference on Cyber-Physical Systems (with CPS-IoT Week 2025)*, pages 1–12, 2025.

[34] A. Chio, D. Jiang, P. Gupta, G. Bouloukakis, R. Yus, S. Mehrotra, and N. Venkatasubramanian. Smartspec: A framework to generate customizable, semantics-based smart space datasets. *Pervasive and Mobile Computing*, 93:101809, 2023.

[35] T. Chondrogiannis, P. Bouros, and W. Emser. Simulation-based evacuation planning for urban areas. In *Proceedings of the 29th International Conference on Advances in Geographic Information Systems*, pages 297–300, 2021.

[36] M. Churchill, J. Monroe, D. Bristow, and C. Crawford. Using electric vehicles to enhance power outage resilience–an agent-based modeling approach. *Heliyon*, 10(11), 2024.

[37] D. J. Cook, A. S. Crandall, B. L. Thomas, and N. C. Krishnan. Casas: A smart home in a box. *Computer*, 46(7):62–69, 2012.

[38] İ. Coşkun, S. Sertok, and B. Anbaroğlu. K-nearest neighbour query performance analyses on a large scale taxi dataset: Postgresql vs. mongodb. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 42:1531–1538, 2019.

[39] S. L. Cutter, B. J. Boruff, and W. L. Shirley. Social vulnerability to environmental hazards. *Social Science Quarterly*, 84(2):242–261, 2003.

[40] Cybersecurity and Infrastructure Security Agency (CISA). Resilient power best practices for critical facilities and sites, 2022.

[41] J. Dagdelen, A. Dunn, S. Lee, N. Walker, A. S. Rosen, G. Ceder, K. A. Persson, and A. Jain. Structured information extraction from scientific text with large language models. *Nature communications*, 15(1):1418, 2024.

[42] G. B. Dantzig and J. H. Ramser. The truck dispatching problem. *Management science*, 6(1):80–91, 1959.

[43] E. Davies. Emergency hospital evacuation as Hurricane Sandy hits new york. *BMJ*, 345:e7357, 2012.

[44] F. de Silva and R. Eglese. Integrating simulation modelling and GIS: Spatial decision support systems for evacuation planning. *Journal of the Operational Research Society*, 51(4):423–430, 2000.

[45] K. Deb, S. Agrawal, A. Pratap, and T. Meyarivan. A fast elitist non-dominated sorting genetic algorithm for multi-objective optimization: Nsga-ii. In *International conference on parallel problem solving from nature*, pages 849–858. Springer, 2000.

[46] P. L. Delamater. Spatial accessibility in suboptimally configured health care systems: A modified two-step floating catchment area (M2SFCA) model. *Health & Place*, 24:30–43, 2013.

[47] D. Dosa, Z. Feng, K. Hyer, L. M. Brown, K. Thomas, and V. Mor. Effects of hurricane katrina on nursing facility resident mortality, hospitalization, and functional decline. *Disaster medicine and public health preparedness*, 4(S1):S28–S32, 2010.

[48] D. Dosa, K. Hyer, K. Thomas, S. Swaminathan, Z. Feng, L. Brown, and V. Mor. To evacuate or shelter in place: implications of universal hurricane evacuation policies on nursing home residents. *Journal of the American Medical Directors Association*, 13(2):190–e1, 2012.

[49] X. Du, Y. Chen, A. Bouferguene, and M. Al-Hussein. An agent-based simulation framework for analysing fall risk among older adults in the evacuation procedures. *Safety science*, 129:104790, 2020.

[50] A. Elkady, J. Hernantes, and L. Labaka. Decision support systems for community resilience and disaster management: A systematic review. *Heliyon*, 10(12):e33116, 2024.

[51] E. J. Emanuel, G. Persad, R. Upshur, B. Thome, M. Parker, A. Glickman, C. Zhang, C. Boyle, M. Smith, and J. P. Phillips. Fair allocation of scarce medical resources in the time of covid-19, 2020.

[52] M. Englert, D. Ozmen, and M. Westermann. The power of reordering for online minimum makespan scheduling. *SIAM Journal on Computing*, 43(3):1220–1237, 2014.

[53] A. T. Ernst, H. Jiang, M. Krishnamoorthy, B. Owens, and D. Sier. Staff scheduling and rostering: A review of applications, methods and models. *European Journal of Operational Research*, 153(1):3–27, 2004.

[54] Esri. What is the ArcGIS Network Analyst extension? Accessed Jan 2026.

[55] R. Fagin and J. H. Williams. A fair carpool scheduling algorithm. *IBM Journal of Research and development*, 27(2), 1983.

[56] C. Fan, C. Zhang, A. Yahja, and A. Mostafavi. Disaster city digital twin: A vision for integrating artificial and human intelligence for disaster management. *International Journal of Information Management*, 56:102049, 2021.

[57] M. Faraji and J. Kiyono. Seismic performance analysis of lifeline systems. *WIT Transactions on The Built Environment*, 119:35–44, 2011.

[58] Federal Emergency Management Agency. National Risk Index. `https://hazards.fema.gov/nri/`, 2026. Interactive map. Accessed: 2026-04-06.

[59] Federal Emergency Management Agency. National Risk Index for Natural Hazards. `https://www.fema.gov/flood-maps/products-tools/national-risk-index`, 2026. Accessed: 2026-04-06.

[60] FEMA. Comprehensive preparedness guide (CPG) 201, third edition: THIRA and SPR guide. `https://bit.ly/4qN5HRh`, 2018.

[61] FEMA. Hazus Earthquake Model Technical Manual. Hazus 6.1, 2024.

[62] G. Ferraro, H.-P. Lam, S. C. Tosatto, F. Olivieri, M. B. Islam, N. van Beest, and G. Governatori. Automatic extraction of legal norms: Evaluation of natural language processing tools. In *JSAI International Symposium on Artificial Intelligence*, pages 64–81. Springer, 2019.

[63] E. H. Field, R. J. Arrowsmith, G. P. Biasi, P. Bird, T. E. Dawson, K. R. Felzer, D. D. Jackson, K. M. Johnson, T. H. Jordan, C. Madden, et al. Uniform california earthquake rupture forecast, version 3 (ucerf3)—the time-independent model. *Bulletin of the Seismological Society of America*, 104(3):1122–1180, 2014.

[64] C. Fikar and P. Hirsch. Home health care routing and scheduling: A review. *Computers & Operations Research*, 77:86–95, 2017.

[65] B. E. Flanagan, E. W. Gregory, E. J. Hallisey, J. L. Heitgerd, and B. Lewis. A social vulnerability index for disaster management. *Journal of Homeland Security and Emergency Management*, 8(1), 2011.

[66] B. E. Flanagan, E. W. Gregory, E. J. Hallisey, J. L. Heitgerd, and B. Lewis. A social vulnerability index for disaster management. *JHSEM*, 8(1), 2011.

[67] S. Fomundam and J. W. Herrmann. A survey of queuing theory applications in healthcare. Technical report, Institute for Systems Research, University of Maryland, 2007.

[68] C. for Disease Control and Prevention. National Risk Index, n.d. Accessed Jun 2024.

[69] A. Fouillet, G. Rey, F. Laurent, G. Pavillon, S. Bellec, C. Guihenneuc-Jouyaux, J. Clavel, E. Jougla, and D. Hémon. Excess Mortality Related to the August 2003 Heat Wave in France. *International Archives of Occupational and Environmental Health*, 80(1):16–24, 2006.

[70] J. Fox, V. Patkar, I. Chronakis, and R. Begent. From practice guidelines to clinical decision support: closing the loop. *Journal of the Royal Society of Medicine*, 102(11):464–473, 2009.

[71] D. Gale and L. S. Shapley. College admissions and the stability of marriage. *The American mathematical monthly*, 69(1):9–15, 1962.

[72] Y. Gao, Y. Xiong, X. Gao, K. Jia, J. Pan, Y. Bi, Y. Dai, J. Sun, Q. Guo, and H. Wang. Retrieval-augmented generation for large language models: A survey. *arXiv preprint arXiv:2312.10997*, 2024.

[73] R. Gayathri. An enhanced resource allocation algorithm to allocate hospital beds during the covid19 pandemic. In *ICIRCA*. IEEE, 2021.

[74] S. Ghaffarian. Rethinking digital twin: Introducing digital risk twin for disaster risk management. *npj Natural Hazards*, 2(1):79, 2025.

[75] A. Ghodsi, M. Zaharia, B. Hindman, A. Konwinski, S. Shenker, and I. Stoica. Dominant resource fairness: Fair allocation of multiple resource types. In *Proceedings of the 8th USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, 2011.

[76] C. A. Glass and R. A. Knight. A critical analysis of nurse rostering research. *Journal of the Operational Research Society*, 61(4):600–605, 2010.

[77] Google. Google distance matrix api, 2005. Accessed Feb 2024.

[78] Gurobi Optimization, LLC. Gurobi Optimizer Reference Manual, 2024.

[79] A. Guttman. R-trees: A dynamic index structure for spatial searching. In *Proceedings of the 1984 ACM SIGMOD International Conference on Management of Data*, pages 47–57, 1984.

[80] Ş. Güven-Koçak, A. Heching, P. Keskinocak, and A. Toriello. Continuity of care in home health care scheduling: a rolling horizon approach. *Journal of Scheduling*, 27(4):375–392, 2024.

[81] M. S. Hagen, J. K. Jopling, T. G. Buchman, and E. K. Lee. Priority queuing models for hospital intensive care units and impacts to severe case patients. In *AMIA annual symposium proceedings*, volume 2013, page 841, 2013.

[82] J. L. Haggerty, R. J. Reid, G. K. Freeman, B. H. Starfield, C. E. Adair, and R. McKendry. Continuity of care: a multidisciplinary review. *Bmj*, 327(7425):1219–1221, 2003.

[83] HCAI. Dep. of Health Care Access and Information, n.d. Accessed Feb 2024.

[84] HHS ASPR. Risk identification and site criticality (RISC) toolkit 2.0. `https://aspr.hhs.gov/RISC/Pages/default.aspx`, 2024.

[85] J. G. Hodge Jr. A legal duty to evacuate patients from healthcare facilities in emergencies. *Health Law.*, 25:20, 2012.

[86] E. Horowitz and S. Sahni. Exact and approximate algorithms for scheduling nonidentical processors. *JACM*, 23(2), 1976.

[87] C. L. Hua, S. Patel, K. S. Thomas, D. J. Jester, C. M. Kosar, L. J. Peterson, D. Dobbs, R. Andel, and D. M. Dosa. Evacuation and health care outcomes among assisted living residents after hurricane irma. *JAMA Network Open*, 7(4):e248572, Apr. 2024.

[88] Institute of Medicine. *Crisis Standards of Care: A Systems Framework for Catastrophic Disaster Response: Volume 1: Introduction and CSC Framework*. The National Academies Press, Washington, DC, 2012.

[89] R. Jain, D.-M. Chiu, and W. R. Hawe. A quantitative measure of fairness and discrimination for resource allocation in shared computer systems. Technical Report DEC-TR-301, Digital Equipment Corporation, 1984.

[90] M. Ji and T. E. Cheng. Parallel-machine scheduling of simple linear deteriorating jobs. *Theoretical Computer Science*, 410(38-40):3761–3768, 2009.

[91] C. Joe-Wong, S. Sen, T. Lan, and M. Chiang. Multi-resource allocation: Fairness-efficiency tradeoffs in a unifying framework. In *Proceedings of IEEE INFOCOM*, 2012.

[92] T. M. John and J. Millum. First come, first served? *Ethics*, 130(2):179–207, 2020.

[93] Kaiser Permanente. Hazard vulnerability analysis (HVA) tool. `https://calhospital.org/updated-hva-tool-kaiser-permanente/`, 2017.

[94] G. Karni, G. N. Rothblum, and G. Yona. On fairness and stability in two-sided matchings. In *Proceedings of the 13th Innovations in Theoretical Computer Science Conference (ITCS)*, Leibniz International Proceedings in Informatics (LIPIcs), 2022.

[95] M. M. Kenne, P. Date, R. T. Eguchi, Z. Hu, J. Rousseau, and N. Venkatasubramanian. iFair: Achieving fairness in the allocation of scarce resources for senior health care. In *2024 IEEE International Conference on Smart Computing (SMARTCOMP)*, pages 22–30. IEEE, 2024.

[96] M. M. Kenne, F. Ventorim, C. Cossey, Z. Hu, J. Rousseau, and N. Venkatasubramanian. RADAR: Resource allocation for disaster resilience in senior health care. In *Proceedings of the 33rd ACM International Conference on Advances in Geographic Information Systems*, pages 959–972, 2025.

[97] S. M. Khalili, M. Mojtahedi, C. Steinmetz-Weiss, and D. Sanderson. A systematic literature review on transit-based evacuation planning in emergency logistics management: Optimisation and modelling approaches. *Buildings*, 14(1):176, 2024.

[98] D. Khorasanian, J. Patrick, and A. Sauré. Dynamic home care routing and scheduling with uncertain number of visits per referral. *Transportation Science*, 58(4):841–859, 2024.

[99] V. A. Knight, P. R. Harper, and L. Smith. Ambulance allocation for maximal survival with heterogeneous outcome measures. *Omega*, 40(6):918–926, 2012.

[100] T. Krityakierne, O. Limphattharachai, and W. Laesanklang. Nurse-patient relationship for multi-period home health care routing and scheduling problem. *Plos one*, 17(5):e0268517, 2022.

[101] M.-P. Kwan and J.-S. Lee. Emergency response after 9/11: The potential of real-time 3d GIS for quick response in micro-spatial environments. *Computers, Environment and Urban Systems*, 29(2):93–113, 2005.

[102] U. Lagap and S. Ghaffarian. Digital post-disaster risk management twinning: A review and improved conceptual framework. *International Journal of Disaster Risk Reduction*, 110:104629, 2024.

[103] T. Lan, D.-M. Chiu, M. Chiang, and A. Sabharwal. An axiomatic theory of fairness in network resource allocation. *arXiv*, 2009.

[104] J. Le Boudec. Rate adaptation, congestion control and fairness: A tutorial. *on line*, 2008.

[105] J. B. Leite, J. R. S. Mantovani, T. Dokic, Q. Yan, P.-C. Chen, and M. Kezunovic. Resiliency assessment in distribution networks using GIS-based predictive risk analytics. *IEEE Transactions on Power Systems*, 34(6):4249–4257, 2019.

[106] P. Lewis, E. Perez, A. Piktus, F. Petroni, V. Karpukhin, N. Goyal, H. Küttler, M. Lewis, W. Yih, T. Rocktäschel, S. Riedel, and D. Kiela. Retrieval-augmented generation for knowledge-intensive NLP tasks. In *Advances in Neural Information Processing Systems*, volume 33, pages 9459–9474, 2020.

[107] H. Li and J. T. Y. Lo. Microscopic simulation of the elderly evacuation in nursing homes. In *EPJ Web of Conferences*, volume 334, page 04004. EDP Sciences, 2025.

[108] B. Liang, D. Yang, X. Qin, and T. Tinta. A risk-averse shelter location and evacuation routing assignment problem in an uncertain environment. *International Journal of Environmental Research and Public Health*, 16(20):4007, 2019.

[109] H. Lim, H. Lee, and J.-H. Hwang. Multi-agent simulation on staff evacuation behavior in elderly nursing home fire emergencies. *Buildings*, 13(2):400, 2023.

[110] Y. Lin, D. Jiang, R. Yus, G. Bouloukakis, A. Chio, S. Mehrotra, and N. Venkatasubramanian. Locater: cleaning wifi connectivity datasets for semantic localization. *arXiv preprint arXiv:2004.09676*, 2020.

[111] S. Liu, A. Vlachokostas, and E. Kontou. Leveraging electric vehicles as a resiliency solution for residential backup power during outages. *Energy*, 318:134613, 2025.

[112] Y. Liu, Y. Yuan, J. Shen, and W. Gao. Emergency response facility location in transportation networks: A literature review. *Journal of traffic and transportation engineering (English edition)*, 8(2):153–169, 2021.

[113] Z. Liu, Z. Wang, L. Chen, T. Tang, and Z. Shen. An agent-based approach to simulating evacuation in elderly care facilities with disabled older adults. *International Journal of Disaster Risk Reduction*, 113:104815, 2024.

[114] A. Loni and S. Asadi. Power system resilience: The role of electric vehicles and social disparities in mitigating the us power outages. *Smart Grids and Sustainable Energy*, 9(1):23, 2024.

[115] L. Luo, J. Li, X. Xu, W. Shen, and L. Xiao. A data-driven hybrid three-stage framework for hospital bed allocation: A case study in a large tertiary hospital in china. *Computational and mathematical methods in medicine*, 2019(1):7370231, 2019.

[116] A. Maranghides, E. Link, C. U. Brown, W. D. Walton, W. Mell, and S. Hawks. A Case Study of the Camp Fire – Notification, Evacuation, Traffic, and Temporary Refuge Areas (NETTRA). NIST Technical Note 2252, National Institute of Standards and Technology, Gaithersburg, MD, 2023.

[117] S. Mehrotra, A. Kobsa, N. Venkatasubramanian, and S. R. Rajagopalan. Tippers: A privacy cognizant iot environment. In *2016 IEEE International Conference on Pervasive Computing and Communication Workshops (PerCom Workshops)*, pages 1–6. IEEE, 2016.

[118] J. Mo and J. Walrand. Fair end-to-end window-based congestion control. *IEEE/ACM Transactions on Networking*, 8(5):556–567, 2000.

[119] A. T. Murray, J. Baik, V. E. Figueroa, D. Rini, M. A. Moritz, D. A. Roberts, S. H. Sweeney, L. M. Carvalho, and C. Jones. Developing effective wildfire risk mitigation plans for the wildland urban interface. *International Journal of Applied Earth Observation and Geoinformation*, 124:103531, 2023.

[120] National Center for Health Statistics. FastStats: Nursing Home Care. `https://www.cdc.gov/nchs/fastats/nursing-home-care.htm`, 2026. Centers for Disease Control and Prevention. Last reviewed February 6, 2026. Accessed: 2026-04-06.

[121] National Center for Health Statistics. FastStats: Residential Care Communities. `https://www.cdc.gov/nchs/fastats/residential-care-communities.htm`, 2026. Centers for Disease Control and Prevention. Last reviewed February 6, 2026. Accessed: 2026-04-06.

[122] NFPA. *Health Care Facilities Code*. NFPA 99, 2021.

[123] NIA. National Institute on Aging, n.d. Accessed Feb 2024.

[124] NOAA National Centers for Environmental Information. U.S. Billion-Dollar Weather and Climate Disasters. `https://www.ncei.noaa.gov/access/billions/state-summary/US`, 2025. Accessed: 2026-04-06.

[125] NOAA National Centers for Environmental Information. U.S. Billion-Dollar Weather and Climate Disasters. `https://www.ncei.noaa.gov/access/billions/`, 2025. Accessed: 2026-04-06.

[126] OCHCA. Orange county health care agency, n.d. Accessed Jan 2024.

[127] S. Pang, C. Liu, C. Zhang, J. Zhang, and H. Zhang. Leveraging electric vehicles for enhancing power system resilience: A review of strategies and challenges. *Current Sustainable/Renewable Energy Reports*, 12(1):11, 2025.

[128] G. Persad, A. Wertheimer, and E. J. Emanuel. Principles for allocation of scarce medical interventions. *The Lancet*, 373(9661):423–431, 2009.

[129] M. D. Petersen, A. M. Shumway, P. M. Powers, C. S. Mueller, M. P. Moschetti, A. D. Frankel, S. Rezaeian, D. E. McNamara, N. Luco, O. S. Boyd, et al. The 2018 update of the US national seismic hazard model: Overview of model and implications. *Earthquake Spectra*, 36(1):5–41, 2020.

[130] G. Pettet, A. Mukhopadhyay, M. J. Kochenderfer, and A. Dubey. Hierarchical planning for dynamic resource allocation in smart and connected communities. *ACM Transactions on Cyber-Physical Systems*, 6(4):1–26, 2022.

[131] M. L. Pinedo. Design and implementation of scheduling systems: More advanced concepts. *Scheduling: Theory, Algorithms, and Systems*, pages 485–508, 2016.

[132] J. Rawls. A theory of justice harvard university press, cambridge, mass. 1971.

[133] ReddiNet. Real-time bed availability data, n.d. Accessed May 2023.

[134] S. M. Rinaldi, J. P. Peerenboom, and T. K. Kelly. Identifying, understanding, and analyzing critical infrastructure interdependencies. *IEEE Control Systems Magazine*, 21(6):11–25, 2001.

[135] S. K. Sahni. Algorithms for scheduling independent tasks. *JACM*, 23(1), 1976.

[136] I. Sattar, M. Shahid, and N. Yasir. Multi-level queue with priority and time sharing for real time scheduling. *International journal of multidisciplinary sciences and engineering*, 5(8):16–17, 2014.

[137] H. Schütze, C. D. Manning, and P. Raghavan. *Introduction to information retrieval*, volume 39. Cambridge University Press Cambridge, 2008.

[138] C. Shahabi and J. P. Wilson. Scalable evacuation routing in a dynamic environment. *Computers, Environment and Urban Systems*, 67:29–40, 2018.

[139] J. F. Shortle, J. M. Thompson, D. Gross, and C. M. Harris. *Fundamentals of Queueing Theory*. John Wiley & Sons, 5 edition, 2018.

[140] G. Snoeren, S. Zlatanova, J. Crompvoets, H. Scholten, and C. Geodan. Spatial data infrastructure for emergency management: the view of the users. In *Proceedings of the 3rd International symposium on Gi4DM*, pages 22–25, 2007.

[141] Y. Song and H. J. Miller. Exploring traffic flow databases using space-time plots and data cubes. *Transportation*, 39:215–234, 2012.

[142] T. Spurlock, K. Sewell, M. M. Sugg, J. D. Runkle, R. Mercado, J. S. Tyson, and J. Russell. A spatial analysis of power-dependent medical equipment and extreme weather risk in the southeastern united states. *International Journal of Disaster Risk Reduction*, 95:103844, 2023.

[143] C. Stone. Machine learning for prediction of hospital length of stay: A systematic review. *PLOS Digital Health*, 2022.

[144] E. Tate, M. A. Rahman, C. T. Emrich, and C. C. Sampson. Flood exposure and social vulnerability in the united states. *Natural Hazards*, 106(1):435–457, 2021.

[145] S. Taylor, A. Rangarajan, N. Rhodes, J. Snodgrass, B. C. Lesieutre, and L. A. Roald. California Test System (CATS): A geographically accurate test system based on the california grid. *IEEE Transactions on Energy Markets, Policy and Regulation*, 2(1):107–118, 2023.

[146] The Associated Press. Power outage at los angeles hospital caused by hilary prompts evacuation of hundreds of patients. *AP News*, Aug 2023.

[147] C. N. T. Tran, M. Heller, A. Berger, and J. Habboushe. Hurricane Sandy: How did we do? assessing a Manhattan hospital's response to the storm. *Frontiers in Public Health*, 2:90, 2014.

[148] U.S. Census Bureau. Older Adults Outnumber Children in 11 States and Nearly Half of U.S. Counties. `https://www.census.gov/newsroom/press-releases/2025/older-adults-outnumber-children.html`, June 2025. Press Release CB25-99. Accessed: 2026-04-06.

[149] U.S. Department of Health and Human Services, ASPR TRACIE. Healthcare facility evacuation/sheltering topic collection, 2026.

[150] USGS. United states geological survey, n.d. Accessed May 2024.

[151] H. R. Varian. Equity, envy, and efficiency. 1973.

[152] N. Venkatasubramanian, C. A. Davis, and R. T. Eguchi. Designing community-based intelligent systems for water infrastructure resilience. In *Proceedings of the 3rd ACM SIGSPATIAL International Workshop on Advances in Resilient and Intelligent Cities*, pages 62–65, 2020.

[153] X. Wang, T.-M. Choi, H. Liu, and X. Yue. A novel hybrid ant colony optimization algorithm for emergency transportation problems during post-disaster scenarios. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 48(4):545–556, 2016.

[154] X. Wang, Y. Yan, L. Huang, X. Zheng, and X.-J. Huang. Hallucination detection for generative large language models by bayesian sequential estimation. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 15361–15371, 2023.

[155] J. L. Ware. Geospatial data fusion: Training gis for disaster relief operations. In *Proc., GIS User Conf*, volume 1097, 2001.

[156] J. M. Ware. Geospatial data fusion: Training GIS for disaster relief operations. In *Proceedings of the ESRI International User Conference*, 2003.

[157] Washington State Department of Health. Heat Wave 2021. `https://doh.wa.gov/emergencies/be-prepared-be-safe/severe-weather-and-natural-disasters/extreme-heat/hot-weather-precautions/heat-wave-2021`, 2021. Final report page with age, sex, and county breakdowns. Accessed: 2026-04-06.

[158] M. Willoughby, C. Kipsaina, N. Ferrah, S. Blau, L. Bugeja, D. Ranson, and J. E. Ibrahim. Mortality in nursing homes following emergency evacuation: a systematic review. *Journal of the American Medical Directors Association*, 18(8):664–670, 2017.

[159] Z. Xu, S. Jain, and M. Kankanhalli. Hallucination is inevitable: An innate limitation of large language models. *arXiv preprint arXiv:2401.11817*, 2024.

[160] Y. Yin, X. Xu, D. Wang, Y. Yu, and T. Cheng. Two-stage recoverable robust optimization for an integrated location–allocation and evacuation planning problem. *Transportation research part B: methodological*, 182:102906, 2024.

[161] M. H. Yousef, Y. N. Alhalaseh, R. Mansour, H. Sultan, N. Alnadi, A. Maswadeh, Y. M. Al-Sheble, R. Sinokrot, K. Ammar, A. Mansour, et al. The fair allocation of scarce medical resources: a comparative study from jordan. *Frontiers in medicine*, 7:603406, 2021.

[162] S. Zlatanova and D. Holweg. 3d geo-information in emergency response: A framework. In *Proceedings of the 4th International Symposium on Mobile Mapping Technology (MMT)*, 2004.