RADAR: Resource Allocation for Disaster Resilience in Senior Health Care

Modeste Mefenya Kenne
Dept. of Computer Science, UC Irvine
Irvine, California, USA
mkenne@uci.edu

Fernanda Ventorim
Dept. of Computer Science, UC Irvine
Irvine, California, USA
f.ventorim@uci.edu

Chad Cossey
Orange County Health Care Agency
Santa Ana, California, USA
ccossey@ochca.com

ZhengHui Hu
ImageCat, Inc
Long Beach, California, USA
zh@imagecatinc.com

Julie Rousseau
School of Medicine, UC Irvine
Irvine, California, USA
jroussea@hs.uci.edu

Nalini Venkatasubramanian
Dept. of Computer Science, UC Irvine
Irvine, California, USA
nalini@uci.edu

ABSTRACT

During disasters, ensuring that resources are efficiently allocated to appropriate locations is essential for minimizing adverse impacts and saving lives. To this end, we present RADAR, a data-driven platform that integrates multisource GIS feeds (USGS earthquake alerts, Cal Fire wildfire perimeters) with facility and transportation data to support proactive planning and real-time recommendations for Emergency Operations Centers. RADAR uses policy-driven stable matching to optimize routing and resource assignment for evacuation planning and resource delivery. The aggregate model allocates across short-term facilities (e.g., hospitals), and a fine-grained extension for long-term senior-care facilities personalizes allocation using resident preferences, medical profiles, and social constraints. RADAR adapts as conditions evolve by utilizing historical data, live traffic, and changing facility status. We validated RADAR's efficacy in many disaster settings, including real events such as the Palisades wildfire and tabletop drills (earthquake and water-contamination scenarios) involving first responders.

CCS CONCEPTS

• **Information systems → Decision support systems**; • **Applied computing → Multi-criterion optimization and decision-making**.

KEYWORDS

emergency management, disaster resilience, resource allocation, senior health care, geographic information system

ACM Reference Format:
Modeste Mefenya Kenne, Fernanda Ventorim, Chad Cossey, ZhengHui Hu, Julie Rousseau, and Nalini Venkatasubramanian. 2025. RADAR: Resource Allocation for Disaster Resilience in Senior Health Care. In *The 33rd ACM International Conference on Advances in Geographic Information Systems (SIGSPATIAL '25), November 3–6, 2025, Minneapolis, MN, USA.* ACM, New York, NY, USA, 14 pages. https://doi.org/10.1145/3748636.3764164

SIGSPATIAL '25, November 3–6, 2025, Minneapolis, MN, USA

ACM ISBN 979-8-4007-2086-4/2025/11
https://doi.org/10.1145/3748636.3764164

1 INTRODUCTION

Natural disasters such as wildfires, floods, and earthquakes are becoming increasingly frequent and severe, significantly impacting communities worldwide. These disasters disrupt critical infrastructure, overwhelm healthcare/emergency systems, and incur substantial economic and social costs. Factors such as climate change, population growth, and aging infrastructure exacerbate these challenges and complicating disaster response efforts [21, 29]. Efficient disaster response involves multiple interconnected components: disaster preparedness ahead of the event, rapid resource allocation and communication coordination during the event, and post-disaster recovery planning. Timely and effective decisions must integrate diverse multimodal data (GIS imagery, weather data, human reports, etc.) to create real-time situational awareness [49]. This is used by decision-makers to rapidly assess infrastructure conditions (e.g., transport), resource availability, and environmental hazards to prioritize actions and allocate scarce resources optimally. Clearly, inefficient allocations can delay relief efforts and exacerbates socioeconomic disparities, often leaving vulnerable populations (e.g. older adults) without adequate support and increasing long-term recovery costs [62]. The effectiveness of the response greatly depends on inter-agency coordination and communication between stakeholders such as Emergency Operations Centers (EOCs), healthcare providers, government agencies, and the affected populations [35].

Older adults represent a notably vulnerable population due to mobility limitations, chronic health conditions, cognitive challenges, and specialized medical requirements – all leading to a reliance on continuous care. During disasters, timely allocation of healthcare resources to this population is critical to mitigate the impact of the disaster [29]. By 2034, the number of older adults in the U.S. is projected to outnumber children, significantly increasing healthcare demands and creating severe staffing shortages in senior care facilities [39]. The complexity of caring for older adults during disasters also extends beyond immediate health impacts; the long-term impacts encompass psychological, social, and emotional effects [41, 42]. Failure to address these multifaceted needs adequately can lead to deterioration in chronic health conditions and ultimately increase mortality rates. Disaster preparedness and response in senior care facilities must involves customized strategies that consider individual medical profiles, personal preferences, and continuity of care, as traditional one-size-fits-all planning models are ineffective for this population.

From a geographic information systems (GIS) perspective, the integration and analysis of spatio-temporal data plays a crucial role in disaster response. Accurate real-time GIS data can improve disaster prediction, streamline evacuation routes, and enhance the coordination of emergency resources by identifying and visualizing hazard zones, traffic disruptions, and demand hotspots dynamically [14, 48, 49]. In addition, storing facility and hazard layers in a spatially enabled PostgreSQL/PostGIS database improves high-performance geospatial indexing (R-tree indices) and sub-second nearest-neighbor queries [14]. By unifying vector feeds (roads, facility footprints) with live traffic streams and dynamic facility status updates, we can enable ad hoc spatial joins and topology-aware route planning that avoid impacted segments in real time.

Building on these spatial foundations, techniques such as AI-driven predictive models, multi-objective optimization, and risk-aware graph algorithms can further enhance disaster resilience. Some approaches use risk-aware pathfinding methods that dynamically update edge weights in response to evolving hazard and traffic conditions [2, 50, 55, 64]. This allows both evacuees and emergency vehicles to be routed via the safest and fastest paths as conditions evolve. Likewise, emerging technologies such as the Internet of Things (IoT), sensor networks, and pervasive computing systems can significantly enhance real-time monitoring and resource management during disasters. IoT-enabled devices can continuously collect and transmit critical data, providing detailed, granular insights into rapidly changing conditions. Thus, there is a need for methods that support situational awareness via the integration of heterogeneous data streams from environmental sensors and smart infrastructure [34]. In addition, incorporating fairness and ethics considerations into the allocation process is also important, as they guide equitable and transparent decision-making, especially under conditions of resource scarcity and intense pressure, ensuring that people receive appropriate care during emergencies [6, 30, 33, 51].

Recent efforts have highlighted the gaps in current emergency planning and response practices. Through the CareDEX disaster resilience effort [38], we conducted wildfire and earthquake drills, as well as tabletop simulations with partner senior care facilities to evaluate evacuation strategies. These exercises revealed that in the aftermath of major disasters, hospitals and senior care facilities often face massive surges in demand for beds, supplies, and staffing, leading to adverse outcomes among vulnerable populations. These findings evidence the need for a unified, spatially informed allocation framework that can anticipate surges, pre-stage critical resources, and coordinate evacuations in real-time.

To ensure seamless integration into operational workflows, we co-designed our platform with key stakeholders, including the Orange County Health Care Agency (OCHCA), ImageCat (hazard analytics experts), and senior healthcare partners. These collaborations ensured that our platform aligns with existing disaster resilience protocols. Based on these insights, we propose RADAR, a data-driven platform designed to enhance disaster response by providing actionable recommendations to EOCs. RADAR provides aggregate decision support, suitable for short-term facilities (e.g., hospitals) and fine-grained allocation for personalized, resident-level assignments in long-term care facilities, such as nursing homes and assisted living facilities. Our key contributions include:

- A GIS-centric platform for disaster planning and mitigation, coupling real-time hazard streams with organizational data and policy-driven stable-matching optimization.
- Customizable resource allocation models: an aggregate model for rapid decision-making, and a fine-grained model for personalized resident-level assignments.
- Integration of multi-source vector feeds, historical data, real-time traffic, and facility status updates in a Spatial Decision Support System (SDSS) for "what-if" scenario analysis.
- Comprehensive emulation-based scenario testing and real-world drills demonstrating improvements in assignment and routing efficiency while ensuring fairness.

Our paper is organized as follows. In §2, we discuss related work and limitations. Our approach is discussed in §3 and the problem is formulated in §4. The decision-making policies, including our model, are described in §5. Then, §6 and §7 present our experimental results, conclusions, and future work.

2 RELATED WORK AND LIMITATIONS

Resource allocation has been extensively studied in various domains such as education [40], economics [31], and transportation [14], where the primary goal is to optimize efficiency by minimizing delays and maximizing resource utilization. Resource scheduling in computing environments, such as CPU and memory allocation in operating systems and cloud platforms, relies on well-defined, deterministic constraints, enabling high-throughput, low-latency policies to be evaluated in closed-loop simulations [19, 28]. By contrast, resource allocation in senior health care with human-in-the-loop (HITL) introduces profound uncertainties: individual compliance, variable mobility impairments, cognitive load, and the need for continuity of medical care. Such HITL scenarios defy the predictability of traditional resource schedulers and require specialized protocols to ensure safety and dignity.

While traditional domains emphasize predictability and efficiency, disaster-oriented health care resource allocation requires fundamentally different approaches. Early efforts utilize mixed-integer programming for multi-period scheduling under capacity constraints [58] and heuristic or genetic algorithms for multi-objective routing and facility assignment [50, 55]. However, systematic reviews call for stronger, data-driven frameworks to overcome inconsistent evidence quality [52]. Recently, large language models have been explored to support decision-making, but their reliability and autonomy in high-stakes disaster contexts remains challenging, e.g., due to hallucinations [57, 61]. Likewise, equity-driven approaches use multi-criteria decision analysis such as Analytic Hierarchy Process (AHP) and Technique for Order of Preference by Similarity to Ideal Solution (TOPSIS) to rank candidate sites by accessibility, capacity, and hazard exposure [1, 30]. However, they struggle to ingest real-time data streams, and static vulnerability indices often fail to reflect rapidly evolving disaster conditions [21].

Geospatial technologies have also become essential to operationalizing decision-making strategies in disasters. These methods offer tools to map, analyze, and act on complex spatial relationships in real time. This forms the backbone of spatial decision support by enabling integrated, multilayered analyses. Here, centralized PostGIS databases deliver high-performance geospatial indexing,

ad-hoc spatial joins, and topology-aware routing that dynamically avoids impacted road segments [14]. However, surveys of GIS in disaster response highlight various challenges in real-time data fusion, and some work points to AI-driven predictive modeling for hazard progression and traffic disruptions [59]. In addition, probabilistic flood-risk models using multi-criteria GIS frameworks show urban vulnerability mapping, while pre-disaster evacuation network design under equilibrium flow conditions informs robust route configuration before congestion occurs [1, 32]. A recent approach has also enabled predictive risk scoring of nursing homes, showing how spatial layers of infrastructure, population density, and hazard exposure can guide pre-positioning of medical resources [46].

To utilize spatial insights in dynamic settings, Spatial Decision Support Systems (SDSS) extend GIS capabilities by uniting geospatial analytics, IoT sensor streams, and simulation for "what-if" scenario exploration [16]. For example, IoT-enabled evacuation platforms dynamically match evacuees to routes based on live sensor and traffic data, while macro-scale traffic flow simulators and agent-based models capture urban evacuation dynamics and potential bottlenecks [3, 16]. Agent-based models for fall risk among older adults during evacuation procedures further highlight the importance of individual-level mobility constraints and facility layouts [20].

Despite these technological advances, the unique needs of healthcare evacuations require specialized consideration [60]. Healthcare-specific evacuation studies emphasize the unique vulnerabilities of frail populations. Retrospective analyses of nursing home evacuations during Hurricanes Gustav and Rita report elevated mortality among cognitively impaired residents and failures in transport provisions such as the use of non-air-conditioned vehicles, which aggravated heat-related risks [45]. Likewise, post-Katrina reviews and qualitative surveys highlight gaps in emergency plans for older adults living independently and in senior care facilities, documenting inadequate preparedness, medication continuity challenges, and ethical dilemmas around prioritization [27].

To bridge these gaps, the RADAR framework builds upon the insights and limitations from prior work. We address some of these challenges by providing a GIS-centric SDSS that couples real-time hazard monitoring, dynamic facility status feeds, and population mobility tracking with policy-driven stable matching and risk-aware, multi-objective routing. By embedding human-in-the-loop constraints, individual health profiles, mobility limitations, and care continuity within our platform, we provide an efficient and fair allocation of resources to older adults in evolving disaster scenarios.

3 THE RADAR APPROACH

In response to these limitations, 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 the EOC. RADAR ties together real-time data integration, dynamic matching, and adaptive routing within a unified, modular architecture. By coupling spatio-temporal GIS streams with predictive hazard forecasts, our system can proactively recompute allocations as conditions evolve, while incremental matching updates preserve allocation stability and bound computational overhead. We embed human oversight at the decision points through

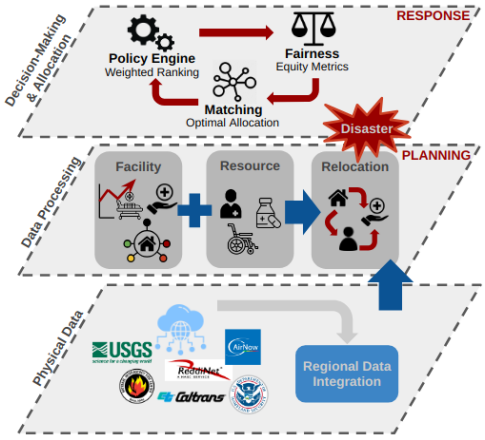

Figure 1: RADAR Architecture

an interactive dashboard that uncovers 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 components: an aggregate model that allocates resources using coarse policies and a fine-grained extension that incorporates resident-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. RADAR's architecture is designed via various modules (see Fig. 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 ideal allocations, aggregating these into concise equity metrics that EOC staff can monitor and balance against efficiency goals.

Thus, by combining dynamic data-driven matching and hazard-aware HITL routing, we can convert the static allocation paradigms into an adaptive and better decision-centric framework for disaster resilience. In the next section, we will discuss the mathematical formulation of this matching and routing problem.

4 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.

4.1 Modeling Key Components

Modeling Facilities. 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). They are classified by a type T_i (Senior Healthcare Facility, Hospital, and Non-Healthcare Facility) and have 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 as established by FEMA [21].

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 [4].

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 [12] 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, social vulnerability, etc.). These scores define the preference lists used in our stable matching procedures. We discuss them in depth in §5. 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 [24, 48] to capture real-world travel conditions. Hence, we integrate them into a unified approach that optimizes both matching and routing efficiency.

4.2 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.

$$\max_{x_{ij}} \quad \sum_{i \in S_{rel}} \sum_{j \in S_{rec}} w_p \, S_p(F_i, F_j) \, x_{ij} \tag{1}$$

$$\text{s.t.} \quad \sum_{j \in S_{rec}} x_{ij} \le 1, \quad \forall \, i \in S_{rel}, \tag{1a}$$

$$\sum_{i \in S_{rel}} x_{ij} \le 1, \quad \forall \, j \in S_{rec}, \tag{1b}$$

$$\sum_{i \in S_{rel}} R_i \, x_{ij} \le C_j, \quad \forall \, j \in S_{rec}, \tag{1c}$$

$$x_{ij} \le S_\Delta(F_i, F_j), \quad \forall \, i \in S_{rel}, \forall \, j \in S_{rec}, \tag{1d}$$

$$D(F_i, F_j) \le D^{\max}, \quad \forall \, i \in S_{rel}, \forall \, j \in S_{rec}, \tag{1e}$$

$$x_{ij} \in \{0, 1\}, \quad \forall \, i \in S_{rel}, \forall \, j \in S_{rec}. \tag{1f}$$

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 then expressed as shown in Eqn. 1.

Constraints (1a) and (1b) ensure a one-to-one matching between facilities relocating-people and receiving-people. Constraint (1c) guarantees that the allocated people do not exceed the resource capacity at the receiving facilities. Constraint (1d) limits assignments to facility pairs connected by safe routes, while constraint (1e) ensures that travel distances remain within a predefined threshold (detailed in Section 5). Together, this forms an integrated framework, simultaneously considering dynamic and static attributes of *facilities*, *people*, *resources*, *hazards*, and *tasks* in a disaster response.

5 DECISION-MAKING POLICIES IN THE RADAR MODEL

In this section, we present the set of policy-driven ranking rules used to guide evacuee relocation (§5.1, §5.2) and introduce our model (§5.3). In addition, we illustrate two extensions of our approach: the application of policy-driven rules for requesting and sending resources (§5.4.1) and a fine-grained extension that incorporates individual-level evacuee preferences (§A.1).

5.1 Policies for Receiving People

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

Triage status: Individuals are prioritized based on their triage category [4], ensuring that limited resources are allocated efficiently. Those in critical condition receive the highest priority to maximize survival rates:

$$S_{\text{tri}}(F_{\text{rel}}, P_i) = \begin{cases} 1, & \text{if } T_i = \text{Red (Immediate)}, \\ 0.75, & \text{if } T_i = \text{Yellow (Delayed)}, \\ 0.5, & \text{if } T_i = \text{Green (Minor)}, \\ 0, & \text{if } T_i = \text{Black (Deceased)}. \end{cases} \tag{2}$$

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

$$S_{\text{res}}(F_{\text{rel}}, F_{\text{rec}}) = \begin{cases} 1, & \text{if } T_{\text{rel}} = T_{\text{rec}}, \\ 0, & \text{otherwise.} \end{cases} \tag{3}$$

Proximity for faster relocation: To minimize delays, facilities that are closer to the receiving facility are prioritized. We quantify this as:

$$S_D(F_{rel}, F_{rec}) = \frac{\max\limits_{F_x \in N_{rel}} D(F_{rel}, F_x)}{D(F_{rel}, F_{rec})}, \tag{4}$$

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. To compute distances and travel times between multiple origins and destinations, we used Google's Distance Matrix API [24], 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 precise, real-world routing data into our decision-making process.

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

$$S_O(F_{rel}, F_{rec}) = \begin{cases} 1, & \text{if } O_{rel} = \text{Red or Black}, \\ 0.75, & \text{if } O_{rel} = \text{Orange}, \\ 0.5, & \text{if } O_{rel} = \text{Yellow}, \\ 0, & \text{if } O_{rel} = \text{Green} \end{cases} \tag{5}$$

5.2 Policies for Relocating People

When relocating evacuees, relocating facilities F_{rel} apply analogous criteria. They prefer receiving facilities of the same residency type (Eqn. 3). In addition, they prioritize facilities that are not approaching their full capacity constraints (e.g., maximum bed capacity). However, only facilities with operational status $O_{rel} = Green$ are eligible for relocation.

Proximity to impacted region: Facilities can receive evacuees if they are located at a safe distance from the impacted region (as shown in Fig. 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 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:

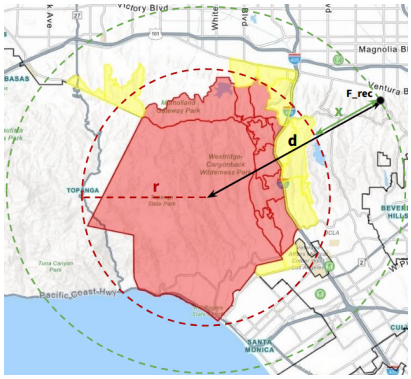

Figure 2: Impacted Region (red) and Mandatory Evacuation Warning Zone (yellow) (Palisades Fire 01/10/25, 7 pm PST)

$$S_{\text{imp}}(F_{\text{rec}}) = \begin{cases} 0, & \text{if } d \le r + x, \\ \frac{r+x}{d}, & \text{if } d > r + x, \end{cases} \tag{6}$$

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:

$$S_\Delta(F_{rec}, F_{rel}) = \begin{cases} 1, & \text{if safe routes exist}, \\ 0, & \text{otherwise.} \end{cases} \tag{7}$$

Social Vulnerability Index (SVI): Facilities in areas with high SVI [22] are also prioritized as shown in Eqn. 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. 9). The range is as follows:

$$SVI = \begin{cases} Very\_Low, & \text{if } 0 < SVI \le 0.25, \\ Low, & \text{if } 0.25 < SVI \le 0.5, \\ High, & \text{if } 0.5 < SVI \le 0.75, \\ Very\_High, & \text{if } 0.75 < SVI \le 1 \end{cases} \tag{8}$$

$$S_{SVI}(F_{rec}, F_{rel}) = \frac{SVI_{rel}}{\max\limits_{F_x \in N_{rec}} SVI_x}. \tag{9}$$

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.

5.3 RADAR Model

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 [23], 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.

Table 1: Relocating-Facilities (F^i_{rel})

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

Table 2: Receiving-Facilities (F^j_{rec})

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

Table 3: Distances $D(F^i_{rel}, F^j_{rec})$

	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

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 $M : A \cup B \rightarrow A \cup B$ 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:

$$\max \sum_{i=1}^{n} \sum_{j=1}^{m} P_{a_i}(b_j) X_{ij}, \tag{10}$$

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

Exploring Stable Matching in RADAR: Let $\mathcal{F}_{rel} = \{F^1_{rel}, F^2_{rel}, \ldots, F^n_{rel}\}$ be the set of relocating-facilities and $\mathcal{F}_{rec} = \{F^1_{rec}, F^2_{rec}, \ldots, F^m_{rec}\}$ the set of receiving-facilities. Each facility $F^i_{rel} \in \mathcal{F}_{rel}$ ranks $F^j_{rec} \in \mathcal{F}_{rec}$ based on their policies. The objective is to maximize preference scores of both F^i_{rel} and F^j_{rec}. 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 Tables 1 and 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. 1. For example, if we consider F^2_{rel}, which needs to relocate people to facilities with available oxygen resources, it prefers F^2_{rec} and F^5_{rec} since only these have oxygen. In addition Table 3 shows that F^2_{rec} is 8 miles away vs. F^5_{rec} which is 10 miles. Assuming there is a safe route to get to those two facilities, F^2_{rel} ranking is $\{F^2_{rec} \succ F^5_{rec} \succ \ldots\}$. 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. 2.

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

Algorithm 1: 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 2: 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;

restriction reduces proposals 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 (Algo. 3). 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 (demand(f)) and safeRoute(f, c, H) test to find an immediate reassignment. If none of those candidates passes (lines 13–15), we dynamically expand

Algorithm 3: Dynamic Rematching under 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';

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 (§5.4.1), and then introduce a fine-grained improvement of our model that integrates resident-specific preferences, including their medical information, into the allocation process (§A.1).

5.4 RADAR Extensions

5.4.1 ***Policies for Requesting and Sending Resources.*** Several of the spatio-temporal policies introduced for relocating people, i.e., proximity (Eqn. 4), safe-route viability (Eqn. 7), operational readiness (Eqn. 5), and social vulnerability (Eqn. 9) also apply in the resource requesting and sending context, since both problems rely on accurate 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) [9] to prevent over-exhaustion of any single source by favoring senders with more available inventory.

$$S_{\text{slack}}(F_{\text{send}}) = \frac{\text{available\_stock}(F_{\text{send}})}{\text{capacity}(F_{\text{send}})} \tag{11}$$

The available stock is the count of uncommitted units (beds, water bottles, etc.), and the capacity represents the maximum stock. Next, the **request urgency** captures the relative criticality of each demand. Drawing on priority-queueing theory [7, 8], we normalize a facility's current demand against the peak regional demand.

$$S_{\text{urg}}(F_{\text{req}}) = \frac{\text{demand}(F_{\text{req}})}{\max_{f' \in \mathcal{F}_{\text{req}}} \text{demand}(f')} \tag{12}$$

Here, demand(F_{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 [17]. We identify clusters of co-located requests $\mathcal{R}_{\text{cluster}}(F_{\text{req}})$ and compare their size to the total local request set $\mathcal{R}_{\text{all}}(F_{\text{req}})$.

$$S_{\text{grp}}(F_{\text{req}}) = \frac{|\mathcal{R}_{\text{cluster}}(F_{\text{req}})|}{|\mathcal{R}_{\text{all}}(F_{\text{req}})|} \tag{13}$$

A higher value indicates a stronger potential for multi-stop deliveries, reducing overall travel distance and vehicle utilization. Finally, the **commitment-ratio** [43] policy ensures senders retain flexibility for future surges by penalizing near-exhaustion of capacity.

$$S_{\text{commit}}(F_{\text{send}}, F_{\text{req}}) = 1 - \frac{\text{allocated}(F_{\text{send}})}{\text{capacity}(F_{\text{send}})} \tag{14}$$

In Eqn. 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. 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.

$$S(F_{req}, F_{send}) = \sum_p w_p \, S_p(F_{req}, F_{send}) \tag{15}$$

5.4.2 ***RADAR Fine-Grained: 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 Appendix A.1.

6 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 (6.1); (2) A 7.8-magnitude **earthquake** drill in Orange County, CA, showcasing hospital bed surge management and safe routing amid infrastructure damage (6.2); and (3) A countywide **water contamination** drill exploring relocation and water resupply for those sheltering-in-place to sustain operations (6.3).

Setup: We obtained facility data (including facility type, licensed bed capacity, location, etc) from California's Health and Human Services Agency [11] and HCAI [26] open datasets (~10409 facilities).

We stored them in a PostgreSQL/PostGIS database to enable efficient geospatial queries. We integrated OpenStreetMap road network data with real-time and historical traffic information with Google's Distance Matrix API [24] 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 [10], USGS ShakeMap for earthquake impact [53], and OCHCA and ReddiNet reports for water-related disruptions [37, 44]. 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.

6.1 Wildfire Use Case: Evacuation Based on the Palisades Fire

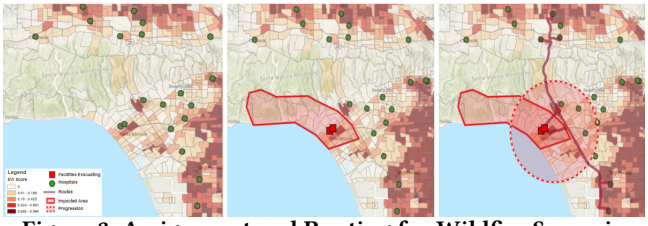

Figure 3: Assignment and Routing for Wildfire Scenario

We simulated evacuations based on real conditions during the Palisades fire (Fig. 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) [18] and (6) Ant Colony Optimization (ACO) [56].

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"[7]), 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). 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. 4a and 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. 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. 4b and 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. 4d and Fig. 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.

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

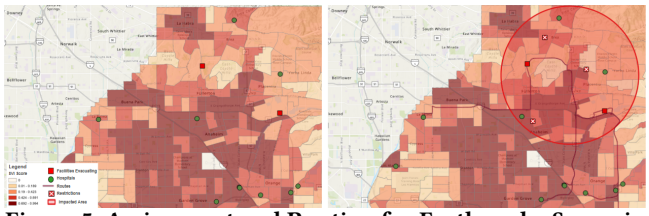

Figure 5: Assignment and Routing for Earthquake Scenario

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). 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. 11) presents a summary of the recommended allocations, including destination options and relevant operational metrics to support decision-making.

Fig. 6 compares estimated travel distances and durations across 20 origin-destination facility pairs, i.e. routes, at four times of day:

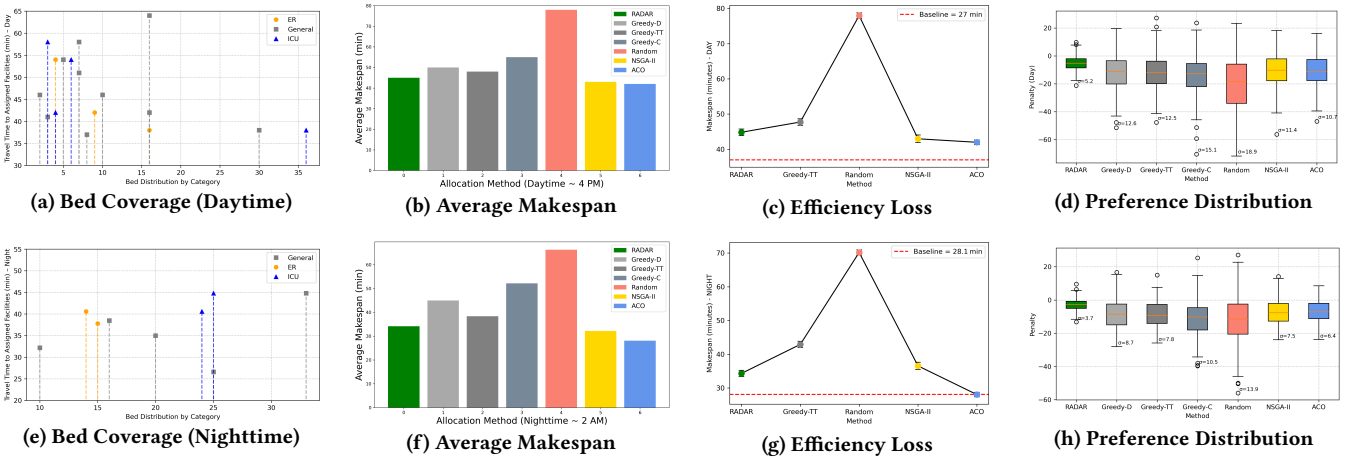

(a) Bed Coverage (Daytime) (b) Average Makespan (c) Efficiency Loss (d) Preference Distribution

(e) Bed Coverage (Nighttime) (f) Average Makespan (g) Efficiency Loss (h) Preference Distribution

Figure 4: Comparison of Daytime (Top row) vs. Nighttime (Bottom row) allocation performance across four core metrics: bed-category coverage (a,e), average makespan (b,f), efficiency loss (c,g), and preference distribution (d,h).

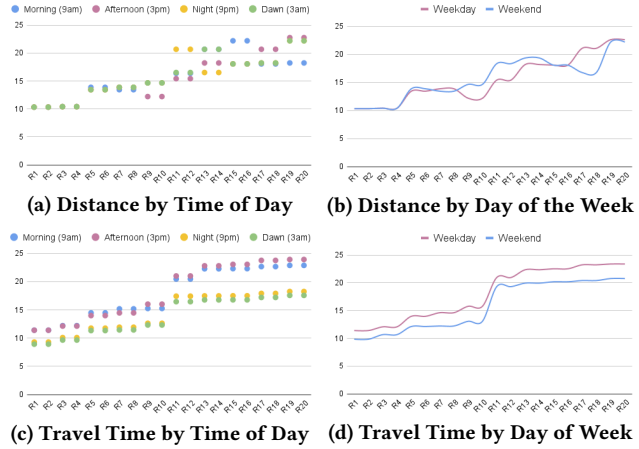

(a) Distance by Time of Day (b) Distance by Day of the Week

(c) Travel Time by Time of Day (d) Travel Time by Day of Week

Figure 6: Earthquake results for Distance (km) and Travel Time Distribution (min) by Time of Day (a, c), and by Time of Week (b, d)

morning (9am), afternoon (3pm), night (9pm), and dawn (3am). As demonstrated in Fig. 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.

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.

6.3 Water Contamination Use Case: Relocation Based on Water Constraints

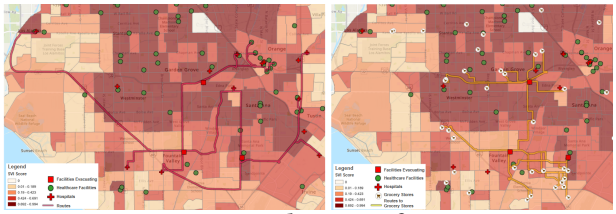

Figure 7: Assignment and Routing for Water Scenario

For this scenario, 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 [15, 54]. We updated facility-level water availability using real-time ReddiNet [44] reports and enabled congestion-aware delivery routing via CalTrans [13] feeds and the Google Distance Matrix API [24]. For re-supply, we identified each facility's ten nearest potable-water vendors and incorporated five official Points-of-Distribution (PODs) from OCHCA [37]. 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 relocations as the scale of water-contamination impacts grows. Fig. 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

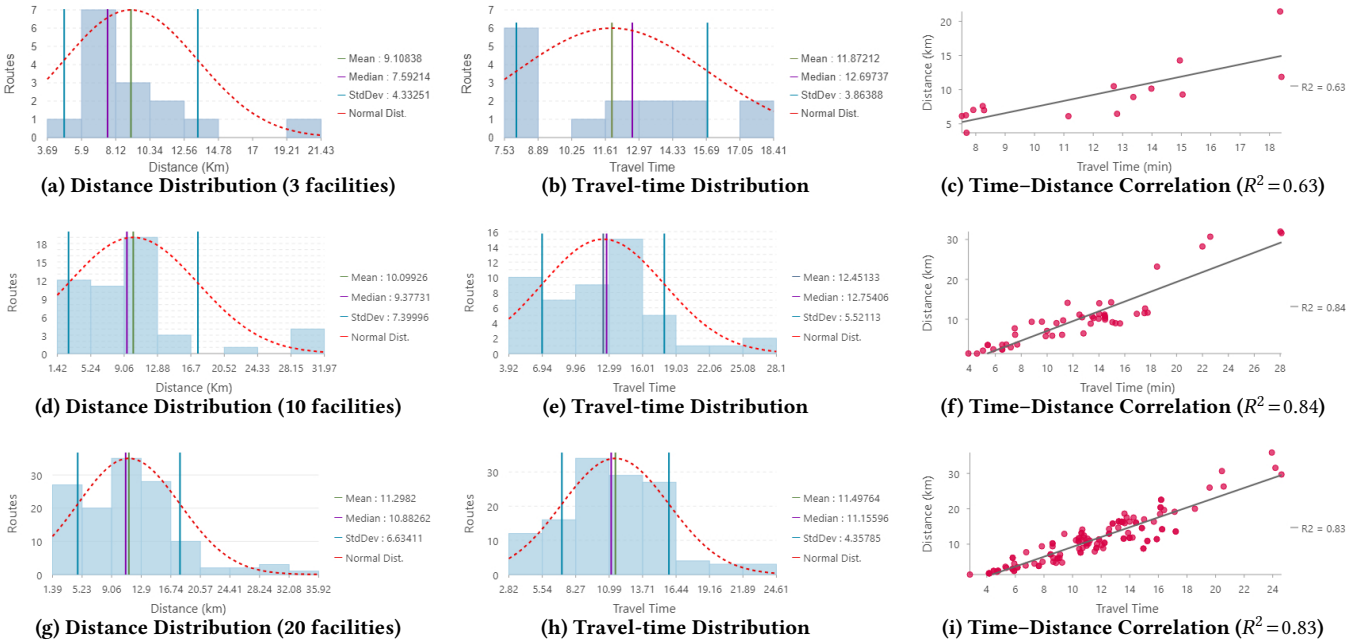

(a) Distance Distribution (3 facilities) (b) Travel-time Distribution (c) Time–Distance Correlation ($R^2 = 0.63$)

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

(1.5–5km) and the extended relocations, highlighting occasional assignments to distant, uncontaminated sites once nearby capacity is exhausted. Travel-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, over 90% of residents/patients reach a safe facility within two standard deviations of these means, ensuring that emergency evacuation remains 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. 10 (Appendices) 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 RADAR 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. 12). 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.

7 CONCLUSION AND FUTURE WORK

In this paper, we presented RADAR, a decision support system that integrates live hazard feeds, real-time traffic, and dynamic facility statuses within a PostGIS-enabled framework to address critical matching and routing challenges in disaster response. By coupling policy-driven stable matching with risk-aware routing, RADAR provides envy-free assignments and adaptive evacuation paths as conditions evolve. Comprehensive simulations and real-world drills with county health agencies and first responders demonstrate significant improvements in both routing efficiency and resource distribution compared to conventional methods. Although our current focus is senior health care, RADAR's modular architecture and policy framework make it extensible to other populations.

Future efforts will strive to enhance the scalability of RADAR and integrate additional real-time data streams. We also plan to cautiously integrate light LLM models as a quick recommendation guide and comparison aid only (given current limitations in accuracy and hallucinations) while keeping the optimization engine authoritative. Ultimately, we aim to deliver more accurate, actionable, and data-driven recommendations to Emergency Operations Centers for better informed decision-making during disasters.

REFERENCES

[1] Mohammad F Abdullah, Sajid Siraj, and Richard E Hodgett. 2021. 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 (2021), 1358.

[2] Jose Almeida, Joao Soares, Fernando Lezama, and Zita Vale. 2022. Robust energy resource management incorporating risk analysis using conditional value-at-risk. *IEEE Access* 10 (2022), 16063–16077.

[3] Joshua R Atencio, Md Sadman Siraj, and Eirini E Tsiropoulou. 2025. EVACUS-CAPE: Internet of Things-enabled emergency evacuation based on matching theory. *Internet of Things* 31 (2025), 101581.

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

[5] Leonardo Bedoya-Valencia and Emre Kirac. 2016. Evaluating alternative resource allocation in an emergency department using discrete event simulation. *Simulation* 92, 12 (2016), 1041–1051.

[6] Nancy Berlinger, Matthew Wynia, Tia Powell, Micah D Hester, Aimee Milliken, Rachel Fabi, and N P Jenks. 2020. Ethical framework for health care institutions responding to novel Coronavirus SARS-CoV-2 (COVID-19) guidelines for institutional ethics services responding to COVID-19. *The Hastings Center* 12 (2020).

[7] Dimitris Bertsimas, Vivek F Farias, and Nikolaos Trichakis. 2011. The price of fairness. *Operations Research* (2011).

[8] Thomas Bonald, Laurent Massoulié, Alexandre Proutiere, and Jorma Virtamo. 2006. A queueing analysis of max-min fairness, proportional fairness and balanced fairness. *Queueing systems* 53 (2006).

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

[10] CALFIRE. n.d.. Palisades fire. https://bit.ly/3F36tYi Accessed Feb 2025.

[11] CalHHS. n.d.. California Health and Human Services Agency. https://data.chhs.ca.gov/dataset/healthcare-facility-locations Accessed Apr 2024.

[12] Caltrans. n.d.. California Department of Transportation. https://quickmap.dot.ca.gov/ Accessed Aug 2025.

[13] Caltrans. n.d.. California Department of Transportation. https://dot.ca.gov/ Accessed Jul 2024.

[14] Changshuai Cao and Yingjuan Su. 2024. Transportation infrastructure and regional resource allocation. *Cities* 155 (2024), 105433.

[15] CDPH. 2023. AFL Emergency Water Requirements. https://bit.ly/44X8W1q Accessed Apr 2025.

[16] Theodoros Chondrogiannis, Panagiotis Bouros, and Winfried Emser. 2021. Simulation-based evacuation planning for urban areas. In *Proceedings of the 29th International Conference on Advances in Geographic Information Systems*. 297–300.

[17] George B Dantzig and John H Ramser. 1959. The truck dispatching problem. *Management science* 6, 1 (1959), 80–91.

[18] Kalyanmoy Deb, Samir Agrawal, Amrit Pratap, and Tanaka Meyarivan. 2000. A fast elitist non-dominated sorting genetic algorithm for multi-objective optimization: NSGA-II. In *International conference on parallel problem solving from nature*. Springer, 849–858.

[19] Thomas E Drabek. 2012. *Human system responses to disaster: An inventory of sociological findings*. Springer Science & Business Media.

[20] Xiaojie Du, Yuan Chen, Ahmed Bouferguene, and Mohamed Al-Hussein. 2020. An agent-based simulation framework for analysing fall risk among older adults in the evacuation procedures. *Safety science* 129 (2020), 104790.

[21] Barry E Flanagan, Edward W Gregory, Elaine J Hallisey, Janet L Heitgerd, and Brian Lewis. 2011. A social vulnerability index for disaster management. *Journal of homeland security and emergency management* 8, 1 (2011).

[22] Centers for Disease Control and Prevention. n.d.. National Risk Index. https://hazards.fema.gov/nri/map Accessed Jun 2024.

[23] David Gale and Lloyd S Shapley. 1962. College admissions and the stability of marriage. *The American mathematical monthly* 69, 1 (1962), 9–15.

[24] Google. 2005. Google Distance Matrix API. https://bit.ly/4boWxo6 Accessed Feb 2024.

[25] Jeannie L Haggerty, Robert J Reid, George K Freeman, Barbara H Starfield, Carol E Adair, and Rachael McKendry. 2003. Continuity of care: a multidisciplinary review. *Bmj* 327, 7425 (2003), 1219–1221.

[26] HCAI. n.d.. Dep. of Health Care Access and Information. https://bit.ly/41ndA5q Accessed Feb 2024.

[27] Tara N Heagele. 2018. *A qualitative study of household emergency preparedness of the elderly and the medically frail living in coastal urban environments*. Ph. D. Dissertation. Rutgers University-Graduate School-Newark.

[28] Zahra J K Abadi and Najme Mansouri. 2024. A comprehensive survey on scheduling algorithms using fuzzy systems in distributed environments. *Artificial Intelligence Review* 57, 1 (2024), 4.

[29] Farid Kadri, Babiga Birregah, and Eric Châtelet. 2014. The impact of natural disasters on critical infrastructures: A domino effect-based study. *Journal of Homeland Security and Emergency Management* 11, 2 (2014), 217–241.

[30] Modeste M Kenne, Prasanna Date, Ronald T Eguchi, ZhengHui Hu, Julie Rousseau, and Nalini Venkatasubramanian. 2024. iFair: Achieving Fairness in the Allocation of Scarce Resources for Senior Health Care. In *2024 IEEE International Conference on Smart Computing (SMARTCOMP)*. IEEE, 22–30.

[31] Leonid Kogan, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman. 2017. Technological innovation, resource allocation, and growth. *The quarterly journal of economics* 132, 2 (2017), 665–712.

[32] Zongmin Li, Xinxin Zhang, Yanfang Ma, Cuiying Feng, and Asaf Hajiyev. 2019. A multi-criteria decision making method for urban flood resilience evaluation with hybrid uncertainties. *International Journal of Disaster Risk Reduction* 36 (2019), 101140.

[33] Julia F Lynch, Isabel M Perera, and Theodore J Iwashyna. 2021. Scarce resource allocation in a pandemic: A protocol to promote equity, timeliness, and transparency. *Critical care explorations* 3, 6 (2021), e0466.

[34] Sharad Mehrotra, Alfred Kobsa, Nalini Venkatasubramanian, and Siva R Rajagopalan. 2016. TIPPERS: A privacy cognizant IoT environment. In *2016 IEEE International Conference on Pervasive Computing and Communication Workshops (PerCom Workshops)*. IEEE, 1–6.

[35] Ayan Mukhopadhyay, Geoffrey Pettet, Sayyed Mohsen Vazirizade, Di Lu, Alejandro Jaimes, Said El S., Hiba Baroud, Yevgeniy Vorobeychik, Mykel Kochenderfer, and Abhishek Dubey. 2022. A review of incident prediction, resource allocation, and dispatch models for emergency management. *Accident Analysis & Prevention* 165 (2022), 106501.

[36] NIA. n.d.. National Institute on Aging. http://bit.ly/3F4hQPL Accessed Feb 2024.

[37] OCHCA. n.d.. Orange County Health Care Agency. https://www.ochealthinfo.com/ Accessed Jan 2024.

[38] University of California Irvine. 2021. CareDEX Workshop. https://sites.uci.edu/caredex/ Accessed May 2024.

[39] Benjamin S Olivari. 2018. CDC grand rounds: promoting well-being and independence in older adults. *MMWR. Morbidity and mortality weekly report* 67 (2018).

[40] Carina Omoeva, Nina M Cunha, and Wael Moussa. 2021. Measuring equity of education resource allocation: An output-based approach. *International Journal of Educational Development* 87 (2021), 102492.

[41] Graziano Onder, Iain Carpenter, Harriet Finne-Soveri, Jacob Gindin, Dinnus Frijters, Jean C. Henrard, Thorsten Nikolaus, Eva Topinkova, Matteo Tosato, Rosa Liperoti, et al. 2012. Assessment of nursing home residents in Europe: the Services and Health for Elderly in Long TERm care (SHELTER) study. *BMC health services research* 12, 1 (2012), 5.

[42] Eric Pike and Merrill R Landers. 2010. Responsiveness of the Physical Mobility Scale in Long-term Care Facility Residents. *Journal of Geriatric Physical Therapy* 33, 2 (2010), 92–98.

[43] Michael L Pinedo. 2016. Design and implementation of scheduling systems: More advanced concepts. *Scheduling: Theory, Algorithms, and Systems* (2016), 485–508.

[44] ReddiNet. n.d.. Real-Time Bed Availability Data. https://www.reddinet.com/ Accessed May 2023.

[45] Carpender S Kay, Paul H Campbell, Barbara J Quiram, Joshua Frances, and Jill J Artzberger. 2006. Urban evacuations and rural America: lessons learned from Hurricane Rita. *Public Health Reports* 121, 6 (2006), 775–779.

[46] Nazmus Sakib, Kathryn Hyer, Debra Dobbs, Lindsay Peterson, Dylan J Jester, Nan Kong, and Mingyang Li. 2022. A GIS enhanced data analytics approach for predicting nursing home hurricane evacuation response. *Health information science and systems* 10, 1 (2022), 28.

[47] Hinrich Schütze, Christopher D Manning, and Prabhakar Raghavan. 2008. *Introduction to information retrieval*. Vol. 39. Cambridge University Press Cambridge.

[48] Ying Song and Harvey J Miller. 2012. Exploring traffic flow databases using space-time plots and data cubes. *Transportation* 39 (2012), 215–234.

[49] Jingran Sun and Zhanmin Zhang. 2020. A post-disaster resource allocation framework for improving resilience of interdependent infrastructure networks. *Transportation Research Part D: Transport and Environment* 85 (2020), 102455.

[50] Chenmei Teng, Poshan Yu, and Liwen Liu. 2024. A cooperative optimization model and enhanced algorithm for guided strategies in emergency mobile facilities. *Humanities and Social Sciences Communications* 11, 1 (2024), 1–11.

[51] Steven Teutsch and Bernd Rechel. 2012. Ethics of resource allocation and rationing medical care in a time of fiscal restraint-US and Europe. *Public Health Reviews* 34, 1 (2012), 15.

[52] Justin W Timbie, Jeanne S Ringel, D Steven Fox, Francesca Pillemer, Daniel A Waxman, Melinda Moore, Cynthia K Hansen, Ann R Knebel, Richard Ricciardi, and Arthur L Kellermann. 2013. Systematic review of strategies to manage and allocate scarce resources during mass casualty events. *Annals of emergency medicine* 61, 6 (2013), 677–689.

[53] USGS. n.d.. United States Geological Survey. https://earthquake.usgs.gov Accessed May 2024.

[54] Nalini Venkatasubramanian, Craig A Davis, and Ronald T Eguchi. 2020. 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*. 62–65.

[55] Feiyue Wang, Ziling Xie, Hui Liu, Zhongwei Pei, and Dingli Liu. 2022. Multiobjective emergency resource allocation under the natural disaster chain with path planning. *International journal of environmental research and public health* 19, 13 (2022), 7876.

[56] Xinyu Wang, Tsan-Ming Choi, Haikuo Liu, and Xiaohang Yue. 2016. 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 (2016), 545–556.

[57] Xiaohua Wang, Yuliang Yan, Longtao Huang, Xiaoqing Zheng, and Xuan-Jing Huang. 2023. Hallucination detection for generative large language models by bayesian sequential estimation. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*. 15361–15371.

[58] Yanyan Wang. 2021. Multiperiod optimal allocation of emergency resources in support of cross-regional disaster sustainable rescue. *International Journal of Disaster Risk Science* 12, 3 (2021), 394–409.

[59] Jared L Ware. 2001. Geospatial data fusion: Training GIS for disaster relief operations. In *Proc., GIS User Conf*, Vol. 1097.

[60] Melissa Willoughby, Chebiwot Kipsaina, Noha Ferrah, Soren Blau, Lyndal Bugeja, David Ranson, and Joseph Elias Ibrahim. 2017. Mortality in nursing homes following emergency evacuation: a systematic review. *Journal of the American Medical Directors Association* 18, 8 (2017), 664–670.

[61] Ziwei Xu, Sanjay Jain, and Mohan Kankanhalli. 2024. Hallucination is inevitable: An innate limitation of large language models. *arXiv preprint arXiv:2401.11817* (2024).

[62] Xuedong Yan, Xiaobing Liu, and Yulei Song. 2018. Optimizing evacuation efficiency under emergency with consideration of social fairness based on a cell transmission model. *PloS one* 13, 11 (2018), e0207916.

[63] Muhannad H Yousef, Yazan N Alhalaseh, Razan Mansour, Hala Sultan, Naseem Alnadi, Ahmad Maswadeh, Yasmeen M Al-Sheble, Raghda Sinokrot, Khawlah Ammar, Asem Mansour, et al. 2021. The fair allocation of scarce medical resources: a comparative study from Jordan. *Frontiers in medicine* 7 (2021), 603406.

[64] Hongzhe Zhang, Xiaohang Zhao, Xiao Fang, and Bintong Chen. 2024. Proactive resource request for disaster response: A deep learning-based optimization model. *Information Systems Research* 35, 2 (2024), 528–550.

A APPENDICES

A.1 RADAR Fine-Grained: 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) [30, 38], 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. 9). 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, key enhancements in RADAR fine-grained are as follows:

Medical compatibility: We adapt the concept of recall from information retrieval [47] 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:

$$S_{\text{med}}(F_{\text{rel}}, F_{\text{rec}}) = \frac{\left|\mathcal{S}_{\text{req}} \cap \mathcal{S}_{\text{avail}}(F_{\text{rec}})\right|}{|\mathcal{S}_{\text{req}}|}, \tag{16}$$

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 [25]. For resident r and candidate facility c:

$$S_{\text{cont}}(r, c) = \begin{cases} 1, & \text{if network}(r) = \text{network}(c), \\ 0, & \text{otherwise}, \end{cases} \tag{17}$$

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:

$$S_{\text{soc}}(r, c) = \sum_{a \in \mathcal{A}} w_a S_a(r, c), \quad \sum_a w_a = 1, \tag{18}$$

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 $\mathcal{A} = \{\text{spouse, pet}\}$ 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 [5, 7, 30, 63].

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 [30]. 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 [36]. 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 §6, 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.

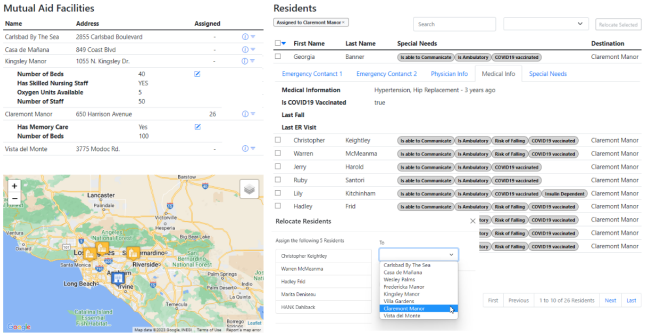

Figure 9: Dashboard with Integrated Medical Profiles

A.2 Additional Experiments and Visualization

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.

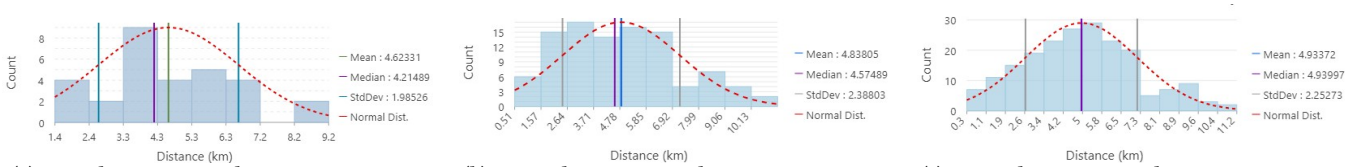

(a) 3-Facilities Impacted: Distance to Stores (b) 10-Facilities Impacted: Distance to Stores (c) 20-Facilities Impacted: Distance to Stores

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

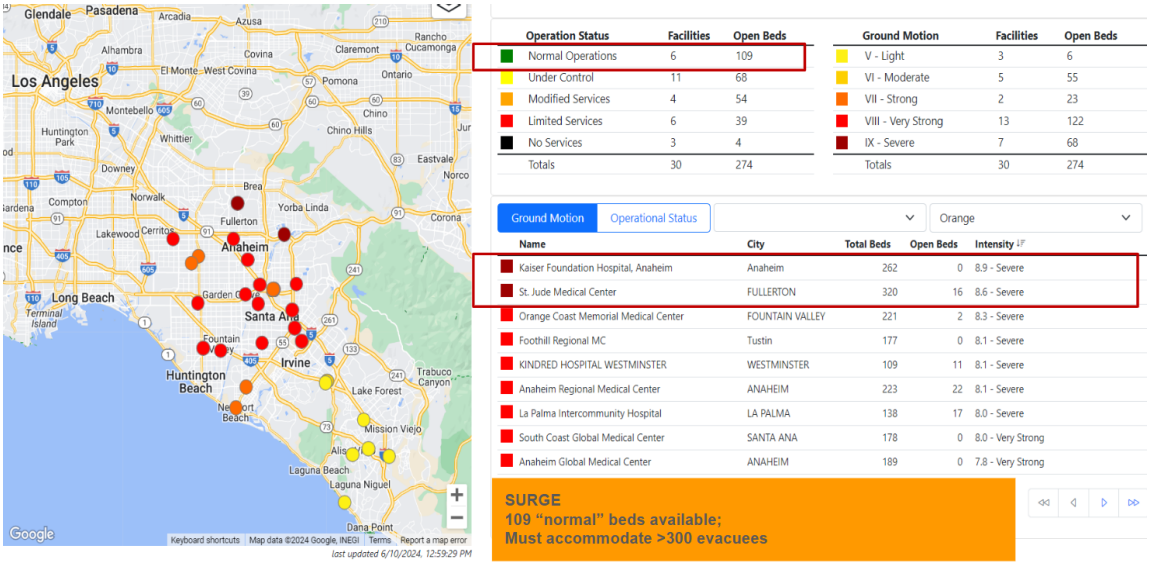

Figure 11: Dashboard of the Earthquake Drill Showing Surge Capacity.

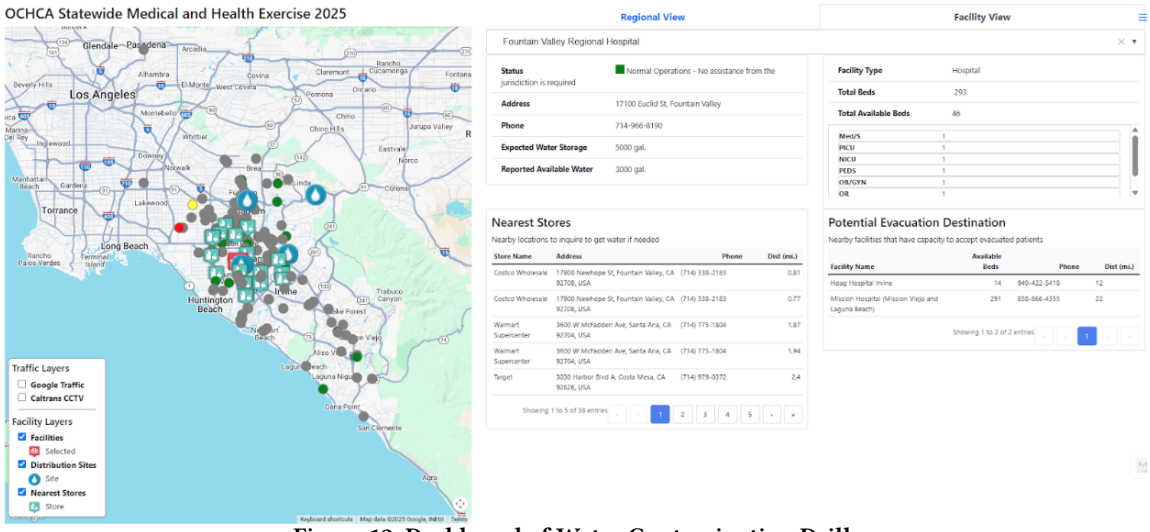

Figure 12: Dashboard of Water Contamination Drill.