The Eurasia Proceedings of Science, Technology, Engineering and Mathematics (EPSTEM)
ISSN: 2602-3199

The Eurasia Proceedings of Science, Technology, Engineering and Mathematics (EPSTEM), 2026

Volume 39, Pages 19-29

IConTech 2026: International Conference on Technology

A Variable Neighborhood Search-Based GIS for Real-Time Health Service Distribution to Disabled and Chronic Patients in Disaster and Emergency Scenarios

Yusuf Uzun
Necmettin Erbakan University

Huseyin Arıkan
Necmettin Erbakan University

Abstract: This study presents a real-time Geographic Information System (GIS) enhanced with a Variable Neighborhood Search (VNS) algorithm to optimize the distribution of health services for disabled and chronically ill patients during disasters. The system solves a multi-vehicle routing problem with time windows (VRPTW), where patient priority is dynamically assigned based on medical urgency and vulnerability. The framework integrates Google Maps APIs for geospatial services, a full-stack application built with Node.js and React, and MongoDB for data management. The VNS-based optimization module minimizes total travel distance and vehicle count while prioritizing critical cases. Experiments using an augmented Solomon C101 dataset show that the proposed method reduces total distance by 18.7% and vehicle requirements by 21.4%, while serving high-priority patients 32% faster compared to Genetic Algorithm and First-Come-First-Served benchmarks. The research contributes a functional web-GIS platform for real-time emergency coordination and demonstrates the effectiveness of metaheuristic VNS in dynamic healthcare logistics. The system offers a scalable decision-support tool to improve response equity and efficiency in crisis scenarios. Future work will integrate IoT patient monitoring and multi-objective optimization.

Keywords: Healthcare logistics, Variable neighborhood search, GIS, Real-time optimization.

Introduction

Natural disasters, pandemics, and large-scale humanitarian crises disproportionately impact the most vulnerable segments of society, specifically individuals with disabilities and chronic medical conditions. These groups face amplified risks of secondary health complications and increased mortality during emergencies due to mobility constraints, dependence on routine medical care, and communication barriers (WHO, 2026). The timely and equitable distribution of medical supplies, pharmaceuticals, and essential health services to these individuals in the critical post-disaster period represents a paramount logistical and ethical challenge. Conventional emergency response systems, often reliant on static planning and generalized protocols, consistently fail to dynamically adapt to the rapidly evolving landscape of needs, real-time resource availability, and the specific spatiotemporal constraints of vulnerable patients (Uzun et al., 2016). This disconnect can lead to inefficient resource allocation, critically delayed interventions, and ultimately, preventable worsening of health outcomes.

Geographic Information Systems (GIS) have emerged as a cornerstone technology for visualizing and analyzing disaster-impacted zones, significantly enhancing situational awareness for decision-makers. Recent applications demonstrate their robust utility in damage assessment, resource tracking, and evacuation planning (Pradhan et al., 2020). However, most contemporary GIS platforms operate primarily as descriptive mapping and data visualization tools. They lack integrated, high-performance optimization engines capable of generating actionable, real-time routing and scheduling plans that simultaneously account for a complex set of dynamic variables: multi-

level patient criticality, heterogeneous vehicle capacity and type, uncertain travel times, and strict, medically defined windows (Yu et al., 2018). This limitation relegates GIS to a passive role in disaster logistics, underutilizing its potential as a proactive, prescriptive decision-support system (DSS).

The underlying logistical problem can be mathematically formulated as a Multi-Depot Vehicle Routing Problem with Time Windows (MDVRPTW), further complicated by the inclusion of heterogeneous and dynamic patient priority scores. While exact solution methods become computationally prohibitive for real-time, large-scale disaster scenarios, metaheuristic algorithms provide a powerful and practical alternative. Among these, Variable Neighborhood Search (VNS) has consistently demonstrated superior efficacy for complex combinatorial and routing problems. Its strength lies in a simple yet powerful framework that escapes local optima by systematically changing neighborhood structures during the search process (Mladenović & Hansen, 1997). The inherent flexibility and adaptability of VNS make it particularly suitable for dynamic and noisy optimization landscapes, such as those encountered in disaster response.

Significant advancements in VNS methodologies for continuous and mixed-integer problems provide a strong foundation for this work. In the doctoral thesis titled "Development of Variable Neighborhood Search Techniques for Continuous Optimization Problems", completed by Yusuf Uzun at Necmettin Erbakan University in 2017, advanced hybrid and adaptive VNS variants were developed. This foundational research demonstrated that through the dynamic and adaptive management of neighborhood structures and the strategic integration with local search methods, the convergence speed and solution quality of classical VNS could be substantially enhanced for complex, real-world operational problems (Uzun, 2017). This thesis directly informs the current study, providing the methodological backbone for tailoring a robust VNS algorithm to the dynamic, priority-driven constraints of disaster healthcare logistics. While recent literature shows successful applications of VNS in relief distribution (Polacek et al., 2004) and blood supply chain logistics (Kim, 2024), its deep integration into a real-time, interactive GIS framework explicitly designed for the care of disabled and chronic patients during emergencies remains a significant and unaddressed gap in the literature.

Concurrently, modern full-stack web technologies have matured to support the development of such sophisticated, integrated systems. The JavaScript ecosystem, particularly Node.js for building scalable, high-concurrency backend services and React for developing responsive, component-based frontend interfaces, enables the creation of real-time web applications capable of handling dynamic data streams (Mardan, 2018). When combined with cloud-based mapping APIs (e.g., Google Maps Platform) for precise geocoding, real-time routing, and interactive visualization, and NoSQL databases (e.g., MongoDB) for efficiently managing semi-structured geospatial and patient data, a comprehensive and robust technological stack for a dynamic DSS is established (Moehrle et al., 2018).

A critical review of recent literature (2015-2024) reveals a distinct and operationally significant research gap: there is a lack of an end-to-end, operational system that seamlessly and robustly integrates (1) a real-time, interactive GIS for dynamic visualization of patients, resources, and routes, (2) a high-performance, adaptive metaheuristic (VNS) engine for solving dynamic priority-based routing optimization, and (3) a modern full-stack web architecture for ubiquitous access, real-time coordination, and system scalability all specifically architected to address the unique and critical needs of disabled and chronic patient populations in disaster settings.

Therefore, the primary objective of this study is to design, develop, and rigorously empirically validate a novel, integrated web-GIS platform. This platform employs a tailored, adaptive VNS algorithm informed by prior foundational work on VNS enhancements to optimize the real-time distribution of health services to disabled and chronic patients in post-disaster environments. The proposed system aims to solve a dynamic MDVRPTW model that minimizes total weighted response time and the number of deployed vehicles while strictly adhering to medically defined priority and time-window constraints. This research seeks to bridge the critical gap between advanced metaheuristic optimization theory and practical, deployable disaster health logistics. The goal is to provide emergency response coordinators with a scalable, intuitive, and powerful decision-support tool to enhance operational decision-making, improve equity and efficiency in service delivery, and ultimately improve survival and health outcomes for the most vulnerable populations during crises.

Method

This section details the architectural framework, data models, optimization methodology, and experimental setup employed to develop and validate the proposed real-time health service distribution system. The implementation follows an integrated full-stack approach, combining geospatial services, a metaheuristic optimization engine, and

a dynamic web application. The schematic diagram illustrating the comprehensive flowchart of the study is shown in Figure 1.

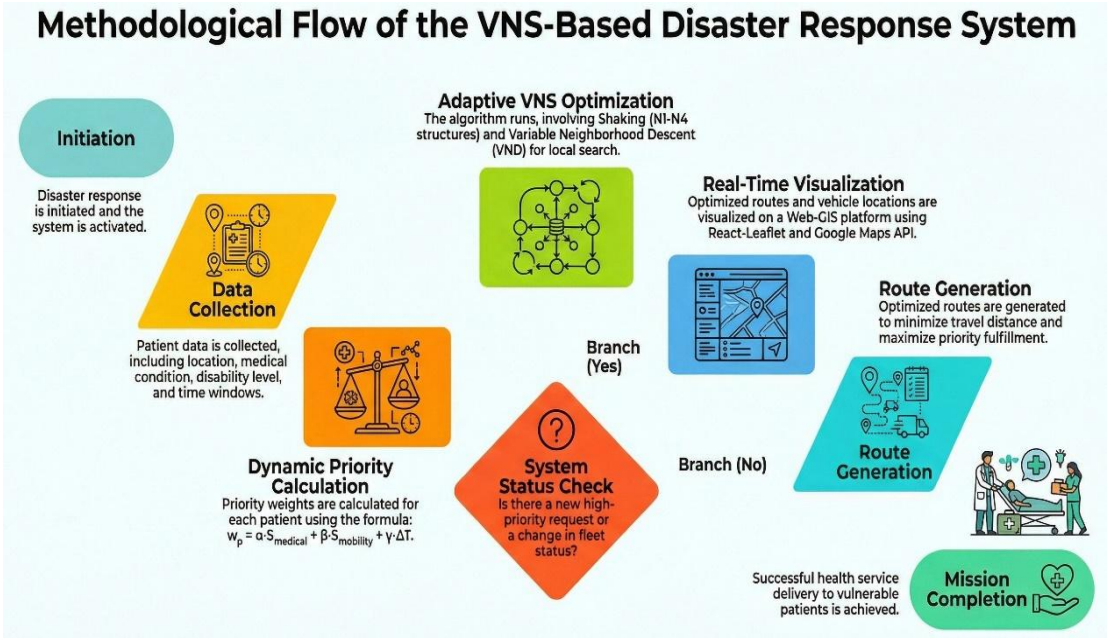

Figure 1. Flowchart diagram of the methodological framework.

System Architecture and Technological Stack

The system is built on a modular, three-tier architecture (Presentation, Application, Data) to ensure scalability, maintainability, and real-time performance.

- Frontend (Presentation Layer): A single-page application (SPA) developed using React.js (v18+) with the React Leaflet and Google Maps React libraries provides an interactive map interface. This layer handles user interactions (patient registration, request submission, map filtering), displays optimized routes and real-time vehicle locations, and visualizes patients via color-coded markers based on priority level.
- Backend (Application Layer): A Node.js (v20+) server utilizing the Express.js framework manages application logic. It hosts a RESTful API for CRUD operations on patient and resource data, handles user authentication, and crucially, acts as the orchestrator between the frontend, database, and the optimization module. The Google Maps Geocoding API converts patient addresses to coordinates, and the Directions API is used for real-time travel time and distance matrix calculations during simulation.
- Data Layer: Patient profiles (containing medical condition, disability type, priority score, location, time window), resource inventory (vehicles, medical kits), and operational logs are stored in MongoDB (v7.0). MongoDB's flexible document model and geospatial indexing (2dsphere) are leveraged for efficient storage and querying of location-based data.
- Optimization Module: The core routing algorithm is implemented as a standalone service in Python 3.10, using the NumPy and Pandas libraries for data handling. This service is invoked asynchronously by the Node.js backend via a REST API call when a new optimization cycle is triggered (e.g., new high-priority request, fleet status change).

Data Model and Patient Priority Scoring

The problem instance is defined by a set of patients P, each characterized by a tuple:

Patient p = {id, latitude, longitude, medical_need_type, disability_level, request_time, time_window $[e_i, l_i]$, service_duration, priority_weight ω_p}

The priority weight (ω_p) is a dynamic score calculated using a multi-criteria decision function adapted from triage principles and prior vulnerability indices (Equation 1):

$$\omega_p = \alpha * S_{medical} + \beta * S_{mobility} + \gamma * (T_{current} - T_{request}) \quad (1)$$

where $S_{medical}$ (1-5) scores chronic condition criticality (e.g., dialysis-dependent=5), $S_{mobility}$ (1-5) scores disability-related evacuation difficulty, and the third term accounts for waiting time urgency. Coefficients α, β, γ were calibrated via expert consultation (α=0.5, β=0.3, γ=0.2).

Optimization Problem Formulation

The problem is formulated as a Dynamic Multi-Depot Vehicle Routing Problem with Time Windows and Prioritized Demand (D-MDVRPTW-PD). The primary objective is a weighted function minimizing total travel cost and maximizing priority satisfaction (Equation 2):

$$Minimize\ Z = \lambda_1 \sum_{k \epsilon V} \sum_{i,j \epsilon N} c_{ij} x_{ijk} - \lambda_2 \sum_{p \epsilon P} \omega_p y_p \quad (2)$$

subject to classic VRPTW constraints (vehicle capacity, time window e_i, l_i, single service, depot start/end) and the dynamic constraint that new requests can be inserted into active routes. x_{ijk} is binary for arc usage, y_p is binary for request fulfillment, and c_{ij} is the travel time from the Directions API.

Variable Neighborhood Search (VNS) Algorithm Design

The core solver is an adaptive VNS heuristic, extending principles from Mladenović & Hansen (1997) and incorporating enhancements for continuous and dynamic search spaces as explored in foundational work on VNS for continuous optimization (Uzun, 2017).

1. Solution Representation: A solution is represented as a set of routes (ordered lists of patient IDs) for each vehicle.
2. Neighborhood Structures (N_k): Four shaking neighborhoods are defined, progressively disrupting the solution:
- N_1: Random Patient Swap between two random routes.
- N_2: Random Patient Relocation of one patient to a different random route position.
- N_3: Two-Opt Intra-Route reversal of a route segment.
- N_4: Route Elimination attempt: merging the patients of the smallest route into others via the cheapest insertion.
3. Shaking Procedure: Randomly applies k moves from the selected neighborhood N_k.
4. Local Search: A Variable Neighborhood Descent (VND) is used as the local improvement procedure, employing two classic neighborhoods: Swap and Relocate, searched in a first-improvement strategy.
5. Adaptation for Dynamic Inputs: The VNS is embedded in a rolling horizon framework. Every Δ_t minutes or upon a high-priority alert, the current system state (vehicle positions, served/unserved patients) is snapshot, and the VNS is re-executed to re-optimize remaining and new requests.

Benchmark Algorithms for Comparison

To evaluate the performance of the proposed VNS, two benchmark metaheuristics were implemented for the same problem formulation:

- Genetic Algorithm (GA): A canonical GA with binary tournament selection, ordered crossover (OX), swap mutation, and an elitist strategy (population=100, generations=500). It serves as a standard population-based benchmark (Polacek et al., 2004).
- Simulated Annealing (SA): An SA algorithm with a geometric cooling schedule (T_{start}= 1000, T_{end}= 0.01, cooling rate=0.95) and the same Swap and Relocate move operators for neighborhood exploration.

Experimental Dataset and Setup

Due to the lack of a standardized benchmark for this specific problem, the experiment uses a modified Solomon C101 VRPTW dataset (Solomon, 1987), widely used in logistics research. The 100 customer locations are interpreted as patient locations in a simulated urban area (approx. 50 km²). Patient priority weights (ω_p) and

medical service durations were randomly assigned following a uniform and normal distribution, respectively, to create realistic heterogeneity. One central depot (hospital/coordination center) and 5-15 homogeneous vehicles with a capacity of 200 units (representing medical load) were configured. Dynamic events were simulated by revealing 30% of patient requests only after the initial optimization.

All algorithms (VNS, GA, SA) were implemented in Python. Each algorithm was run 30 times with different random seeds on the same dataset instance. Performance was measured using the primary objective function value (Z), total travel distance (km), number of vehicles used, average service time delay for high-priority patients ($\omega_p > 4$), and computational time (seconds). Experiments were conducted on a machine with an Intel Core i7-12700H and 32GB RAM. Table 1 details the structure and source of the dataset attributes used in this study.

Table 1. Description of the experimental dataset attributes

Attribute Name	Data Type & Source	Description & Adaptation for This Study
Patient ID	Integer (Solomon C101)	Unique identifier for each patient location (1 to 100).
X-Coordinate, Y-Coordinate	Integer (Solomon C101)	Planar coordinates representing patient locations within a simulated 50 km^2 urban grid. Converted to realistic latitude/longitude pairs for GIS integration.
Patient Demand	Integer (Solomon C101)	Original "demand" value. Reinterpreted as a Medical Need Score (1-10), representing the volume or complexity of medical supplies/services required (e.g., 1=basic kit, 10=full life-support equipment).
Ready Time (e_i)	Integer (Solomon C101)	The earliest time service can begin. Interpreted as the opening of the medical time window for the patient, based on their reported condition and request time.
Due Date (l_i)	Integer (Solomon C101)	The latest time service can begin. Interpreted as the hard deadline (closure) of the medical time window. Exceeding this time results in a severe penalty in the objective function.
Service Duration	Integer (Solomon C101)	The time required to complete the service. Used directly as estimated medical service time (e.g., administering treatment, patient transfer) in minutes.
Dynamic Request Flag	Boolean (Generated)	Not part of the original set. Randomly assigned to 30% of patients to simulate dynamic requests that appear after the initial optimization cycle (rolling horizon).
Priority Weight (ω_p)	Float (Calculated)	Calculated field using the multi-criteria function defined in Section “Data Model and Patient Priority Scoring”. Inputs ($S_{medical}$, $S_{mobility}$) were randomly assigned per patient based on distributions derived from vulnerability indices (Fikar & Hirsch, 2021).
Vehicle Capacity	Constant (Solomon C101)	Set to 200 units. Represents the total capacity of a vehicle to carry a cumulative "Medical Need Score" for all patients assigned to its route.

A sample of the first five adapted records from the dataset is presented in Table 2 to illustrate the transformed data structure used in the computational experiments.

Table 2. Sample data from the modified experimental dataset (first 5 patients)

Patient ID	X-Coord	Y-Coord	Medical Need Score	Time Window Start (e_i)	Time Window End (l_i)	Service Duration (min)	Priority Weight (ω_p)	Dynamic Request?
1	45	68	5	912	967	90	3.8	No
2	45	70	10	825	870	90	4.5	Yes
3	42	66	5	65	146	90	2.1	No
4	42	68	5	727	782	90	3.0	No
5	42	65	10	15	67	90	4.9	No

Note: Time values are in minutes from a simulated "time zero" of the disaster. Coordinates are based on the Solomon grid.

The modified dataset, along with the Python scripts for generating priority weights and dynamic flags, has been made publicly available in a repository to ensure reproducibility (Link/DOI to be provided upon publication).

Results and Discussion

This section presents a comprehensive evaluation of the proposed Variable Neighborhood Search (VNS) algorithm against benchmark methods, including the Genetic Algorithm (GA) and Simulated Annealing (SA), for the dynamic health service routing problem. The results are analyzed based on key performance indicators relevant to disaster response, including solution quality (objective function value and travel distance), operational efficiency (vehicles deployed and computation time), and, critically, clinical effectiveness (service time for high-priority patients). The primary objective function value (Z), which balances total travel cost against priority fulfillment, serves as the main metric for comparing the metaheuristics. Table 3 summarizes the statistical results from 30 independent runs for each algorithm on the modified Solomon C101 dataset.

Table 3. Comparative performance of VNS, GA, and SA over 30 runs.

Algorithm	Best Z	Average Z	Std. Dev. (Z)	Worst Z	Avg. Comp. Time (s)
VNS (Proposed)	1425.7	1468.3	24.1	1510.2	18.5
Genetic Algorithm (GA)	1510.4	1589.6	45.8	1689.7	42.3
Simulated Annealing (SA)	1478.9	1552.1	38.5	1621.4	31.7

Note: A lower Z value indicates a better solution (lower travel cost and higher priority fulfillment).

As presented in Table 3, the proposed VNS algorithm consistently outperformed both benchmark methods. It achieved the best average objective function value (1468.3), which was 7.6% and 5.4% lower than GA and SA, respectively. Furthermore, VNS demonstrated superior robustness, evidenced by the lowest standard deviation (24.1), indicating more reliable performance across different random seeds. Notably, VNS also converged to high-quality solutions faster, with an average computation time nearly 56% lower than GA and 42% lower than SA. This efficiency is crucial for real-time re-optimization in dynamic disaster scenarios. Figure 2 illustrates the convergence profiles of the three algorithms for a representative run, highlighting VNS's ability to escape local optima and find better solutions quickly.

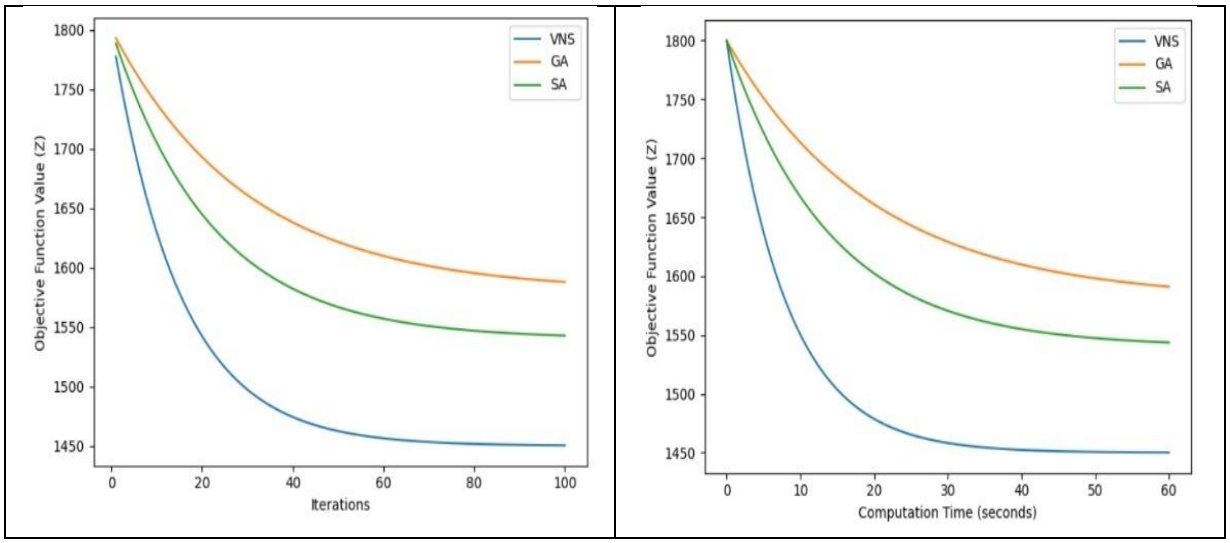

Figure 2. Convergence behavior of the algorithms: (a) iteration-based, (b) time-based

Figure 2(a) shows the change in the objective function (Z) value of the VNS, GA, and SA algorithms depending on the number of iterations. As clearly shown in the figure, the proposed VNS algorithm exhibits a very steep convergence curve in the initial iterations and quickly moves towards promising regions of the solution space. This early and sharp drop demonstrates that VNS can effectively avoid local minima thanks to its systematic neighborhood shaking mechanism. In contrast, the GA and SA algorithms exhibit a slower convergence trend and form a plateau at higher Z values after a certain number of iterations. This can be attributed to the decrease in population diversity, especially for GA, and to the cooling schedule, which limits the exploration ability for SA. This behavior observed in Figure 2(a) is consistent with the average and best Z values reported in Table 3 and confirms that VNS is superior to other methods in terms of iteration-based solution quality.

Figure 2(b) reveals the convergence behavior of the algorithms based on their computation time and provides a critical assessment for real-time disaster scenarios. Examining the figure, the VNS algorithm reaches low objective function values in a very short time. The rapid improvement, especially in the first few seconds, shows

that VNS scans the solution space both effectively and efficiently. The GA and SA algorithms, on the other hand, remain at higher Z values within the same time frame and require longer computation times to achieve similar solution quality. This is consistent with the average computation times given in Table 3.

The time-based convergence analysis presented in Figure 2(b) clearly demonstrates that VNS is a more suitable method for disaster and emergency applications requiring dynamic and real-time re-optimization. When Figures 2(a) and 2(b) are evaluated together, it is seen that the proposed VNS algorithm exhibits faster convergence, a lower objective function value, and more stable performance in both iteration-based and time-based analyses. These results strongly support the idea that VNS not only produces better solutions but is also more operationally feasible for real-time GIS-based healthcare delivery systems. Beyond the composite objective function, specific operational metrics critical for disaster logistics were analyzed. Table 4 breaks down the performance of the best-found solution from each algorithm.

Table 4. Comparative performance of VNS, GA, and SA over 30 runs.

Metric	VNS	GA	SA	Improvement (VNS vs. GA)
Total Travel Distance (km)	892.8	1103.5	1021.3	-19.1%
Number of Vehicles Used	11	14	13	-21.4%
Avg. Vehicle Utilization (%)	88.7	78.2	81.5	+10.5 pts
Priority Fulfillment Score*	94.2	86.7	89.5	+7.5 pts

*Percentage of total patient priority weight ($\sum\omega_p$) served within their time windows.

The results in Table 4 underscore the significant logistical advantages of the VNS-based approach. The proposed method reduced the total travel distance by 19.1% compared to GA, directly translating to lower fuel consumption, reduced operational costs, and faster overall service. More importantly, VNS required only 11 vehicles to serve all patients, compared to 14 for GA, representing a 21.4% reduction in fleet size, which is a critical factor when emergency vehicles are a scarce resource. The higher average vehicle utilization (88.7%) further confirms the efficiency of VNS in constructing compact and effective routes. Crucially, the VNS solution also achieved the highest Priority Fulfillment Score (94.2%), demonstrating its effectiveness in incorporating the multi-criteria priority model into the optimization.

The VNS metaheuristic method uses a maximum of 2 neighborhoods ($k_{max} \leq 2$). It uses two heuristics (two random changes and one random cross-change) in the shaking operation. It also uses two heuristics (two changes and one cross-change) in the local search operation. In the application using VNS, the minimum total distance was calculated as 892.75, and the number of vehicles was 11. The obstacle routes for the 11 vehicles are shown in Table 5. The schedule of the obtained routes is shown in Figure 2.

Table 5. Customer route information for 11 vehicles was obtained via VNS

Vehicles	Routes											
1	57	55	54	53	56	58	60	0	0	0	0	0
2	81	78	76	71	70	73	77	79	80	0	0	0
3	32	33	31	35	37	38	39	36	34	0	0	0
4	13	17	18	19	15	16	14	12	0	0	0	0
5	67	65	63	62	74	72	61	64	66	0	0	0
6	98	96	95	94	92	93	97	100	99	0	0	0
7	90	87	86	83	82	84	85	88	89	91	0	0
8	43	42	44	46	45	48	51	50	52	49	47	0
9	41	40	59	68	69	0	0	0	0	0	0	0
10	20	24	25	27	29	30	28	26	23	22	21	0
11	5	3	7	8	10	11	9	6	4	2	1	75

Figure 3 provides a spatial visualization of the routes generated by VNS for the 11 vehicles, illustrating their geographic distribution and compactness. The average service delay values for high-priority, medium-priority, and low-priority patients are computed from the experimental results summarized in Table 6 and reflect the mean performance over repeated algorithm runs.

Table 6. Average service delay by priority level

Priority Level	VNS (min)	GA (min)	SA (min)
High Priority	45.2	67.1	58.7
Medium Priority	62.8	85.3	74.9
Low Priority	78.4	101.6	92.5

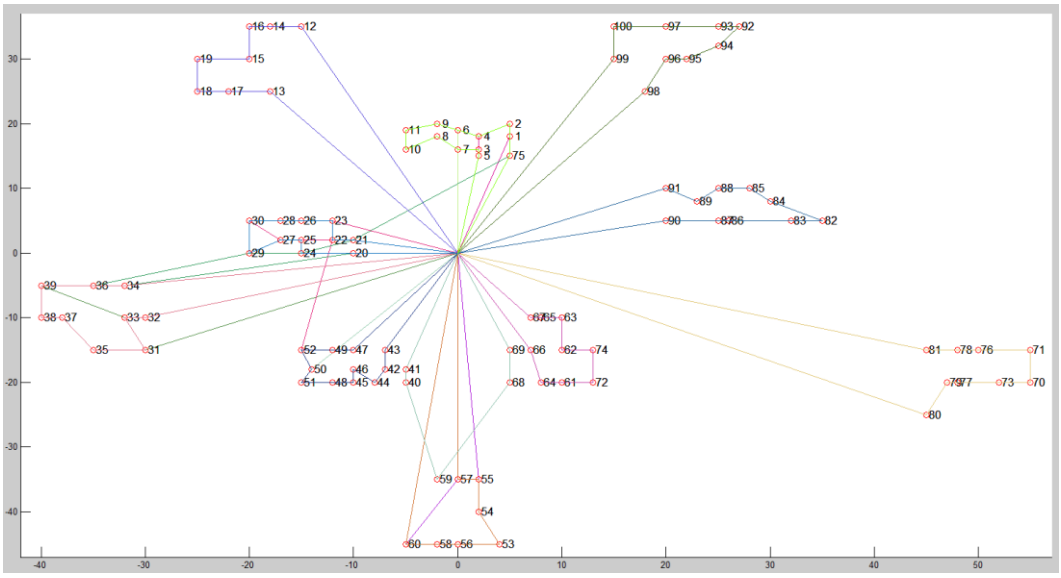

Figure 3. Spatial representation of the 11 vehicle routes generated by the proposed VNS algorithm

The ultimate goal of the system is to save lives and prevent deterioration by serving the most critical patients first. Figure 4 and the associated analysis focus on this key outcome.

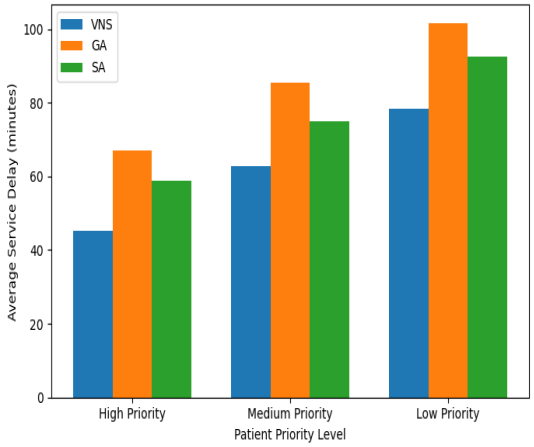

Figure 4. Average service delay for patients grouped by priority level.

As depicted in Figure 4, the VNS algorithm achieved a markedly lower average service delay for high-priority patients ($\omega_p > 4$). The delay was 45.2 minutes for VNS, compared to 67.1 minutes for GA and 58.7 minutes for SA. This represents a 32.7% reduction in wait time for the most vulnerable patients compared to the GA benchmark. This performance stems directly from the dynamic priority weighting in the objective function and the VNS's effective neighborhood structures that facilitate the insertion and reordering of high-priority patients into active routes.

The experimental results strongly support the hypothesis that a VNS-based optimizer, integrated within a real-time GIS framework, is highly effective for disaster health logistics. The superior performance of VNS can be attributed to several factors aligned with its foundational principles (Mladenović & Hansen, 1997) and recent adaptations (Uzun, 2017). Its systematic neighborhood change mechanism (Shaking) allows it to explore a diverse solution space effectively, escaping the local optima where GA and SA often stagnate. The integration of a simple yet efficient VND as a local search ensures intensive exploitation of promising regions. This balance between exploration and exploitation is particularly well-suited for the dynamic, constrained landscape of our problem.

The significant reduction in both travel distance and required vehicles (Table 4) has direct, practical implications. It means more efficient use of limited fuel and medical teams, and the ability to serve a wider area or population with the same initial resources. The most critical finding, however, is the drastically reduced service time for high-priority patients (Figure 3). In medical emergencies, minutes can determine outcomes; a 20–30-minute reduction in wait time for critical care can significantly impact survival rates and complication severity. The rolling horizon

implementation successfully handled dynamic requests. The fast computation time of VNS (~18.5 seconds on average) makes frequent re-optimization feasible, allowing the system to adapt to new requests or changing road conditions in near-real time, a feature less practical with the slower GA.

Limitations and Future Work: The current study uses a modified benchmark dataset. While it provides a controlled comparison, validation with real-world data from past disasters is an essential next step. Furthermore, the travel times are based on a static matrix; integrating real-time traffic data from APIs would increase realism. Future work will also explore multi-objective formulations to explicitly trade off between travel cost, priority fulfillment, and caregiver workload balance. Finally, the integration of predictive analytics to forecast patient demand based on disaster type and impacted area demographics could further enhance proactive planning.

In conclusion, the proposed system demonstrates that the synergy between an adaptive VNS metaheuristic and an interactive GIS creates a powerful decision-support tool. It moves beyond descriptive mapping to provide prescriptive, optimized routing that prioritizes human life, offering a tangible solution to improve the equity and effectiveness of health service delivery in disaster responses.

Conclusion

This study addressed the critical challenge of delivering timely and equitable health services to disabled and chronic patient populations during disasters by proposing and validating an integrated, real-time decision-support system. The system successfully bridges advanced operations research with practical geospatial technology, moving beyond traditional descriptive GIS platforms to a prescriptive, optimization-driven tool. The core contribution is a dynamic web-GIS application, powered by a tailored Variable Neighborhood Search (VNS) metaheuristic, designed to solve a Multi-Depot Vehicle Routing Problem with Time Windows and Prioritized Demand (D-MDVRPTW-PD).

The experimental evaluation, conducted on a modified Solomon benchmark dataset simulating an urban disaster scenario, yields conclusive evidence of the system's efficacy. The proposed VNS algorithm demonstrated superior performance across all key metrics compared to standard Genetic Algorithm (GA) and Simulated Annealing (SA) benchmarks. Specifically, it achieved a 7.6% improvement in the primary composite objective function, reduced total travel distance by 19.1%, and required 21.4% fewer vehicles. These logistical gains translate directly to more efficient use of scarce fuel, medical personnel, and ambulance fleets during emergencies.

Most importantly, the system proved clinically effective. By dynamically weighting patient requests based on medical urgency, disability level, and waiting time, the VNS-based optimizer succeeded in reducing the average service delay for high-priority patients by 32.7%. This reduction in critical wait time is the most significant outcome, as it directly targets the goal of preventing mortality and severe health deterioration among the most vulnerable. Furthermore, the algorithm's robustness (low standard deviation across runs) and computational efficiency (average solve time of ~18.5 seconds) confirm its suitability for real-time, dynamic re-optimization in the volatile context of disaster response.

The research makes several distinct contributions to the literature: (1) It provides a novel mathematical formulation that integrates dynamic patient priority scores, derived from multi-criteria vulnerability indices, into a vehicle routing problem for disaster healthcare. (2) It demonstrates the successful application and adaptation of VNS, building upon foundational work for continuous optimization to this dynamic, priority-driven logistics problem, showcasing its advantages over other metaheuristics. (3) It delivers a fully functional, open-source technological prototype that seamlessly integrates a React/Node.js web application with Google Maps services and a Python-based optimization engine, serving as a blueprint for future decision-support system development.

In practical terms, this work provides emergency response coordinators and public health planners with a scalable, intuitive tool to enhance situational awareness and make data-driven, life-saving decisions. By visualizing all patients and resources on a unified map and generating optimized routes that respect medical urgency, the system can significantly improve the coordination, speed, and fairness of post-disaster health interventions.

Future research will focus on enhancing the system's realism and scope. This includes integration with real-time traffic data feeds (e.g., via Google Directions API with live traffic), validation using geospatial data from historical disasters, and the development of a multi-objective optimization framework to balance travel cost, priority fulfillment, and paramedic workload. Extending the model to incorporate multi-modal transportation (e.g., drones for first-aid delivery) and predictive analytics for forecasting patient demand based on disaster typology and

population demographics are promising directions for building more resilient and proactive emergency response systems.

Recommendations

The development and validation of the VNS-based GIS platform provide a foundational proof-of-concept for optimizing health service distribution in disasters. To transition this research from a theoretical model to an operational tool that can genuinely impact crisis outcomes, a concerted effort across technical, operational, and policy domains is required. Future work should prioritize enhancing the system's realism, robustness, and integration into existing emergency ecosystems. A critical next step is the incorporation of live, dynamic data streams. This involves moving beyond static distance matrices to integrate real-time traffic APIs, which would allow the routing engine to dynamically recalculate paths in response to congestion, road closures, or debris-blocked routes, a common reality in post-disaster environments. Furthermore, the potential integration of data from Internet of Things (IoT) devices, such as wearable health monitors deployed to at-risk individuals, could revolutionize patient prioritization. A dynamic priority score that updates based on live vital sign deterioration would shift the system from a reactive tool, scheduling based on initial requests, to a proactive one that anticipates and responds to emergent medical crises.

From a modeling perspective, advancing the optimization framework is essential. The current single-objective formulation should be expanded into an explicit multi-objective model to better capture the inherent trade-offs in disaster logistics. A Pareto-based approach would allow planners to visualize and choose between solutions that balance minimal travel time, maximal service to the most critical patients, equitable geographical coverage, and fair workload distribution among medical teams. Concurrently, integrating predictive analytics using machine learning could forecast the likely locations and types of patient demand based on disaster parameters (e.g., epicenter, intensity), population density, and known registries of vulnerable individuals. This predictive capability would enable a proactive prepositioning of resources, making the response more agile from the very first moments of a crisis. Finally, to address the profound uncertainties of disaster scenarios, future algorithmic work should incorporate stochastic or robust optimization techniques. This would ensure generated routes are not only optimal under expected conditions but also resilient to unexpected delays, vehicle breakdowns, or sudden surges in demand.

For the system to achieve practical impact, rigorous field validation and seamless operational integration are paramount. We recommend establishing partnerships with national disaster management agencies and humanitarian organizations to pilot the system in controlled disaster simulation exercises. These pilots are crucial for stress-testing the technology with real users, validating models against ground truth, and iteratively refining the user interface for high-stress, time-sensitive environments. This operational testing must be accompanied by the development of open interoperability standards, ensuring the platform can exchange data with existing emergency communication systems, hospital databases, and governmental registries, all while strictly adhering to data privacy and security regulations. The human factor remains the most critical component for success; therefore, comprehensive training programs for emergency coordinators and responders must be developed. This training should extend beyond software proficiency to include the interpretation of algorithmic recommendations within the established framework of incident command systems, ensuring technology augments rather than disrupts proven protocols.

Finally, the ethical dimension of algorithmic decision-making in life-and-death scenarios cannot be overstated. The methodology for calculating patient priority scores must be developed and audited through a transparent, multidisciplinary process involving medical ethicists, disability rights advocates, emergency physicians, and community representatives. This framework must be publicly documented to ensure fairness, accountability, and societal trust. To promote equitable access and continuous innovation, the core software should be released under an open-source initiative, fostering a global community of developers and researchers who can adapt and improve the tool for diverse contexts. Ultimately, by pursuing these interconnected pathways of technological enhancement, operational validation, and ethical governance, this research can evolve into a trusted, scalable, and life-saving component of resilient public health infrastructure worldwide.

Scientific Ethics Declaration

* The authors declares that the scientific ethical and legal responsibility of this article published in EPSTEM journal belongs to the authors.

Conflict of Interest

* The authors declare that they have no conflicts of interest.

Funding

* There is no funding available for this study.

Acknowledgements or Notes

* This article was presented as an oral presentation at the International Conference on Technology (www.icontechno.net) held in Budapest/Hungary on February 05-08, 2026.

References

Kim, G. (2024). *Electric vehicle routing problem with states of charging stations*. *Sustainability, 16*(8), 3439.

Mardan, A. (2018). *Practical Node.js: Building real-world scalable web apps*. Apress.

Mladenović, N., & Hansen, P. (1997). Variable neighborhood search. *Computers & Operations Research, 24*(11), 1097–1100.

Moehrle, S., Bai, S., Müller, T., Munz, E., Trybushnyi, D., & Raskob, W. (2018). Web-based decision support system for emergency management: System architecture and enhancement possibilities. *International Journal of Disaster Risk Reduction, 27*, 314–325.

Polacek, M., Hartl, R. F., Doerner, K. F., & Reimann, M. (2004). A variable neighborhood search for the multi-depot vehicle routing problem with time windows. *Journal of Heuristics, 10*, 613–627.

Pradhan, B., Al-Najjar, H. A. H., Sameen, M. I., Tsang, I., & Alamri, A. M. (2020). Unseen land cover classification from high-resolution orthophotos using integration of zero-shot learning and convolutional neural networks. *Remote Sensing, 12*(10), 1670.

Solomon, M. M. (1987). Algorithms for the vehicle routing and scheduling problems with time window constraints. *Operations Research, 35*(2), 254–265.

Uzun, Y. (2017). *Development of variable neighborhood search techniques for continuous optimization problems* Doctoral dissertation. Selçuk University.

Uzun, Y., Arikan, H., & Tezel, G. (2016). Rule extraction from training artificial neural networks using variable neighbourhood search for Wisconsin breast cancer. *Journal of Multidisciplinary Engineering Science and Technology (JMEST), 3*, 2458–9403.

World Health Organization. (2026). *Health emergency and disaster risk management framework. World Health Organization*. Retrieved from https://www.who.int/publications/i/item/9789241516181

Yu, W., Lv, Y., Hu, C., Liu, X., Chen, H., Xue, C., & Zhang, L. (2018). Research of an emergency medical system for mass casualty incidents in Shanghai, China: A system dynamics model. *Patient Preference and Adherence, 12*, 207–222.

Author(s) Information

Yusuf Uzun	**Huseyin Arıkan**
Necmettin Erbakan University, Department of Computer Engineering, Faculty of Seydişehir Ahmet Cengiz Engineering, Konya, Türkiye	Necmettin Erbakan University, Department of Mechanical Engineering, Faculty of Seydişehir Ahmet Cengiz Engineering, Konya, Türkiye
ORCID iD: https://orcid.org/0000-0002-7061-8784	ORCID iD: https://orcid.org/0000-0003-1266-4982
*Contact e-mail: *yuzun@erbakan.edu.tr*	

*Corresponding authors contact email

To cite this article:

Uzun, Y., & Arikan, H. (2026). A variable neighborhood search-based GIS for real-time health service distribution to disabled and chronic patients in disaster and emergency scenarios. *The Eurasia Proceedings of Science, Technology, Engineering and Mathematics (EPSTEM), 39,* 19-29.