

MORSO for Multi-Objective Fire Station Location on Urban Road Networks: The Mosul Case

Ziadoon Mohand Khaleel^{1,*}, Safa Jawad Abed², Jalal Abdulkareem Sultan², Noor Marwan Ahmeed³

¹*Department of Petroleum Reservoir Eng., College of Petroleum and Mining Engineering, University of Mosul, Iraq*

²*Nineveh Agriculture Directorate, Iraq*

³*Department of Construction and Projects, University of Mosul, Iraq*

Abstract The effectiveness of urban fire response is heavily reliant on the fire station location strategy on a realistic road network. This paper presents a two-stage GIS framework for the additive fire station location problem on an urban road graph. Stage 1 relies on an economic location criterion to identify the number of new stations (N^*) needed. Stage 2 employs a reinforced multi-objective particle swarm optimizer (MORSO) to explore potential locations and simultaneously optimize (i) the maximum and (ii) the average nearest station network distance, computed using multi-source Dijkstra algorithms with careful treatment of disconnected components. Feasibility constraints (such as minimum distance and equity/coverage criteria) are used to guarantee the existence of implementable solutions. A test example on Mosul, Iraq, illustrates that adding three stations to the existing nine (total 12) leads to a significant improvement in accessibility: the best Pareto solution decreases the average distance from 3,474.67 m to 2,823.09 m (–18.8%) and the maximum distance from 12,595.07 m to 8,752.15 m (–30.5%), with further tail optimization (P90). Distances are also reported as travel-time proxy bands using an average speed of 35 km/h ($\approx 3.0/6.0/7.5$ km for 4/8/10 minutes): coverage increases from 52.60/89.02/93.64% to 64.16/93.06/97.11% for the $\leq 4/\leq 8/\leq 10$ -minute bands, respectively. Multi run bench marking and sensitivity analysis further support the robustness and practical usability of the proposed planning workflow. The resulting Pareto front and mapped layouts enable transparent efficiency risk trade offs for deployment.

Keywords Emergency-service accessibility; Fire-station siting; GIS road-network distances; Multi-objective optimization; Reinforced particle swarm optimization

AMS 2010 subject classifications 90B80, 90C29

DOI: 10.19139/soic-2310-5070-3297

1. Introduction

Siting emergency service facilities, such as fire stations or ambulance depots, is a classic optimization problem that directly affects citizen safety. Traditional models of facility location include the maximal covering location problem (MCLP), which aims to maximize the demand covered within a given time or distance budget [1], and the p -median model, which aims to minimize the average distance between demand points and facilities [2]. These classic models provided the theoretical foundations for facility location analysis. However, in modern facility location planning, planners are faced with multiple, sometimes competing, criteria such as minimizing emergency response times while keeping construction and operating expenses under control, leading to the development of multi-objective models that can clearly articulate trade-offs to planners [3]. For instance, He et al. developed a bi-objective model that simultaneously sought to minimize construction costs and maximize fire risk coverage, and used an epsilon-constraint approach together with a tailored genetic algorithm to obtain Pareto-optimal solutions that presented planners with a set of highly effective alternatives instead of a single optimal solution [4].

*Correspondence to: ziadoon Mohand Khaleel (Email: ziadoon.khaleel@uomosul.edu.iq). Department of Petroleum Reservoir Eng., University of Mosul. Mosul, Iraq.

Because of the NP-hardness of the multi-objective facility location problem, metaheuristics have been usually used for obtaining high-quality solutions within reasonable processing time. [5]. Of these, evolutionary algorithms have gained significant attention. Although introduced by Deb et al., [6], NSGA-II has been used in a wide range of applications, including land use planning [7] and post-disaster logistics and hybridized versions have improved convergence properties for difficult instances [8]. Decomposition algorithms, such as the multiobjective evolutionary algorithm based on decomposition (MOEA/D), have reformulated vector-valued objectives into sets of scalar subproblems that are solved collaboratively, often providing well-distributed approximations of the Pareto set [9]. Experiments have shown that dominance-based algorithms (such as NSGA-II) and decomposition-based algorithms (such as MOEA/D) are complementary; many recent studies have evaluated both to measure robustness [10]. In the context of emergency services, these evolutionary approaches remain popular benchmarks because of their ability to offer informed selection and diversity criteria while allowing standardized performance comparisons.

Apart from these evolutionary approaches, population-based swarm intelligence algorithms, such as particle swarm optimization (PSO), have also been explored for multi-objective facility location problems because of their simplicity, few hyperparameters, and fast convergence properties [11]. In multi-objective PSO algorithms, a set of candidate solutions and an external archive of non-dominated solutions are maintained. Xiang et al. [12] applied discrete multi-objective particle swarm optimization (MOPSO) to fire safety layout planning in construction mega-projects, addressing coverage, distance, response time, and cost simultaneously. In competitive scenarios, population-based evolutionary algorithms generally outperform simpler heuristics, and logistics research has also found NSGA-II and multi-objective differential evolution to offer better Pareto fronts than basic MOPSO variants [13]. Single-trajectory heuristics have also been explored; for example, Chen and Lai [14] proposed a multi-objective simulated annealing framework for rural emergency medical services (EMS) location-allocation and found it efficient. For example, Bolouri et al. [15] found that a genetic algorithm (GA) outperformed simulated annealing (SA) on a multi-objective fire station location-allocation problem in Tehran. These findings all indicate that the choice of algorithm is context-dependent: while evolutionary paradigms can offer robust global Pareto approximations, PSO-family algorithms can be considered a good choice when constraint-aware deployment planning and fast scenario exploration are important.

An increasing number of recent studies have utilized hybrids that either combine search methods or use multi-criteria decision making (MCDM) to cope with real-world complexity. A framework for a multi-objective genetic algorithm for finding urban gas stations under planning and safety considerations was provided by Yang et al., illustrating the capability of evolutionary search to handle competing risk and cost criteria [16]. To more effectively explore the extreme trade-off areas in disaster relief, Ransikarbum and Mason combined local search and repair operators with NSGA-II [8]. MCDM combinations are also prevalent; for example, Ransikarbum and Pitakaso combined fuzzy AHP with a multi-objective GA for sustainable biofuel network design [17], while Vargas-Santiago et al. developed a facility location problem using adaptable Pareto sets, which allow for post-optimization modifications through interactive analysis [18]. Hybrid swarm methods have also progressed; for example, Tanoumand et al. located emergency water reservoirs for fire-following-earthquake fire suppression using a metaheuristic-based method that combined hazard mapping and logistics considerations [19].

Meanwhile, advances in data and modeling have also increased realism. More recent uses include traffic, population, and hazard data to simulate operational realities [20]. Geographic Information Systems (GIS) allow analysis on road networks instead of Euclidean spaces, and early work showed how network travel times can lead to significant differences in coverage calculations [21, 22, 23]. GIS models have been employed for emergency facility location problems for many years; for instance, Habibi et al. proposed an analytic hierarchy process (AHP) and GIS approach for urban fire station location [22], while Brotcorne et al. examined ambulance relocation on realistic road networks [23]. Current applications involve high-resolution road networks and multiple spatial data sources. However, two traditional challenges remain in applying this technology: (i) road network planning with potential bottlenecks or disconnects, and (ii) generating alternatives that represent trade-offs between average accessibility and worst-case vulnerability.

With this context in mind, this paper proposes a two-stage, GIS-informed strategy for the fire station location problem in Mosul. In Stage 1, a cost benefit rule of thumb is used to estimate the number of additions. In Stage 2,

we propose a reinforced multi-objective PSO (MORSO) algorithm, which extends traditional velocity updates with a reinforcement component that kicks in whenever personal bests are first discovered, to promote more aggressive search around promising areas while preserving diversity. The algorithm works directly with the real road network (harmonized to UTM (Universal Transverse Mercator) Zone 38N) and uses multi-source Dijkstra search to compute shortest path distances, with a conservative fallback to disconnected components. Additionally, an optional spatial equity guard clause is provided to ensure a minimum coverage threshold for specified sub-regions. In a case study for Mosul with an additive model (optimizing a few new locations within the existing infrastructure), the findings show significant improvements in average and early time network accessibility, as well as how road network topology can control worst-case distances an operational insight that is consistent with GIS location theory [21, 22, 23].

This paper offers the following contributions. First, it presents a clear, GIS-informed two-stage procedure for additive fire-station planning on real-world urban road networks, decoupling the questions of how many stations to build and where to build them. Second, it presents a reinforced multi-objective PSO variant (MORSO) with deployability-focused characteristics, such as diversity-protecting archiving and optional feasibility/equity guards enabling policy-compliant designs. Third, it offers a case study of the city of Mosul, where the benefits of improved accessibility and proxy coverage times are quantified, although road network connectivity is identified as a constraint on worst-case accessibility

The paper is organized as follows. Section 2 defines the multi-objective fire station location problem on a road network and introduces the equity guard constraint. Section 3 describes the MORSO algorithm and its implementation. Section 4 describes the case study of the city of Mosul and the computational results, including comparisons with existing baselines. Section 5 discusses planning implications and limitations. Section 6 provides conclusions and suggestions for future research.

2. Problem Statement and Model Formulation

2.1. Operational context and motivation

The city of Mosul, in Iraq, is currently covered by a limited number of fire stations compared to the size of the city and the road network. A recent GIS analysis audit has shown that there are coverage gaps and an unbalanced distribution of services, with a high concentration of services in the city center and a lack of accessibility in the surrounding areas, especially in the west bank [24]. With respect to response times, the current standard requires a response time of 10 minutes, while international standards require more accurate first-due targets; for example, NFPA 1710 requires a travel time of 4 minutes for the first-due engine in certain conditions [25]. These realities motivate a network-based planning model that (i) evaluates accessibility using road travel on the network graph rather than Euclidean distance, (ii) explicitly quantifies the trade-off between typical accessibility and worst-case exposure, and (iii) supports policy constraints (e.g., minimum coverage floors for under-served subareas) to counter documented spatial imbalance [24, 27].

2.2. Road-network representation, sets, and distances

Let the navigable road network be represented by a weighted graph (Eq. 1)

$$G = (V, E) \quad (1)$$

where V is the set of road nodes and E is the set of road segments. Each edge $(u, v) \in E$ is assigned a nonnegative weight w_{uv} equal to its length in meters after projecting GIS layers to a metric coordinate reference system (CRS), e.g., UTM Zone 38N [32, 33].

Demand is represented by a set I of demand points (e.g., neighborhood representative-point centroids). Each demand $i \in I$ is snapped to the closest road node $c(i) \in V$ so that travel distances are evaluated on the traversable network [32]. Candidate locations for new stations are restricted to a feasible subset $J \subseteq V$ (navigable nodes), optionally pre-screened to reduce search complexity (e.g., nodes near demand clusters and/or high-degree hubs)

[32, 33]. Let $B \subseteq V$ denote the set of baseline (existing) station nodes, which remain fixed throughout the optimization.

Shortest-path distance. For any demand $i \in I$ and facility node $s \in (J \cup B)$, define the network travel distance (Eq.2)

$$d_{is}^{\text{net}} = \text{dist}_G(c(i), s) \quad (2)$$

computed as the shortest-path distance on G (e.g., Dijkstra) [32, 33].

Disconnected-network handling (fallback penalty). If s is unreachable from $c(i)$ due to graph disconnection, we use a conservative fallback distance to avoid optimistic bias and to reflect missing or disconnected links in GIS road data. Let d_{is}^{uc} denote the Euclidean (straight-line) distance between i and s in the projected CRS. The effective distance is defined as (Eq.3)

$$d_{is} = \begin{cases} d_{is}^{\text{net}}, & \text{if a path exists in } G, \\ \gamma d_{is}^{\text{uc}} + \delta, & \text{otherwise.} \end{cases} \quad (3)$$

where $\gamma > 1$ inflates straight-line distance to reflect detours and $\delta \geq 0$ is an additive penalty (implementation details are provided in Section3.4) [32, 33]. This definition ensures unreachable pairs are not treated as artificially ‘‘close’’ and prevents disconnected areas from being under penalized.

Standards-based reporting bands. To make sense of distances, distances are converted to travel time bands based on an average emergency speed v (km/h) (Section 4). If d is in meters, the travel time in minutes is (Eq.4)

$$t(d) = \frac{d}{1000 v} \times 60 \quad (4)$$

We report coverage within 4-, 8-, and 10-minute bands using distance thresholds $\tau_4, \tau_8, \tau_{10}$ obtained from the same speed assumption (e.g., $\tau_k = 1000 v k/60$) [25, 28]. In the Mosul environment, a reasonable approximation of the 4-minute belt can be considered as $\tau_4 \approx 3000$ m for common speed ranges, and the values of τ_8, τ_{10} are obtained consistently [24, 25].

2.3. Decision variables and additive siting structure

The problem is formulated as an additive problem, where the initial stations are not changed, and a few new stations are added. For each siting option $j \in J$, a binary decision variable $y_j \in \{0, 1\}$, where $y_j = 1$ represents the addition of a new station at siting option j . The set of active stations is given by (Eq.5).

$$S = B \cup \{j \in J : y_j = 1\} \quad (5)$$

Denote by N^* the number of new stations to be added, determined in Stage 1 (Section2.5). The constraint on the number of new stations is given by (Eq.6).

$$\sum_{j \in J} y_j = N^*, \quad y_j \in \{0, 1\} \forall j \in J \quad (6)$$

For each demand point $i \in I$, the minimum network distance to the nearest station is defined as (Eq.7)

$$r_i(S) = \min_{s \in S} d_{is} \quad (7)$$

2.4. Objective functions (max, mean) and constraint options

We adopt a bi-objective formulation that minimizes both worst-case and average accessibility (Eq.8):

$$\min(f_1(S), f_2(S)) = \left(\max_{i \in I} r_i(S), \frac{1}{|I|} \sum_{i \in I} r_i(S) \right) \quad (8)$$

subject to Eqs.(5)-(7) and feasibility constraints. The first objective f_1 (minimax) captures extreme-case risk (p -center logic), protecting the most delayed neighborhoods, while the second objective f_2 minimizes the mean distance to improve overall system efficiency [27]. Optimizing the pair reveals a Pareto front of deployable plans and avoids the brittleness of single-threshold coverage models, where small network changes near the threshold can cause large discontinuities in performance.

Feasibility and spacing (optional). Candidate locations are limited to feasible locations within the administrative boundary and on the traversable road network. To avoid infeasible clustering, a minimum spacing constraint can optionally be imposed between any two chosen additions (Eq.9):

$$\|x_j - x_k\|_2 \geq \delta_s, \quad \forall j \neq k \text{ (selected additions)} \quad (9)$$

where $x_j \in \mathbb{R}^2$ represents the projected location of candidate j and $\delta_s > 0$ is planner-specified separation distance. This constraint can optionally be applied between new locations and baseline stations as well.

Equity/coverage floor (optional). To account for known imbalances (e.g., west bank under-service), the model can optionally enforce a policy-specified minimum coverage fraction for a specified subset $W \subseteq I$. For instance, requiring a floor on the coverage fraction within the 8-minute zone requires (Eq.10)

$$\frac{1}{|W|} \sum_{i \in W} \mathbf{1}\{r_i(S) \leq \tau_8\} \geq \theta, \quad \theta \in [0, 1] \quad (10)$$

where $\mathbf{1}\{\cdot\}$ is the indicator function. This type of direct equity constraint is well understood in facility location modeling and offers a controllable way to balance efficiency and equity considerations across sub-areas [27].

2.5. Stage-1 economic sizing to determine N^*

To avoid arbitrary choices for the number of additions, Stage 1 employs a clear economic sizing rule that trades off an annualized per-station cost SC with a citywide loss proxy TLC (e.g., marginal losses due to coverage inadequacies). The required number of additions is estimated as (Eq.11)

$$N^* = \text{round} \left(\sqrt{\frac{\alpha \cdot TLC}{SC}} \right) \quad (11)$$

where $\alpha > 0$ is a policy lever controlling the cost risk trade-off. The scaling $N^* \propto \sqrt{TLC/SC}$ stabilizes planning against extreme values and provides a simple scenario-analysis knob: increasing α corresponds to more conservative planning (higher value placed on reducing losses), while decreasing α reflects tighter budget emphasis. Once N^* is set, Stage 2 explores the Pareto set of spatial layouts for that fixed addition count, and sensitivity analysis can be performed by running multiple α (or N^*) scenarios. Based on consultation with the Mosul fire department, we used the scenario parameters $SC = 50,000$ USD, $TLC = 10,000$ USD, and $\alpha = 50$, which yield $N^* = 3$.

2.6. Modeling scope and alignment with standards

This study focuses on the (max, mean) distance objectives because they provide an interpretable and policy-relevant balance between worst-case protection and average efficiency on urban road networks [27]. While the present formulation captures accessibility and supports optional equity constraints, reliability effects (busy fraction and unit availability) can be incorporated when dispatch and workload data are available, using established availability-aware models (e.g., Maximum Expected Covering Location Problem (MEXCLP)-type formulations) [26]. For reporting and international comparability, we include NFPA-oriented 4-minute first-due travel-time guidance and 8- and 10-minute bands commonly used in operational policy discussions [25, 28]. This reporting does not replace the distance-based optimization; rather, it provides a consistent, scenario-based interpretation layer for decision makers.

3. MORSO Algorithm and Implementation

3.1. Two-stage workflow and design rationale

We adopt a two-stage workflow tailored to additive fire-station planning on a real urban road network. Stage 1 determines the required number of new stations N^* using the economic sizing rule (Section 2.5). Stage 2 then solves the spatial siting problem on the road graph using a reinforced multi-objective particle swarm optimizer (MORSO) to minimize both (i) the worst-case (maximum) nearest-station network distance and (ii) the mean nearest-station network distance. This bi-objective design explicitly reveals the efficiency risk trade-off and returns a Pareto set of deployable plans rather than a single solution [27]. Network (not Euclidean) distances are used to reflect operational travel constraints and enable standards-oriented reporting bands (Section 2.2) [25]. The proposed MORSO follows the archive based MOPSO family (external archive, dominance ranking, diversity maintenance, leader selection) [6], and introduces a reinforcement term that intensifies successful particle moves; when the reinforcement factor is set to zero, the method reduces to a standard archive-based MOPSO baseline [29, 30].

3.2. GIS preprocessing and network-distance evaluation

The road layer is converted into a weighted undirected graph $G = (V, E)$, where each edge $(u, v) \in E$ has a weight w_{uv} equal to the segment length (meters) in a metric CRS (e.g., UTM Zone 38N). Demand is represented by a set of points I (e.g., neighborhood representative centroids), each snapped to its nearest graph node $c(i) \in V$. Candidate facility locations are restricted to a filtered subset $J \subseteq V$ (e.g., nodes near demand clusters and/or high-degree hubs) to keep the search discrete and computationally tractable [32, 33]. Existing (baseline) stations form a fixed set $B \subseteq V$.

For any demand $i \in I$ and facility node $s \in (B \cup J)$, the network distance d_{is}^{net} is computed as the shortest-path distance on G from $c(i)$ to s , using multi-source Dijkstra during objective evaluation [32]. If s is unreachable from $c(i)$ due to graph disconnection, we apply the conservative fallback distance defined in Eq. 3 to avoid optimistic bias (Section 2.2) [33]. Standards-oriented reporting uses time-proxy distance thresholds $\tau_4, \tau_8, \tau_{10}$ derived from the assumed average response speed (Section 2.2) [25, 28].

3.3. Additive siting representation and objective evaluation

Let $B \subseteq V$ denote baseline (existing) station nodes (fixed). Decision variables are binary indicators $y_j \in \{0, 1\}$ for each candidate $j \in J$, where $y_j = 1$ means a new station is added at node j . The total number of additions is fixed to N^* from Stage 1 (Eq. 12):

$$\sum_{j \in J} y_j = N^*, \quad y_j \in \{0, 1\} \forall j \in J \quad (12)$$

The active station set under an additive plan is (Eq. 13)

$$S = B \cup \{j \in J : y_j = 1\} \quad (13)$$

For each demand $i \in I$, the nearest-station distance is (Eq. 14):

$$r_i(S) = \min_{s \in S} d_{is} \quad (14)$$

We optimize the bi-objective vector (Eq. 15)

$$\min(f_1(S), f_2(S)) = \left(\max_{i \in I} r_i(S), \frac{1}{|I|} \sum_{i \in I} r_i(S) \right) \quad (15)$$

The first objective f_1 is a minimax (p -center) criterion that protects the worst-served neighborhood, while f_2 improves system-wide efficiency by reducing the average distance [27]. When reliability/availability data are available, the formulation can be extended with availability-aware coverage or double-coverage criteria (e.g., MEXCLP-type models) [26].

Practical constraints. Candidate nodes are restricted to feasible locations inside the administrative boundary and on the traversable road graph. Optionally, a minimum spacing constraint is enforced to avoid impractical clustering of new sites. If $x_j \in \mathbb{R}^2$ denotes the coordinate of candidate j , then for any two selected additions $j \neq k$ (Eq.16):

$$\|x_j - x_k\|_2 \geq \delta \quad (16)$$

where $\delta > 0$ is a planner chosen separation distance. The same spacing rule can be enforced between new sites and baseline stations.

Equity/coverage floors (optional). Equity objectives can be imposed via an ε -constraint. For a designated subset $W \subseteq I$ (e.g., west-bank demands), requiring that at least a fraction $\theta \in [0, 1]$ be served within the 8-minute band yields (Eq.17):

$$\frac{1}{|W|} \sum_{i \in W} \mathbf{1}\{r_i(S) \leq \tau_8\} \geq \theta \quad (17)$$

where $\mathbf{1}\{\cdot\}$ is the indicator function. This turns an equity goal into an operational constraint and limits purely efficiency-driven plans from under-serving disadvantaged or peripheral subareas [27].

3.4. Reinforced multi-objective PSO (MORSO)

3.4.1. Continuous encoding and Top- N^* projection

The fixed cardinality selection is encoded using a continuous score vector $x \in [0, 1]^{|J|}$. A binary selection $y(x)$ is decoded by choosing the indices of the Top- N^* components of x (highest scores). Let π be the permutation of indices that sorts x in descending order, i.e., $x_{\pi(1)} \geq x_{\pi(2)} \geq \dots$. The decoded decision is (Eq.18):

$$y_{\pi(k)}(x) = \begin{cases} 1, & k = 1, \dots, N^*, \\ 0, & \text{otherwise.} \end{cases} \quad (18)$$

A repair operator $\mathcal{R}(\cdot)$ is then used to enforce feasibility (boundary and spacing). This projection is a standard part of PSO-style discrete selection: the swarm searches in continuous space, but the decoded solution is always a fixed-size set.

3.4.2. Archive-based multi-objective search and leader (“gbest”) definition

Unlike single-objective PSO, multi-objective optimization has no unique global best. Therefore, we maintain an external archive \mathcal{A} of non-dominated solutions, updated by Pareto dominance after each evaluation. Diversity within \mathcal{A} is maintained using a spread measure such as crowding distance; leaders are then sampled from \mathcal{A} with a bias toward less-crowded regions to balance convergence and diversity [6]. In this paper, “gbest” is defined as a sampled archive leader, not a single fixed point, consistent with standard MOPSO practice [35].

3.4.3. Velocity and position updates with reinforcement

For particle i , let $x_i^{(t)}$ and $v_i^{(t)}$ be position and velocity at iteration t . Let $p_i^{(t)}$ denote the particle’s personal-best position (according to the constraint-aware dominance rule in (Section3.4.4)), and let $g^{(t)}$ be the leader sampled from the archive \mathcal{A} . The reinforced update is

Stage (1) (economic sizing). We estimate $N^* = \text{round}\left(\sqrt{\alpha \cdot TLC/SC}\right)$ from annualized station cost SC , a city-wide loss proxy TLC , and a policy lever $\alpha > 0$ (Section 2.5; Eq.11). This provides transparent, scenario-testable counts prior to spatial optimization.

Stage (2) (MORSO search). A solution is encoded as a continuous score vector $x \in [0, 1]^{|J|}$; the Top- N^* components define the selected nodes (yielding a fixed-cardinality mask), with repair to enforce spacing/boundary. The swarm evolves via Eqs.(19–20):

$$v_i^{(t+1)} = \omega v_i^{(t)} + c_1 r_1^{(t)} (p_i^{(t)} - x_i^{(t)}) + c_2 r_2^{(t)} (g^{(t)} - x_i^{(t)}) + \eta \mathbf{1}\left\{f\left(x_i^{(t)}\right) \prec f\left(p_i^{(t)}\right)\right\} (p_i^{(t)} - x_i^{(t)}) \quad (19)$$

$$x_i^{(t+1)} = \Pi_{[0,1]} \left(x_i^{(t)} + v_i^{(t+1)} \right) \quad (20)$$

Here ω is the inertia weight; c_1, c_2 are cognitive and social gains; $r_1^{(t)}, r_2^{(t)} \sim U(0, 1)^{|J|}$ are random vectors; $\Pi_{[0,1]}(\cdot)$ clips components into $[0, 1]$; $\eta \geq 0$ is the reinforcement coefficient; $\mathcal{S}(x)$ denotes the decoded and repaired plan obtained from x ; and \prec denotes Pareto dominance on the objective vector (under feasibility handling). When $\eta = 0$, (Eq.19) reduces to the standard archive-based MOPSO update [29]. Optionally, a low-rate uniform mutation can be applied to a small subset of components of x to preserve exploration.

3.4.4. Constraint handling (spacing and equity)

Decoded selections are repaired to satisfy fixed cardinality and spacing. Policy constraints such as (Eq.17) are handled using an ε -constraint principle: solutions are compared first by total constraint violation, then by objectives. Let $v(S) \geq 0$ denote total violation, for example (Eq.21)

$$v(S) = v_{\text{eq}}(S) + v_{\text{sp}}(S) \quad (21)$$

where

$$v_{\text{eq}}(S) = \max \left(0, \theta - \frac{1}{|W|} \sum_{i \in W} \mathbf{1}\{r_i(S) \leq \tau_8\} \right), \quad v_{\text{sp}}(S) = (\text{aggregate spacing violation}).$$

A solution S_a is preferred to S_b if

$$S_a \prec_{\varepsilon} S_b \iff [v(S_a) < v(S_b)] \text{ or } [v(S_a) = v(S_b) \text{ and } (f_1(S_a), f_2(S_a)) \prec (f_1(S_b), f_2(S_b))]. \quad (22)$$

This ensures feasible solutions dominate infeasible ones while preserving Pareto ranking among feasible solutions [35].

3.4.5. Pseudocode (Stage 2: MORSO)

Algorithm 1. MORSO for additive siting on a road network

Input: candidate set J , baseline stations B , demand set I , additions N^* , parameters ω, c_1, c_2, η , mutation rate p_m , spacing δ , optional policy constraints.

Output: archive \mathcal{A} of non-dominated solutions.

1. Initialize swarm positions $\{x_i\} \subset [0, 1]^{|J|}$ and velocities $\{v_i\}$. Set archive $\mathcal{A} \leftarrow \emptyset$.
2. For each particle i :
 - (a) Unpack x_i via Top- N^* projection (Eq.18) to obtain y_i , construct S_i via(Eq.13), and apply repair $\mathcal{R}(\cdot)$.
 - (b) Evaluate $f_1(S_i), f_2(S_i)$ (Eq.15) and constraint violation $v(S_i)$.
 - (c) Set personal best $p_i \leftarrow x_i$.
 - (d) Update archive \mathcal{A} by inserting non-dominated solutions and removing dominated ones; update diversity measure (e.g., crowding distance).
3. For $t = 1, \dots, T$:
 - (a) For each particle i :
 - i. Sample leader $g^{(t)}$ from \mathcal{A} using diversity-biased selection.
 - ii. Update $v_i^{(t+1)}$ and $x_i^{(t+1)}$ using Eqs.19–20. Apply mutation with probability p_m .
 - iii. Decode/repair $x_i^{(t+1)} \mapsto S_i$. Evaluate $f_1(S_i), f_2(S_i), v(S_i)$.
 - iv. Update personal best using the constraint-aware dominance rule (Eq.22).
 - v. Update archive \mathcal{A} .
4. Return \mathcal{A} .

3.5. Output set, knee selection, and reporting metrics

The proposed algorithm provides an archive of non-dominated plans approximating the Pareto front. To obtain a single feasible plan for reporting, a knee-point solution selection method is used to determine a compromise solution with the maximum perpendicular distance to the line connecting the two most distant points in objective space [34]. Let $z(S) = (f_1(S), f_2(S))$. Let z^{\min} and z^{\max} the two most extreme archive points (minimizing each objective). For each non dominated $z \in \mathcal{A}$, the knee score is calculated as the point to line distance (Eq.23):

$$k(z) = \frac{\|(z^{\max} - z^{\min}) \times (z^{\min} - z)\|}{\|z^{\max} - z^{\min}\|} \quad (23)$$

and the corresponding knee solution is selected as $z^* = \arg \max_{z \in \mathcal{A}} \kappa(z)$. (In 2D, the numerator is simply the absolute value of the scalar cross-product magnitude.)

The performance metrics reported in Section 4 include f_1 (maximum distance), f_2 (mean distance), distribution summaries (e.g., median and $P90$), and coverage metrics for τ_4 τ_8 τ_{10} , allowing direct operational interpretation and comparison with standards [25, 28]. For multi-run comparisons, the 2D hypervolume (HV) indicator with respect to a reference point is also reported as a standard Pareto indicator [31].

3.6. Parameter settings, computational complexity, and reproducibility

Default search parameters follow common practice in PSO/MOPSO studies and were validated by preliminary trials: swarm size 50–60, iterations 150–250, $\omega \in [0.60, 0.75]$, $c_1 = c_2 \in [1.4, 1.8]$, reinforcement $\eta \in [0.15, 0.30]$, and mutation $p_m \approx 0.01$ – 0.03 [6, 29]. All experiments enforce a metric CRS, snapping demands to the road graph, and fixed-cardinality decoding to ensure reproducible and operationally interpretable outputs [32].

Let $|V|$ and $|E|$ denote the number of nodes and edges in the road graph, and $|J|$ the number of candidate nodes. For each decoded plan S , objective evaluation requires multi-source shortest-path computation from the selected stations to all demand nodes, with time complexity $O(|E| \log |V|)$ using a standard priority-queue implementation (multi-source Dijkstra can be implemented by pushing all sources initially with distance 0) [33]. The Top- N^* projection requires sorting $|J|$ scores, i.e., $O(|J| \log |J|)$. Therefore, the per-iteration cost is approximately (Eq.24)

$$O(\text{swarm} \cdot (|E| \log |V| + |J| \log |J|)) \quad (24)$$

which is practical for city scale road graphs when $|J|$ is filtered and N^* is small.

4. Case Study, Results, Algorithmic Comparison, and Sensitivity Analysis

4.1. Study area and datasets

The case study focuses on Mosul, Iraq. Demand points are constructed from representative-point centroids of a neighborhood (administrative areas) polygon layer. The road network is derived from a road center line GIS layer and converted into a weighted urban road graph. Existing fire stations are taken from a point facilities layer. To ensure metric consistency, all GIS layers are projected into UTM Zone 38N (EPSG:32638), which is a common practice in network-based facility location studies [21, 22, 23].

The resulting road graph contains approximately 64k nodes and 76k edges and includes multiple disconnected components an artifact that can arise in real urban GIS data, especially in cities affected by disruption or incomplete mapping. Network distances are evaluated using multi-source shortest paths (multi-source Dijkstra) on the road graph. For any demand that becomes unreachable from the selected set of stations due to topological disconnection, a conservative fallback rule is applied to prevent optimistic bias in reporting (as discussed in Section 3).

4.2. Experimental setup and reporting protocol

The proposed framework follows two stages (Section 3).

Stage 1 (economic sizing). The number of additional stations N^* is estimated using the economic sizing rule (Eq.11), based on the ratio of long-term loss cost to station cost parameters. Using the scenario parameters $SC = 50,000$ USD, $TLC = 10,000$ USD, and $\alpha = 50$ (obtained via consultation with the Mosul fire department), (Eq.11) yields $N^* = 3$, i.e., three new stations are added.

Stage 2 (additive multi-objective optimization). In additive mode, the optimizer selects N^* new station locations from a discrete set of candidate nodes while keeping the baseline stations fixed. The objectives are minimized jointly as follows:

1. Maximum nearest-station network distance (tail-risk / worst-case protection).
2. Mean nearest-station network distance (system-wide efficiency).

To ensure deployable solutions, feasibility constraints are added (e.g., minimum station separation and equity floor constraint to ensure minimum level of service presence in key partitions of the city). For operational purposes, distances are also expressed in travel-time proxy bands based on assumed mean emergency vehicle speed of 35 km/h, resulting in distance thresholds of about 3.0 km/6.0 km/7.5 km for 4/8/10 minutes, respectively (Table 1). This distance-to-time conversion is used for reporting purposes only and does not change the objective of optimization, which is still based on network distance [21, 22, 23].

Table 1. Main experimental settings (Stage 2, additive optimization)

Setting	Value
Baseline stations (fixed)	9
Additions N^*	3
Total stations after plan	12
Candidate nodes ($ J $)	99
Objectives	Minimize (max distance, mean distance)
Constraints (examples)	Minimum spacing; equity floor
Reporting speed proxy	35 km/h
Time proxy bands	$\leq 4 / \leq 8 / \leq 10$ minutes ($\approx 3.0/6.0/7.5$ km)

4.3. Baseline vs. proposed additive plan (9 + 3)

(Table 2) presents a comparison between the baseline solution (9 stations) and the proposed additive solution (9 + 3 = 12). All distances are measured in network meters. Coverage is measured as the percentage of demands for which the nearest station distance is within the proxy thresholds for 4/8/10 minutes.

Table 2. Baseline vs. proposed plan (9 + 3) on the Mosul road network

Metric	Baseline	Proposed (selected)	Change (Proposed – Baseline)
Max distance (m)	12,595.07	8,752.15	-3,842.92
Mean distance (m)	3,474.67	2,823.09	-651.58
Median distance (m)	2,870.11	2,373.68	-496.43
P90 distance (m)	6,336.03	5,188.96	-1,147.07
Coverage ≤ 4 min (%)	52.60	64.16	+11.56 pp
Coverage ≤ 8 min (%)	89.02	93.06	+4.05 pp
Coverage ≤ 10 min (%)	93.64	97.11	+3.47 pp

The proposed solution results in a systematic change in the distance distribution to smaller values. The average nearest station distance is reduced by about 18.8% (3,474.67 m \rightarrow 2,823.09 m), and tail accessibility is enhanced, with the 90th percentile (P90) reduced by about 1.15 km. The early response distance band (≤ 4 minutes), which is operationally important, is increased by 11.56 percentage points, indicating that a much larger proportion of

neighborhoods will be accessible within a short travel time after the opening of three new stations. Notably, the maximum nearest station distance is reduced by 30.5% (12,595.07 m \rightarrow 8,752.15 m), indicating that at least one of the new stations is placed in a region that was hitherto not well-covered in network terms.

4.4. Pareto front and selection of a deployable compromise plan

The optimizer provides a Pareto set of non-dominated solutions that capture the trade-off between mean distance and maximum distance (Fig. 1). This approach is well-suited to public safety planning, where different actors may have different priorities regarding risk scenarios, potentially preferring either enhanced “typical” accessibility or explicit support for the most underserved neighborhoods.

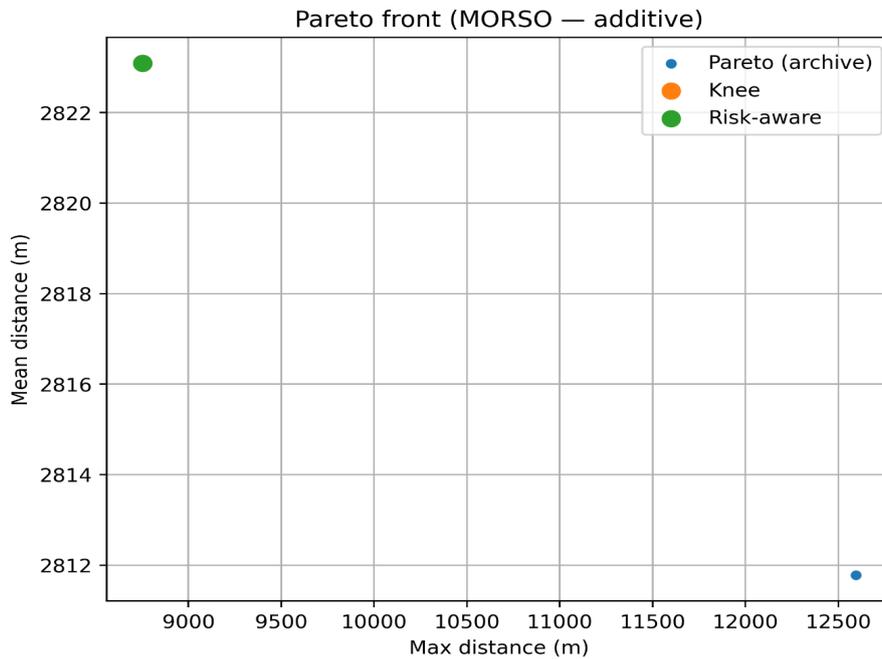


Figure 1. Pareto front (MORSO—additive).

One way to proceed with solution selection is to consider a compromise solution (such as a knee solution) that provides a large improvement in mean/median accessibility without compromising worst-case risk exposure, as common in multi-objective decision support [34]. In the Mosul problem, the selected solution provides a simultaneous improvement in both maximum and mean accessibility metrics (Table 2), thus representing a policy-feasible compromise solution.

4.5. Spatial layout and coverage effects

The spatial layout of the baseline and proposed additions is examined at the administrative level and the road network (Fig 2). The selected additions are not concentrated around existing stations; instead, they are spread out to minimize coverage gaps in underserved corridors while satisfying spatial constraints. This layout supports the noted improvement in early-time coverage (≤ 4 minutes) and the decrease in the P90 metric (Table 2).

Coverage improvements by time-proxy band are summarized visually (Fig. 3):

4.6. Numerical comparison with established multi-objective algorithms

To measure the performance of the algorithms over multiple runs, we compare the proposed MORSO algorithm with two well-known multi-objective baselines: a traditional MOPSO variant (obtained by setting the reinforcement

Layouts: administrative areas, roads, baseline and selected additions

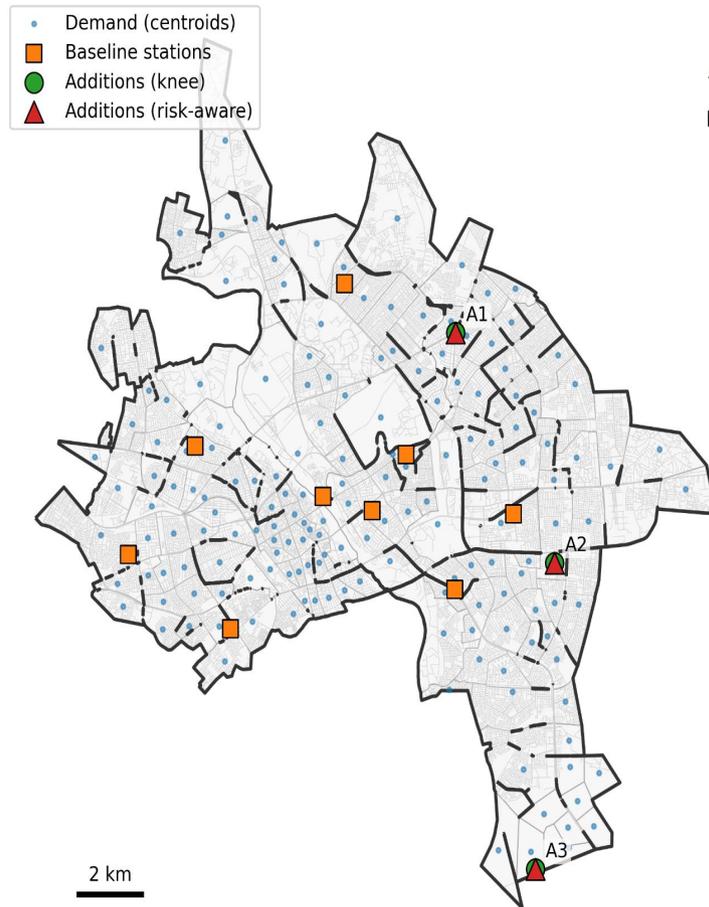


Figure 2. Layouts: administrative areas, roads, baseline stations, and selected additions.

weight $\eta = 0$) and NSGA-II. All algorithms were compared over the same candidate solutions, constraints, and demand/network distance calculations, with ten independent runs for each setting. Table 3 shows the mean and standard deviation of the hypervolume (HV), inverted generational distance (IGD), knee solution objectives, and runtime.

MORSO provides competitive Pareto solutions with the additional benefit of a fast operational solution that is always deployable and constraint-compliant at very low computational cost, allowing for fast scenario analysis and sensitivity studies on realistic road graphs. Note that the knee solutions of the PSO-family algorithms (MOPSO/MORSO) provide better worst-case protection (minimum maximum distance) under the given setting, reflecting the public-safety planning focus. NSGA-II provides better global front indicators (HV/IGD) but with much higher runtimes and knee solutions that improve worst-case distance at the cost of marginally better mean distance. For the case study, MORSO performs equally well as the baseline MOPSO algorithm on aggregate quality metrics, indicating that the reinforcement does not harm Pareto quality while maintaining the proposed mechanisms for feasibility and decision support for deployable solutions. All experiments were run on a laptop computer with an Intel Core i7-6820HQ @ 2.70 GHz CPU and 16 GB RAM under 64-bit Windows. The code was written in Python and run from the PyCharm IDE, with runtimes as reported in Table 3.

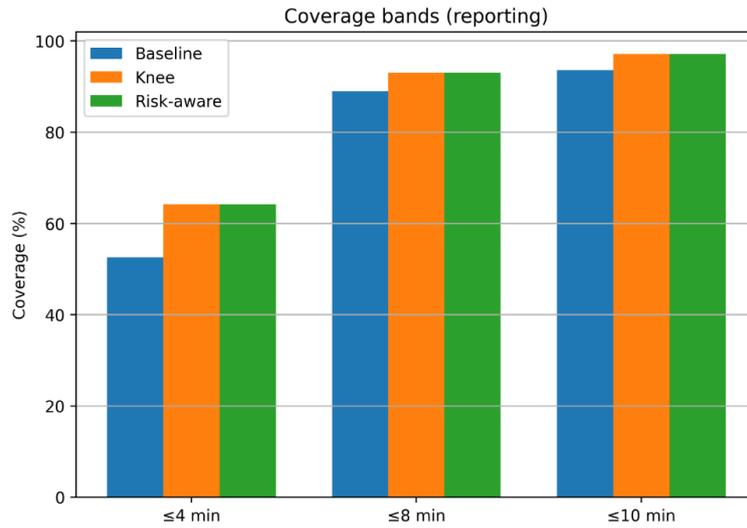


Figure 3. Coverage bands (reporting): baseline vs. selected solutions.

In this case, the reinforcement term does not affect the Pareto set obtained under the fixed evaluation budget significantly, and thus the similarity between MORSO and the $\eta = 0$ baseline. However, MORSO has a slightly smaller average runtime (Table 3), due to implementation details such as early feasibility checks and smaller constraint and archive update overheads, while retaining the same evaluation pipeline and solution feasibility.

Table 3. Algorithm comparison over repeated runs (mean \pm std).

Algorithm	Runs	HV (mean \pm std)	IGD (mean \pm std)	Knee max-distance (m) mean \pm std	Knee mean-distance (m) mean \pm std	Runtime (s) mean \pm std
MOPSO	10	1.077e6 \pm 6.873e4	81.21 \pm 59.94	7,754.38 \pm 583.97	2,826.40 \pm 69.75	6.25 \pm 0.55
MORSO	10	1.077e6 \pm 6.873e4	81.21 \pm 59.94	7,754.38 \pm 583.97	2,826.40 \pm 69.75	6.13 \pm 0.61
NSGA-II	10	1.115e6 \pm 6.063e4	30.79 \pm 50.47	8,178.35 \pm 2.05	2,765.72 \pm 20.25	49.57 \pm 1.74

4.7. Sensitivity analysis (speed, additions N^* , and reinforcement/constraints)

A sensitivity analysis assesses the validity of the results against important modeling assumptions (Table 4):

1. **Speed assumption (reporting only).** Because optimization is based on network distances, the travel speed assumption impacts only the reported coverage percentages in time bands, not the optimized placement location. (Figure 4) illustrates that the ≤ 4 -minute coverage is very sensitive to the assumed average speed.
2. **Number of additions N^* (captures economic sizing uncertainty).** Varying N^* offers a convenient proxy for uncertainty in Stage-1 economic sizing parameters (SC , TLC , etc.). (Figure 5) illustrates a monotonic gain in the median mean-distance as N^* varies from 2 to 4.
3. **Reinforcement weight and feasibility controls (η , spacing, equity floor).** For the explored ranges, the medians are robust, and thus moderate policy constraints do not reverse the overall results.

The results confirm that (i) the addition of more stations always leads to an improvement in typical accessibility (mean distance), and (ii) the reported time-band coverage is highly dependent on speed; hence, time proxies should be considered as a form of scenario reporting rather than actual travel times. The major results, based on improvements in network distances, are robust to variations in parameters.

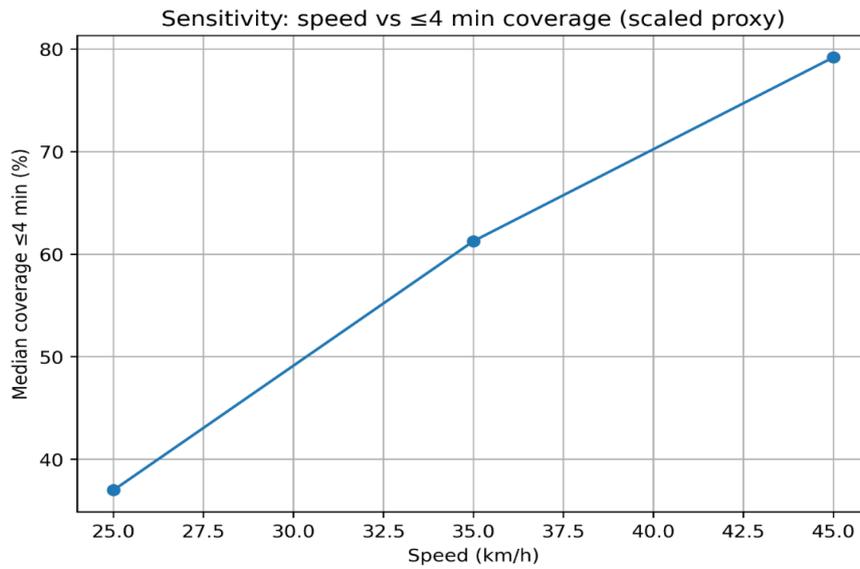


Figure 4. Sensitivity: speed vs. median coverage within ≤ 4 minutes (scaled proxy).

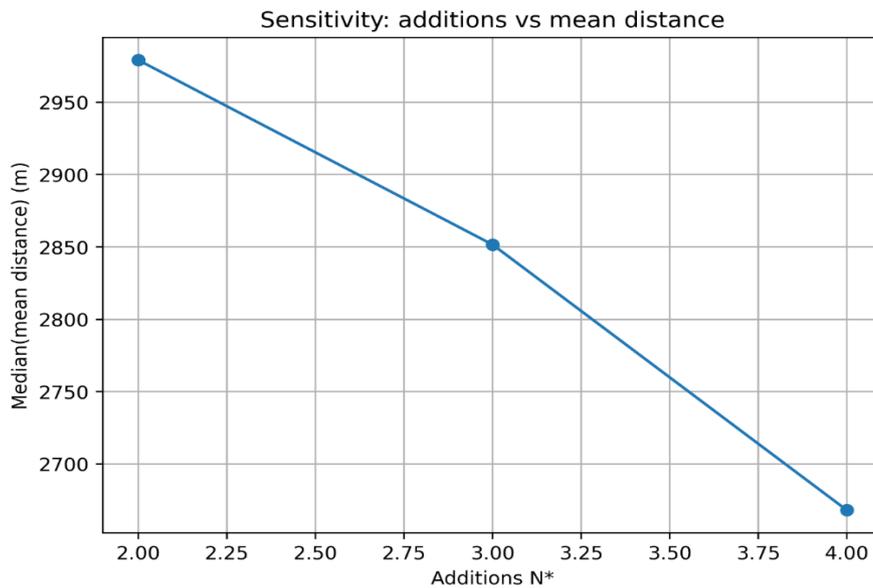


Figure 5. Sensitivity: additions N^* vs. median(mean distance).

5. Discussion and Practical Implications

5.1. Why the additive plan improves early-time coverage and mean accessibility

The difference between the baseline and the proposed plan (Table 2) shows that adding $N^* = 3$ stations makes a big difference in the service area, as it reduces the mean and median nearest station distances by about 652 m and 496 m, respectively, while it also increases the area of ≤ 4 -minute coverage by 11.56 percentage points. This is a

Table 4. Sensitivity summary (median indicators).

Factor	Setting	Reported indicator (median)
Additions N^* ($v = 35$ km/h)	$N^* = 2$	Mean distance $\approx 2,979$ m
	$N^* = 3$	Mean distance $\approx 2,852$ m
	$N^* = 4$	Mean distance $\approx 2,668$ m
Speed v ($N^* = 3$)	25 km/h	Coverage ≤ 4 min $\approx 37.0\%$
	35 km/h	Coverage ≤ 4 min $\approx 61.3\%$
	45 km/h	Coverage ≤ 4 min $\approx 79.2\%$

common property of additive facility location problems in urban road networks, where a few stations, located in a carefully chosen strategic way, have a strong impact on the “short-trip” part of the distribution, as most of the demands are located along corridors and local collectors, and thus the majority of the neighborhoods have improved access to at least one station in terms of network distance, resulting in a big improvement in mean accessibility and early-time coverage [21, 22, 23].

The geographic distribution (Figure 2) helps to understand this phenomenon. The new stations are not placed around existing stations but rather fill in uncovered corridors while maintaining spacing and equity constraints. This reduces detours and increases the proportion of neighborhoods that lie within strict early-time coverage bounds. From an operational perspective, the improvement in the ≤ 4 -minute area is particularly significant because early-time coverage is essential for fire growth suppression and loss prevention, and best-practice fire service delivery aims to achieve short first-due distances [25, 36, 37].

5.2. Worst-case distance (max), tail risk, and topology effects

The proposed plan decreases the maximum nearest station distance from 12,595.07 m to 8,752.15 m (Table 2), i.e., by 30.5%. This shows that at least one new station aims at an originally most poorly served demand pocket in network terms. Nevertheless, the new maximum distance is still above the 10-minute proxy threshold (7.5 km at 35 km/h). This illustrates an important aspect of network-based station location: accessibility in the worst-case scenario not only depends on the number and location of stations but also on the road network topology. Long paths can remain due to sparse corridors, bridge-like bottlenecks, or partial disconnections in the traversable subnetwork conditions that can be met in real-world urban GIS data [21, 22, 23].

In practice, the remaining tail risk can be mitigated by two different but complementary approaches. First, capacity scenarios: an increased N^* (from 3 to 4) may enhance typical accessibility and can also enhance tail risk if an additional station can be located close to the remaining hard-to-serve pockets. Second, topology and data modifications: road graph validation and modification (e.g., closing digitization gaps or reconnecting edges) may help to decrease structural barriers that cannot be overcome by station location alone. Hence, minimax performance should be considered as a common function of station location and road network connectivity.

5.3. Why (max, mean) is a policy-relevant objective pair and how it relates to planning decisions

The objective pair (maximum distance, mean distance) offers an explicit policy-relevant trade-off between worst-case protection and system-wide efficiency. Minimization of mean distance enhances typical service levels and resource efficiency, whereas minimization of maximum distance offers an explicit protection of the most remote or endangered neighborhoods (p-center logic) and supports equity-oriented planning objectives [27, 40].

The Pareto analysis (Figure 1) illustrates these trade-offs: solutions that slightly improve the mean can increase worst-case exposure, and vice versa. The selection of a compromise (knee) solution is therefore a justifiable decision-support approach, as it identifies large marginal improvements without moving the solution towards less desirable regions, as in knee-based selection for multi-objective decision support [34, 43, 44]. In the Mosul scenario, the selected solution simultaneously improves mean and max values (Table 2), justifying its selection as a compromise solution.

5.4. Algorithmic considerations: what the comparison indicates in this case study

The multi-run comparison (Table 3) indicates two important aspects. First, NSGA-II demonstrates better global Pareto performance (better HV and IGD values) for the given setting, indicating a good global approximation of the trade-off front. [6]. Second, this global performance is achieved at the cost of much higher runtime, and its knee solutions trade off slightly better mean distance for worse worst-case exposure. In contrast, the PSO-family solutions (MOPSO/MORSO) offer knee solutions with lower maximum distance and extremely low runtime, making them preferable when worst-case protection and fast scenario exploration are of primary interest [29, 35].

In this particular problem setting, MORSO performs on par with baseline MOPSO on global indicators (Table 3), indicating that reinforcement does not harm Pareto quality while maintaining the feasibility and decision-support mechanisms integral to the proposed planning process. In practice, the primary strength of the proposed solution is not the optimizer, but the entire constraint-aware and network-based decision-support process yielding deployable plans and revealing efficiency-risk trade-offs. [31].

5.5. Speed assumption: what it changes and what it does not

Speed of travel is used in this study to transform network distances into meaningful response time proxy bands. Sensitivity analysis indicates that the reported ≤ 4 -minute coverage is highly speed-dependent (Figure 5; Table 4): the median coverage varies significantly over realistic speeds (e.g., 25-45 km/h). Thus, time-band percentages should be considered scenario-reported data that vary according to operational conditions such as congestion, road type, signal delay, and time of day.

Notably, because the optimization objectives are network distances, the key finding improved accessibility under the proposed siting plan is independent of a specific speed proxy. Distance-based results remain the key planning tool, and time bands offer an operational insight with which to intuitively interpret common emergency response reporting [21, 22, 23, 25].

5.6. Practical implications for Mosul and transferable lessons

For Mosul, the additive plan improves the system from 9 to 12 stations and provides concrete improvements in early response accessibility, typical travel burden, and worst-case exposure (Table 2; Figures 2-3). These results support an incremental implementation strategy: add a small number of stations to realize immediate improvements in ≤ 4 -minute coverage; then investigate remaining worst-case neighborhoods to decide whether the problem is a lack of stations (evaluate higher N^* scenarios) or network topology (repair connectivity/data).

Two general lessons are transferable. First, network modeling is critical because Euclidean distance can be misleading about road detours and barriers; this can lead to biased siting and coverage analysis compared to road network distances [21, 22, 23, 41, 42]. Second, multi-objective models are more accurate decision tools than single-threshold models because they explicitly illuminate the efficiency-risk trade-off and allow for transparent identification of policy-relevant compromise solutions [27].

5.7. Limitations and extensions

Some simplifications have been made to ensure the development of a deployable planning tool. Demand is modeled using neighborhood centroids instead of incident-level call data; using past incident locations and risk-weighted demand can enhance realism. Travel time is modeled using a uniform speed proxy; road-class-specific or time-specific travel time models can enhance operational realism [21, 22, 23]. Busy fraction (reliability) is not modeled and can be incorporated using availability-formulation models (e.g., coverage and failure cost views) if dispatch/workload information is available [26, 43, 44]. Lastly, more comprehensive benchmarking on closely comparable evaluation budgets and other MOEAs can further enhance generality, as is common practice in multi-objective optimization evaluation [31, 35]. Other accessibility metrics, such as those based on floating catchment area, can also be investigated in future extensions [45, 46].

6. Conclusions

This work introduced a GIS and road network two-stage additive fire station planning framework on urban road graphs and applied it to the case of Mosul, Iraq. The framework distinguishes between the number of stations to add (using an economic sizing criterion to obtain N^*) and where to add them (using multi-objective optimization on the road graph). The Stage-2 optimization focuses on policy-relevant objectives minimizing maximum and average nearest-station network distances thus balancing tail-risk protection and system-wide efficiency. The method explicitly works with shortest-path network distances instead of Euclidean distances, and it provides an interpretable Pareto set of deployable solutions under realistic conditions.

For the Mosul case study, the addition of three new stations to the existing nine (total of 12 stations) leads to significant improvements. The chosen solution decreases the average nearest station distance from 3,474.67 m to 2,823.09 m ($\approx -18.8\%$) and the maximum distance from 12,595.07 m to 8,752.15 m ($\approx -30.5\%$). Standards compliant reporting with a speed proxy of 35 km/h indicates improvements in coverage from 52.60% \rightarrow 64.16% for ≤ 4 minutes, 89.02% \rightarrow 93.06% for ≤ 8 minutes, and 93.64% \rightarrow 97.11% for ≤ 10 minutes. These improvements suggest that a few well-located additions can lead to immediate improvements in early response accessibility as well as minimize worst-case exposure.

A numerical multi-run comparison with respect to pre-defined multi-objective baselines demonstrates that various optimizers can focus on different parts of the efficiency-risk trade-off. For the given setup, NSGA-II outperforms in terms of global Pareto front quality metrics (HV/IGD), but at the cost of significantly higher runtimes and knee solutions that balance worst-case exposure in favor of marginally better mean distance. The proposed MORSO-based workflow offers a useful decision support benefit by enabling the rapid generation of constraint-satisfying solutions and enabling rapid scenario exploration and sensitivity analysis on realistic road graphs. In this case, MORSO equals the baseline MOPSO on aggregate metrics, suggesting that reinforcement does not degrade solution quality while maintaining the feasibility handling and deployability mechanisms that are integral to the proposed framework.

Sensitivity analysis verifies that the more additions N^* result in a monotonic increase in typical accessibility (median mean distance reduces from $\approx 2,979$ m for $N^* = 2$ to $\approx 2,668$ m for $N^* = 4$). It is also clear from the sensitivity analysis that the median time-band coverage is highly sensitive to the choice of speed proxy (median time-band coverage ≤ 4 minutes rises from $\approx 37\%$ at 25 km/h to $\approx 79\%$ at 45 km/h), thus reiterating that time-band outcomes need to be considered in the context of scenario reporting. In summary, the results clearly suggest that the network distance benefits delivered by the proposed planning framework are robust, while the remaining gaps in extreme cases can be filled by higher-capacity scenarios and/or specific efforts to enhance road network connectivity.

REFERENCES

1. R. L. Church, and C. S. ReVelle, *The maximal covering location problem*, Papers in Regional Science, vol. 32, no. 1, pp. 101–118, 1974. doi: 10.1007/BF01942293.
2. C. S. ReVelle, and H. A. Eiselt, *Location analysis: A synthesis and survey*, Eur. J. Oper. Res., vol. 165, no. 1, pp. 1–19, 2005. doi: 10.1016/j.ejor.2003.11.032.
3. G.-H. Tzeng, and Y.-W. Chen, *The optimal location of airport fire stations: A fuzzy multi-objective programming and revised genetic algorithm approach*, Transp. Plan. Technol., vol. 23, no. 3, pp. 225–251, 1999. doi: 10.1080/03081069908717638.
4. J. He, X. Guan, H. Lu, and J. Yang, *Research on multi-objective hierarchical site selection coverage of fire stations*, PLOS ONE, vol. 19, no. 12, Art. no. e0309731, 2024. doi: 10.1371/journal.pone.0309731.
5. I. Landa-Torres et al., *A multi-objective grouping harmony search algorithm for the optimal distribution of 24-hour medical emergency units*, Expert Syst. Appl., vol. 40, no. 6, pp. 2343–2349, 2013. doi: 10.1016/j.eswa.2012.10.051.
6. K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan, *A fast and elitist multiobjective genetic algorithm: NSGA-II*, IEEE Trans. Evol. Comput., vol. 6, no. 2, pp. 182–197, 2002. doi: 10.1109/4235.996017.
7. P. Gao et al., *Sustainable land-use optimization using NSGA-II: theoretical and experimental comparisons of improved algorithms*, Landscape Ecol., vol. 36, pp. 1877–1892, 2021. doi: 10.1007/s10980-020-01051-3.
8. K. Ransikarbum, and S. J. Mason, *A bi-objective optimisation of post-disaster relief distribution and short-term network restoration using a hybrid NSGA-II algorithm*, Int. J. Prod. Res., vol. 60, no. 19, pp. 5769–5793, 2022. doi: 10.1080/00207543.2021.1970846.
9. Q. Zhang, and H. Li, *MOEA/D: A multiobjective evolutionary algorithm based on decomposition*, IEEE Trans. Evol. Comput., vol. 11, no. 6, pp. 712–731, 2007. doi: 10.1109/TEVC.2007.892759.

10. A. Zhou et al., *Multiobjective evolutionary algorithms: A survey of the state of the art*, Swarm Evol. Comput., vol. 1, no. 1, pp. 32–49, 2011. doi: 10.1016/j.swevo.2011.03.001.
11. Y. Shi, and R. Eberhart, *A modified particle swarm optimizer*, in Proc. IEEE Int. Conf. Evol. Comput., 1998, pp. 69–73. doi: 10.1109/ICEC.1998.699146.
12. Y. Xiang et al., *Optimization research on the site selection of fire safety for mega-project sites based on multi-objective particle swarm*, Evol. Intell., vol. 15, no. 4, pp. 2455–2471, 2022. doi: 10.1007/s12065-020-00520-7.
13. Y. Zheng, B. Guo, and Y. Zhou, *Multi-objective White Shark Optimizer for global optimization and rural sports-facilities location problem*, Biomimetics, vol. 10, no. 8, Art. 537, 2025. doi: 10.3390/biomimetics10080537.
14. Y. Chen, and Z. Lai, *A multi-objective optimization approach for emergency medical service facilities location-allocation in rural areas*, Risk Manag. Healthc. Policy, vol. 15, pp. 473–490, 2022. doi: 10.2147/RMHP.S332215.
15. S. Bolouri et al., *The ordered capacitated multi-objective location-allocation problem for fire stations using spatial optimization*, ISPRS Int. J. Geo-Inf., vol. 7, no. 2, Art. 44, 2018. doi: 10.3390/ijgi7020044.
16. Y. Zhu, M. Zhu, L. Chen, and Y. Lin, *Multi-Objective Optimization of Urban Gas Station Site Selection under Territorial Spatial Planning Constraints*, ISPRS Int. J. Geo-Inf., vol. 13, no. 11, Art. 375, 2024. doi: 10.3390/ijgi13110375.
17. K. Ransikarbum, and R. Pitakaso, *Multi-objective optimization design of sustainable biofuel network with integrated fuzzy analytic hierarchy process*, Expert Syst. Appl., vol. 240, Art. 122586, 2024. doi: 10.1016/j.eswa.2023.122586.
18. M. Vargas-Santiago et al., *The adaptable Pareto set problem for facility location: a video game approach*, Expert Syst. Appl., vol. 186, Art. 115682, 2021. doi: 10.1016/j.eswa.2021.115682.
19. A. Tanoumand et al., *A metaheuristic-based approach for optimizing the allocation of emergency water reservoirs for fire-following-earthquake suppression*, Results Eng., vol. 27, Art. 105925, 2025. doi: 10.1016/j.rineng.2025.105925.
20. J. Li et al., *Multi-objective optimization of EMS facilities using multi-source data: A case study in Dangtu, China*, Front. Archit. Res., vol. 14, no. 4, pp. 1090–1107, 2025. doi: 10.1016/j.foar.2024.10.011.
21. F. Delgado-Enríquez et al., *A framework to improve urban accessibility and environmental conditions in age-friendly cities using graph modeling and multi-objective optimization*, Comput. Environ. Urban Syst., vol. 102, Art. 101966, 2023. doi: 10.1016/j.compenvurbysys.2023.101966.
22. M. Habibi, S. Lotfi, and M. J. Koohsari, *Spatial Analysis of Urban Fire Station Locations by Integrating AHP Model and IO Logic Using GIS: A Case Study of Zone 6 of Tehran*, Journal of Applied Sciences, vol. 8, no. 19, pp. 3302–3315, 2008. doi: 10.3923/jas.2008.3302.3315.
23. L. Brotcorne, G. Laporte, and F. Semet, *Ambulance location and relocation models*, Eur. J. Oper. Res., vol. 147, no. 3, pp. 451–463, 2003. doi: 10.1016/S0377-2217(02)00364-8.
24. L. H. Omar, and A. I. A. Sulaiman, *Evaluation of the efficiency of the performance of the service of civil defense centers in the city of Mosul using geographic information systems*, Special Issue of the Fourth International Scientific Conference / Historical and Geographical Studies, Univ. of Mosul, pp. 1201–1234, 2022. doi: 10.33899/berj.2023.177194.
25. IAFF, *NFPA 1710: Summary Sheet—Standard for the Organization and Deployment of Fire Suppression Operations, Emergency Medical Operations, and Special Operations to the Public by Career Fire Departments*, 2016. [Online].
26. M. S. Daskin, *A Maximum Expected Covering Location Model: Formulation, Properties, and Heuristic Solution*, Transportation Science, vol. 17, no. 1, pp. 48–70, 1983. doi: 10.1287/trsc.17.1.48.
27. M. T. Marsh, and D. A. Schilling, *Equity Measurement in Facility Location Analysis: A Review and Framework*, Eur. J. Oper. Res., vol. 74, no. 1, pp. 1–17, 1994. doi: 10.1016/0377-2217(94)90200-3.
28. J. A. P. III, *Running head: COMPLYING WITH ISO RESPONSE TIME CRITERIA. Complying with ISO Response Time Criteria*, Unpublished document, n.d.
29. C. A. C. Coello Coello, and M. S. Lechuga, *MOPSO: A proposal for multiple objective particle swarm optimization*, in Proc. CEC, 2002, pp. 1051–1056. doi: 10.1109/CEC.2002.1004388.
30. R. Poli, J. Kennedy, and T. Blackwell, *Particle swarm optimization: An overview*, Swarm Intelligence, vol. 1, no. 1, pp. 33–57, 2007. doi: 10.1007/s11721-007-0002-0.
31. L. While, P. Hingston, L. Barone, and S. Huband, *A faster algorithm for calculating the hypervolume indicator*, IEEE Trans. Evol. Comput., vol. 10, no. 2, pp. 230–234, 2006. doi: 10.1109/TEVC.2005.851275.
32. S. L. Hakimi, *Optimum location of switching centers and the absolute centers and medians of a graph*, Operations Research, vol. 12, no. 3, pp. 450–459, 1964. doi: 10.1287/opre.12.3.450.
33. G. Boeing, *OSMnx: New methods for acquiring, constructing, analyzing, and visualizing complex street networks*, Comput. Environ. Urban Syst., vol. 65, pp. 126–139, 2017. doi: 10.1016/j.compenvurbysys.2017.05.004.
34. J. Branke, K. Deb, K. Miettinen, and R. Slowinski (eds.), *Finding knees in multi-objective optimization*, in *Multiobjective Optimization: Interactive and Evolutionary Approaches*, LNCS 5252, pp. 283–318, 2008. doi: 10.1007/978-3-540-30217-9_73.
35. C. A. C. Coello Coello, G. T. Pulido, and M. S. Lechuga, *Handling multiple objectives with particle swarm optimization*, IEEE Trans. Evol. Comput., vol. 8, no. 3, pp. 256–279, 2004. doi: 10.1109/TEVC.2004.826067.
36. NFPA 1710, *Standard for the Organization and Deployment of Fire Suppression Operations, Emergency Medical Operations, and Special Operations to the Public by Career Fire Departments*, National Fire Protection Association (NFPA), latest ed.
37. IAFF, *Understanding NFPA 1710: Response time and staffing benchmarks*, International Association of Fire Fighters, technical brief.
38. E. Talen, *Visualizing fairness: Equity maps for planners*, Journal of the American Planning Association, vol. 64, no. 1, pp. 22–38, 1998. doi: 10.1080/01944369808975954.
39. W. Li, Q. Zhang, and Y. Jin, *Knee point-guided multiobjective optimization*, Complexity, vol. 2020, Art. ID 8877008, 2020. doi: 10.1155/2020/8877008.
40. O. Cuate, J. Velázquez-Iturbide, and E. Alba, *Pareto explorer for finding the knee in multiobjective optimization*, Mathematics, vol. 8, no. 10, p. 1651, 2020. doi: 10.3390/math8101651.
41. P. Tatit, R. N. Rodrigues, and D. Silva, *Navigating the Maps: Euclidean vs. road network distances in spatial queries: Accuracy and implications*, Algorithms, vol. 17, no. 1, p. 29, 2024. doi: 10.3390/a17010029.

42. A. Okabe, and K. Sugihara, *Spatial Analysis Along Networks: Statistical and Computational Methods*, Wiley, Chichester, UK, 2012. doi: 10.1002/9781119967101.
43. L. V. Snyder, and M. S. Daskin, *Reliability models for facility location: The expected failure cost case*, *Transportation Science*, vol. 39, no. 3, pp. 400–416, 2005. doi: 10.1287/trsc.1040.0107.
44. R. Batta, *The maximal expected covering location problem: Revisited*, *Transportation Science*, vol. 23, no. 4, pp. 277–287, 1989. doi: 10.1287/trsc.23.4.277.
45. M. F. Guagliardo, *Spatial accessibility of primary care: concepts, methods and challenges*, *International Journal of Health Geographics*, vol. 3, Art. 3, 2004. doi: 10.1186/1476-072X-3-3.
46. W. Luo, and Y. Qi, *An enhanced two-step floating catchment area (E2SFCA) method for measuring spatial accessibility to primary care physicians*, *International Journal of Health Geographics*, vol. 8, p. 58, 2009. doi: 10.1186/1476-072X-8-58.