MORAF: A Multi-Objective Framework for Resource Allocation with Fairness for Efficient Disaster Management

V Shreya Sivani
PES University, Bengaluru, India
vshreyasivani@gmail.com

Vandana J
PES University, Bengaluru, India
vandanaj0110@gmail.com

Rachit Bhatnagar
PES University, Bengaluru, India
rachitbhatnagar04@gmail.com

Mandadi Sai Likesh Reddy
PES University, Bengaluru, India
reddylikesh28@gmail.com

Arti Arya
PES University, Bengaluru, India
artiarya@pes.edu

ABSTRACT

After large-scale disasters, emergency responders must allocate limited resources across affected areas within critical time windows. This creates a fundamental challenge: concentrating resources maximizes total lives saved but leaves some areas underserved, while distributing resources equally ensures coverage but reduces overall effectiveness. This work presents MORAF (Multi-Objective Resource Allocation with Fairness), an optimization framework that balances both objectives simultaneously. MORAF incorporates diminishing returns (the first ambulance saves more lives than the tenth), context-specific effectiveness (concrete breakers work for earthquakes but not floods), and resource synergy (teams, vehicles, and equipment work better together). Since this optimization problem is non-convex, the framework employs Sequential Least Squares Programming with multiple starting points. MORAF is validated on a hurricane scenario using the Hurricane Michael road network across 100 zones with 12 resource types. Results demonstrate 19.4% higher rescue utility and 69% improved fairness (Fairness score: 0.910 versus 0.537) compared to baseline approaches.

Keywords

Disaster response, resource allocation, multi-objective optimization, fairness, humanitarian logistics.

INTRODUCTION

Post-disaster resource allocation effectiveness depends not only on the quantity of available resources but also on how allocation objectives are defined. Humanitarian logistics models often prioritize utilitarian formulations to maximize aggregate welfare or total lives saved. These approaches frequently use objective functions based on social or deprivation costs (Holguín-Veras et al., 2013). However, alternative distribution strategies that emphasize broader coverage have also been explored. This is particularly relevant in last-mile humanitarian operations, where ensuring service reach often conflicts with efficiency-oriented objectives (Ghahremani-Nahr et al., 2024). This tension highlights the difficult balance between overall effectiveness and equitable access in disaster response planning.

The urgency of this challenge is visible in recent large-scale events. Hurricane Michael (2018) caused over $25 billion in damages across Florida, while the 2023 Turkey–Syria earthquake affected more than 15 million people. In such events, emergency responders must allocate limited resources such as ambulances, rescue teams, and heavy machinery across many affected zones under strict time constraints. Early decisions in the response process play a critical role in shaping overall rescue outcomes and system performance (Y. Wang & Sun, 2018).

Current state-of-the-art approaches for emergency resource allocation frequently employ Mixed-Integer Linear Programming (MILP) or Deep Reinforcement Learning (DRL) techniques (Zhao & Wang, 2023). While effective for large-scale optimization, these formulations often rely on simplified utility structures. They implicitly assume that adding more resources yields proportional gains, which is rarely true in the field. Moreover, many existing models treat resources as homogeneous units, overlooking the fact that operational effectiveness depends strongly on disaster type and local conditions. As a result, solutions may be mathematically optimal yet operationally impractical. For instance, models might suggest deploying fractions of a rescue team (indivisible resources) leading to a phenomenon referred to as micro-allocation fragmentation.

At the same time, modern disaster response increasingly benefits from AI-driven perception systems that enhance situational awareness. UAV-based imagery analysis enables semantic segmentation of damaged infrastructure (Rahnemoonfar et al., 2023), victim detection using thermal and near-infrared imagery (Cruz Ulloa et al., 2021; Shao, 2024), disaster type classification (Kyrkou & Theocharides, 2020), and debris material identification (Rahnemoonfar et al., 2023). It also aids in predicting secondary risks such as landslides (Arabameri et al., 2022). Advanced routing algorithms further determine feasible deployment strategies under realistic constraints, including time limits, service duration, and uncertainty in victim conditions (Rabbani et al., 2022; Shiri et al., 2024). Despite these advances, a gap remains in translating this rich perceptual information into resource allocation decisions that are both operationally realistic and sensitive to competing optimization objectives.

This work presents MORAF, a Multi-Objective Resource Allocation Framework designed to address these challenges. It jointly models diminishing returns, context-aware resource effectiveness, resource synergy, and fairness-related constraints within a single optimization formulation. The primary contributions are:

- **Operational Realism via Nonlinear Modeling**: The framework uses a modified Cobb-Douglas function ($\alpha = 0.6$) to model the reality that adding more resources eventually yields diminishing returns. It also applies a viability threshold ($\tau = 0.5$) to prevent the system from suggesting operationally insignificant allocations.
- **Context-Aware Effectiveness Matrix**: A heterogeneity matrix dynamically scores how well a resource performs based on the specific disaster and debris type. These classifications align with the Federal Emergency Management Agency (FEMA) Resource Typing Library Tool (RTLT).
- **Synergy-Driven Optimization**: The model rewards the deployment of complete response units combining Personnel, Vehicles, and Equipment rather than isolated assets. A tiered synergy bonus (up to 1.15×) is applied to reflect this real-world operational efficiency.
- **Pareto-Optimal Validation**: Ablation studies demonstrate that MORAF outperforms standard Greedy (efficiency-focused) and Proportional (equity-focused) baselines. The framework consistently achieves a superior balance of high utility and fairness across diverse operational scenarios.

RELATED WORK

Disaster Resource Allocation Models

Traditional disaster resource allocation has been extensively studied through operations research perspectives. Early formulations focused on minimizing delivery time or cost through MILP approaches, while recent work has shifted toward multi-objective strategies. Wang et al. (F. Wang et al., 2022) proposed a model for natural disaster chains that integrates path planning with periodic supply strategies. Similarly, Sun et al. (H. Sun et al., 2021) developed a bi-objective robust optimization framework for facility location under demand uncertainty. However, both approaches share a critical limitation: they assume uniform resource effectiveness. By failing to model context-specific variations such as the suitability of specific equipment for earthquake debris versus flood conditions, these models overlook the diminishing returns and operational viability thresholds essential for realistic planning.

Fairness in Humanitarian Logistics

The concept of fairness has gained significant traction in recent literature. Eriskin et al. (Eriskin et al., 2024) developed a robust multi-objective model for pandemic healthcare resource allocation, explicitly incorporating equity constraints through budgeted uncertainty sets. While effective for pandemics, their framework assumes homogeneous medical resources and does not address the physical heterogeneity of disaster debris or the non-linear utility of rescue assets. More broadly, fairness has emerged as a multidimensional challenge in humanitarian logistics, encompassing both the equitable distribution of relief items and the robustness of preparedness decisions under uncertainty. (Erbeyoğlu & Bilge, 2020) demonstrated that integrating fairness objectives into disaster preparedness network design significantly improves equitable coverage across demand zones during response operations. Huang et al. (Huang et al., 2012) further showed that equity, efficiency, and efficacy represent three distinct and often conflicting objectives in relief routing, highlighting the need for multi-objective formulations that explicitly model all three dimensions simultaneously. In a different approach, Rezaei-Malek et al. (Rezaei-Malek et al., 2016) incorporated fairness through max-min objectives. While this prioritizes the worst-off zones, it often leads to inefficient utilization by sacrificing overall system effectiveness for local equity. The limitations of relying solely on Max-Min objectives have been documented in broader humanitarian logistics contexts, where such approaches can lead to arbitrarily inequitable outcomes across beneficiary zones. Gini-based and distributional metrics have been proposed as more balanced alternatives that simultaneously account for the full distribution of outcomes rather than only the worst-off zone (Alem et al., 2022; Gutjahr & Fischer, 2018).

Theoretical Frameworks of Algorithmic Fairness

Beyond specific logistic models, translating ethical fairness into mathematical constraints remains a fundamental challenge. Recent literature distinguishes between *individual fairness* (treating similar zones similarly) and *group fairness* (ensuring equitable outcomes across different regions). Traditional utilitarian models often violate both principles. By focusing strictly on maximizing total survivors, these models tend to prioritize "easy-to-save" zones while neglecting difficult or remote areas. Anaya-Arenas et al. (Anaya-Arenas et al., 2018) argue that such cost-minimization objectives fail to capture the true "social cost" of suffering, which grows non-linearly with deprivation time.

Prior research has attempted to solve this using **Rawlsian Max-Min** fairness. This approach tries to maximize the utility of the worst-off zone. While ethically sound, Max-Min can be inefficient in mass-casualty scenarios. It often leads to the "drowning swimmer" paradox, where excessive resources are diverted to save a single person in a hard-to-reach area, sacrificing the chance to save ten people in a more accessible zone. Conversely, simple demand-based allocation ignores whether resources are actually effective in specific environments. MORAF diverges from these extremes by adopting a Gini-based fairness measure computed over fulfillment ratios rather than raw resource quantities. Unlike Max-Min objectives, the Gini coefficient accounts for the full distribution of outcomes across all zones simultaneously, avoiding the inefficiency of diverting excessive resources to a single worst-off zone (Gutjahr & Fischer, 2018). Unlike demand-proportional metrics, Gini-based fairness measured over fulfillment ratios captures whether zones receive resources commensurate with actual need. Critically, the Gini coefficient produces a continuous, differentiable fairness signal that integrates naturally into gradient-based optimization via SLSQP, whereas lexicographic or risk-based alternatives would require integer programming incompatible with real-time operational requirements. This combination of distributional sensitivity, outcome-based measurement, and computational compatibility makes the Gini coefficient the appropriate fairness metric for MORAF's joint optimization setting (Alem et al., 2022; Gutjahr & Fischer, 2018). While prior work has applied Gini-based metrics to supply chain equity (Alem et al., 2022) and deprivation cost minimization (Gutjahr & Fischer, 2018), MORAF is the first to apply it over fulfillment ratios within a joint efficiency–fairness nonlinear formulation for heterogeneous, context-aware disaster resource allocation.

Context-Aware Resource Effectiveness

Recent studies have begun to recognize that resource utility is context-dependent. Sun and Zhang (J. Sun & Zhang, 2020) demonstrated that repair crew effectiveness varies significantly across interdependent infrastructure networks (power-transportation-water) during earthquake recovery. However, their agent-based framework focuses on sequential repair prioritization. It does not provide a systematic method to integrate heterogeneous resource effectiveness into the kind of multi-resource, multi-zone optimization required for immediate response operations.

Optimization Challenges in Disaster Response

The resource allocation problem addressed here is non-convex, meaning the mathematical "landscape" has many local optima. The complexity arises from the synergy bonuses and the non-linear utility function. Traditional **Mixed-Integer Linear Programming (MILP)** solvers, which are standard in logistics, struggle with these non-linearities. As noted by Sun et al. (H. Sun et al., 2021), approximating these curves with straight lines often leads to significant errors in estimating resource effectiveness under uncertainty.

While recent studies have explored Deep Reinforcement Learning (DRL) approaches, such as the Deep Q-Networks proposed by Zhao and Wang (Zhao & Wang, 2023), developed in the context of urban public health emergencies,

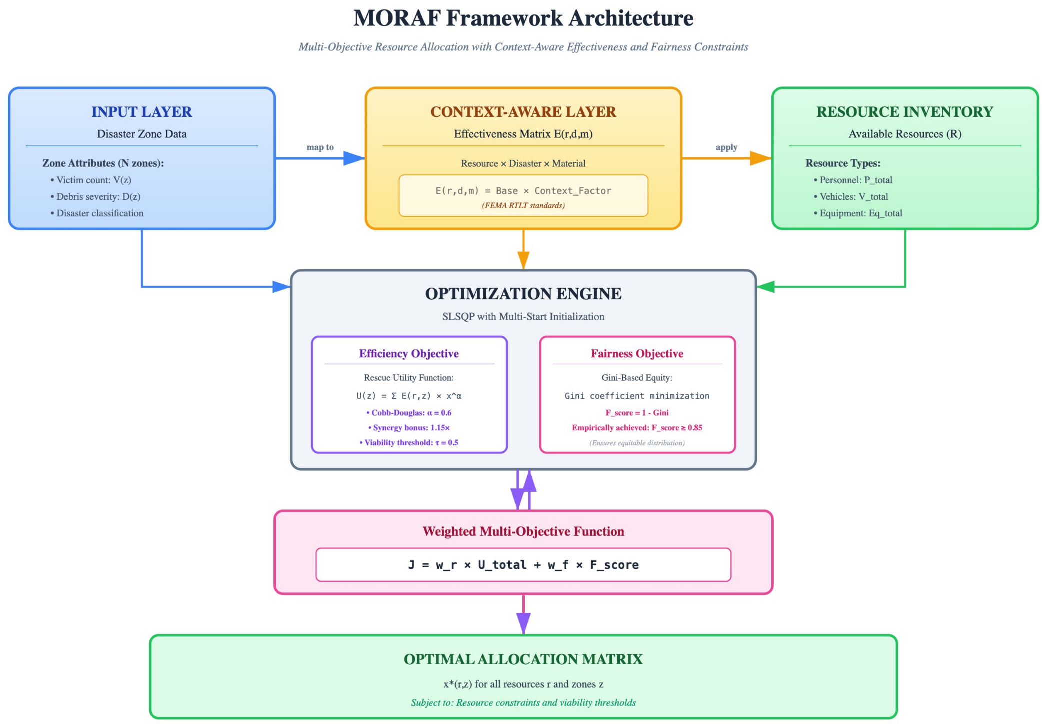

Figure 1. MORAF Framework Architecture: End-to-end system overview showing input layer (zone-level disaster data), context-aware effectiveness layer (FEMA-derived multipliers), resource inventory, optimization engine with dual objectives (efficiency and fairness), and optimal allocation output matrix.

these methods can be computationally demanding, particularly when rapid re-optimization is required. In real-time response, where decision windows are measured in minutes, waiting for a neural network to converge or requiring massive training data presents operational challenges. This study utilizes **Sequential Least Squares Programming (SLSQP)** as a strategic compromise, using gradient-based search with multi-start initialization to navigate the utility surface efficiently.

Research Gaps

Despite these advances, existing approaches exhibit three critical limitations. First, most models employ linear utility functions that fail to capture diminishing marginal returns; in reality, allocating a fifth ambulance to a zone yields less incremental benefit than the first. This limitation overlooks well-documented operational dynamics in disaster response environments. As multiple rescue units operate within confined debris zones, physical congestion reduces individual team productivity. Coordination and command overhead also increase with team size, as communication and supervision demands grow faster than effective rescue capacity. Moreover, early-arriving resources typically reach the most accessible victims, while subsequent deployments face progressively more complex extraction conditions. Empirical studies in post-disaster logistics establish that both deprivation costs and resource productivity exhibit diminishing marginal behavior, supporting nonlinear benefit structures in disaster response contexts (Holguín-Veras et al., 2013, 2016). These mechanisms justify modeling rescue utility using a concave functional form rather than assuming proportional gains. Second, resources are typically treated as homogeneous units, ignoring the operational reality that effectiveness varies dramatically based on disaster type and debris characteristics. Third, current multi-objective formulations often produce theoretical Pareto frontiers without respecting viability thresholds or synergistic resource combinations. MORAF addresses these gaps through nonlinear utility modeling, context-aware effectiveness matrices, and constraint formulations that ensure operational feasibility.

PROBLEM FORMULATION

Input Data and Notation

The framework builds upon an integrated disaster response pipeline, utilizing zone-level situational data generated by upstream perception and routing components. Specifically, the allocation model leverages outputs from the perception module which provides detailed site characterization including victim counts, disaster classifications, debris material compositions, and secondary risk scores alongside a prioritized subset of accessible zones within specific operational time windows (12h, 24h, 48h) established by the routing module.

While upstream modules for semantic segmentation, risk prediction, and dynamic route optimization constitute distinct contributions from the broader project, they establish the essential operational foundation for this work. Consequently, this work focuses specifically on the subsequent challenge: the optimal allocation of heterogeneous resources across these identified priority zones.

Table 1. Summary of Notation

Symbol	Description
Sets and Indices	
$i \in \{1, \dots, n\}$	Index over prioritized disaster zones
$j \in \{1, \dots, m\}$	Index over available resource types
C	Set of resource categories: $\{$Personnel, Vehicle, Equipment$\}$
Zone Parameters	
V_i	Number of detected victims in zone i
D_i	Disaster type classification of zone i
M_i	Debris material composition of zone i
δ_i	Normalized debris severity score, $\delta_i \in [0, 1]$
R_i	Predicted secondary risk score, $R_i \in [0, 1]$
P_i	Priority score used by upstream routing module to select zone subset
Resource Parameters	
β_j	Base effectiveness of resource j
K_j	Total available capacity of resource j
C_j	Category of resource j, $C_j \in C$
$\phi_D(j, D_i)$	Disaster-type effectiveness multiplier for resource j in zone i
$\phi_M(j, M_i)$	Debris-material effectiveness multiplier for resource j in zone i
E_{ij}	Context-aware effectiveness of resource j in zone i
Decision Variables and Derived Quantities	
x_{ij}	Amount of resource j allocated to zone i
x_{ij}^{eff}	Effective allocation after viability threshold (Eq. 5)
S_i	Synergy multiplier for zone i (Eq. 6)
$U_i(X)$	Rescue utility of zone i given allocation X (Eq. 4)
f_i	Fulfillment ratio of zone i: $U_i(X)/V_i$
$F(X)$	Gini-based fairness score over fulfillment ratios (Eq. 7)
Scalar Parameters	
α	Diminishing returns exponent ($\alpha = 0.6$)
τ	Viability threshold ($\tau = 0.5$)
k	Sigmoid steepness parameter ($k = 5$)
w_r, w_f	Policy weights for rescue utility and fairness ($w_r + w_f = 1$)

Table 1 summarizes all notation used throughout the formulation. For each zone $i \in \{1, 2, \dots, n\}$ in the prioritized set, the following attributes are defined:

- Number of detected victims: V_i
- Disaster type classification: D_i (e.g., earthquake, flood, cyclone, fire)
- Debris material composition: M_i (e.g., metal, rubber, wood, organic, electronic, plastic, cement, ceramic)
- Debris severity score: δ_i (normalized index $\in [0, 1]$ derived from debris extent)
- Predicted secondary risk score: $R_i \in [0, 1]$

The priority score P_i is calculated as a weighted sum of these critical factors:

$$P_i = w_v V_i + w_d \delta_i + w_\rho R_i \tag{1}$$

where w_v, w_d, w_ρ are learned weights that balance the importance of saving lives, clearing obstructions, and mitigating future risks. The priority score P_i serves as the zone-ranking criterion for the upstream routing module, which determines the subset of accessible zones within each operational time window. This prioritized subset constitutes the input to the allocation optimization; P_i itself does not appear as a decision variable within the MORAF formulation.

The routing system determines which zones can be visited within operational time constraints (12h, 24h, or 48h) and outputs a prioritized subset of zones for resource allocation. Additionally, a set of available resources $j \in \{1, 2, \ldots, m\}$ is defined, where each resource is characterized by:

- Resource type: ambulances, rescue teams, equipment, etc.
- Resource category: $C_j \in C$, where $C = \{\text{Personnel}, \text{Vehicle}, \text{Equipment}\}$
- Total capacity: K_j (total units available for deployment)
- Base effectiveness: β_j (intrinsic capability measure)

The complete MORAF architecture integrating these components is illustrated in Figure 1.

Decision Variables

The primary decision variable is the allocation matrix:

$$X = [x_{ij}] \in \mathbb{R}^{n \times m} \tag{2}$$

where x_{ij} represents the amount of resource j allocated to zone i.

Context-Aware Resource Effectiveness

A key innovation in MORAF is the context-aware effectiveness matrix $E \in \mathbb{R}^{n \times m}$, where:

$$E_{ij} = \beta_j \cdot \phi_D(j, D_i) \cdot \phi_M(j, M_i) \tag{3}$$

Here, $\phi_D(j, D_i)$ is the disaster-type multiplier and $\phi_M(j, M_i)$ is the debris-material multiplier, both derived from FEMA Resource Typing Library Tool (RTLT) standards. Representative values for these multipliers are provided in Table 2.

Table 2. Sample Effectiveness Multipliers calibrated to FEMA RTLT Standards

Resource	**Operational Context**	$\phi_\mathbf{D}$	$\phi_\mathbf{M}$
Concrete Breakers	Earthquake (Concrete debris)	3.0	3.2
USAR (Urban Search and Rescue) Teams	Earthquake (Concrete/Steel)	2.8	2.5
Boats	Flood zones	3.0	1.0
Helicopters	Flood / Hurricane zones	2.8 / 2.3	1.0
Chainsaws	Tornado / Hurricane (Wood debris)	2.4 / 2.2	3.0

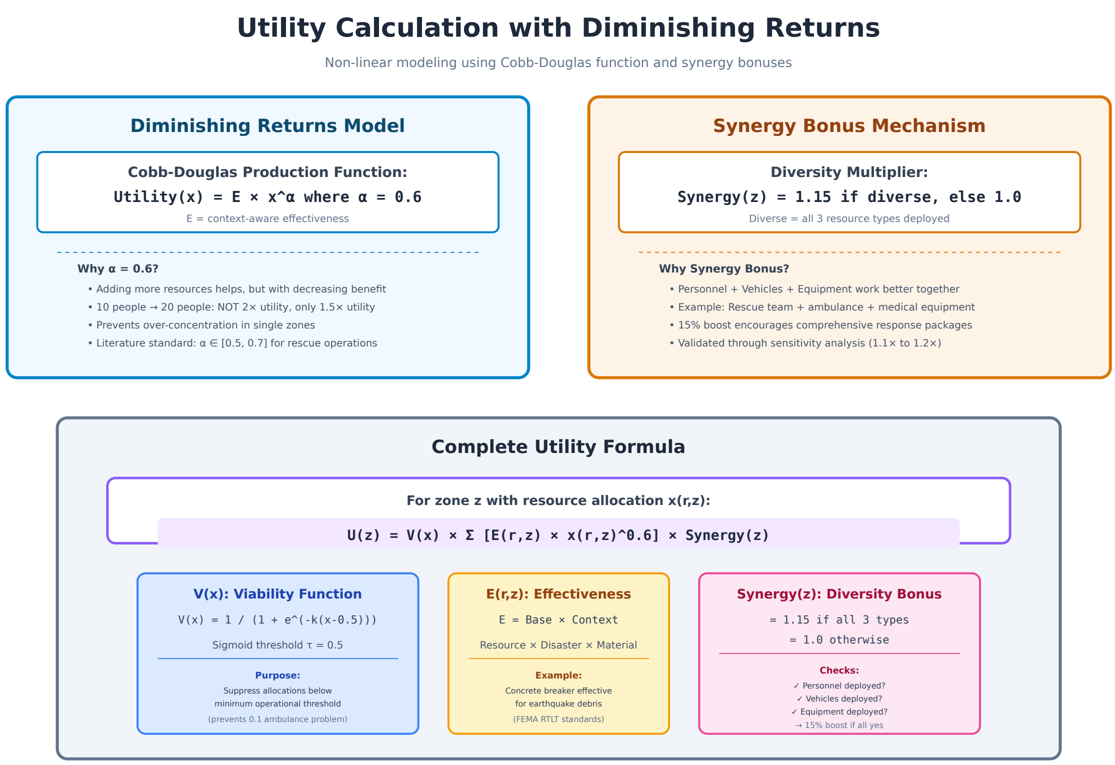

Figure 2. Utility Calculation with Diminishing Returns: Complete rescue utility formulation showing (left) Cobb-Douglas diminishing returns ($\alpha = 0.6$) preventing over-concentration, (center) viability threshold ($\tau = 0.5$) suppressing micro-allocations, and (right) synergy bonus mechanism ($1.15\times$ for full diversity) rewarding comprehensive resource packages combining Personnel, Vehicles, and Equipment.

Rescue Utility Function

Rescue utility $U_i(X)$ is defined as the expected number of victims that can be effectively rescued in zone i given the allocated resources X. Unlike a simple victim count, rescue utility accounts for resource effectiveness in specific disaster contexts, diminishing marginal returns, synergistic benefits of diverse resource deployment, and operational viability constraints.

Drawing from the foundational structure of the Cobb-Douglas production function (Cobb & Douglas, 1928; Colther & Doussoulin, 2025), the rescue utility for zone i is modeled using a modified formulation that captures diminishing returns:

$$U_i(X) = \min\left(\sum_{j=1}^{m} E_{ij} \cdot \left(x_{ij}^{\text{eff}}\right)^{\alpha} \cdot S_i,\ V_i\right) \tag{4}$$

where x_{ij}^{eff} represents the effective allocation after applying the viability threshold (defined in Eq. 5), $\alpha = 0.6$ captures diminishing returns (validated through sensitivity analysis showing flat performance for $\alpha \in [0.5, 0.8]$), V_i caps utility at the victim count, and S_i is the synergy multiplier. The complete utility calculation incorporating these three components diminishing returns, viability threshold, and synergy bonus is illustrated in Figure 2.

Viability Threshold

To prevent micro-allocation fragmentation, a soft viability threshold using a sigmoid function is used:

$$x_{ij}^{\text{eff}} = x_{ij} \cdot \sigma(x_{ij}) = x_{ij} \cdot \frac{1}{1 + e^{-k(x_{ij} - \tau)}} \tag{5}$$

where $\tau = 0.5$ is the threshold value and $k = 5$ controls the steepness. This formulation smoothly suppresses allocations below the operational minimum while preserving differentiability for gradient-based optimization.

Synergy Bonus

The synergy multiplier S_i rewards diverse resource deployment:

$$S_i = \begin{cases} 1.15 & \text{if } \left|\{C_j : x_{ij} > 0.5\}\right| \geq 3 \\ 1.10 & \text{if } \left|\{C_j : x_{ij} > 0.5\}\right| = 2 \\ 1.00 & \text{otherwise} \end{cases} \tag{6}$$

This encourages allocations that combine personnel, vehicles, and equipment for comprehensive response capability. Sensitivity analysis validated that a maximum multiplier of 1.15× for full diversity provides the optimal balance between incentivizing comprehensive deployments and maintaining realistic utility scores. The synergy bonus operates on the aggregate utility term in Eq. 4, which is weighted by the context-aware effectiveness matrix E_{ij} (Eq. 3). Resources with near-zero effectiveness in a given zone – such as flood boats in earthquake debris or concrete breakers in flood zones – contribute negligibly to zone utility regardless of synergy, making their deployment unattractive to the optimizer. The framework therefore naturally avoids deploying contextually inappropriate resources, preserving inventory for zones where they are operationally effective.

Theoretical Justification of Utility Components

We selected the **Cobb-Douglas** form ($U \propto x^{\alpha}$) because it captures a simple reality of disaster response: diminishing returns. In any rescue operation, the first few units sent to a zone make the biggest difference. However, sending too many units to the same small area causes congestion and coordination issues, reducing their added value. By setting the parameter $\alpha = 0.6$, the model naturally penalizes "hoarding" resources in a single zone. This forces the system to spread resources out, which improves overall coverage even before we apply specific fairness rules. The range $\alpha \in [0.5, 0.7]$ referenced in Figure 2 reflects the sensitivity-validated plateau demonstrated in Figure 5(a), where less than 2% performance variation was observed across this interval, rather than a fixed external standard. While diminishing returns reduce the marginal utility of each additional resource unit, they do not prevent high-demand zones from receiving large allocations. The $\min(\cdot, V_i)$ operator in Eq. 4 bounds utility at the victim count rather than at any resource limit, so a zone with fifty victims remains eligible for sufficient resources to achieve full coverage. The role of $\alpha = 0.6$ is to discourage over-concentration in already-served zones – not to restrict allocation to critical areas. Sensitivity analysis confirms this: across all tested values of α, high-demand zones consistently received the largest allocation shares with no systematic under-service observed (Figure 5a).

Similarly, we used a sigmoid function for the **Viability Threshold** instead of a sharp "step" function. A sharp cut-off makes the math discontinuous, which stops standard optimization solvers (like SLSQP) from working. The sigmoid (Eq. 5) creates a smooth curve that mimics the operational rule: "you need a minimum team size to be effective." The steepness parameter ($k = 5$) ensures the transition is strict enough to prevent tiny, useless allocations (like 0.1 ambulances) but smooth enough for the solver to calculate the best solution without getting stuck.

Fairness Metric

Allocation fairness is measured using a Gini-based coefficient:

$$F(X) = 1 - G(X) \tag{7}$$

where $G(X)$ is the Gini coefficient computed over the fulfillment ratios

$$f_i = \frac{U_i(X)}{V_i}. \tag{8}$$

Fairness is measured over fulfillment ratios rather than raw allocated quantities because the operational concern in disaster response is not whether zones receive equal resources, but whether they receive resources proportional to their actual need. A zone with fifty victims requiring ten ambulances and a zone with five victims requiring one are treated equitably when both achieve similar fulfillment ratios, even if their absolute allocations differ substantially.

A fairness score of 1.0 indicates perfect equity, while 0.0 indicates maximal inequality. The benchmark of $F \geq 0.85$ referenced in Figure 1 reflects an empirically observed performance level across the experimental scenarios reported in this work. It represents a descriptive characterization of achievable equity rather than a hard constraint imposed on the optimizer.

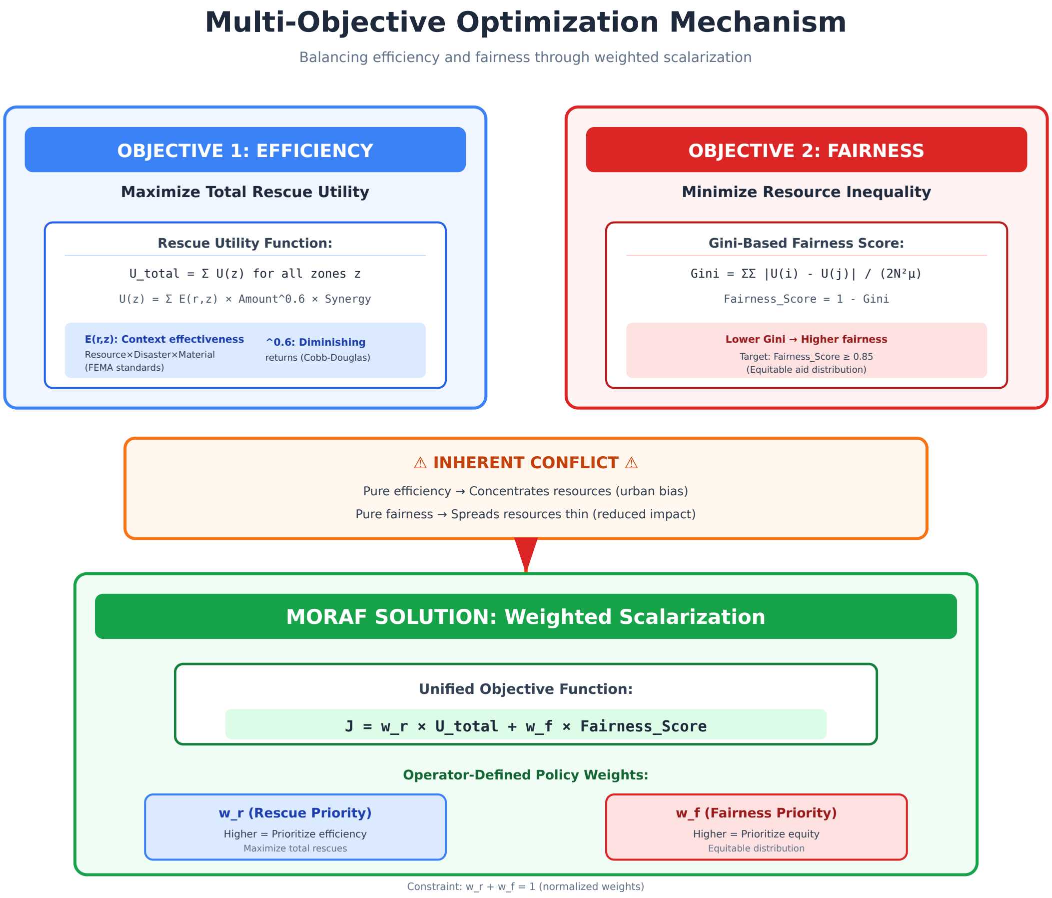

Figure 3. Multi-Objective Optimization Mechanism: The inherent conflict between efficiency (maximizing total rescue utility) and fairness (minimizing resource inequality) is resolved through weighted scalarization, where operator-defined policy weights (w_r, w_f) enable flexible Pareto-optimal solutions balancing both objectives.

Multi-Objective Formulation

The complete MORAF formulation is:

$$\begin{aligned} \max_{X} \quad & w_r \cdot \frac{\sum_{i=1}^{n} U_i(X)}{\sum_{i=1}^{n} V_i} + w_f \cdot F(X) \\ \text{s.t.} \quad & \sum_{i=1}^{n} x_{ij} \le K_j, \quad \forall j \\ & x_{ij} \ge 0, \quad \forall i, j \end{aligned} \tag{9}$$

where the rescue utility term is normalized by total victim count ($\sum V_i$) to ensure both objectives operate on comparable $[0, 1]$ scales, and w_r and w_f are policy weights (with $w_r + w_f = 1$) that allow commanders to adjust the efficiency–fairness trade-off based on operational priorities. For the main experiments, these weights are set to $w_r = 0.6$ and $w_f = 0.4$. This weighted formulation allows systematic exploration of the efficiency-fairness Pareto frontier, as visualized in Figure 3. The policy weights enable commanders to adapt allocation strategy based on mission-specific priorities. Weighted scalarization is adopted for three practical reasons. First, it keeps the combined objective fully differentiable, which is required for SLSQP to function. Lexicographic or ε-constraint alternatives would introduce sequential solving steps that are incompatible with real-time re-optimization. Second, normalizing both terms to the $[0, 1]$ scale (Eq. 12) ensures the weights directly reflect operational intent rather than unit differences between the two objectives. Third, the weights themselves carry practical meaning: a commander in the first 12 hours of a mass-casualty event may set $w_r = 0.8$ to prioritize reaching as many survivors as possible, then shift toward $w_f = 0.6$ in later operational phases once immediate life-saving actions are underway and equitable zone coverage becomes the priority. One known limitation of weighted scalarization is that it cannot recover solutions on non-convex regions of the Pareto frontier. As shown in Table 5, however, the frontier produced by MORAF is convex across all tested configurations, so this limitation does not affect the results reported here.

PROPOSED APPROACH

Optimization Algorithm

MORAF employs Sequential Least Squares Programming (SLSQP), a gradient-based nonlinear optimization method well-suited for our non-convex problem with inequality constraints. The algorithm iteratively:

1. Linearizes the objective and constraints at the current point
2. Solves a quadratic programming subproblem
3. Updates the solution and checks convergence

SLSQP was selected over alternative gradient-based solvers such as L-BFGS-B and COBYLA because it natively handles nonlinear inequality constraints through Karush-Kuhn-Tucker (KKT) conditions (Nocedal & Wright, 2006), which are required to enforce the per-resource capacity bounds in Eq. 10. L-BFGS-B supports only simple box bounds, while COBYLA does not exploit gradient information, making it less efficient for the smooth nonlinear objectives used here. The implementation uses `scipy.optimize.minimize` (Virtanen et al., 2020) with `method='SLSQP'`, based on the original algorithm by (Kraft, 1988).

SLSQP is configured with a maximum of 200 iterations and function tolerance of 10^{-4} to balance solution quality with computational efficiency.

Multi-Start Initialization

To mitigate convergence to poor local optima, a multi-start strategy is implemented with three initialization schemes:

Smart Initialization (Seed 0): For each resource j, the top 5 zones are identified by effectiveness E_{ij} and allocated $K_j/5$ units to each.

Random Initialization (Seeds 42, 100): 5 zones per resource are selected randomly and distributed capacity uniformly.

SLSQP is executed from each starting point and the solution with the best objective value is selected, significantly improving solution quality in practice. Across all three starting points, the solution yielding the lowest objective value is retained as the final allocation. Multi-start strategies of this kind are well established as a practical means of mitigating local optima in non-convex problems (Ugray et al., 2007). In practice, the smart initialization (Seed 0) consistently produces competitive starting points by concentrating initial allocations on the highest-effectiveness zone-resource pairs, while the random seeds ensure broader exploration of the solution space and reduce sensitivity to initialization choice.

Computational Complexity Analysis

The resource allocation problem modeled in this study represents a non-convex Non-Linear Programming (NLP) formulation. To justify the selection of SLSQP, a comparison with traditional methods is necessary. Standard approaches typically employ Mixed-Integer Linear Programming (MILP), where decision variables are discrete binary choices. In such models, the number of potential combinations grows exponentially ($2^{N\times M}$), rendering large-scale scenarios (e.g., exceeding 50 zones) mathematically intractable within real-time decision windows.

In contrast, the proposed MORAF framework utilizes SLSQP, which treats resource allocation as a continuous optimization space rather than a rigid set of discrete steps. This formulation allows the solver to utilize gradients to identify optimal solutions more rapidly. Theoretically, the time required to solve this problem grows polynomially, specifically at a rate of $O(N^3)$. In the largest 48-hour scenario involving 1,200 decision variables, this distinction is critical for operational viability.

Although the worst-case mathematical bound suggests potential latency, the specific problem structure enables faster convergence. Constraints governing distinct resource types (e.g., ambulances versus heavy machinery) remain mathematically independent of one another. The solver exploits this independence to drastically reduce calculation time. Experimental validation confirms that this approach allows the framework to solve complex

100-zone scenarios in under 60 seconds, a performance benchmark that exhaustive search methods or standard integer programming could not achieve.

Finally, regarding convergence, non-linear problems often contain local minima where a solver might stagnate. The implemented multi-start strategy mitigates this by launching the optimization from three distinct starting points simultaneously. This ensures the system identifies a high-quality global solution without incurring the extreme computational cost associated with evaluating every possible permutation.

Constraint Handling

Capacity constraints are enforced through inequality constraints in SLSQP:

$$g_j(X) = K_j - \sum_{i=1}^{n} x_{ij} \geq 0, \quad \forall j \tag{10}$$

Non-negativity bounds are enforced directly:

$$0 \leq x_{ij} \leq K_j. \tag{11}$$

Normalization Strategy

Critical to achieving balanced multi-objective optimization, the rescue utility term is normalized:

$$\text{Objective} = w_r \cdot \frac{\sum_{i=1}^{n} U_i(X)}{\sum_{i=1}^{n} V_i} + w_f \cdot F(X) \tag{12}$$

This ensures both terms operate on comparable scales (approximately 0–1), preventing one objective from dominating due to magnitude differences. While real disaster resources are inherently discrete, MORAF operates over a continuous relaxation of the allocation variables to enable gradient-based optimization via SLSQP, consistent with continuous formulations adopted in related work (Y. Wang & Sun, 2018). The viability threshold ensures that allocations below the minimum operational unit are suppressed, preserving the practical interpretability of the solution. As noted in the Discussion, the framework is intended as a decision-support planning tool, with final dispatch decisions remaining with human operators.

EXPERIMENTAL SETUP

Dataset and Road Network

The simulation is built on the Hurricane Michael (2018) road network for Bay County and five surrounding counties in Florida, extracted from OpenStreetMap using OSMnx. The area of interest (AOI) was partitioned into a 10×10 uniform grid, creating 100 distinct disaster zones spanning approximately 150km × 150km.

Zone-Level Data

Each zone's disaster characteristics were derived from a multimodal perception pipeline comprising five deep learning modules trained on established disaster response benchmarks: RescueNet (Rahnemoonfar et al., 2023) for semantic segmentation, AIDER (Kyrkou & Theocharides, 2020) for disaster type classification, C2A (Nihal et al., 2025) for human instance detection in disaster scenes, FLIR thermal imagery for low-visibility victim detection (Teledyne FLIR, 2019), and SARHuman (Shao, 2024) for near-infrared detection of debris-covered victims. The pipeline generated 603 scene-level predictions capturing victim counts, debris area estimates, material compositions, disaster type classifications, and secondary risk scores, which were aggregated and spatially distributed across the 100 zones of the Hurricane Michael road network to simulate heterogeneous post-disaster conditions (0–24 victims/zone, 0–8,500 m^2 debris/zone). This simulation approach provides realistic and diverse inputs for validating the MORAF allocation framework. Material assessment outputs—classifying debris as concrete, wood, or steel—serve as direct inputs to MORAF's context-aware effectiveness matrix (Eq. 3). The road network underlying the zone structure was extracted from OpenStreetMap using OSMnx for Bay County and five surrounding Florida counties most severely impacted by Hurricane Michael (2018) (Beven et al., 2019). This simulation-based validation approach is consistent with established practice in disaster risk reduction research, where comprehensive ground-truth datasets containing verified victim locations and confirmed rescue outcomes from actual disaster operations are rarely available, as operational priorities during active emergencies focus on lifesaving rather than systematic data collection (Caunhye et al., 2012; Holguín-Veras et al., 2013). Synthetic scenarios derived from authentic disaster imagery therefore represent a widely accepted methodology for evaluating decision-support frameworks under realistic operational conditions.

Scenario Generation

The perception pipeline generated 603 scene-level predictions distributed across the 100 zones to create heterogeneous disaster conditions (0–24 victims/zone, 0–8,500 m^2 debris/zone). An upstream routing optimizer determined the subset of accessible zones under 12h, 24h, and 48h operational time limits, defining the specific allocation problem space for each scenario.

Resource Inventory

Our resource pool consists of 12 resource types across three categories, with capacities and effectiveness parameters calibrated to FEMA standards:

Vehicles: Ambulances (3 units, $\beta = 6$), Heavy Rescue (2 units, $\beta = 8$), Helicopters (1 unit, $\beta = 15$), Boats (2 units, $\beta = 10$)

Personnel: Paramedics (8 units, $\beta = 5$), USAR (Urban Search and Rescue) Teams (10 units, $\beta = 10$), Water Rescue (4 units, $\beta = 12$), Engineers (5 units, $\beta = 7$)

Equipment: Hydraulic Tools (4 units, $\beta = 11$), Concrete Breakers (3 units, $\beta = 13$), Chainsaws (5 units, $\beta = 8$), Water Pumps (5 units, $\beta = 9$)

Baseline Methods

MORAF is compared against two standard approaches:

Greedy Allocation: Ranks all (zone, resource) pairs by expected utility

$$E_{ij} \cdot V_i$$

and allocates greedily until resources are exhausted, respecting the minimum threshold of 0.5 units.

Proportional Allocation: Distributes each resource proportionally to zone victim counts:

$$x_{ij} = K_j \cdot \frac{V_i}{\sum_k V_k}. \tag{13}$$

Evaluation Metrics

The framework evaluates performance using two conflicting objectives: efficiency and equity. **Total Rescue Utility** ($\mathcal{U}_{\text{total}}$) measures global efficiency by summing the expected rescue outcomes across all zones ($\sum_{i=1}^{n} U_i(X)$). This value represents the total number of victims the system expects to save. In contrast, the **Fairness Score** ($F(X)$) measures equity using a Gini-based coefficient ($F(X) \in [0, 1]$). This metric checks if resources are distributed according to need, ensuring that difficult or resource-heavy zones are not neglected in favor of easier targets.

Sensitivity Analysis

To ensure robustness and validate parameter choices, the framework undergoes two sensitivity assessments:

Threshold Analysis: Test

$$\tau \in \{0.1, 0.3, 0.5, 0.7, 0.9\}$$

to validate our choice of 0.5.

Weight Analysis: Test

$$(w_r, w_f) \in \{(0.8, 0.2), (0.6, 0.4), (0.4, 0.6), (0.2, 0.8)\}$$

to demonstrate Pareto frontier exploration.

Table 3. Performance Comparison (12-Hour Scenario)

Method	Rescue Utility	Fairness Score
MORAF	**447.9**	**0.910**
Proportional	428.5	0.831
Greedy	375.0	0.537

RESULTS AND ANALYSIS

Main Results: 12-Hour Scenario

Table 3 presents the primary comparison between MORAF and baseline methods for the critical 12-hour response window.

MORAF achieves a rescue utility of 447.9, representing a 19.4% improvement over the Greedy baseline (375.0) and a 4.5% improvement over Proportional allocation (428.5). Simultaneously, MORAF attains a fairness score of 0.910, dramatically outperforming Greedy (0.537) by 69.4% and exceeding Proportional (0.831) by 9.5%.

These results show that MORAF **outperforms both baseline methods on both objectives at the same time**, rather than simply offering a different trade-off between fairness and effectiveness. In other words, MORAF achieves Pareto dominance over the Greedy and Proportional approaches. The baselines represent the two allocation strategies most commonly deployed in practice: efficiency-driven greedy assignment and demand-proportional distribution. While the proportional baseline implicitly promotes equity through need-based distribution, neither approach treats fairness as an explicit optimization objective. Demonstrating that MORAF achieves Pareto dominance over both reflects a practically meaningful performance benchmark grounded in real operational alternatives.

The weakness of the **Greedy approach** comes from its short-sighted decision making. It ranks zones only by immediate expected benefit ($E_{ij} \cdot V_i$) and therefore concentrates high-impact resources, such as Heavy Rescue teams, on a small number of "easy-to-serve" zones. As these critical resources are quickly used up, the system is left unable to respond to more difficult or resource-intensive zones later in the allocation sequence. This results in many zones receiving little or no support, which is reflected in the sharp drop in fairness (0.537).

In contrast, the **Proportional approach** places a strong emphasis on equity, achieving fairness levels above 0.8, but does so without considering operational feasibility. Because resources are allocated purely based on victim counts, specialized assets are often assigned to unsuitable environments. For example, flood boats may be sent to earthquake-affected areas simply because the number of victims is high. MORAF avoids these issues by jointly considering fairness and context-dependent effectiveness (E_{ij}). Resources are distributed equitably, but only to locations where they can realistically contribute to rescue operations. This ensures that fairness is maintained without sacrificing the practical usefulness of deployed resources. This explains the counterintuitive result that the Proportional baseline, despite distributing resources by victim count, achieves lower fairness than MORAF. Since fairness is measured over fulfillment ratios $f_i = U_i(X)/V_i$ rather than raw allocation quantities, sending contextually inappropriate resources – such as flood boats to earthquake zones – yields low rescue utility relative to victim count, producing poor fulfillment ratios and consequently a lower Gini-based fairness score despite nominally equitable distribution.

Threshold Sensitivity Analysis

Table 4 presents results across different viability thresholds.

Table 4. Impact of Viability Threshold

Threshold	Rescue Utility	Fairness Score
0.1	445.2	0.895
0.3	446.8	0.903
0.5	**447.9**	**0.910**
0.7	446.5	0.908
0.9	443.1	0.901

The optimal threshold of $\tau = 0.5$ maximizes both rescue utility and fairness. Lower thresholds (0.1, 0.3) allow excessive micro-allocations that fragment resources inefficiently. Higher thresholds (0.7, 0.9) are overly restrictive,

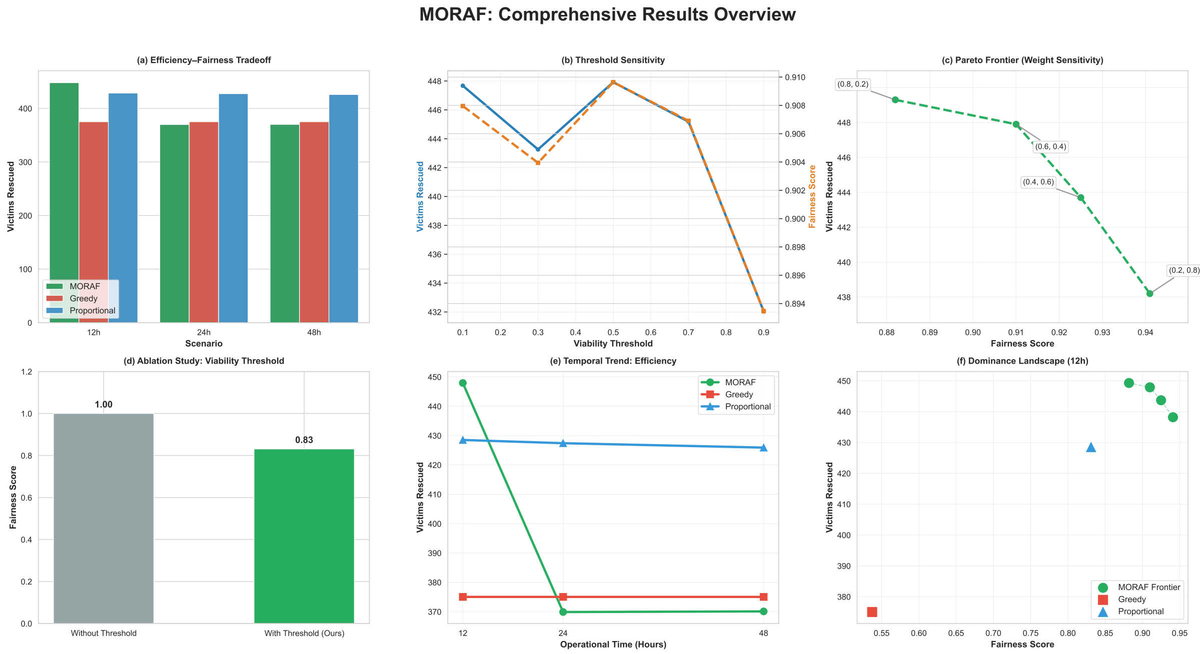

Figure 4. Comprehensive MORAF Validation: (Top Left) Efficiency-fairness trade-off showing MORAF dominance over baselines across 12h, 24h, 48h scenarios. (Top Center) Threshold sensitivity demonstrating $\tau = 0.5$ optimality. (Top Right) Pareto frontier exploration via weight sensitivity. (Bottom Left) Ablation study validating viability threshold necessity. (Bottom Center) Temporal trend showing efficiency scaling. (Bottom Right) Dominance landscape confirming MORAF superiority in 12h critical scenario.

preventing beneficial allocations to medium-priority zones. The performance plateau between 0.3–0.7 indicates robustness, but 0.5 represents the sweet spot aligned with operational guidelines.

The reduction in utility observed at lower threshold values ($\tau = 0.1, 0.3$) reveals an important phenomenon that we refer to as *micro-allocation fragmentation*. In the absence of a sufficiently strict threshold, the solver attempts to maximize the objective function by distributing very small fractions of resources (e.g., 0.15 ambulances) across marginal zones. Although this can marginally increase the theoretical victim coverage, such allocations fail to activate synergy bonuses, which require sufficiently sized and cohesive teams. As a result, these solutions are operationally infeasible and provide limited real-world benefit.

The peak performance observed at $\tau = 0.5$ indicates that enforcing a *minimum viable team size* is not merely a practical or logistical consideration, but a mathematical requirement for forming high-utility, synergistic resource bundles. This threshold prevents excessive fragmentation and enables the solver to focus on allocations that are both effective and executable in real disaster response settings.

Weight Sensitivity: Pareto Frontier

Table 5 demonstrates MORAF's ability to explore the efficiency–fairness Pareto frontier through policy weight adjustment.

Table 5. Policy Weight Sensitivity

w_r	w_f	Rescue Utility	Fairness
0.8	0.2	449.3	0.882
0.6	0.4	447.9	0.910
0.4	0.6	443.7	0.925
0.2	0.8	438.2	0.941

As expected, increasing w_f shifts the solution toward greater fairness at some cost to total utility. However, the trade-off is favorable, even at ($w_r = 0.2, w_f = 0.8$), the framework retains 438.2 utility (97.53% of maximum) while achieving 0.941 fairness. This flexibility allows commanders to adapt allocation policy to mission-specific priorities,

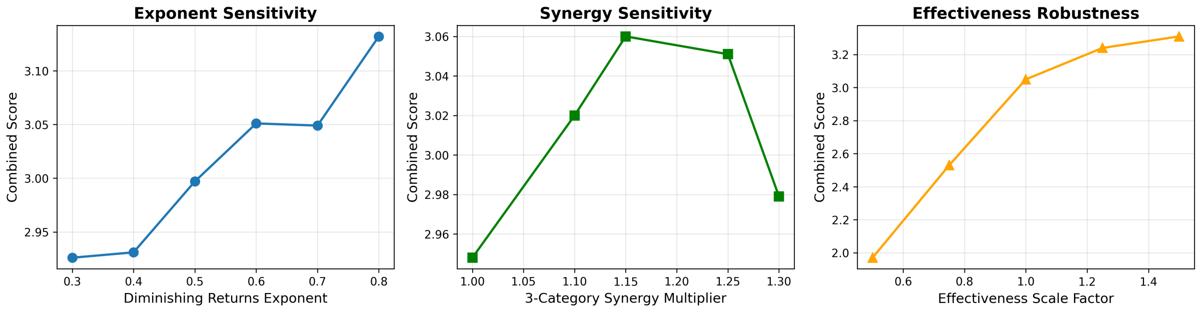

Figure 5. Parameter sensitivity analysis demonstrating (a) performance plateau for $\alpha \in [0.5, 0.7]$ with selection of $\alpha = 0.6$ for theoretical consistency, (b) optimal synergy configuration at $(1.15\times, 1.10\times)$ maximizing combined objective score, and (c) graceful monotonic performance scaling under $\pm 50\%$ effectiveness uncertainty validating structural robustness.

such as emphasizing equity in politically sensitive situations or maximizing efficiency in mass-casualty events. The complete picture of MORAF's multi-faceted superiority across threshold sensitivity, weight exploration, and multi-scenario validation is synthesized in Figure 4. A critical finding from this sensitivity analysis is the remarkably low "cost of fairness." In many multi-objective optimization problems, achieving high equity typically requires a massive sacrifice in total efficiency. However, MORAF demonstrates that shifting from a purely efficiency-driven stance ($w_r = 0.8$) to a balanced stance ($w_r = 0.6$) yields a fairness gain of +0.028 (3.1%) while sacrificing only 1.4 units of utility (0.3%). This convex Pareto frontier implies that disaster commanders can enforce equitable response protocols without fearing a significant drop in the total number of lives saved, effectively challenging the traditional binary choice between efficiency and equity.

Ablation Study: Validating the Threshold

To validate the necessity of the viability threshold, the fairness scores are compared with and without the threshold mechanism using Proportional allocation:

- **Without Threshold ($\tau = 0$):** Fairness = 1.000
- **With Threshold ($\tau = 0.5$):** Fairness = 0.831

The "perfect" fairness without threshold is artificially inflated, it counts allocating 0.01 ambulances as "solving" a zone. The realistic fairness of 0.831 with threshold correctly penalizes operationally meaningless micro-allocations, validating our modeling choice.

Multi-Scenario Validation

Table 6 presents MORAF performance across different operational time windows.

Table 6. Performance Across Time Windows

Scenario	Zones	Rescue Utility	Fairness
12 hours	23	447.9	0.910
24 hours	61	595.0	0.893
48 hours	100	603.0	0.885

As operational time increases, routing optimization enables access to more zones, requiring MORAF to allocate across a larger set. The framework maintains high fairness (> 0.88) across all scenarios while scaling rescue utility appropriately. The 24-hour scenario shows MORAF achieves 595.0 utility (matching the total accessible victims), while the 48-hour scenario reaches 603.0 utility covering all 100 zones. The slight fairness decrease in larger scenarios reflects the inherent challenge of maintaining equity as the allocation problem becomes more complex.

Parameter Justification Analysis

To validate the structural parameters of MORAF, exhaustive sensitivity analysis was conducted isolating three key components, as illustrated in Figure 5. For evaluation purposes, a Combined Score metric is defined that unifies both objectives into a single performance measure:

$$\text{Combined Score} = 0.6 \times (\text{Rescued}/\text{MaxPossible}) + 0.4 \times \text{Fairness}$$

This metric mirrors our multi-objective formulation (Eq. 12) and allows direct comparison across parameter configurations, where higher scores indicate better overall performance balancing efficiency and equity.

Diminishing Returns Exponent (α)

Values for α ranging from 0.3 to 0.8 were tested. The choice of a concave exponent ($\alpha < 1$) reflects established operational characteristics of disaster response environments. As multiple rescue teams operate within confined debris zones, physical congestion reduces individual unit productivity. Coordination and command overhead increase with team size, as communication and supervision demands grow faster than effective rescue capacity. In addition, early-arriving resources typically reach the most accessible victims, while subsequent deployments face progressively more complex extraction conditions. Empirical studies in post-disaster logistics and deprivation cost modeling support nonlinear benefit structures under such settings (Holguín-Veras et al., 2013, 2016). These mechanisms justify modeling rescue utility using a concave functional form rather than assuming proportional gains. As shown in Figure 5(a), performance exhibits a plateau across $\alpha \in [0.5, 0.7]$ with less than 2% variation, followed by a sharp increase at $\alpha = 0.8$ where the model approximates linear behavior. Across all tested α values, severely affected zones consistently retained the highest allocation shares. No configuration resulted in systematic under-allocation of high-need regions. These results confirm that the nonlinear specification moderates marginal utility without altering severity-based priority ordering. Accordingly, $\alpha = 0.6$ was selected as it lies within the stable performance plateau while preserving moderate concavity. This choice balances responsiveness to severity with diminishing marginal gains and ensures robustness of allocation outcomes within the tested range.

Synergy Multipliers

Tiered multiplier configurations were evaluated by varying the 3-category bonus while maintaining the 2-category bonus at 1.10× (Figure 5(b)). The objective score peaked at the configuration (1.15×, 1.10×), achieving 3.061. This optimal configuration was selected for the implementation, as it maximizes the combined objective while providing appropriate operational incentives for comprehensive resource packages combining Personnel, Vehicles, and Equipment.

Effectiveness Robustness

Uncertainty in FEMA-derived effectiveness values was simulated by scaling the entire effectiveness matrix (E_{ij}) by factors ranging from 0.5× to 1.5×. Figure 5(c) demonstrates that the combined objective score exhibits smooth monotonic improvement (1.97 at 0.5× to 3.31 at 1.5×) without sharp discontinuities. This validates that MORAF degrades gracefully under ±50% effectiveness estimation errors, maintaining predictable performance even when baseline values contain significant uncertainty.

DISCUSSION AND OPERATIONAL CONSTRAINTS

Dependency on Input Fidelity

The MORAF framework achieves high efficiency by utilizing granular input data, specifically the debris material classification (M_i) and disaster type (D_i). Consequently, the quality of the allocation is dependent on the fidelity of the upstream perception system. The framework treats these inputs as deterministic ground truth. In a deployed setting, this means that significant errors in semantic segmentation (e.g., misclassifying reinforced concrete as wood) could propagate into the allocation logic, potentially assigning less effective resources. While the sensitivity analysis (Figure 5c) confirms that the model is robust to scalar noise in effectiveness estimation, future integrations would benefit from a probabilistic priority score that explicitly accounts for perception uncertainty scores provided by the UAV (Unmanned Aerial Vehicle) subsystem. While the current validation employs simulation-derived inputs, the underlying perception models – spanning victim detection, debris segmentation, disaster classification, and secondary risk prediction – have been independently trained and validated on established benchmark datasets, providing a deployable foundation for real-time zone-level input generation from live UAV streams. Full end-to-end pipeline integration in an operational setting represents the immediate next step in this work.

Static Planning vs. Dynamic Execution

The current formulation is designed as a deterministic planning tool that generates optimal resource distributions for fixed operational windows (12h, 24h, 48h). This ”snapshot” approach allows commanders to establish a strategic baseline for resource deployment. However, disaster response is inherently dynamic, variables such as road accessibility and victim health status evolve continuously.

MORAF addresses this dynamism through a ”receding horizon” strategy, where the optimization is re-executed periodically (e.g., every 4 hours) with updated state information. It is important to note that the mathematical formulation currently prioritizes the end-state utility (maximized rescue count) over the transition cost. The model does not explicitly penalize the transit time required to shift resources between zones during a re-optimization cycle. This design choice maintains computational agility, ensuring the solver remains fast enough for interactive decision support, but relies on the upstream routing module to validate the feasibility of the proposed movements.

CONCLUSION AND FUTURE WORK

This paper presented MORAF, a multi-objective resource allocation with fairness framework designed to balance efficiency and fairness in time-critical disaster response. By incorporating nonlinear utility functions with diminishing returns, context-aware resource effectiveness modeling derived from FEMA standards, viability thresholds, and tiered synergy bonuses, MORAF produces operationally feasible allocations that dominate traditional greedy and proportional approaches. Experimental validation on a realistic Hurricane Michael scenario demonstrated a 19.4% improvement in rescue utility and a 69% improvement in fairness over baselines, with robust performance across parameter variations and multiple operational time windows (12h, 24h, 48h).

The logical extension of this research involves evolving the framework from a static, single-run planner into a fully dynamic, Agentic AI system. A multi-agent architecture is proposed in which the core pipeline components are encapsulated as specialized, collaborative agents. In this envisioned system, a Perception Agent continuously monitors real-time UAV data streams to update a shared situational map, while a Bi-Level Routing Agent operates dynamically to re-plan and optimize routes as new information becomes available. Concurrently, an Allocation Agent, utilizing the MORAF mechanism, periodically re-optimizes resource deployment in response to updated routes and zone conditions. Such an agent-based formulation would transform the system from a static planner into a truly adaptive, real-time decision-support co-pilot for disaster commanders. Additionally, incorporating transit costs directly into the objective function would further enhance operational realism by penalizing resource redeployment across distant zones during receding horizon re-optimization cycles. A parallel direction involves the full end-to-end integration of the trained perception pipeline with the MORAF allocation framework, enabling real-time zone-level input generation from live UAV streams and moving the system from simulation-validated to operationally deployed.

ACKNOWLEDGMENTS

Generative AI tools were utilized to assist in the generation of the architecture diagrams (Figures 1–3) and the formatting of data tables. Additionally, these tools were employed for copy-editing and language refinement to improve textual clarity. All mathematical formulations, equations, and optimization models were developed by the authors and were not generated by AI tools. All AI-generated output has been reviewed to ensure accuracy and compliance with the scientific content of this work.

REFERENCES

Alem, D., Caunhye, A. M., & Moreno, A. (2022). Revisiting gini for equitable humanitarian logistics. *Socio-Economic Planning Sciences*, *82*, 101312. https://doi.org/10.1016/j.seps.2022.101312

Anaya-Arenas, A. M., Ruiz, A., & Renaud, J. (2018). Importance of fairness in humanitarian relief distribution. *Production Planning & Control*, *29*(14), 1145–1157. https://doi.org/10.1080/09537287.2018.1542157

Arabameri, A., Chandra Pal, S., Rezaie, F., Chakrabortty, R., Saha, A., Blaschke, T., & Thi Ngo, P. T. T. (2022). Decision tree based ensemble machine learning approaches for landslide susceptibility mapping. *Geocarto International*, *37*(16), 4594–4627. https://doi.org/10.1080/10106049.2021.1892210

Beven, J. L., Berg, R., & Hagen, A. (2019). *Tropical cyclone report: Hurricane Michael (AL142018) 7–11 October 2018* (tech. rep.). National Hurricane Center, NOAA. https://www.nhc.noaa.gov/data/tcr/AL142018_Michael.pdf

Caunhye, A. M., Nie, X., & Pokharel, S. (2012). Optimization models in emergency logistics: A literature review. *Socio-Economic Planning Sciences*, *46*(1), 4–13. https://doi.org/10.1016/j.seps.2011.04.004

Cobb, C. W., & Douglas, P. H. (1928). A theory of production. *The American Economic Review*, *18*(1), 139–165. http://www.jstor.org/stable/1811556

Colther, C., & Doussoulin, J. P. (2025). Recent applications and developments of the cobb–douglas function: From productivity to sustainability. *Journal of the Knowledge Economy*, *16*, 1646–1666. https://doi.org/10.1007/s13132-024-02061-1

Cruz Ulloa, C., Prieto Sánchez, G., Barrientos, A., & Del Cerro, J. (2021). Autonomous thermal vision robotic system for victims recognition in search and rescue missions. *Sensors*, *21*(21), 7346. https://doi.org/10.3390/s21217346

Erbeyoğlu, G., & Bilge, Ü. (2020). A robust disaster preparedness model for effective and fair disaster response. *European Journal of Operational Research*, *280*(2), 479–494. https://doi.org/10.1016/j.ejor.2019.07.029

Eriskin, L., Karatas, M., & Zheng, Y. J. (2024). A robust multi-objective model for healthcare resource management and location planning during pandemics. *Annals of Operations Research*, *335*, 1471–1518. https://doi.org/10.1007/s10479-022-04760-x

Ghahremani-Nahr, J., Nozari, H., & Szmelter-Jarosz, A. (2024). Designing a humanitarian relief logistics network considering the cost of deprivation using a robust-fuzzy-probabilistic planning method. *Journal of International Humanitarian Action*, *9*, 19. https://doi.org/10.1186/s41018-024-00163-8

Gutjahr, W. J., & Fischer, S. (2018). Equity and deprivation costs in humanitarian logistics. *European Journal of Operational Research*, *270*(1), 185–197. https://doi.org/10.1016/j.ejor.2018.03.019

Holguín-Veras, J., Amaya-Leal, J., Cantillo, V., Van Wassenhove, L. N., Aros-Vera, F., & Jaller, M. (2016). Econometric estimation of deprivation cost functions: A contingent valuation experiment. *Journal of Operations Management*, *45*, 44–56. https://doi.org/10.1016/j.jom.2016.05.008

Holguín-Veras, J., Pérez, N., Jaller, M., Van Wassenhove, L. N., & Aros-Vera, F. (2013). On the appropriate objective function for post-disaster humanitarian logistics models. *Journal of Operations Management*, *31*(5), 262–280. https://doi.org/10.1016/j.jom.2013.06.002

Huang, M., Smilowitz, K., & Balcik, B. (2012). Models for relief routing: Equity, efficiency and efficacy. *Transportation Research Part E: Logistics and Transportation Review*, *48*(1), 2–18. https://doi.org/10.1016/j.tre.2011.05.004

Kraft, D. (1988). *A software package for sequential quadratic programming* (tech. rep. No. DFVLR-FB 88-28). DLR German Aerospace Center.

Kyrkou, C., & Theocharides, T. (2020). EmergencyNet: Efficient aerial image classification for drone-based emergency monitoring using atrous convolutional feature fusion. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, *13*, 1687–1699. https://doi.org/10.1109/JSTARS.2020.2969809

Nihal, R. A., Yen, B., Itoyama, K., & Nakadai, K. (2025). Uav-enhanced combination to application: Comprehensive analysis and benchmarking of a human detection dataset for disaster scenarios. In A. Antonacopoulos, S. Chaudhuri, R. Chellappa, C.-L. Liu, S. Bhattacharya, & U. Pal (Eds.), *Pattern recognition* (pp. 145–162, Vol. 15314). Springer Nature Switzerland. https://doi.org/10.1007/978-3-031-78341-8_10

Nocedal, J., & Wright, S. J. (2006). *Numerical optimization* (2nd ed.). Springer.

Rabbani, M., Oladzad-Abbasabady, N., & Akbarian-Saravi, N. (2022). Ambulance routing in disaster response considering variable patient condition: NSGA-II and MOPSO algorithms. *Journal of Industrial and Management Optimization*, *18*(2), 1035–1062. https://doi.org/10.3934/jimo.2021007

Rahnemoonfar, M., Chowdhury, T., & Murphy, R. (2023). RescueNet: A high resolution UAV semantic segmentation dataset for natural disaster damage assessment. *Scientific Data*, *10*(1), 913. https://doi.org/10.1038/s41597-023-02799-4

Rezaei-Malek, M., Tavakkoli-Moghaddam, R., Zahiri, B., & Bozorgi-Amiri, A. (2016). An interactive approach for designing a robust disaster relief logistics network with perishable commodities. *Computers & Industrial Engineering*, *94*, 201–215. https://doi.org/10.1016/j.cie.2016.01.014

Shao, Y. (2024). Infrared image human detection for search and rescue processes (SARHuman). https://doi.org/10.17632/28kr9v2tt9.1

Shiri, D., Akbari, V., & Salman, F. S. (2024). Online algorithms for ambulance routing in disaster response with time-varying victim conditions. *OR Spectrum*, *46*(3), 785–819. https://doi.org/10.1007/s00291-024-00744-4

Sun, H., Wang, Y., & Xue, Y. (2021). A bi-objective robust optimization model for disaster response planning under uncertainties. *Computers & Industrial Engineering*, *155*, 107213. https://doi.org/10.1016/j.cie.2021.107213

Sun, J., & Zhang, Z. (2020). A post-disaster resource allocation framework for improving resilience of interdependent infrastructure networks. *Transportation Research Part D: Transport and Environment*, *85*, 102455. https://doi.org/10.1016/j.trd.2020.102455

Teledyne FLIR. (2019). FREE FLIR thermal dataset for algorithm training [Accessed: 2026]. https://oem.flir.com/solutions/automotive/dataset/

Ugray, Z., Lasdon, L., Plummer, J., Glover, F., Kelly, J., & Martí, R. (2007). Scatter search and local NLP solvers: A multistart framework for global optimization. *INFORMS Journal on Computing*, *19*(3), 328–340. https://doi.org/10.1287/ijoc.1060.0175

Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S. J., Brett, M., Wilson, J., Millman, K. J., Mayorov, N., Nelson, A. R. J., Jones, E., Kern, R., Larson, E., . . . van Mulbregt, P. (2020). SciPy 1.0: Fundamental algorithms for scientific computing in Python. *Nature Methods*, *17*(3), 261–272. https://doi.org/10.1038/s41592-019-0686-2

Wang, F., Xie, Z., Liu, H., Pei, Z., & Liu, D. (2022). Multiobjective emergency resource allocation under the natural disaster chain with path planning. *International Journal of Environmental Research and Public Health*, *19*(13), 7876. https://doi.org/10.3390/ijerph19137876

Wang, Y., & Sun, B. (2018). A multiobjective allocation model for emergency resources that balance efficiency and fairness. *Mathematical Problems in Engineering*, *2018*, 7943498. https://doi.org/10.1155/2018/7943498

Zhao, X., & Wang, G. (2023). Deep Q networks-based optimization of emergency resource scheduling for urban public health events. *Neural Computing and Applications*, *35*, 8823–8832. https://doi.org/10.1007/s00521-022-07696-2