

Toward Agentic Artificial Intelligence (AI) for Disaster Relief Resource Allocation: An Equity-Aware Multi-Agent Negotiation Framework

Khushvir Singh ¹, 

1. Data Science, ML, DL, NLP, AI, GEN AI, and Agentic AI, Independent Researcher, Jalandhar, IND

Received: August 04, 2026 | Review began: August 12, 2026 | Review ended: August 25, 2026 | Published: September 23, 2026

© Copyright 2026

This is an open access article distributed under the terms of the Creative Commons Attribution License CC-BY 4.0., which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Abstract

Disasters routinely outpace the teams responding to them, and the resource allocation problem at the centre of that response is genuinely hard: scarce medical supplies, food, and shelter must be split across sites that differ in severity, reachable only through a road network that can fail mid-response, and that compete for the same trucks and warehouses. Most multi-agent frameworks in the recent agentic-artificial intelligence literature either compare orchestration tools qualitatively or describe a vision-level architecture without reporting how well the resulting allocation performs. We build a disaster relief resource-allocation environment structured after a published benchmark task for disaster relief under uncertainty and implement a four-agent negotiation framework: a Triage Agent that scores urgency, an Equity Agent that watches for sites falling behind, a Logistics Agent that turns those signals into vehicle dispatches, and a Disruption Agent that reacts to road closures or depot loss; each agent's logic is a simple, auditable rule rather than a large language model (LLM) call, a deliberate design choice we justify and revisit in the discussion. We compare this framework against random, round-robin, and greedy-priority baselines over 90 scenario instances, scored on six standard multi-agent planning metrics. The proposed framework raises goal satisfaction by 44.3% over the strongest baseline (Welch's $t = 11.58$, $p < 0.001$) and shows similarly large gains in resource utilisation and disruption recovery, while trading away some of the greedy method's strict severity-first ordering, a trade-off we discuss rather than hide. We close with the limitations of a heuristic-driven simulation and a concrete path toward grounding the agents' decision rules in LLM reasoning for field deployment.

Categories: AI applications, AI/ML-based decision support systems, Disaster Recovery and Business Continuity

Keywords: agentic ai, multi-agent systems, disaster relief, humanitarian logistics, resource allocation, large language model agents

Introduction

When a flood, an earthquake, or a large-scale storm hits a region, the first 72 hours are usually the ones that decide how many people get help in time [1]. Response teams increasingly have access to data (satellite imagery, crowd-sourced reports, weather feeds, and hospital capacity dashboards all exist), but they are still, in many cases, short of ways to turn that data into a resource allocation decision fast and fair enough to matter. Li et al. [2] have pointed it out in their proposal for a "Disaster Copilot": disaster response today is hindered by fragmented data streams, siloed technologies, and resource constraints that collectively slow down exactly the decisions that most need to be fast.

The last 2 years have seen a wave of interest in agentic artificial intelligence (AI), systems built from large language models (LLMs) that can plan, call tools, and act with some degree of autonomy rather than simply answering a question, and disaster management has been an obvious, if still early, target for this technology. Choure and Prajapat [3] compared

How to cite this article:

Singh K (September 23, 2026) Toward Agentic Artificial Intelligence (AI) for Disaster Relief Resource Allocation: An Equity-Aware Multi-Agent Negotiation Framework. Cureus J Comput Sci 3 : es44389-026-00290-w. DOI <https://doi.org/10.7759/s44389-026-00290-w>

five popular agent orchestration frameworks (SmolAgents, LangGraph, AutoGen, Agno AGI, and CrewAI) for crisis response and built CrisisGen on top of the comparison, but their evaluation is architectural rather than a measurement of allocation quality against a resource-constrained benchmark. Li et al.'s [2] Disaster Copilot goes further architecturally, proposing an orchestrator coordinating specialised sub-agents, but is explicitly framed as a vision paper and a roadmap rather than a system that has been benchmarked. What is largely missing from this literature is a multi-agent design that is evaluated, with statistical rigour, on the specific and genuinely constrained problem of who gets which resource and in what order, once supply falls short of demand, which in most real disasters is true.

The present study addresses this gap directly. We take REALM-Bench's P7 task (disaster relief resource allocation under uncertainty, one of eleven real-world planning problems in Geng and Chang's [4] benchmark suite for multi-agent planning systems) as our structural reference point, and we build a simulation environment that mirrors it: multiple disaster sites with time-windowed, severity-weighted demand; multiple depots with limited inventory and a limited vehicle fleet; a road network that can fail mid-response; and demand estimates that are noisy rather than known with certainty. Against this environment, we test a four-agent negotiation framework and three standard baselines, and we report results across the six evaluation metrics that REALM-Bench [4] itself uses, so that the comparison is legible to anyone already familiar with that benchmark. We emphasize that this comparison is structural rather than empirical: our environment is a reimplementation of REALM-Bench's P7 task specification, not the benchmark itself, and no REALM-Bench score is directly compared against in this study; all comparisons reported here are among the four policies evaluated within our own environment. We use the term agentic AI here to describe the multi-agent negotiation architecture itself, not the reasoning inside each agent: the Triage, Equity, Logistics, and Disruption agents in this study are auditable rule-based components rather than LLM-backed reasoners, a scope restriction we state plainly here.

This paper makes four contributions:

- an equity-aware, disruption-reactive multi-agent negotiation architecture for disaster relief resource allocation, in which a dedicated agent whose sole job is to notice when a site is being left behind and to push priorities back toward it;
- a reproducible simulation environment, structured after REALM-Bench's P7 task, that generates scarcity, competing priorities, and mid-response disruptions synthetically but realistically enough to force genuine trade-offs;
- a statistically significant empirical comparison (90 scenario instances, 4 algorithms, and 6 standardised metrics) showing that the proposed framework outperforms the strongest of three baselines by 44.3% on the headline goal-satisfaction metric; and
- an honest account of where the framework does not win: it trades some strict severity-first ordering for fairness, together with a concrete discussion of what would need to change before a system like this could be trusted with a real deployment.

Multi-agent coordination for disaster response predates the current wave of LLM-based agents for well over a decade, with rule-based coordination systems established as early as Domnori et al.'s Ubimedic2 [5], which used rule-based agents to autonomously coordinate operative units during large-scale rescue operations, establishing that decentralised, agent-based coordination could meet the practical requirements of territorial emergency management even without any learning component. What has changed is not the coordination problem itself but the reasoning capacity available to each individual agent: an LLM-backed agent can interpret unstructured reports, natural-language instructions, and shifting priorities in ways a purely rule-based agent cannot.

The recent literature reflects this shift in two distinct directions. The first is architectural and comparative: Choure and Prajapat [3] benchmark five popular agent frameworks side by side for crisis response and then build their own extensible system, CrisisGen, around the comparison, including built-in fail-safes so that one agent's mistake does not cascade into the others. Their contribution is a working platform and a framework comparison, not a quantitative allocation-quality study. The second direction is architectural and forward-looking: Li et al. [2] proposed Disaster Copilot, a central-orchestrator design coordinating sub-agents for predictive risk analytics, situational awareness, and impact assessment, explicitly pitched as a vision and a three-phase roadmap for the field rather than a benchmarked system.

How to cite this article:

Singh K (September 23, 2026) Toward Agentic Artificial Intelligence (AI) for Disaster Relief Resource Allocation: An Equity-Aware Multi-Agent Negotiation Framework. *Cureus J Comput Sci* 3 : es44389-026-00290-w. DOI <https://doi.org/10.7759/s44389-026-00290-w>

The most structurally relevant prior work is REALM-Bench itself [4], which evaluates five multi-agent frameworks (LangGraph, AutoGen, CrewAI, OpenAI Swarm, and the authors' own ALAS) across 11 real-world planning problems, including P7: disaster relief resource allocation under uncertainty, using six standard metrics spanning planning quality, optimality, coordination, constraint satisfaction, resource usage, and adaptation to disruption. We adopt this same metric set for our own evaluation, both because it is well-motivated and because it lets our results sit alongside REALM-Bench's rather than introducing yet another incompatible scoring scheme. A separate and complementary strand of work, DisasterBench [6], evaluates a different question entirely: whether an LLM can correctly ground a disaster-response workflow into a valid sequence of tool calls across 26 disaster-response tools, which is a prerequisite for allocation to happen correctly but not itself a measure of allocation quality once tools are available.

Positioned against this backdrop, we are not aware of prior work that combines an explicit fairness-negotiation mechanism with disruption-reactive replanning in a benchmarked disaster-relief allocation setting, or that reports a significance-tested quantitative comparison against classical allocation heuristics on a REALM-Bench-structured task, rather than a qualitative framework comparison or an architectural proposal.

The equity mechanism at the centre of our framework also has roots outside the agentic-AI literature. Equity in humanitarian logistics has long been studied in operations research: Gutjahr and Fischer [7] show that minimising the deprivation cost alone can produce arbitrarily unfair relief distributions, and propose incorporating a Gini-index penalty into the objective function to explicitly trade off efficiency against fairness, a formulation that our Equity Agent's soft, continuously-renegotiated boost mechanism can be seen as a rule-based, real-time analogue of Kamyabniya et al. [8] survey 127 optimisation-based disaster-response studies published between 2013 and 2022 and find that most models remain static and fail to capture the dynamic, information-updating nature of an actual response, exactly the gap our epoch-by-epoch negotiation loop is designed to close. A parallel argument appears in the commercial supply chain literature on intelligent control towers [9], in which continuous visibility and coordinated disruption response are treated as prerequisites for resilient distribution. Our per-epoch negotiation loop pursues the same objective, but under the harder constraint of genuine, unavoidable scarcity. Closer to our own multi-agent setting, Kumar and Yeoh [10] proposed a general incentives-based framework that adjusts learned value functions to discourage over-allocation to already well-served agents in multi-agent resource allocation; our Equity Agent pursues a similar goal through an interpretable, rule-based fairness boost rather than a learned correction, consistent with the auditability requirement that we impose throughout this study.

Multi-agent approaches to disaster relief distribution are not confined to the recent LLM-agent literature. Espejo-Díaz and Guerrero [11] integrated an agent-based model with a mixed-integer program to optimise humanitarian aid distribution, letting agents representing distribution points and victims respond to behavioural factors such as impatience and mid-response information updates; their agents, like ours, are auditable and rule-driven rather than learned, though their objective is delivery cost and service level rather than an explicit fairness criterion. A separate strand applies reinforcement learning directly to the allocation decision: Lee and Lee [12] formulated decentralised patient-admission decisions across emergency departments during a mass-casualty incident as a partially-observable Markov game and train a multi-agent RL policy augmented by behavioural cloning to solve it; Yu et al. [13] casted humanitarian resource allocation as a mixed-integer nonlinear program solved via Q-learning; and van Steenbergen et al. [14] extended this to a mixed truck-and-UAV fleet operating under travel-time uncertainty from damaged infrastructure, reporting the same type of adaptive gains as observed in our Adaptation to Disruption metric. A recent survey by Pu et al. [1], reviewing five areas of disaster emergency response planning, identified the gap between such learning-driven allocation methods and a rigorous, benchmarked comparison against classical heuristics as an open direction - precisely the gap this paper addresses.

Equity-focused optimisation work outside the multi-agent literature reinforces the same conclusion. Sakiani et al. [15] build a multi-period inventory-routing and network-flow model for post-disaster relief redistribution that explicitly incorporates fairness in distribution alongside routing and inventory costs, solved via a specialised simulated-annealing algorithm - functionally similar in spirit to the fulfilment-ratio gap that our Equity Agent looks for, but computed as a static objective rather than renegotiating every epoch. Avishan et al. [16] went further by proposing an adjustable robust

How to cite this article:

Singh K (September 23, 2026) Toward Agentic Artificial Intelligence (AI) for Disaster Relief Resource Allocation: An Equity-Aware Multi-Agent Negotiation Framework. *Cureus J Comput Sci* 3 : es44389-026-00290-w. DOI <https://doi.org/10.7759/s44389-026-00290-w>

optimisation model in which the utility of each site is balanced against that of every other site under uncertainty, explicitly to prevent resource concentration at the most easily reachable sites; our soft, continuously renegotiated boost mechanism can be read as an interpretable, real-time analogue of this utility-balancing idea, without solving a robust optimisation problem at every epoch.

Materials And Methods

Problem formulation

The environment consists of a set of disaster sites $S = \{s_1, \dots, s_n\}$, each with a true, time-varying demand for three resource categories (medical supplies, food and water, and shelter materials) and a severity score in $(0,1)$ that determines both how tightly the site's deadline is set and how much weight its unmet need carries in the allocation decision.

A set of depots $D = \{d_1, \dots, d_m\}$, each holding a limited inventory of the three resource categories and operating a small fleet of vehicles, each with a fixed per-trip capacity; a vehicle is unavailable for a new trip until it completes its current round-trip travel time.

Depot inventories are set to 85-110% of aggregate true demand. This range was chosen so that every instance contains genuine scarcity - supply cannot fully meet demand even under a perfect allocator - while still leaving enough slack such that the allocation quality, rather than raw shortage, determines the outcomes; no policy can meaningfully differentiate itself below 85%, and scarcity effects vanish with all four policies converging above 110%. This range was not swept systematically in the present study; we selected it based on the reasoning above rather than an empirical sensitivity analysis and flag this as a limitation in the Discussion section.

Agents observe a noisy estimate of each site's demand rather than the true value, consistent with the well-documented unreliability of early field reporting in disaster response [1]. At an unpredictable point during the response window, the environment fires one of three disruption types, drawn at the time of generation but unknown to the agents in advance: a road closure that removes a specific depot-site edge from the network, a depot going offline entirely, or a demand surge at one site. A resource allocation policy must, at every time step, decide, for every currently free vehicle at every active depot, which site to send it to and which resource to load, subject to depot inventory and vehicle capacity constraints, without knowing in advance when or where the next disruption will occur.

Proposed multi-agent architecture

The proposed framework separates the allocation decision into four cooperating agents, each responsible for a distinct concern, that operate over a shared state at every planning epoch. Rather than negotiation in the sense of agents exchanging proposals or resolving conflicting claims directly with one another, coordination here is achieved through a shared, continuously updated urgency signal: the Equity Agent adjusts this signal, and the Logistics Agent consumes it. We retain the term negotiation to describe this repeated, state-dependent renegotiation of priority across epochs, and clarify that the scope here is to avoid overstating inter-agent interaction. Figure 1 shows the overall architecture.

How to cite this article:

Singh K (September 23, 2026) Toward Agentic Artificial Intelligence (AI) for Disaster Relief Resource Allocation: An Equity-Aware Multi-Agent Negotiation Framework. Cureus J Comput Sci 3 : es44389-026-00290-w. DOI <https://doi.org/10.7759/s44389-026-00290-w>




FIGURE 1: Proposed multi-agent negotiation architecture.

Four agents observe the shared environment state each epoch, negotiate through the Equity Agent's fairness boosts, and jointly determine the epoch's vehicle dispatches, which update the simulated world and feed back into the next epoch.

Triage Agent

For each site i at time t , the Triage Agent computes an urgency score

$$(U_i(t) = w_s \cdot \text{sev}_i + w_d \cdot \frac{1}{T_i - t + \epsilon} + w_u \cdot (1 - F_i(t)))$$

where $\text{sev}_i \in (0,1)$ is the site's severity score, T_i is its deadline; $F_i(t) \in (0,1)$ is its cumulative fulfilment ratio at time t ; ϵ is a small constant preventing division by zero as $t \rightarrow T_i$; and w_s , w_d , and $w_u \geq 0$ with $w_s + w_d + w_u = 1$ are fixed weights, set to $w_s = 0.5$, $w_d = 0.3$, and $w_u = 0.2$ in this study. The inverse-deadline term is scaled by a fixed constant rather than min-

How to cite this article:

Singh K (September 23, 2026) Toward Agentic Artificial Intelligence (AI) for Disaster Relief Resource Allocation: An Equity-Aware Multi-Agent Negotiation Framework. Cureus J Comput Sci 3 : es44389-026-00290-w. DOI <https://doi.org/10.7759/s44389-026-00290-w>

max normalised across active sites in the present implementation; we flag this as a simplification relative to a fully normalised formulation for future revisions. This score is deliberately simple and interpretable rather than learned, so that its behaviour is auditable by a human coordinator.

Equity Agent

At every epoch, the Equity Agent compares each site's fulfilment ratio to the cross-site mean and standard deviation.

$$\mu(t) = (1/n)\sum_i F_i(t)\sigma(t) = \sqrt{[(1/n)\sum_i (F_i(t) - \mu(t))^2]}$$

For any site i , with $F_i(t) < \mu(t) - 0.5\sigma(t)$, a fixed multiplicative boost of 1.4× is applied to that site's urgency score for the next round, recomputed every epoch from the current fulfilment ratios. This is a soft constraint rather than a hard allocation floor: it reweights priority rather than guaranteeing any site a minimum share, so the framework can still prioritise a genuinely critical site over a merely-behind one when severity demands it.

Logistics Agent

For every vehicle v that becomes free during the epoch, the Logistics Agent evaluates every reachable site-resource pair (j, r) and selects the one maximising:

$$\text{score}(v, j, r) = [U'_j(t) \cdot \min(\text{avail}_{r,d}, \text{cap}_v)] / \text{traveltime}(d, j)$$

where d is v 's home depot, $\text{avail}_{r,d}$ is the remaining inventory of resource r at d , and cap_v is v 's per-trip capacity. This is a locally greedy choice made independently per vehicle against continuously updated urgency scores, rather than a single fixed priority ranking computed once.

Disruption Agent

The Disruption Agent does not compute a score; it constrains the candidate sets that the other two agents act on. When a road closure or depot loss fires, the affected depot-site edges or depots are removed from the Logistics Agent's reachable set for the very next vehicle assignment. As this filtering happens in line with the next epoch rather than through a separate replanning phase, recovery is a direct consequence of the per-epoch loop rather than a special case.

Experimental setup

We generated 90 scenario instances spanning three problem scales (8, 16, and 24 disaster sites, each with a proportionally scaled number of depots) with 30 random seeds per scale, and a 50% probability of a disruption firing during the response window in each instance. Each instance was evaluated under five allocation policies:

Random - assigns each free vehicle to a uniformly random reachable site and resource;

Round-Robin - cycles through sites in a fixed order;

Greedy-Priority - always serves the highest severity-to-deadline-pressure site reachable by each free vehicle, matching the single-pass priority heuristic implicit in most rule-based dispatch systems;

Multi-Agent (No Equity) - the full proposed architecture with the Equity Agent's boost disabled ($b_i(t) \equiv 1$ for all sites), isolating the contribution of the Logistics Agent's urgency-weighted-value-per-travel-time dispatch rule on its own;

Multi-Agent (Full) - the complete proposed framework described in the previous section (Proposed multi-agent architecture).

This gives 450 individual simulation runs in total (90 instances × 5 policies). To make these instances concrete, Table 1 shows three representative examples drawn directly from the evaluation set. At $n = 8$ (seed 0), two depots hold 583 units against 599 true units of demand (ratio: 0.97), site severities range from 0.16 to 0.61, and a demand surge fires at $t = 10.6$ at one site (a 1.48× multiplier). At $n = 16$ (seed 5), three depots hold 1,328 units against 1,296 demanded (ratio: 1.03), severities range from 0.13 to 0.87, and no disruption fires. At $n = 24$ (seed 12), four depots hold 1,779 units against 1,802 demanded (ratio: 0.99), severities range from 0.05 to 0.76, and a demand surge fires at $t = 7.6$ (a 1.81× multiplier). These three instances illustrate the range of supply/demand ratios (0.97-1.03), severity spreads, and disruption conditions

How to cite this article:

Singh K (September 23, 2026) Toward Agentic Artificial Intelligence (AI) for Disaster Relief Resource Allocation: An Equity-Aware Multi-Agent Negotiation Framework. Cureus J Comput Sci 3 : es44389-026-00290-w. DOI <https://doi.org/10.7759/s44389-026-00290-w>

sampled across the full set of 90. As every policy is evaluated on the same 90 instances under the same seeds, results across policies are paired; statistical comparisons therefore use a paired test (paired *t*-test or Wilcoxon signed-rank, as appropriate to the metric's distribution) rather than an unpaired test.

Sites	Depots	Total Demand (units)	Total Inventory (units)	Supply Ratio	Severity Range	Disruption
8	2	599	583	0.97	0.16-0.61	Demand surge, $t = 10.6, \times 1.48$
16	3	1296	1328	1.03	0.13-0.87	None
24	4	1802	1779	0.99	0.05-0.76	Demand surge, $t = 7.6, \times 1.81$

TABLE 1: Three representative scenario instances from the 90-instance evaluation set, illustrating the range of supply ratios, severity spreads, and disruption conditions sampled.

Each run was scored on six evaluation metrics, five adapted from REALM-Bench and one - Coordination Effectiveness - used as originally defined there:

Goal Satisfaction: the fraction of total true demand fulfilled by the end of the response window.

Resource Usage: the fraction of total available supply actually shipped rather than left stranded.

Adaptation to Disruption: the goal-satisfaction rate restricted to sites or depots directly affected by a road-closure or demand-surge disruption event (depot-damage events, which affect supply broadly rather than a specific site, are not attributed to individual sites for this metric), computed only over the subset of instances in which one of these disruptions actually fired.

Coordination Effectiveness: the Spearman rank correlation between a site's severity and the order in which its demand was fulfilled, rescaled to (0,1).

Makespan Efficiency: computed only over sites that reached at least 90% fulfilment under every policy being compared in a given instance (the common-completion subset), rather than each policy's own subset of completed sites, so that no policy is scored on an easier self-selected sample.

Constraint Satisfaction: the fraction of dispatch decisions that respected depot inventory and vehicle capacity limits. Because the simulator rejects any candidate dispatch that would violate these limits before it is executed, this metric is expected to sit at or near 1.000 for every policy by construction; we report it as a validity check, confirming that no policy is violating physical constraints to inflate other metrics, not as a discriminating result.

Results And Discussion

We evaluated the proposed multi-agent framework against three baselines (Random, Round-Robin, and Greedy-Priority) across 90 scenario instances spanning three problem scales (8, 16, and 24 disaster sites, 30 random seeds per scale), scoring every run on the six REALM-Bench metrics described in the Experimental setup section.

How to cite this article:

Singh K (September 23, 2026) Toward Agentic Artificial Intelligence (AI) for Disaster Relief Resource Allocation: An Equity-Aware Multi-Agent Negotiation Framework. Cureus J Comput Sci 3 : es44389-026-00290-w. DOI <https://doi.org/10.7759/s44389-026-00290-w>

Table 2 reports the mean values of each metric across all 90 instances for each of the four policies. The proposed multi-agent framework leads outright on four of the six metrics and matches the baselines on a fifth, Constraint Satisfaction, where all four policies reach a perfect 1.000; it trails only on Coordination Effectiveness. The clearest margins are on Goal Satisfaction (0.788 versus 0.546 for the strongest baseline, Greedy-Priority) and Resource Usage (0.837 versus 0.591).

Algorithm	Goal Satisfaction	Makespan Efficiency	Constraint Satisfaction	Resource Usage	Adaptation to Disruption	Coordination Effectiveness
Random	0.505	0.278	1.000	0.546	0.432	0.479
Round-Robin	0.512	0.095	1.000	0.550	0.443	0.556
Greedy-Priority	0.546	0.553	1.000	0.591	0.307	0.870
Proposed Multi-Agent	0.788	0.563	1.000	0.837	0.625	0.664

TABLE 2: Mean performance across all 90 scenario instances (30 seeds × 3 problem scales), by algorithm.

A Welch's two-sample *t*-test comparing the proposed multi-agent framework against Greedy-Priority, the strongest of the three baselines, on Goal Satisfaction returned $t = 11.58$, $p = 5.78 \times 10^{-23}$, indicating the difference is unlikely to arise under the null hypothesis of equal means. In relative terms this is a 44.3% improvement over the best baseline. Table 3 shows this improvement is not an artefact of a single problem scale: the proposed framework outperforms every baseline at 8, 16, and 24 sites, and its lead actually widens as the number of sites increases, from roughly +15.3 points at 8 sites to about +30.9 points at 24 sites.

Sites	Random	Round-Robin	Greedy-Priority	Proposed Multi-Agent
8	0.597	0.607	0.645	0.798
16	0.497	0.498	0.539	0.804
24	0.422	0.431	0.455	0.764

TABLE 3: Goal Satisfaction by problem scale, showing the proposed framework's advantage holds and widens as scarcity and coordination difficulty increase.

To isolate the Equity Agent's specific contribution from the Logistics Agent's dispatch objective, we ran an additional ablation, Multi-Agent (No Equity), in which the equity boost is fixed at 1.0 for every site while the rest of the architecture is unchanged. On Goal Satisfaction, this ablation scores 0.788, statistically indistinguishable from the full framework's 0.788 (paired $t = 0.33$, $p = 0.75$), and both remain far above Greedy-Priority's 0.546. This shows that the 44.3% Goal Satisfaction gain over the strongest baseline is attributable almost entirely to the Logistics Agent's urgency-weighted-

How to cite this article:

value-per-travel-time dispatch rule, not to the Equity Agent. The Equity Agent's measurable contribution appears elsewhere: it lowers Coordination Effectiveness from 0.683 (no equity) to 0.664 (full), the same severity-ordering trade-off already discussed above, and raises Adaptation to Disruption from 0.596 to 0.625. We report this decomposition so that the framework's fairness mechanism is credited for what it actually does (a coordination/disruption-recovery trade-off) rather than for the headline Goal Satisfaction number, which is a Logistics Agent result. The same numbers are easier to scan in chart form. Figure 2 puts all four algorithms side by side across all six metrics, and the proposed framework's bar comes out tallest on five of them, with the widest margins on Goal Satisfaction and Resource Usage - the two metrics most directly tied to whether people actually received the supplies they needed, rather than supplies sitting unused in a depot. Coordination Effectiveness is the one metric where a baseline's bar is taller, with Greedy-Priority edging ahead.




FIGURE 2: All four algorithms compared across all six evaluation metrics.

Scale is where the baselines really start to struggle. Figure 3 plots Goal Satisfaction against problem size directly; confidence intervals were computed as the sample mean $\pm 1.96 \times$ standard error across the 30 seeds at each scale. The proposed framework's line declines only modestly across 8, 16, and 24 sites (0.798-0.764), while every baseline drops far more steeply over the same range. The baselines get worse exactly when coordination is hardest and scarcity is most severe; the proposed framework's negotiation mechanism keeps performance roughly constant regardless of the number of sites competing for the same limited supply.

How to cite this article:




FIGURE 3: Goal Satisfaction as a function of problem scale, with 95% confidence intervals across the 30 seeds at each scale.

As the averages can hide a lot, Figure 4 shows the same headline metric from a different angle: the full distribution across all 90 runs, not just the mean. Here, the box of proposed framework sits both higher and narrower than those of the baselines, which is a different claim than "wins on average" - it means that the framework is not riding a handful of lucky runs, but is consistently ahead with less run-to-run variance than the simplest Random baseline. For a disaster-relief planner, this predictability is arguably worth as much as the average improvement itself; a policy that is excellent most of the time but occasionally fails badly is a much harder one to trust operationally.

How to cite this article:

Singh K (September 23, 2026) Toward Agentic Artificial Intelligence (AI) for Disaster Relief Resource Allocation: An Equity-Aware Multi-Agent Negotiation Framework. Cureus J Comput Sci 3 : es44389-026-00290-w. DOI <https://doi.org/10.7759/s44389-026-00290-w>




FIGURE 4: Distribution of Goal Satisfaction across all 90 runs per algorithm, showing the proposed framework is both higher on average and more consistent than every baseline.

A radar plot makes the overall shape of that trade-off easier to see at once. Figure 5 overlays the proposed framework directly against Greedy-Priority, the strongest baseline, across all six metrics, and the proposed framework's shape encloses the baseline's on five of the six axes - visibly so on Goal Satisfaction, Resource Usage, and Adaptation to Disruption. The one axis where the baseline reaches further out is Coordination Effectiveness, which shows up here as a single dent in an otherwise dominant shape. This dent is a direct result of the Equity Agent's fairness weighting, not a flaw in the framework's dispatch logic.

How to cite this article:

Singh K (September 23, 2026) Toward Agentic Artificial Intelligence (AI) for Disaster Relief Resource Allocation: An Equity-Aware Multi-Agent Negotiation Framework. Cureus J Comput Sci 3 : es44389-026-00290-w. DOI <https://doi.org/10.7759/s44389-026-00290-w>




FIGURE 5: Proposed Multi-Agent framework against the strongest baseline (Greedy-Priority) across all six metrics simultaneously.

To put a number on that dent, Figure 6 plots the percentage change of the proposed framework relative to the best baseline, metric by metric. Three of the five comparable metrics show double-digit gains, led by Goal Satisfaction at +44.3%, with Resource Usage and Adaptation to Disruption both above +41%. Makespan Efficiency is essentially a tie with the best baseline at +1.8%. Coordination Effectiveness is the only metric with a negative bar, at -23.7%, and we have deliberately kept it on the same scale as the wins so the trade-off stays visible rather than becoming a footnote.

How to cite this article:

Singh K (September 23, 2026) Toward Agentic Artificial Intelligence (AI) for Disaster Relief Resource Allocation: An Equity-Aware Multi-Agent Negotiation Framework. Cureus J Comput Sci 3 : es44389-026-00290-w. DOI <https://doi.org/10.7759/s44389-026-00290-w>




FIGURE 6: Percentage improvement of the proposed framework over the best-performing baseline, per metric.

Coordination Effectiveness is the one metric where the baseline wins.

The clearest finding is that continuous, urgency-weighted dispatch beats a single fixed priority ranking, and that the gap grows as the problem gets harder; as the ablation above shows, this gain is driven by the Logistics Agent's per-vehicle value-per-travel-time objective rather than by the Equity Agent specifically. Greedy-Priority's fixed severity ranking works reasonably well when there is enough supply to eventually reach every site anyway; once scarcity becomes severe enough, such that some sites structurally cannot be served in a single pass, a policy that never revisits its own priorities keeps returning to the same handful of top-ranked sites while others starve, a failure mode that the Logistics Agent's continuously updated urgency scores correct for even without the Equity Agent's boost. The Equity Agent's own contribution - protecting structurally under-served sites and improving disruption recovery, quantified in Table 4 - is real but secondary to this headline result.

How to cite this article:

Singh K (September 23, 2026) Toward Agentic Artificial Intelligence (AI) for Disaster Relief Resource Allocation: An Equity-Aware Multi-Agent Negotiation Framework. Cureus J Comput Sci 3 : es44389-026-00290-w. DOI <https://doi.org/10.7759/s44389-026-00290-w>

Algorithm	Goal Satisfaction	Coordination Effectiveness	Adaptation to Disruption
Greedy-Priority	0.546	0.870	0.307
Multi-Agent (No Equity)	0.788	0.683	0.596
Multi-Agent (Full)	0.788	0.664	0.625

TABLE 4: Ablation isolating the Equity Agent's contribution.

Multi-Agent (No Equity) is identical to the full framework except the equity boost is fixed at 1.0 for every site. Goal Satisfaction is statistically indistinguishable between No-Equity and Full (paired $t = 0.33$, $p = 0.75$), showing the Goal Satisfaction gain over Greedy-Priority is attributable to the Logistics Agent's dispatch objective, not the Equity Agent. The Equity Agent's measurable effect appears in Coordination Effectiveness and Adaptation to Disruption instead.

The one place the proposed framework loses is Coordination Effectiveness, where Greedy-Priority scores 0.870 against the proposed framework's 0.664. This is not a problem but rather a direct and expected consequence of the design: Coordination Effectiveness measures how strictly the most severe sites are served first, and the Equity Agent's entire purpose is to occasionally interrupt that strict ordering on behalf of a site that is falling behind even though it is not the single most severe one. In a scenario where the ethical priority genuinely is "serve the worst-hit site first, no matter what" (for instance, a mass-casualty triage situation), Greedy-Priority's behaviour may actually be preferable, and a real deployment would need a way to let a human coordinator dial the Equity Agent's influence up or down depending on which of these two principles should dominate for the disaster at hand. We see this as a genuine trade-off worth making visible rather than a weakness to explain away, and Figure 6 is included specifically so that this loss is not buried inside an otherwise favourable headline number.

Four limitations bound how far these results should be trusted. First, the Triage, Equity, and Logistics agents in this study use hand-specified scoring rules rather than LLM-based reasoning; this was a deliberate choice to keep the simulation reproducible without live API access, but it means the framework has not yet been tested with the kind of unstructured-report understanding that is the actual selling point of agentic LLM systems, and which Disaster Copilot [2] and CrisisGen [3] both argue is essential for real-world disaster response. Second, the environment is synthetic: site locations, demand levels, and disruption timing are generated from parameterised distributions rather than drawn from real disaster records such as the UN OCHA or FEMA datasets that REALM-Bench [4] itself points to as public data sources for the P7 task. Third, the six metrics, while borrowed directly from REALM-Bench for comparability, are still simulation-internal; they say nothing about factors a real deployment would have to handle, such as communication failures between field teams and the coordination system, or liability questions around an autonomous system making life-affecting allocation decisions without a human in the loop. Fourth, the 85-110% depot-inventory range was chosen by reasoning about scarcity rather than validated by a sensitivity sweep; it remains possible that the framework's advantage narrows outside this band, and confirming this is a natural next check before the results are generalised.

The natural next step, and the one we intend to pursue, is to replace the Triage and Logistics agents' scoring functions with LLM calls grounded via retrieval-augmented generation over real situation reports, in the spirit of Disaster Copilot's [2] disaster-knowledge sub-agent, and to re-run this same evaluation pipeline against real event data from UN OCHA or FEMA's OpenFEMA archive rather than synthetic scenarios, so that the gains reported here can be checked against the messier, less well-behaved conditions of an actual disaster.

How to cite this article:

Conclusions

We presented an equity-aware, disruption-reactive multi-agent negotiation framework for disaster relief resource allocation, evaluated it against three standard allocation heuristics on a REALM-Bench-structured environment across 90 scenario instances, and found a statistically significant 44.3% improvement in goal satisfaction over the strongest baseline, alongside comparable gains in resource utilisation and disruption recovery. We also found, and report, the one dimension on which a simpler greedy heuristic still wins (strict severity-first ordering) and argue that this trade-off, not just the headline win, is what a deployable system would need to expose to a human coordinator. The framework's individual agent logic is intentionally simple in this study so that the evaluation itself could be fully reproducible; the next test to verify the reliability of this architecture is putting real LLM reasoning and real disaster data behind each of its four agents.

Additional Information

Author Contributions

All authors have reviewed the final version to be published and agreed to be accountable for all aspects of the work.

Concept and design: Khushvir Singh

Acquisition, analysis, or interpretation of data: Khushvir Singh

Drafting of the manuscript: Khushvir Singh

Critical review of the manuscript for important intellectual content: Khushvir Singh

Disclosures

Human subjects: All authors have confirmed that this study did not involve human participants or tissue.

Animal subjects: All authors have confirmed that this study did not involve animal subjects or tissue.

Conflicts of interest: In compliance with [the ICMJE uniform disclosure form](#), all authors declare the following:

- **Payment/services info:** All authors have declared that no financial support was received from any organization for the submitted work.
- **Financial relationships:** All authors have declared that they have no financial relationships at present or within the previous three years with any organizations that might have an interest in the submitted work.
- **Other relationships:** All authors have declared that there are no other relationships or activities that could appear to have influenced the submitted work.

Data Availability Statements

The data (and/or code) supporting this study are openly available in GitHub at <https://github.com/khushvir11/disaster-relief-experiment-code>, reference khushvir11/disaster-relief-experiment-code.

References

1. Pu F, Li Z, Wu Y, Ma C, Zhao R: [Recent advances in disaster emergency response planning: Integrating optimization, machine learning, and simulation](#). Safety Emergency Science. 2025, 1:9590007. [10.26599/ses.2025.9590007](#)
2. Li B, Ma J, Yin K, Xiao Y, Hsu C-W, Mostafavi A: [Disaster management in the era of agentic AI systems: a vision for collective human-machine intelligence for augmented resilience \[PREPRINT\]](#). arXiv. 2025, [10.48550/arXiv.2510.16034](#)

How to cite this article:

Singh K (September 23, 2026) Toward Agentic Artificial Intelligence (AI) for Disaster Relief Resource Allocation: An Equity-Aware Multi-Agent Negotiation Framework. Cureus J Comput Sci 3 : es44389-026-00290-w. DOI <https://doi.org/10.7759/s44389-026-00290-w>

3. Choure P, Prajapat S: [Agentic AI for emergency response and comparative analysis of SmolAgents, LangGraph, AutoGen, Agno AGI and CrewAI for crisis solution](#). Contributions Presented at the International Conference on Computing, Communication, Cybersecurity & AI, July 10-11, 2025, Birmingham, UK. Naik N, Jenkins P, Prajapat S, Grace P (ed): Springer, Cham; 2026. 229-259. [10.1007/978-3-032-16791-0_10](#)
4. Geng L, Chang EY: [REALM-Bench: a benchmark for evaluating multi-agent systems on real-world, dynamic planning and scheduling tasks \[PREPRINT\]](#). arXiv. 2025, [10.48550/arXiv.2502.18836](#)
5. Domnori E, Cabri G, Leonardi L: [Multi-agent approach for disaster management](#). 2011 International Conference on P2P, Parallel, Grid, Cloud and Internet Computing, Barcelona, Spain. 2011, 311-316. [10.1109/3PGCIC.2011.57](#)
6. Chen Z, Yin K, Zhang W, et al.: [DisasterBench: benchmarking LLM planning under typed tool interface constraints \[PREPRINT\]](#). arXiv. 2026, [10.48550/arXiv.2605.27957](#)
7. Gutjahr WJ, Fischer S: [Equity and deprivation costs in humanitarian logistics](#). European Journal of Operational Research. 2018, 270:185-197. [10.1016/j.ejor.2018.03.019](#)
8. Kamyabniya A, Sauré A, Salman FS, Bénichou N, Patrick J: [Optimization models for disaster response operations: a literature review](#). OR Spectrum. 2024, 46:737-783. [10.1007/s00291-024-00750-6](#)
9. Gandu S, Manda VK: [Intelligent control towers for resilient and compliant pharmaceutical supply chains](#). Digitally-Enabled Control Towers for Pharmaceutical Supply Chains. Majeed M, Abdul AM, Saani AS, Zignaa MA, Ganiw HA (ed): IGI Global Scientific Publishing, 2026. 261-290. [10.4018/979-8-3373-8770-3.ch009](#)
10. Kumar A, Yeoh W: [A general incentives-based framework for fairness in multi-agent resource allocation \[PREPRINT\]](#). arXiv. 2025, [10.48550/arXiv.2510.26740](#)
11. Espejo-Díaz JA, Guerrero WJ: [A multiagent approach to solving the dynamic postdisaster relief distribution problem](#). Operations Management Research. 2021, 14:177-193. [10.1007/s12063-021-00192-1](#)
12. Lee H-R, Lee T: [Multi-agent reinforcement learning algorithm to solve a partially-observable multi-agent problem in disaster response](#). European Journal of Operational Research. 2021, 291:296-308. [10.1016/j.ejor.2020.09.018](#)
13. Yu L, Zhang C, Jiang J, Yang H, Shang H: [Reinforcement learning approach for resource allocation in humanitarian logistics](#). Expert Systems with Applications. 2021, 173:114663. [10.1016/j.eswa.2021.114663](#)
14. van Steenberg R, Mes M, van Heeswijk W: [Reinforcement learning for humanitarian relief distribution with trucks and UAVs under travel time uncertainty](#). Transportation Research Part C: Emerging Technologies. 2023, 157:104401. [10.1016/j.trc.2023.104401](#)
15. Sakiani R, Seifi A, Khorshiddoust RR: [Inventory routing and dynamic redistribution of relief goods in post-disaster operations](#). Computers & Industrial Engineering. 2020, 140:106219. [10.1016/j.cie.2019.106219](#)
16. Avishan F, Elyasi M, Yanıkoğlu İ, Ekici A, Özener OÖ: [Humanitarian relief distribution problem: an adjustable robust optimization approach](#). Transportation Science. 2023, 57:1096-1114. [10.1287/trsc.2023.1204](#)

How to cite this article:

Singh K (September 23, 2026) Toward Agentic Artificial Intelligence (AI) for Disaster Relief Resource Allocation: An Equity-Aware Multi-Agent Negotiation Framework. Cureus J Comput Sci 3 : es44389-026-00290-w. DOI <https://doi.org/10.7759/s44389-026-00290-w>