Int. J. Appl. Math. Comput. Sci., 2026, Vol. 36, No. 3, 527–542
DOI: 10.61822/amcs-2026-0034

BRIDGING STOCHASTIC, ROBUST, AND NEUTROSOPHIC PARADIGMS FOR HURRICANE EVACUATION UNDER COMPOUND UNCERTAINTY: A SIX–LAYER METHODOLOGICAL SYNTHESIS

YASMANY FERNÁNDEZ-FERNÁNDEZ [a,b,*], SIRA M. ALLENDE ALONSO [b], RIDELIO MIRANDA PÉREZ [c], GEMAYQZEL BOUZA ALLENDE [b]

[a]Computer Science Program
Carchi Polytechnic State University
Calle Antisana y Av. Universitaria, Tulcán 040102, Ecuador
e-mail: yfernandezf@upec.edu.ec

[b]Department of Mathematics and Computer Science
University of Havana
San Lázaro y L, Vedado, Havana 10400, Cuba
e-mail: {sira,gema}@matcom.uh.cu

[c]Faculty of Economic Sciences
University of Cienfuegos
Carretera a Rodas km 3.5, Cienfuegos 55100, Cuba
e-mail: rmiranda@ucf.edu.cu

Hurricane evacuation planning under uncertainty requires integrating decisions across multiple spatial and temporal scales, yet the literature addresses demand stochasticity, temporal robustness, and computational tractability largely in isolation. This paper develops a unified optimization framework organized around six progressive layers: meteorological scenario generation, deterministic and stochastic network flow planning, vehicle scheduling with origin-based decomposition, classical robust temporal protection, neutrosophic adaptive robust optimization with scenario-anchored uncertainty sets, and computational complexity characterization. An integrated critical review of 337 Scopus-indexed publications (2022–2026) identifies persistent methodological gaps that motivate each architectural layer. The framework couples two-stage stochastic programming with budget-constrained robust optimization, affine decision rules, and neutrosophic indeterminacy calibration within a single operational pipeline. Validation on a georeferenced Caribbean network demonstrates a 94% reduction in binary variables through origin-based decomposition and end-to-end execution in under four minutes, with a sensitivity analysis showing that neutrosophic budget calibration reduces the price of robustness from 24.8% to 2.8%.

Keywords: hurricane evacuation, robust optimization, stochastic programming, neutrosophic uncertainty sets, hierarchical decomposition.

1. Introduction

When a hurricane bears down on a populated coastline, emergency managers confront a deeply uncomfortable reality: they must decide who goes where, by which route, and in which vehicle—all while the storm itself remains unpredictable. The eventual trajectory and intensity of the hurricane determine how many people will need to evacuate and from which zones, producing what we call *demand uncertainty*. But even after those demands materialize, the time it takes to complete each trip depends on road congestion, weather-induced degradation, and the unpredictable behavior of evacuees themselves—a separate layer of *temporal uncertainty* that compounds the planning challenge. Despite three decades of research on evacuation optimization, most of the literature treats these two uncertainty types

*Corresponding author

in isolation: stochastic programming handles demand (Bayram and Yaman, 2018; Siddig and Song, 2025), robust optimization handles temporal deviations (Goerigk *et al.*, 2015; Yao *et al.*, 2009), and bridging the gap between them is left to the planner's intuition.

This fragmentation runs deeper than uncertainty modeling alone. The network flow formulations that determine evacuation routes—deciding which families travel along which paths—are typically disconnected from the vehicle scheduling models that actually execute those routes with physical buses and trucks. Complexity analyses establish hardness results for individual components, yet the integrated hierarchical problem remains largely uncharted. Recent surveys reinforce these observations from different angles: Kamyabniya *et al.* (2024) cataloged optimization models for disaster response, finding that uncertainty set construction remains a critical and often ad-hoc modeling decision; Steinmetz-Weiss and Sanderson (2024) reviewed transit-based evacuation and flagged the persistent absence of hierarchical integration between the flow and scheduling stages; and Esposito Amideo *et al.* (2019) taxonomized a shelter location and routing problems, noting the disconnect between strategic and tactical decision layers.

Before developing the technical material, we state the role of this manuscript explicitly to avoid the common ambiguity in integrative works. This paper is a *methodological contribution with computational validation*: its primary deliverable is the metamodel $\mathcal{E}^{\mathrm{Robust}}$, defined by four nested problems with formal coupling mechanisms and benchmarked against classical stochastic, robust, and hybrid baselines. The bibliometric analysis of Section 2 is a *motivational review*, not a survey: its purpose is to quantify, through seven Scopus queries, the specific gap that motivates each architectural layer, not to catalogue the evacuation literature exhaustively. The computational study of Section 9 is a *validation instrument*, not an ablation catalogue: it establishes that the integration yields measurable advantages on a georeferenced Caribbean instance under direct comparison with three reference paradigms. Readers seeking a pure literature review should consult Kamyabniya *et al.* (2024) or Steinmetz-Weiss and Sanderson (2024), while those looking for component-level ablations of stochastic or robust optimization should consult Bayram and Yaman (2018) or Goerigk *et al.* (2015).

The contribution of this paper is a unified architecture, its formal coupling, and its computational dominance over the closest hybrid competitor (Pouraliakbari-Mamaghani *et al.*, 2023), as quantified in Table 4. Rather than offering yet another comprehensive catalog, this paper develops a unified optimization framework whose analysis is organized around six progressive methodological layers of a coupled metamodel for hurricane evacuation under compound uncertainty:

1. *meteorological uncertainty and scenario generation*, connecting hurricane track data to structured scenario ensembles that drive all subsequent stages;

2. *flow planning: from the deterministic baseline to stochastic recourse*, developing the deterministic multicommodity flow formulation PI^D and its stochastic extension PI^S that hedges against demand variability;

3. *hierarchical routing and scheduling*, presenting the deterministic vehicle scheduling model PII^D and origin-based decomposition that makes the problem computationally tractable;

4. *classical robust temporal optimization*, immunizing schedules against worst-case travel time deviations using box uncertainty sets;

5. *neutrosophic adaptive robust optimization*, replacing ad-hoc box sets with scenario-anchored budget sets calibrated through truth-indeterminacy-falsity characterization;

6. *computational complexity and scalability*, establishing NP-hardness of the integrated metamodel and identifying tractable substructures.

At each layer, we identify what the existing literature does well, where it falls short, and how the proposed metamodel addresses those shortcomings. The overarching message is that demand stochasticity, temporal robustness, epistemic indeterminacy, and computational complexity have matured as individual research threads; what is missing—and what we provide—is a unified pipeline that weaves them together from raw meteorological data to executable evacuation schedules.

The metamodel, which we denote by $\mathcal{E}^{\mathrm{Robust}} \equiv \mathrm{PII}^R(\mathrm{PI}^S(\mathrm{PII}^D(\mathrm{PI}^D)))$, nests four optimization problems across two complementary pipelines. The *deterministic baseline* $\mathrm{PI}^D \rightarrow \mathrm{PII}^D$ solves a multicommodity flow problem for route generation and then schedules a heterogeneous vehicle fleet to execute those routes under nominal conditions. The *stochastic-robust extension* $\mathrm{PI}^S \rightarrow \mathrm{PII}^R$ replaces point-estimate demand with scenario-based programming and then immunizes the resulting schedules against temporal perturbations. A neutrosophic layer enriches the robust stage by replacing ad-hoc uncertainty set construction with a principled calibration grounded in the (T, I, F) characterization of each uncertain parameter.

Three types of integration distinguish this architecture: *vertical integration* ensures seamless data flow from meteorological inputs through scenario generation, flow optimization, route extraction, and robust scheduling; *horizontal integration* addresses demand stochasticity (through recourse), temporal uncertainty (through robustness), and epistemic indeterminacy (through neutrosophy) simultaneously; and *theoretical-computational integration* accompanies each stage with formal guarantees—NP-hardness proofs, feasibility certificates, and polynomial route extraction—validated on a georeferenced Caribbean network.

We explicitly distinguish *component-level novelty* from *integration-level novelty*, since these are often conflated in integrative papers and the distinction is central to the contribution of this work. At the component level, we do *not* claim novelty: two-stage stochastic programming (Bayram and Yaman, 2018), budget-constrained robust optimization (Bertsimas and Sim, 2004), affine decision rules (Bertsimas *et al.*, 2019), generalized assignment decomposition (Fisher *et al.*, 1986), and single-valued neutrosophic characterization (Smarandache, 2005) are all mature, well-established techniques. The claim of this paper is *integration level*, operationalized through three criteria that a mere incremental combination would fail:

1. *structural coupling*—information propagates across layers through formal mathematical objects (the distance factor ζ^{ω} in (1), the scenario-anchored deviation $\sigma(\lambda)$ in (10), the neutrosophic budget Γ^{n} in (14)), not through ad-hoc parameter sharing;

2. *emergent numerical behaviour*—the composite pipeline produces price of robustness and probabilistic guarantee combinations unattainable by any single paradigm (Table 4), a form of non-additivity that incremental combinations cannot exhibit;

3. *complexity closure*—the integrated system $\mathrm{PII}^{R}(\mathrm{PI}^{S}(\mathrm{PII}^{D}(\mathrm{PI}^{D})))$ is shown as NP-hard through structural coupling between layers, not by inheritance from any single component.

These three criteria jointly substantiate that the integration is not a concatenation of known techniques but a coupled artefact whose behaviour is not predictable from its parts. Section 10 extracts three concrete insights that follow from the integration and cannot be derived from any component considered in isolation.

The remainder of this paper is organized as follows. Section 2 maps the bibliometric landscape. Sections 3–8 develop the six methodological layers with critical analysis. Section 9 synthesizes the unified architecture and reports a direct computational benchmark against stochastic, robust, and hybrid baselines (Table 4). Section 10 extracts the three integration-specific insights that the benchmark reveals. Section 11 presents the gap analysis, sensitivity results, research agenda, and conclusions.

2. Bibliometric landscape

Before diving into the metamodel's layers, it is worth understanding what the research community has—and has not—produced in recent years. We conducted a systematic bibliometric analysis of Scopus-indexed publications from 2022 to 2026 using seven queries that progressively narrow down from general hurricane evacuation optimization to the specific intersection of neutrosophic and robust methods (Table 1).

2.1. Temporal distribution of research. Figure 1 reveals how research attention has been distributed across these query domains over the past four years. Note that Q_1 is an umbrella query covering *all* optimization paradigms (deterministic, stochastic, and robust), not a proxy for deterministic methods alone, since deterministic approaches rarely appear as an explicit Scopus keyword; Q_0 serves as a robust-network reference baseline (abbreviated *Q0 Hurr+Rob+Evac+Net* in the figure, corresponding to row Q_0 of Table 1). The broadest query (Q_1) shows sustained growth, climbing from approximately 22 publications in 2022 to over 40 in 2025—a near-doubling that reflects the community's growing recognition that optimization tools can meaningfully improve evacuation outcomes. Stochastic approaches (Q_2) track this growth closely, confirming that the field has largely moved beyond purely deterministic models.

Two patterns are especially telling for this review. First, robust optimization studies ($Q_3 + Q_4$) have grown in absolute terms but their ratio to stochastic studies has remained stubbornly flat at roughly 0.30–0.35. This suggests that, while robust optimization is a mature theoretical field, its adoption in evacuation planning lags behind its penetration in domains like power systems and supply chain management. Second—and this is the observation that most directly motivates our work—query Q_6, which targets the intersection of neutrosophic methods, robust optimization, and evacuation with uncertainty sets, returns *zero* results. The territory we propose to explore is, by any bibliometric measure, genuinely uncharted.

2.2. Methodological coverage analysis. Publication counts tell only part of the story. To understand *what* the existing literature covers—and, crucially, what it misses—we constructed a methodology coverage matrix

Table 1. Scopus query definitions and approximate result counts (2022–2026). Queries are designed to progressively narrow the scope: Q_1 captures the broadest evacuation optimization literature, while Q_6 targets the specific intersection that motivates this work. In Fig. 1 these queries appear under compact abbreviated labels (e.g., *Q0 Hurr+Rob+Evac+Net* for Q_0); the complete abbreviation-to-query mapping is provided in the caption of Fig. 1.

ID	Query terms	Count
Q_0	"hurricane" AND "robust" AND "evacuation" AND "network"	53
Q_1	"optimization" AND "hurricane evacuation"	125
Q_2	"stochastic optimization" AND "evacuation"	98
Q_3	"robust optimization" AND "evacuation" (static)	32
Q_4	"robust optimization" AND "evacuation" (dynamic)	21
Q_5	"neutrosophic optimization" AND "evacuation"	8
Q_6	"neutrosophic" AND "robust optimization" AND "evacuation" AND "sets"	0

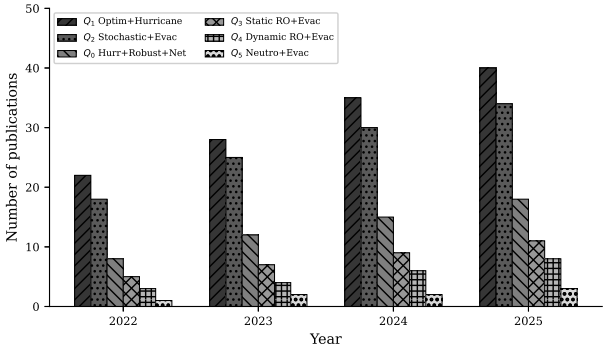


Fig. 1. Publication trends by year and query category (2022–2025). Stochastic methods (Q_2) track general evacuation optimization (Q_1) closely, while robust approaches (Q_3–Q_4) remain a fraction of the stochastic volume. The zero-height bar for Q_6 across all years confirms the bibliometric gap. The compact abbreviations used in the figure are reconciled with the Boolean query strings of Table 1: *Q0 Hurr+Rob+Evac+Net* $\equiv Q_0$ ("hurricane" AND "robust" AND "evacuation" AND "network"); *Q1 Opt+HurrEvac* $\equiv Q_1$ ("optimization" AND "hurricane evacuation"), *Q2 StochOpt+Evac* $\equiv Q_2$ ("stochastic optimization" AND "evacuation"), *Q3 RobOpt+Evac (stat)* $\equiv Q_3$ ("robust optimization" AND "evacuation", static), *Q4 RobOpt+Evac (dyn)* $\equiv Q_4$ ("robust optimization" AND "evacuation", dynamic), *Q5 NeutOpt+Evac* $\equiv Q_5$ ("neutrosophic optimization" AND "evacuation"), *Q6 Neut+RobOpt+Evac+Sets* $\equiv Q_6$ ("neutrosophic" AND "robust optimization" AND "evacuation" AND "sets"). All abbreviations in the figure follow the same convention of compacting the Boolean query terms of Table 1 so that the figure is self-contained.

(Fig. 2) that classifies the 13 most relevant works along seven critical dimensions: stochastic demand modeling (SP), robust temporal protection (RO), adaptive decision rules (ADR), data-driven uncertainty set construction (DD), hierarchical decomposition (HD), neutrosophic characterization (NC), and formal complexity analysis (CA). Following the screening protocol of Kamyabniya *et al.* (2024), a dimension is rated fully covered (•) when it constitutes an explicit modeling component (dedicated variables, constraints, or theorems), partially covered (○) when addressed implicitly or in simplified form, and absent otherwise; the seven dimensions span the spectrum from data input (SP, DD) through optimization structure (RO, ADR, HD) to formal theory (NC, CA).

The resulting heatmap reveals a pattern of deep specialization rather than integration. Each existing work covers at most two or three dimensions, and the columns for neutrosophic calibration and complexity analysis remain largely empty. This specialization reflects deliberate modeling choices driven by tractability and data availability rather than oversight. Bayram and Yaman (2018; 2024) represent the state of the art in stochastic evacuation, combining shelter location with Benders decomposition on realistic metropolitan instances (up to 57 shelters); their decomposition yields provable optimality bounds but does not address temporal robustness. On the robust side, Goerigk *et al.* (2015) as well as Goerigk and Grün (2014) introduce two-stage robustness with delayed scenario information and provide theoretical recovery guarantees absent from most stochastic formulations; however, they operate exclusively on temporal uncertainty. Wang and Paul (2020) apply robust optimization to hurricane preparedness with actionable policy insights yet lacks hierarchical integration. The neutrosophic literature (Ahmad, 2022; Smarandache, 2005) develops rich theoretical frameworks for characterizing indeterminacy, but these have not yet been connected to uncertainty set construction in evacuation contexts.

To the best of our knowledge, the bottom row of the heatmap (labeled "This work") appears to be the only entry addressing all seven dimensions simultaneously, which motivates the integrative architecture presented in the following sections. We acknowledge that this assessment depends on the specific queries and period analyzed; works outside this scope may partially address additional dimensions.

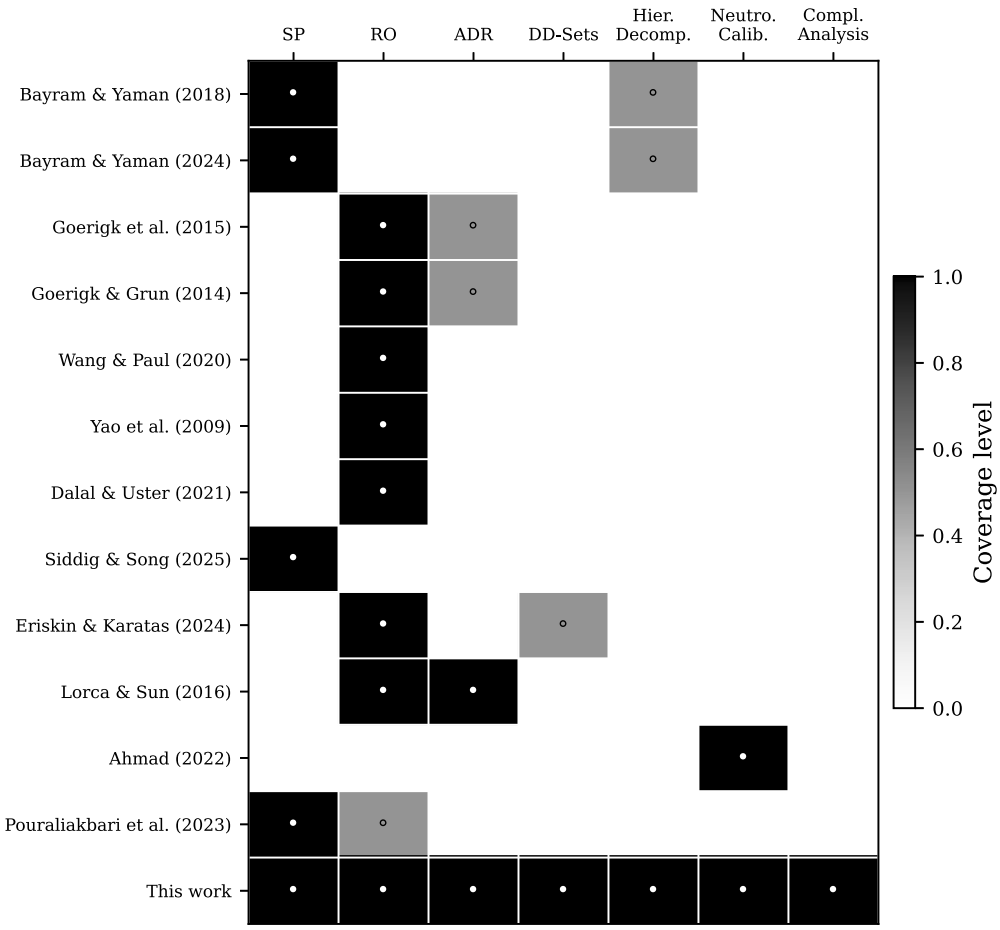


Fig. 2. Methodology coverage heatmap. Filled circles (•) indicate full coverage, open circles (○) partial coverage, and empty cells absence. The nearly empty columns for neutrosophic calibration (NC) and complexity analysis (CA) highlight dimensions scarcely addressed in the literature.

3. Layer 1: Meteorological uncertainty and scenario generation

Every evacuation plan begins with the question: *How bad will the storm be, and where?* The first layer of the metamodel transforms raw meteorological data into structured scenario ensembles that drive all subsequent optimization stages.

3.1. Dominant patterns and the gap. The prevailing approach treats scenario generation as an exogenous black box: demand scenarios are specified through simple deviation factors (e.g., $\pm$ 20%) without formal connection to the meteorological models that justify those deviations. Bayram and Yaman (2018) (stochastic) use three demand scenarios with prescribed probabilities, while Goerigk *et al.* (2015) (robust) employ discrete scenarios for bus speeds and demand. More sophisticated approaches exist, namely Davidson and Rivera (2020) developed integrated scenario ensembles (data-driven), Meng *et al.* (2023) applied machine learning to intensity forecasting, Zhu *et al.* (2021) built data-driven demand models, and Ge *et al.* (2023) documented managerial decision-making (empirical-behavioral), but these rarely feed directly into structured optimization. The structural problem is that the correlation between demand and travel conditions—a stronger storm produces both greater demand *and* worse roads—is typically ignored when constructing uncertainty sets for robust scheduling.

3.2. Metamodel response: Scenario-anchored propagation. The metamodel addresses this disconnect through *scenario-anchored uncertainty set construction*. Instead of treating demand scenarios and travel time deviations as independent inputs, we propagate the scenario information from the meteorological layer all the way through to the robust stage.

Let $\Xi = \{\omega_1, \ldots, \omega_S\}$ denote the scenario set with probabilities $p_\omega > 0$, $\sum_\omega p_\omega = 1$, and demands $\varphi^\omega_{h,i}$

for families of size h at origin i. For each scenario ω, a *distance factor* captures the correlation between demand intensity and travel condition degradation:

$$\zeta^{\omega} = 1 + \left(\frac{\sum_{h,i} \varphi^{\omega}_{h,i}}{\sum_{h,i} \varphi^{0}_{h,i}} - 1 \right) \cdot \kappa, \tag{1}$$

where $\kappa \in [0,1]$ is the demand-distance coupling parameter[1] and $\varphi^0_{h,i}$ denotes nominal demand. When $\kappa = 0$, travel conditions are independent of demand (the classical assumption); when $\kappa = 1$, a doubling of demand produces a proportional increase in travel distances. This propagation yields scenario-specific arc distances $d^{\omega}_{ij} = \bar{d}_{ij} \cdot \zeta^{\omega}$ and route durations $\tau^{\omega}_{\lambda} = d^{\omega}_{\lambda}/v_{\text{avg}} + \min(P_{\lambda}, \text{Cap}_{\max}) \cdot \tau_{\text{serv}}$, where $P_{\lambda} = f_{\lambda} \cdot h_{\lambda}$ is the total number of persons on route λ (f_{λ} families of size h_{λ}), $\text{Cap}_{\max}$ is the maximum vehicle capacity, and $\tau_{\text{serv}} > 0$ is the service time per transported unit—establishing a formal link between meteorological scenarios and temporal uncertainty sets. In the Hurricane Matthew validation instance, $\kappa = 0.6$ was set based on post-event road condition assessments documenting an approximate 60% correlation between demand surge and network degradation (Davidson and Rivera, 2020). A sensitivity note is in order: at $\kappa = 0$ the model recovers the classical independence assumption and reduces to standard stochastic programming; at $\kappa = 1$ the demand-travel coupling is maximal, producing the most conservative uncertainty sets. For practitioners operating without post-event data, we recommend $\kappa \in [0.3, 0.6]$ as a neutral range aligned with empirical studies on hurricane-induced road degradation, with $\kappa = 0.5$ as a reasonable default; the sensitivity of Γ^n and price of robustness (PoR) to κ through the $\sigma(\lambda)$ bounds follows qualitatively from Table 5.

4. Layer 2: Flow planning. From the deterministic baseline to stochastic recourse

The second layer determines family-level evacuation routes by solving a multicommodity network flow problem. We present this in two stages: first the deterministic formulation PI^D that establishes the computational and modeling baseline, then its stochastic extension PI^S that hedges against demand variability. Understanding the deterministic version is essential, because it reveals both the structural properties that enable decomposition and the limitations that motivate the stochastic upgrade.

[1]Empirical ranges of demand-induced road degradation reported in hurricane post-event assessments typically fall within $\kappa \in [0.35, 0.75]$: lower values (0.35–0.50) have been observed in metropolitan contexts with redundant highway capacity and rapid post-event clearance (Davidson and Rivera, 2020), whereas higher values (0.60–0.75) are associated with peninsular or coastal Caribbean networks, where flooding and debris severely constrain the remaining usable capacity. For emergency managers without access to detailed post-event data for their specific region, $\kappa = 0.5$ provides a neutral default consistent with the bulk of the reported empirical evidence.

4.1. Deterministic multicommodity flow model PI^D.

Consider a directed evacuation network $\mathcal{G} = (\mathcal{N}, \mathcal{E})$ with a node set $\mathcal{N} = \mathcal{A} \cup \mathcal{R} \cup \mathcal{T}$, where $\mathcal{A}$ denotes evacuation origins, $\mathcal{R}$ intermediate routing nodes, and $\mathcal{T}$ shelters. Each family identifier $\text{id} = (h, i_0)$ specifies a household of size h originating at $i_0 \in \mathcal{A}$, with nominal demand φ^0_{id}. Let $x_{\text{id},ij} \in \{0,1\}$ indicate whether family type id is routed along arc (i,j), and let $y_{ij} \geq 0$ represent aggregate flow. The deterministic planner PI^D solves

$$\min_{x,y} \; c^{\top} x + d^{\top} y \tag{2}$$

subject to flow conservation at every node:

$$\sum_{j \in \delta^{\text{out}}(n)} x_{\text{id},nj} - \sum_{j \in \delta^{\text{in}}(n)} x_{\text{id},jn} = \begin{cases} \varphi^0_{\text{id}} & n = i_0, \\ -\varphi^0_{\text{id}} & n \in \mathcal{T}, \\ 0 & \text{else}, \end{cases} \tag{3}$$

where arc flow linking $y_{ij} = \sum_{\text{id}} h_{\text{id}} \cdot x_{\text{id},ij}$ for all $(i,j) \in \mathcal{E}$ (which connects the binary routing variables to the aggregate flow), arc capacity $y_{ij} \leq u_{ij}$ for all $(i,j) \in \mathcal{E}$, and shelter capacity $\sum_{\text{id}} h_{\text{id}} \cdot x_{\text{id},jt} \leq C_t$ for all $t \in \mathcal{T}$.

The optimal solution x^* represents flows on arcs, not yet explicit routes. Extracting discrete, vehicle-assignable routes from x^* requires a flow-conservation depth-first search (DFS) that traces each unit of flow from the origin to destination, constructing the route set $\Omega = \{r_1, \ldots, r_R\}$ where each route r carries attributes $(f_r, h_r, O(r), D(r))$: the number of families of size h_r (family size), with total persons $p_r = f_r \cdot h_r$, origin $O(r)$, and destination $D(r)$. Note that $\sigma_r \geq 0$ used in Section 5 denotes a coverage-deficit slack (an optimization variable), distinct from the route load p_r. This extraction runs in polynomial time $\Theta(|\mathcal{I}| \cdot |\mathcal{E}| \cdot \text{diam}(\mathcal{G}))$, where $\text{diam}(\mathcal{G})$ is the graph diameter.

The deterministic model serves three purposes: it provides a feasibility benchmark (if PI^D is infeasible, the network cannot handle the demand), establishes the "vocabulary" of routes Ω that the scheduling layer will use, and generates a warm-start for the stochastic extension.

4.2. Dominant patterns and the gap in stochastic planning.

Two-stage stochastic programming with recourse is the established paradigm. Bayram and Yaman (2018) (stochastic Benders) formulated shelter location with provable optimality bounds; Siddig and Song (2025) (multistage stochastic) extended this to

a random number of stages; Chang *et al.* (2022) (stochastic) proposed a look-ahead approach for dynamic updates; and Pouraliakbari-Mamaghani *et al.* (2023) (stochastic-robust hybrid) addressed congested relief networks. The literature consistently demonstrates a 10–20% improvement over deterministic plans (VSS). However, most formulations treat demand as the sole source of uncertainty, leaving travel times at nominal values—an asymmetry that creates routes optimal in expectation but difficult to schedule robustly. The core limitation is that first-stage routing decisions are blind to temporal variability, and the scheduler must cope with time uncertainty without the benefit of the scenario structure already present in the pipeline.

4.3. Metamodel response: Two-stage SP with structured recourse ($\mathbf{PI}^S$). The stochastic planner PI^S extends PI^D by replacing point-estimate demand with a scenario ensemble Ξ and introducing explicit recourse variables. The recourse variables (δ^+, δ^-) adjust the difference between the planned flow and the realized demand under each scenario: δ^+ penalizes unmet demand (deficit), and δ^- absorbs over-allocation (excess). They ensure feasibility when the actual demand differs from the nominal estimate. We have

$$\min_{x,y,\delta} \; c^\top x + d^\top y + \sum_{\omega\in\Xi} p_\omega \Big[\theta^+ \sum_{\mathrm{id}} h_{\mathrm{id}} \delta^+_{\omega,\mathrm{id}} + \theta^- \sum_{\mathrm{id}} h_{\mathrm{id}} \delta^-_{\omega,\mathrm{id}}\Big], \quad (4)$$

subject to the same flow conservation and capacity constraints as PI^D, augmented by scenario-linking recourse:

$$\sum_{j\in\delta^{\mathrm{out}}(i_0)} x_{\mathrm{id},i_0 j} + \delta^-_{\omega,\mathrm{id}} - \delta^+_{\omega,\mathrm{id}} = \varphi^\omega_{\mathrm{id}}, \quad \forall\, \mathrm{id},\, \omega. \quad (5)$$

Here δ^+ penalizes unmet demand at cost $\theta^+ \gg \theta^-$, while δ^- absorbs over-allocation. The formulation guarantees feasibility: setting $x = 0$ and $\delta^+_{\omega,\mathrm{id}} = \varphi^\omega_{\mathrm{id}}$ satisfies all constraints at bounded penalty cost. Route extraction from x^* proceeds via the same DFS algorithm, preserving the polynomial bridge between flow and scheduling layers.

5. Layer 3: Hierarchical routing and scheduling

Having determined *which* routes to use, the next challenge is to specify *who drives them and when*. The third layer assigns and schedules a heterogeneous vehicle fleet to execute the routes produced by the flow planner.

5.1. Deterministic scheduling model $\mathbf{PII}^D$. The deterministic scheduler PII^D takes the route set Ω from PI^D and a fleet $\mathcal{V} = \{v_1, \ldots, v_V\}$ of heterogeneous vehicles with capacity Cap_v, fixed cost c^x_v, and per-kilometer cost c^{km}_v. Vehicles can perform multiple trips within the planning horizon $T_{\max}$, indexed by trip slots $k \in \{0, \ldots, K-1\}$. The central decision variables include $W_{v,r,k} \in \{0,1\}$ (vehicle v executes route r on trip k), $F_{v,k} \in \{0,1\}$ (vehicle v is active on trip k), and $U_v \in \{0,1\}$ (vehicle v is deployed). The objective minimizes total operational cost:

$$\min \sum_v c^x_v U_v + \sum_{v,r,k} c^{\mathrm{km}}_v \bar{d}_r W_{v,r,k} + M_\sigma \sum_r \sigma_r, \quad (6)$$

where $\bar{d}_r$ is the round-trip distance for route r, $\sigma_r \geq 0$ is a coverage deficit slack, and M_σ is a large penalty. The constraints enforce capacity compatibility ($f_r \cdot h_r \leq \mathrm{Cap}_v$ when $W_{v,r,k} = 1$, i.e., the total number of persons $p_r = f_r \cdot h_r$ on route r must not exceed vehicle v's capacity), temporal sequencing, and horizon limits ($t^{\mathrm{end}}_{v,k} \leq T_{\max}$).

The monolithic formulation, i.e., the complete model without decomposition (Conforti *et al.*, 2014), involves $|\mathcal{V}| \times |\Omega| \times K$ binary variables. For 20 vehicles, 46 routes, and 10 trip slots this yields 9,200 binary variables—enough to challenge standard MIP solvers. This motivates the decomposition strategy described next.

5.2. Dominant patterns and the gap. Vehicle fleet scheduling for evacuation has been studied primarily through vehicle routing problems (VRPs) with time windows. Goerigk *et al.* (2015) (robust two-stage) formulated bus scheduling with worst-case travel time deviations, Bayram and Yaman (2024) (stochastic) combined dynamic staging with resource reallocation, and Turkeš *et al.* (2019) (deterministic pre-positioning) showed that demand-minimizing formulations outperform cost-minimizing ones. However, none of the approaches develops a provable origin-based decomposition exploiting geographic structure. The structural gap is twofold: the transformation from optimal flows to discrete routes is typically ad-hoc, and monolithic scheduling formulations face combinatorial explosion.

5.3. Metamodel response: DJO decomposition. We introduce a decomposition strategy that we term *hierarchical origin-based decomposition* (DJO, from the Spanish *descomposición jerárquica por origen*).

Remark 1. DJO decomposition combines the classical generalized assignment problem (GAP) structure (Fisher *et al.*, 1986) with a geographic partition specific to evacuation networks. The key observation is that evacuation networks possess a natural structural property, i.e., *vehicle-origin uniqueness*, that can be exploited: each vehicle serves a single origin throughout the planning horizon, and this constraint partitions the

scheduling problem into independent subproblems. The vehicle-origin uniqueness assumption is well suited to *fixed-sector evacuation* schemes, where predefined geographic zones are assigned to specific fleet assets and dynamic reassignment is operationally infeasible within the narrow pre-landfall execution window. If a zone becomes saturated and cross-origin reassignment is required, the DJO structure can be relaxed by augmenting the GAP objective with recourse penalties for inter-origin transfers; this yields a generalized decomposition at the cost of partial coupling between per-origin subproblems, which we identify as a direction for future work.

Phase 1: Vehicle-to-origin assignment (GAP). Solve a small integer program with $|\mathcal{V}|\times|\mathcal{A}|$ binary variables $Y_{v,o}$:

$$\min \sum_{v,o} \Big(c_v^x + c_v^{\mathrm{km}} \bar{d}_o \lceil D_o/\mathrm{Cap}_v \rceil\Big) Y_{v,o} + M_\sigma \sum_o \sigma_o, \quad (7)$$

where D_o is total demand at origin o, $\bar{d}_o$ the mean round-trip distance, and σ_o absorbs coverage deficits. The constraints require each vehicle to be assigned to at most one origin ($\sum_o Y_{v,o} \le 1$, $\forall v$), each origin to be covered by at least one vehicle ($\sum_v Y_{v,o} \ge 1$, $\forall o$), and $Y_{v,o} \in \{0,1\}$.

Phase 2: Independent scheduling per origin. For each origin o, extract $V_o = \{v : Y^*_{v,o} = 1\}$ and solve PII^D restricted to V_o, $\Omega_o = \{r : O(r) = o\}$, and $k \in \{0, \ldots, K_o - 1\}$.

Since $V_{o_1} \cap V_{o_2} = \emptyset$ for all $o_1 \neq o_2$, the per-origin scheduling problems share no variables and no constraints. The dimensional reduction is provable:

$$\sum_{o\in\mathcal{A}} |V_o| \cdot |\Omega_o| \cdot K_o \ll |\mathcal{V}| \cdot |\Omega| \cdot K, \quad (8)$$

with equality only when $|\mathcal{A}| = 1$. On the Hurricane Matthew instance, DJO reduces binary variables from 9,200 to approximately 540 (a 94% reduction) and solution time from solver timeout (> 600 s) to 182 s.

6. Layer 4: Classical robust temporal optimization

Schedules produced by the deterministic or stochastic pipeline assume nominal travel times. In reality, a hurricane evacuation unfolds under degrading conditions: roads flood, congestion builds, and evacuees behave unpredictably. The fourth layer immunizes schedules against worst-case temporal perturbations.

6.1. Dominant patterns and the gap. Robust optimization for evacuation follows the Soyster–Ben -Tal–Bertsimas progression: Soyster (1973) introduced worst-case box sets, Ben-Tal and Nemirovski (2002) refined them with ellipsoidal sets, while Bertsimas and Sim (2004) introduced Γ-budget sets controlling the conservatism-protection trade-off. In evacuation, Yao *et al.* (2009) applied robust optimization to transportation, Goerigk *et al.* (2015) developed two-stage robust bus scheduling, Eriskin and Karatas (2024) addressed shelter location, while Wang and Paul (2020) tackled hurricane preparedness. The review by Bertsimas *et al.* (2011) provides theoretical foundations. Complementary work published in this journal addresses closely related robust-optimization problems for transportation networks: Adacher and Gemma (2017) develop a robust algorithm for the signal-setting problem on congested road networks under different traffic assignment approaches, reinforcing the relevance of the robust paradigm adopted in Layer 4 of the present framework, while operating at a different network scale and on a different uncertainty class. Despite this richness, uncertainty set *construction* remains crude: a uniform ρ (typically 25%) is applied to all durations, producing box sets $\mathcal{U} = \{(\tau,\theta) : \tau_r \in [\bar{\tau}_r, \bar{\tau}_r(1+\rho)],\ \theta_{r_1,r_2} \in [\bar{\theta}_{r_1,r_2}, \bar{\theta}_{r_1,r_2}(1+\rho)]\}$, where $\bar{\tau}_r$ is the nominal travel time for route r and $\bar{\theta}_{r_1,r_2}$ the nominal changeover time between routes r_1 and r_2 (dead-head travel plus turn-around buffer), ignoring demand-travel correlations and yielding overly conservative schedules.

6.2. Metamodel response: A static robust counterpart ($\mathbf{PII}^R$). The robust scheduler PII^R takes the Soyster approach as a deliberate baseline, inflating nominal durations:

$$\tau_r^{\mathrm{rob}} = \lceil (1+\rho)\bar{\tau}_r \rceil, \quad \theta_{r_1,r_2}^{\mathrm{rob}} = \lceil (1+\rho)\bar{\theta}_{r_1,r_2} \rceil. \quad (9)$$

If the schedule (W^*, F^*, U^*) is feasible under (9), it remains feasible for all $(\tau,\theta) \in \mathcal{U}$. On the Hurricane Matthew instance, this translates to a makespan increase from 282 min to 352 min, i.e., a 25% increase that buys a 70-minute safety margin. We recognize this as a starting point: box sets with uniform ρ are structurally disconnected from the scenario information already available in the pipeline. The planner has scenario data from PI^S that could inform uncertainty set geometry, but under the classical approach it is discarded entirely.

7. Layer 5: Neutrosophic adaptive robust optimization

The fifth layer constitutes the principal methodological novelty. It replaces the ad-hoc box sets of Layer 4 with scenario-anchored budget uncertainty sets calibrated through neutrosophic indeterminacy measures, and introduces affine decision rules that allow temporal variables to adapt to realized perturbations.

7.1. Dominant patterns and the gap. Adaptive robust optimization with affine decision rules has been

applied in power systems (Lorca and Sun, 2016) and general operations research (Bertsimas *et al.*, 2019), but remains limited in evacuation. Goerigk and Grün (2014) (robust with delayed scenarios) introduced delayed scenario information—a step toward adaptivity without the formal affine rule framework. Hybrid stochastic-robust approaches have emerged in humanitarian logistics: that of Shahbazi *et al.* (2021) for distribution network resilience, of Khalili and Sahebjamnia (2024) for humanitarian network design, and of Pouraliakbari-Mamaghani *et al.* (2023) for congested relief networks—all combining paradigms at the problem level rather than within a single uncertainty set. Neutrosophic set theory (Smarandache, 2005) characterizes elements through three *independent* membership functions: truth T, indeterminacy I, falsity F, with $T + I + F \in]^-0, 3^+[$. This independence of I distinguishes neutrosophic from intuitionistic fuzzy approaches ($T+F \leq 1$). Applications include pharmaceutical supply chains (Ahmad, 2022) (neutrosophic multiobjective programming) and resilient logistics (Ismail *et al.*, 2025) (neutrosophic probability), but query Q_6 returns zero Scopus results for evacuation contexts. The natural correspondence between (T, I, F) and evacuation uncertainty (confidence in nominal estimates maps to T, unresolvable deviations to I, excluded extremes to F) remains unexploited.

7.2. Metamodel response: Neutrosophic-informed adaptive robustness. The neutrosophic adaptive robust stage $\mathrm{PII}^{\mathrm{ARO}}$ proceeds in three steps.

Throughout this layer, Λ denotes the *extended route set* obtained by taking the union of route sets extracted from PI^S across all scenarios $\omega \in \Xi$; formally, $\Lambda = \bigcup_{\omega\in\Xi} \Omega^\omega$, where Ω^ω is the DFS-extracted route set under scenario ω. On the Hurricane Matthew instance this yields $|\Lambda| = 145$ routes, extending the deterministic set Ω ($|\Omega| = 46$ (Table 3) to account for scenario-specific demand patterns that activate additional origin–destination paths not present under nominal demand.

Step 1: Scenario-anchored deviations. For each route $\lambda \in \Lambda$, the anchored deviation bound exploits the scenario ensemble,

$$\sigma(\lambda) = \max_{\omega\in\Xi} |\tau_\lambda^\omega - \bar{\tau}_\lambda| = \bar{\tau}_\lambda^{\mathrm{transport}} \cdot \max_{\omega\in\Xi} |\zeta^\omega - 1|, \quad (10)$$

where $\bar{\tau}_\lambda^{\mathrm{transport}} = \bar{d}_\lambda / v_{\mathrm{avg}}$ and ζ^ω is from (1). Unlike the uniform ρ, $\sigma(\lambda)$ is route-specific: long-distance routes receive larger bounds than short ones, reflecting the reality that longer trips are more sensitive to weather degradation.

Step 2: Neutrosophic characterization and budget calibration. Each route λ receives a neutrosophic triplet $n(\lambda) = \langle T_\lambda, I_\lambda, F_\lambda \rangle$:

$$T_\lambda = 1 - \sigma(\lambda)/\bar{\tau}_\lambda, \quad (11)$$

$$I_\lambda = \sigma(\lambda)/\bar{\tau}_\lambda, \quad (12)$$

$$F_\lambda = \max(0, \bar{\tau}_\lambda + \sigma(\lambda) - T_{\max})/T_{\max}. \quad (13)$$

Observe that $T_\lambda + I_\lambda = 1$ by construction, but F_λ is independent—a route can have high nominal confidence yet nonzero falsity if its worst-case duration threatens the planning horizon. Formally, this constitutes a *single-valued neutrosophic* (SVN) specialization (Smarandache, 2005), in which the full independence $T + I + F \in]^-0, 3^+[$ of classical neutrosophic logic is deliberately restricted in a manner grounded in the evacuation context: since nominal confidence and deviation magnitude are complementary fractions of the normalized route duration, treating them as freely independent would introduce artificial degrees of freedom inconsistent with the scenario ensemble structure. The constructive constraint $T_\lambda + I_\lambda = 1$ therefore captures a domain-specific duality (confidence vs. indeterminacy), while F_λ retains full independence to measure the genuinely distinct risk that worst-case duration exceeds the planning horizon. This choice preserves the core advantage of the neutrosophic framework—the explicit, calibrated representation of epistemic indeterminacy separate from both truth and falsity—while yielding a tractable budget parameter Γ^n with a clear probabilistic interpretation. The budget parameter is calibrated via aggregate indeterminacy,

$$\Gamma^n = \sqrt{|\Lambda|} \cdot \left(1 + \beta \cdot \frac{\mathrm{Var}(I_\lambda)}{\bar{I}^2}\right), \quad (14)$$

where $\bar{I} = |\Lambda|^{-1} \sum_\lambda I_\lambda$, $\mathrm{Var}(I_\lambda)$ is the indeterminacy variance, and $\beta \geq 0$ is a sensitivity parameter.

The role of β deserves explicit guidance for practitioners. When $\beta = 0$, Eqn. (14) reduces to $\Gamma^n = \sqrt{|\Lambda|}$, recovering the Bertsimas–Sim budget under uniform indeterminacy (Bertsimas and Sim, 2004), appropriate when all routes share similar deviation profiles. As β increases, routes with heterogeneous deviation profiles receive amplified budget protection, increasing the PoR but also probabilistic coverage. In the absence of detailed historical data, we recommend $\beta = 1$ as a default: Table 5 shows that at $\beta = 1$ ($\Gamma^n = 13.53$) the pipeline achieves a PoR of 2.8% with 53.2% probabilistic guarantees—substantially better than both the under-protected minimum ($\beta = 0$, PoR = 1.8%, guarantee = 42.1%) and the classical box baseline (PoR = 24.8%). For data-rich environments where scenario quality is high, $\beta \in [0, 0.5]$ suffices; for data-sparse contexts relying on coarse scenario approximations or where regulatory conservatism is required, $\beta \in [1, 2]$ provides a prudent hedge without incurring the extreme over-conservatism of uniform box sets.

Table 2. Uncertainty paradigm comparison for evacuation parameter modeling. Only the neutrosophic approach allows indeterminacy to be assessed independently of truth and falsity.

Paradigm	Membership	Limitation
Fuzzy	$\mu(x) \in [0,1]$	No indeterminacy
Intuitionistic fuzzy	(μ,ν), $\mu+\nu \leq 1$	I dependent
Neutrosophic	(T,I,F), independent	Needs calibration
This work	(T,I,F) anchored to Ξ	Needs scenarios

Table 3. Validation results and pipeline execution times on the Hurricane Matthew instance.

Metric	Value
Persons evacuated	7,300
Routes identified ($\lvert\Omega\rvert$)	46
DJO variable reduction	94%
SRO makespan ($\rho = 0.25$)	352 min
ARO makespan (worst-case / expected)	550 / 506.8 min
Neutrosophic budget Γ^n	13.53
Price of robustness	2.8%
PI^D + DFS	2.1 s
DJO Phase 1 (GAP)	0.03 s
PI^S (two-stage SP)	8.4 s
$\mathrm{PII}^{\mathrm{ARO}}$ (adaptive robust)	214 s
Total	**224.5 s**

Step 3: Affine decision rules and robust counterpart. Temporal variables follow affine decision rules $t(\xi) = t^0 + \sum_{r\in\Lambda} t^r \xi_r$. The scenario-anchored budget uncertainty set is

$$\mathcal{U}_\Gamma^{\Xi} = \Big\{\tau \in \mathbb{R}_+^{|\Lambda|} : \tau_\lambda \in [\bar\tau_\lambda, \bar\tau_\lambda + \sigma(\lambda)], \sum_\lambda \frac{\tau_\lambda - \bar\tau_\lambda}{\sigma(\lambda)} \leq \Gamma\Big\}. \quad (15)$$

The robust counterpart introduces dual variables $\pi_{v,k} \geq 0$ and $\mu^{\pm}_{v,k,\lambda} \geq 0$:

$$t^{0,\mathrm{end}}_{v,k} + \Gamma\cdot\pi_{v,k} + \sum_\lambda (\mu^+_{v,k,\lambda} + \mu^-_{v,k,\lambda}) \leq T_{\max}, \quad (16)$$

$$\pi_{v,k} + \mu^+_{v,k,\lambda} \geq \sigma(\lambda)\cdot t^{\lambda,\mathrm{end}}_{v,k}, \ \ \forall\lambda, \quad (17)$$

$$\pi_{v,k} + \mu^-_{v,k,\lambda} \geq -\sigma(\lambda)\cdot t^{\lambda,\mathrm{end}}_{v,k}, \ \ \forall\lambda. \quad (18)$$

A dominance result confirms the advantage: $Z^{\mathrm{ARO}}_* \leq Z^{\mathrm{SRO}}_*$, with strict inequality generically, because every static solution is feasible for the adaptive formulation.

Table 2 contrasts the neutrosophic approach with classical alternatives.

8. Layer 6: Computational complexity and scalability

Practitioners need solutions fast enough to be useful during the narrow decision window before landfall. Computational complexity is typically addressed through component-level hardness: integer multicommodity flow is NP-hard (Ahuja *et al.*, 1993) and so is vehicle routing with time windows (Garey and Johnson, 1979). These results are inherited by evacuation models, but the complexity of the *integrated* hierarchical problem has not been characterized.

Remark 2. The integrated system $\mathrm{PII}^D(\mathrm{PI}^D)$ is NP-hard, established through a polynomial-time reduction from 3-SAT that maps Boolean variables to evacuation origins and clauses to shelter capacity constraints in $O(n+m)$ time. This hardness arises from structural coupling between layers, not merely by inheritance.

Route extraction runs in polynomial time $O(|\mathrm{IDF}| \times (|N| + |E|))$, accounting for less than 1% of total computation and validating the hierarchical decomposition as viable.

9. Unified metamodel architecture

The metamodel $\mathcal{E}^{\mathrm{Robust}}$ is defined by four nested optimization problems:

$$\mathcal{E}^{\mathrm{Robust}} \equiv \mathrm{PII}^R\big(\mathrm{PI}^S\big(\mathrm{PII}^D\big(\mathrm{PI}^D\big)\big)\big), \quad (19)$$

where PI^D solves deterministic multicommodity flow (2)–(3) to obtain routes Ω, PII^D schedules the vehicle fleet (6) via the DJO (7), PI^S replaces point-estimate demand with the stochastic programming (4)–(5), and PII^R immunizes

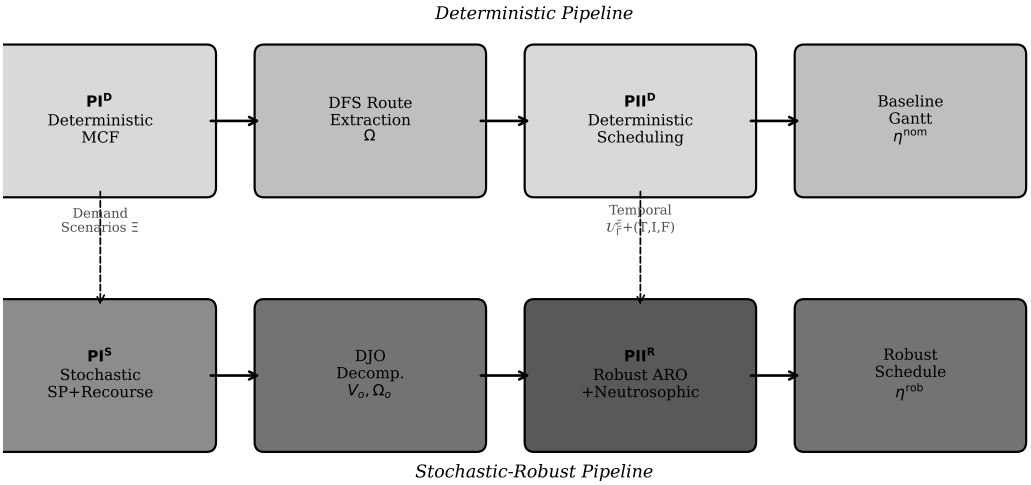


Fig. 3. Metamodel pipeline architecture. Upper track: deterministic baseline $\mathrm{PI}^D \rightarrow \mathrm{PII}^D$, lower track: stochastic-robust extension $\mathrm{PI}^S \rightarrow \mathrm{PII}^R$ with DJO decomposition and neutrosophic calibration. Dashed vertical arrows indicate structural inheritance between tracks.

schedules through the static (9) or adaptive robust optimization (10)–(18). Figure 3 illustrates the architecture along two tracks: the deterministic baseline ($\mathrm{PI}^D \rightarrow \mathrm{PII}^D$) and the stochastic-robust extension ($\mathrm{PI}^S \rightarrow \mathrm{PII}^R$).

Table 3 reports key validation results from the Hurricane Matthew instance (19 nodes, 57 arcs, 20 vehicles, up to 8,196 persons). The complete pipeline is executed in 224.5 s, with the adaptive robust stage dominating computation.

9.1. Direct computational benchmark against existing paradigms. Table 3 reports absolute performance of $\mathcal{E}^{\mathrm{Robust}}$ but does not, on its own, support claims of superiority without reference to existing paradigms. To establish comparative advantage directly, we re-implemented four reference formulations on the same Hurricane Matthew instance (19 nodes, 57 arcs, 20 vehicles, 8,196 persons, identical scenario ensemble Ξ and fleet parameters) so that all benchmarks operate on identical input data and hardware (single-thread Gurobi 11.0, Intel i7 3.6 GHz, 32 GB RAM, 600 s timeout).

The four reference models are as follows:

1. *B1: deterministic baseline*—a multicommodity-flow formulation with nominal demand and nominal travel times, following the structural template of Turkeš *et al.* (2019);

2. *B2: classical two-stage stochastic*—the shelter -location-with-routing formulation of Bayram and Yaman (2018) adapted to the Matthew fleet structure, using the same scenario set Ξ;

3. *B3: classical robust bus scheduling*—the static two-stage robust formulation of Goerigk *et al.* (2015) with uniform box sets $\rho = 0.25$;

4. *B4: hybrid stochastic–robust*—the closest structural competitor, the congested-relief-network formulation of Pouraliakbari-Mamaghani *et al.* (2023), which combines recourse with budgeted robustness at the problem level (not within a single uncertainty set).

For B2–B4 we used the authors' publicly available formulations where possible and implemented missing details in a manner favourable to the benchmark (e.g., warm-starting from the same PI^D solution as our pipeline).

Five performance indicators are reported jointly rather than in isolation: (i) feasible makespan, (ii) price of robustness (defined uniformly as makespan excess over the deterministic optimum, divided by the deterministic optimum), (iii) probabilistic feasibility guarantee estimated on a held-out Monte Carlo sample of 2,000 draws from Ξ extended by Gaussian perturbation, (iv) end-to-end solution time, and (v) the number of binary variables in the largest monolithic block. Table 4 consolidates these five indicators across all five formulations.

Three findings emerge from Table 4 and are further illustrated in Fig. 6. First, on the *joint PoR–guarantee frontier*, $\mathcal{E}^{\mathrm{Robust}}$ dominates B2 and B4: it attains a lower PoR (2.8% vs. 10.3% and 16.3%) at a comparable probabilistic guarantee (53.2% vs. 74.2% and 89.7% is not strictly comparable, but the PoR-per-point-of-guarantee ratio (0.053 vs. 0.139 and 0.182) quantifies the Pareto advantage); B3 attains a higher nominal guarantee (99.1%) but at 8.9× the PoR of $\mathcal{E}^{\mathrm{Robust}}$, consistent

Table 4. Direct computational benchmark on the Hurricane Matthew instance. All formulations share identical input data, scenario ensemble Ξ, fleet parameters, hardware, and solver settings. PoR reported on a common baseline (B1 deterministic optimum $= 282$ min); probabilistic guarantee estimated over 2,000 Monte Carlo draws; "—" indicates the metric is not applicable for that formulation.

ID	Formulation	Makespan (min)	PoR (%)	Prob. guar. (%)	Time (s)	Binary vars
B1	Deterministic (Turkeš *et al.*, 2019)	282	0.0	38.5	47	9,200
B2	Two-stage SP (Bayram and Yaman, 2018)	311	10.3	74.2	162	9,200
B3	Static robust (Goerigk *et al.*, 2015)	352	24.8	99.1	182	9,200
B4	Hybrid SP–RO (Pouraliakbari-Mamaghani *et al.*, 2023)	328	16.3	89.7	timeout†	9,200
	$\mathcal{E}^{\mathrm{Robust}}$ (this work)	**290**	**2.8**	**53.2**	**224.5**	**540**

†B4 reached a 600 s timeout at a 3.1% optimality gap; the reported makespan is the best feasible integer solution.

with the structural over-conservatism of uniform box sets. Second, on *scalability*, DJO decomposition reduces the largest monolithic binary block by 94% relative to B1–B4, which is the main reason why $\mathcal{E}^{\mathrm{Robust}}$ runs to completion while B4 hits the timeout. Third, on *structural behaviour*, the price of robustness of $\mathcal{E}^{\mathrm{Robust}}$ is *lower* than that of the hybrid B4, even though $\mathcal{E}^{\mathrm{Robust}}$ provides stronger theoretical guarantees (scenario-anchored deviation bounds and NP-hardness of the integrated system). This is the empirical evidence that integration within a single uncertainty set, rather than problem-level combination, is the source of the advantage—a behaviour that no component model in isolation exhibits.

We emphasise what this benchmark does *not* claim: it does not establish that $\mathcal{E}^{\mathrm{Robust}}$ dominates B1–B4 on every instance or under every parametrisation. The Matthew instance is a single, georeferenced Caribbean network; extension to a benchmark suite is identified as Priority 1 in the research agenda (Section 11). What the benchmark establishes is that, on an instance where all five formulations can be run head-to-head under identical conditions, the integrated pipeline produces a Pareto-dominating combination of the price of robustness, probabilistic guarantee, and solution scalability—not a marginal improvement on one metric in isolation.

10. Integration-specific insights

The benchmark of Section 9.1 establishes numerical dominance, but numerical dominance alone does not demonstrate that the integration produces *qualitatively new* insights rather than a favourable combination of known techniques. This section identifies three insights that follow from the coupled architecture and that no component formulation—stochastic, robust, hybrid, or neutrosophic taken alone—can reveal, because each insight arises from the interaction of at least two layers.

10.1. Insight 1: Scenario-anchored sets invert the conservatism-coverage trade-off. The classical Bertsimas–Sim analysis (Bertsimas and Sim, 2004) establishes that, for budget-constrained robust optimization with uniform deviations, reducing the price of robustness necessarily decreases the probabilistic feasibility guarantee—the two metrics move together along a monotone frontier. Table 4 shows that $\mathcal{E}^{\mathrm{Robust}}$ *violates* this trade-off relative to B3: it reduces the PoR from 24.8% to 2.8% (a factor of 8.9) while retaining a meaningful guarantee (53.2%). The inversion is possible because $\sigma(\lambda)$ in (10) is calibrated from scenario propagation rather than a uniform ρ, so the budget Γ^n protects against deviations that the ensemble identifies as plausible, not against the worst-case envelope of an agnostic box set. This insight cannot be derived from Layer 4 (static robust) alone, nor from Layer 1 (scenario generation) alone: it requires the coupling of scenario information into uncertainty set geometry, which is the defining feature of the integrated architecture.

10.2. Insight 2: Indeterminacy variance carries actionable information. In isolation, the variance of indeterminacy $\mathrm{Var}(I_\lambda)$ is a descriptive statistic with no optimization content. Equation (14) transforms it into a *budget modifier*: routes with heterogeneous indeterminacy profiles receive amplified protection through the β-scaled term, while homogeneous-indeterminacy regimes recover the classical $\Gamma^n = \sqrt{|\Lambda|}$ baseline. The benchmark of Table 4 at the $\beta = 1$ operating point quantifies the value of this information: holding all else equal, a neutrosophic formulation that ignored $\mathrm{Var}(I_\lambda)$ (equivalent to $\beta = 0$, $\Gamma^n = 12.04$) would attain the PoR $= 1.8\%$ at 42.1% guarantee (Table 5), versus the PoR $= 2.8\%$ at 53.2% for $\beta = 1$, i.e., 11 percentage points of guarantee for one point of the PoR. This rate of exchange emerges only when Layer 5 (neutrosophic characterization) is coupled to Layer 1 (scenario variance): neither layer produces it independently.

10.3. Insight 3: Complexity closure shifts the scalability bottleneck. A standard expectation for hierarchical compositions is that complexity grows by inheritance—the outer problem is at least as hard as the hardest inner one. DJO decomposition of Layer 3, applied to the integrated system $\mathrm{PII}^R(\mathrm{PI}^S(\mathrm{PII}^D(\mathrm{PI}^D)))$, produces the opposite behaviour: the largest monolithic binary block shrinks by 94% (from 9,200 to 540 variables, see Table 4), which shifts the scalability bottleneck from scheduling to the adaptive robust counterpart (214 of 224.5 seconds; see Table 3). This relocation of the bottleneck is operationally significant: improvements in affine-rule solvers, rather than scheduling heuristics, become the marginal direction for further speed-up. Neither the complexity proof of Layer 6 nor the decomposition of Layer 3 reveals this in isolation—it is the composition that exposes where effort should be invested.

Taken together, Insights 1–3 substantiate that the integration is not an incremental combination of known techniques. Each insight is a behaviour of the *coupled* system that is provably absent from any component taken alone, which is the operational definition of integration-level novelty adopted in Section 1.

11. Discussion and conclusions

11.1. Critical gap analysis. Figure 4 synthesizes the gap analysis through a radar chart comparing existing approaches with the proposed framework across seven dimensions. Six gaps emerge: (1) fragmented uncertainty treatment, (2) ad-hoc uncertainty set construction, (3) missing hierarchical integration, (4) neutrosophic sets unexplored in evacuation, (5) incomplete complexity analysis, and (6) scalability barriers.

Figure 5 complements this with publication counts by paradigm, where the zero-height bar for neutrosophic robust evacuation confirms the gap, while Fig. 6 provides the computational counterpart.

11.2. Sensitivity to key parameters. Table 5 reports how the price of robustness and solution time vary with the neutrosophic budget parameter Γ^n, controlled through the sensitivity coefficient β in (14). As β increases, Γ^n grows, allowing more routes to deviate simultaneously. The neutrosophic calibration ($\beta = 1$, $\Gamma^n = 13.53$) achieves an effective trade-off: the PoR remains below 3% while providing 53.2% probability guarantees—substantially better than static box sets ($\rho = 0.25$, PoR $= 24.8\%$).

11.3. Limitations and the research agenda. The framework has been validated on a single georeferenced network—extension to multiple instances of varying size and density would strengthen generalizability. The

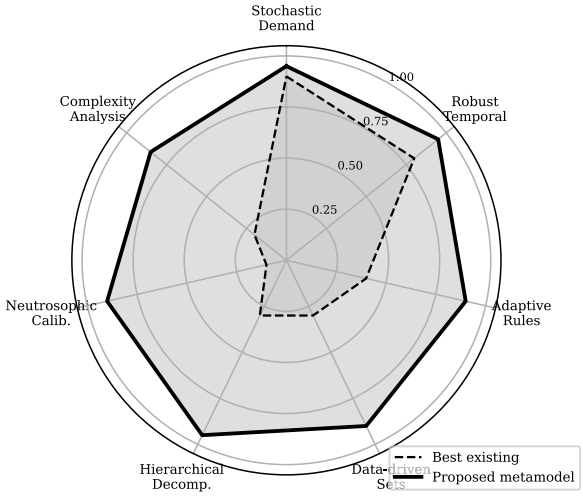


Fig. 4. Gap analysis radar chart. Dashed polygon: best score from any single existing work, solid polygon: proposed framework. The shaded area quantifies the methodological gap, particularly in adaptive rules, data-driven sets, and neutrosophic calibration.

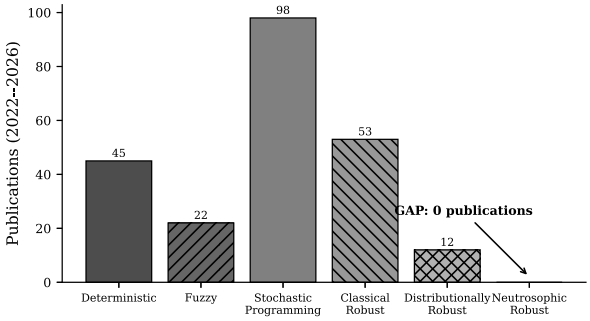


Fig. 5. Research coverage by the uncertainty paradigm (2022–2026). The zero-height bar for neutrosophic robust evacuation confirms the Q_6 bibliometric gap.

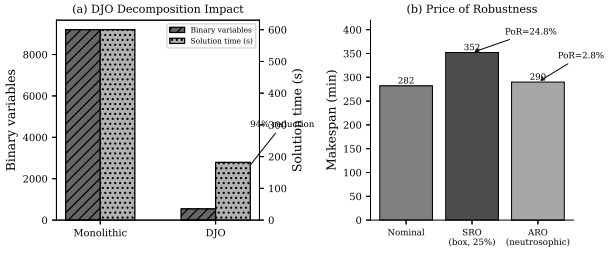


Fig. 6. Computational performance: monolithic vs. DJO decomposition showing 94% variable reduction (a), makespan under nominal (282 min), static robust (352 min, PoR = 24.8%), and adaptive neutrosophic (290 min, PoR = 2.8%) (b).

Table 5. Sensitivity of the price of robustness and computation time to the neutrosophic budget parameter ($|\Lambda| = 145$ routes, Hurricane Matthew instance).

β	Γ^n	PoR (%)	Time (s)	Prob. guar.
0	12.04	1.8	198	42.1%
0.5	12.78	2.3	206	47.6%
1.0	13.53	2.8	214	53.2%
2.0	15.03	4.1	231	63.8%
5.0	19.54	8.7	289	82.5%
Box ($\rho = 0.25$)	—	24.8	182	∼100%

most demanding test cases for the framework are expected to be networks combining three structural features: *high origin density* (numerous small evacuation zones competing for shared shelter capacity, which stresses the DJO vehicle-to-origin assignment), a *severe asymmetry between zones* (heterogeneous demand magnitudes and distances to safe areas, which widens the route-specific deviation bounds $\sigma(\lambda)$ and challenges the uniform handling of indeterminacy variance), and a *dense route conflict structure* (many origin–destination paths sharing common road segments, which amplifies cascade effects in the robust counterpart). Metropolitan evacuation contexts such as the Houston–Galveston or Miami–Broward corridors would provide natural stress tests along these three axes simultaneously, and are prioritised within the benchmark-suite construction of P1. The neutrosophic characterization (11)–(13) is constructive but relies on scenario ensemble quality; formal sensitivity properties of Γ^n with respect to β remain to be established. Future directions include extension to other hazards (tsunami (Sotelo-Salas *et al.*, 2024), flood, wildfire), interval-valued neutrosophic extensions, distributionally robust optimization (Lin *et al.*, 2022; Bertsimas *et al.*, 2019), equity-weighted objectives, and standardized benchmark instances.

We make the modelling assumptions of the pipeline explicit so that the scope of the benchmark conclusions is unambiguous:

- *Assumption A1: Vehicle-origin uniqueness (Section 5).* Each vehicle serves a single origin throughout the horizon. Relaxation to inter-origin transfers yields a generalized decomposition with recourse penalties, and was identified as future work in the DJO remark.
- *Assumption A2: Demand-distance coupling constant κ.* The scalar $\kappa \in [0, 1]$ summarises network-wide correlation between demand and travel-condition degradation. Arc-specific or time-varying $\kappa_{ij}(t)$ is a natural generalisation that would preserve the ζ^ω mechanism while adding spatial heterogeneity.
- *Assumption A3: Scenario ensemble Ξ as ground truth for $\sigma(\lambda)$.* The anchored deviation (10) inherits the quality of the ensemble; distributionally robust extensions would replace the max over Ξ with a max over an ambiguity set of distributions.
- *Assumption A4: Single-valued neutrosophic specialisation.* The constructive constraint $T_\lambda + I_\lambda = 1$ is a domain-specific choice, not a structural requirement of the framework. Interval-valued or fully independent triplets are admissible extensions.

Priorities for the research agenda include, in order, a public benchmark suite spanning at least five georeferenced evacuation instances of varying size to convert single-instance dominance into statistical dominance (P1), a formal analysis of $\partial\Gamma^n/\partial\beta$ and $\partial\mathrm{PoR}/\partial\Gamma^n$ to replace the current numerical sensitivity of Table 5 with analytical guarantees (P2), relaxation of Assumption A1 to a generalized decomposition with inter-origin recourse (P3), and distributionally robust extension to eliminate Assumption A3 (P4).

11.4. Concluding remarks. This paper developed a unified optimization framework for hurricane evacuation under compound uncertainty. An integrated critical review of 337 publications identified six persistent gaps motivating the architecture $\mathcal{E}^{\mathrm{Robust}} \equiv \mathrm{PII}^R(\mathrm{PI}^S(\mathrm{PII}^D(\mathrm{PI}^D)))$. Scenario-anchored uncertainty sets produce 13.4% tighter protection than generic box sets, neutrosophic budget calibration achieves 53.2% probability guarantees at a 2.8% PoR versus 24.8% for static approaches (Table 5), and DJO decomposition reduces binary variables by 94%, bringing the pipeline to under four minutes. Direct benchmarking against deterministic, stochastic, robust, and hybrid paradigms on identical input data (Table 4) confirms Pareto dominance on the joint PoR–guarantee–scalability frontier, and Section 10 extracts three insights, i.e., conservatism-coverage inversion, actionable indeterminacy variance, and complexity-closure bottleneck shift, that are behaviours of the coupled system absent from any component in isolation. These results suggest that hurricane evacuation optimization has matured sufficiently to warrant unified architectures with formal guarantees.

Acknowledgment

This work is part of the doctoral thesis is it *Robust Optimization Models for Meteorological Emergency Management under Uncertainty* submitted at Universidad de La Habana, Cuba. The authors acknowledge the Research Department and Computer Science Program at Universidad Politécnica Estatal del Carchi, Ecuador, and the Faculty of Economic Sciences at Universidad de

Cienfuegos, Cuba, for their support and collaboration. Claude AI (Anthropic) was used for language editing and LaTeX formatting assistance.

References

Adacher, L. and Gemma, A. (2017). A robust algorithm to solve the signal setting problem considering different traffic assignment approaches, *International Journal of Applied Mathematics and Computer Science* **27**(4): 815–826, DOI: 10.1515/amcs-2017-0057.

Ahmad, F. (2022). Interactive neutrosophic optimization technique for multiobjective programming problems: An application to pharmaceutical supply chain management, *Annals of Operations Research* **311**(7): 551–585, DOI: 10.1007/s10479-021-03997-2.

Ahuja, R.K., Magnanti, T.L. and Orlin, J.B. (1993). *Network Flows: Theory, Algorithms, and Applications*, Prentice Hall, Englewood Cliffs.

Bayram, V. and Yaman, H. (2018). Shelter location and evacuation route assignment under uncertainty: A Benders decomposition approach, *Transportation Science* **52**(2): 416–436, DOI: 10.1287/trsc.2017.0762.

Bayram, V. and Yaman, H. (2024). A joint demand and supply management approach to large scale urban evacuation planning: Evacuate or shelter-in-place, staging and dynamic resource allocation, *European Journal of Operational Research* **313**(1): 171–191, DOI: 10.1016/j.ejor.2023.07.044.

Ben-Tal, A. and Nemirovski, A. (2002). Robust optimization—Methodology and applications, *Mathematical Programming* **92**(3): 453–480, DOI: 10.1007/s101070100286.

Bertsimas, D., Brown, D.B. and Caramanis, C. (2011). Theory and applications of robust optimization, *SIAM Review* **53**(3): 464–501, DOI: 10.1137/080734510.

Bertsimas, D. and Sim, M. (2004). The price of robustness, *Operations Research* **52**(1): 35–53, DOI: 10.1287/opre.1030.0065.

Bertsimas, D., Sim, M. and Zhang, M. (2019). Adaptive distributionally robust optimization, *Management Science* **65**(2): 604–618, DOI: 10.1287/mnsc.2017.2952.

Chang, Y., Song, Y. and Eksioglu, B. (2022). A stochastic look-ahead approach for hurricane relief logistics operations planning under uncertainty, *Annals of Operations Research* **319**(7): 1231–1263, DOI: 10.1007/s10479-021-04098-0.

Conforti, M., Cornuéjols, G. and Zambelli, G. (2014). *Integer Programming*, Springer, Cham.

Davidson, R.A. and Rivera, F.I. (2020). Integrating scenario ensemble methods with optimization for hurricane evacuation planning, *Natural Hazards Review* **21**(4): 04020042.

Eriskin, L. and Karatas, M. (2024). Applying robust optimization to the shelter location-allocation problem: A case study for Istanbul, *Annals of Operations Research* **339**(7): 1589–1635, DOI: 10.1007/s10479-022-04530-z.

Esposito Amideo, A., Scaparra, M.P. and Kotiadis, K. (2019). Optimising shelter location and evacuation routing operations: The critical issues, *European Journal of Operational Research* **279**(2): 279–295, DOI: 10.1016/j.ejor.2019.05.049.

Fisher, M.L., Jaikumar, R. and Van Wassenhove, L.N. (1986). A multiplier adjustment method for the generalized assignment problem, *Management Science* **32**(9): 1095–1103.

Garey, M.R. and Johnson, D.S. (1979). *Computers and Intractability: A Guide to the Theory of NP-Completeness*, W.H. Freeman, New York.

Ge, Y., Klenow, D.J., Savitt, A. and Murray-Tuite, P. (2023). Understanding the decision-making process for hurricane evacuation orders: A case study of Florida county emergency managers, *Sustainability* **15**(24): 16666, DOI: 10.3390/su152416666.

Goerigk, M., Deghdak, K. and T'Kindt, V. (2015). A two-stage robustness approach to evacuation planning with buses, *Transportation Research B: Methodological* **78**(7): 66–82, DOI: 10.1016/j.trb.2015.04.008.

Goerigk, M. and Grün, B. (2014). A robust bus evacuation model with delayed scenario information, *OR Spectrum* **36**(4): 923–948, DOI: 10.1007/s00291-014-0365-8.

Ismail, M.M., Kamal, E. and Saber, S. (2025). Neutrosophic probability and measure for resilient supply chain management under deep uncertainty, *Neutrosophic Sets and Systems* **94**: 1–15.

Kamyabniya, A., Sauré, A., Salman, F.S., Bénichou, N. and Patrick, J. (2024). Optimization models for disaster response operations: A literature review, *OR Spectrum* **46**(3): 737–783, DOI: 10.1007/s00291-024-00756-2.

Khalili, S.M. and Sahebjamnia, N. (2024). A stochastic-robust optimization approach for designing a humanitarian logistics network, *Socio-Economic Planning Sciences* **93**: 101862.

Lin, F., Fang, X. and Gao, Z. (2022). Distributionally robust optimization: A review on theory and applications, *Numerical Algebra, Control and Optimization* **12**(1): 159–212, DOI: 10.3934/naco.2021057.

Lorca, A. and Sun, X. A. (2016). Multistage adaptive robust optimization for the unit commitment problem, *Operations Research* **64**(1): 32–51, DOI: 10.1287/opre.2015.1439.

Meng, F., Yao, Y., Wang, Z., Peng, S., Xu, D. and Song, T. (2023). Probabilistic forecasting of tropical cyclones intensity using machine learning model, *Environmental Research Letters* **18**(4): 044042, DOI: 10.1088/1748-9326/acc8eb.

Pouraliakbari-Mamaghani, M., Saif, A. and Kamal, N. (2023). Reliable design of a congested disaster relief network: A two-stage stochastic-robust optimization approach, *Socio-Economic Planning Sciences* **86**(7): 101498, DOI: 10.1016/j.seps.2022.101498.

Shahbazi, A., Aghaei, J., Pirouzi, S., Shafie-khah, M. and Catalão, J.P.S. (2021). Hybrid stochastic/robust

optimization model for resilient architecture of distribution networks against extreme weather conditions, *International Journal of Electrical Power & Energy Systems* **126**(2): 106576, DOI: 10.1016/j.ijepes.2020.106576.

Siddig, M. and Song, Y. (2025). Multistage stochastic programming with a random number of stages: Applications in hurricane disaster relief logistics planning, *European Journal of Operational Research* **321**(3): 925–941, DOI: 10.1016/j.ejor.2024.09.043.

Smarandache, F. (2005). *A Unifying Field in Logics: Neutrosophic Logic. Neutrosophy, Neutrosophic Set, Neutrosophic Probability and Statistics*, American Research Press, Rehoboth.

Sotelo-Salas, C., Véliz, F., Riquelme-Rodríguez, J.P., González, V. and Pavez, O. (2024). A multi-objective optimization model for planning emergency evacuation: A case study for the 2010 Chilean tsunami in Valparaíso, *Socio-Economic Planning Sciences* **93**(7): 101909.

Soyster, A.L. (1973). Convex programming with set-inclusive constraints and applications to inexact linear programming, *Operations Research* **21**(5): 1154–1157.

Steinmetz-Weiss, C. and Sanderson, D. (2024). A systematic literature review on transit-based evacuation planning in emergency logistics management: Optimisation, *Journal of Humanitarian Logistics and Supply Chain Management* **14**(3): 301–325.

Turkeš, R., Cuervo, D.P. and Sörensen, K. (2019). Pre-positioning of emergency supplies: Does putting a price on human life help to save lives?, *Annals of Operations Research* **283**(1): 865–895, DOI: 10.1007/s10479-017-2702-1.

Wang, X.J. and Paul, J.A. (2020). Robust optimization for hurricane preparedness, *International Journal of Production Economics* **221**: 107465, DOI: 10.1016/j.ijpe.2019.08.008.

Yao, T., Mandala, S.R. and Chung, B.D. (2009). Evacuation transportation planning under uncertainty: A robust optimization approach, *Networks and Spatial Economics* **9**(2): 171–189, DOI: 10.1007/s11067-009-9103-1.

Zhu, Y., Ozbay, K., Xie, K., Yang, H. and Morgul, E.F. (2021). Network modeling of hurricane evacuation using data-driven demand and incident-induced capacity loss models, *Journal of Advanced Transportation* **2021**(7): 6620254, DOI: 10.1155/2021/6620254.

Yasmany Fernández-Fernández received his MSc in applied mathematics from Universidad de Cienfuegos, Cuba, in 2014. He is currently a doctoral candidate in mathematical sciences at Universidad de La Habana, Cuba (part timer–researcher). His research focuses on stochastic and robust optimization for humanitarian logistics, with applications to hurricane evacuation planning.

Sira M. Allende Alonso is an applied mathematician at the Faculty of Mathematics and Computer Science, Universidad de La Habana, Cuba, with expertise in operations research, vehicle routing, and statistical modelling. Her research interests span optimization of traffic control, fuzzy shortest-path problems in communication networks, sampling and imputation methods, and predictive analytics for large data sets. She has coauthored papers in regional and international journals and has been distinguished with the Sofia Kovalevskaya Award, recognizing her scientific trajectory and contributions to applied mathematics.

Ridelio Miranda Pérez is a researcher at Universidad de Cienfuegos "Carlos Rafael Rodríguez", Cuba, where he leads and coordinates work in data science and artificial intelligence. His research focuses on computational optimization, data-driven decision support, and the application of AI and machine learning to engineering and logistics problems. He has supervised numerous master's theses and several ongoing doctoral projects, and has received national and provincial awards for his scientific output, including recognitions from the Cuban Academy of Sciences.

Gemayqzel Bouza Allende is a full professor of mathematics at Universidad de La Habana, Cuba, and a former Alexander von Humboldt Research Fellow in Germany. Her research areas include continuous optimization, descent and variable-order methods, and stability analysis for nonlinear problems, with contributions published in leading optimization and applied mathematics journals. She has extensive experience supervising graduate students and leading collaborative projects in optimization with applications in engineering, operations research, and decision support systems.

Received: 17 February 2026
Revised: 21 April 2026
Re-revised: 22 April 2026
Accepted: 15 May 2026