

Context Matters: Optimizing Crisis Response Through Contextual Stochastic Decision-Making

Nafe Moradkhani*

KEDGE Business School, France
nafe.moradkhani@kedgebs.com

Walid Klibi

KEDGE Business School, France
walid.klibi@kedgebs.com

Frederick Benaben

IMT Mines Albi, France
frederick.benaben@mines-albi.fr

ABSTRACT

The growing complexity and uncertainty of crisis response demand innovative approaches that integrate predictive intelligence with decision-making under dynamic conditions. Contextual optimization, sitting at the confluence of machine learning and operations research, offers a powerful framework for aligning data-driven predictions with real-world objectives. This paradigm excels in addressing the dual challenges of accurate demand forecasting and cost-efficient resource allocation, crucial in disaster response scenarios. By iteratively refining predictions and optimizing downstream outcomes, contextual optimization bridges the gap between abstract models and actionable strategies. This study evaluates and demonstrates the effectiveness of contextual optimization frameworks using a simulation-based methodology in earthquake response scenarios, highlighting their potential to significantly reduce unmet demand and total operational costs. The results underscore its pivotal role as a cornerstone for the next generation of efficient and responsive disaster management systems, emphasizing its adaptability to evolving conditions and ability to prioritize critical outcomes.

Keywords

Crisis Management, Stochastic Optimization, Decision-Making, Risk Mitigation, Disaster Operations

INTRODUCTION

In an era marked by frequent and diverse crises—ranging from natural disasters to pandemics and supply chain disruptions—decision-makers face unprecedented challenges. Crisis management is evolving beyond traditional reactive measures, which often rely on post-event responses and recovery strategies (Boin & McConnell, 2007; Boin et al., 2016; McConnell, 2011). Instead, it demands proactive and dynamic decision-making that integrates predictive analytics and real-time data to address the complexities of real-world situations. Despite advancements in technology and data analytics, one fundamental question remains: how can decision-makers tailor their strategies to the unique and evolving context of each crisis to achieve better outcomes?

Traditional crisis management systems often rely on static or generalized frameworks, designed to handle crises under broad assumptions (Alshoubaki, 2020; Hale et al., 2005). While these approaches provide robustness, they frequently fail to adapt to the variability and complexity of real-world scenarios. For example, an earthquake in a densely populated urban center presents vastly different challenges than one in a sparsely populated rural area. Similarly, resource constraints, demographic factors, and the dynamic nature of crises require strategies that go beyond a “one-size-fits-all” solution.

This paper focuses on **contextual optimization**—a decision-making paradigm that integrates context-specific information into optimization frameworks to enhance the effectiveness of crisis management. Linear programming

*Corresponding author

(LP) assumes a fixed and deterministic set of constraints and objective functions, limiting its adaptability to dynamic conditions (Dantzig, 1998). Stochastic programming incorporates randomness but relies on fixed probabilistic distributions of uncertainty, making it less effective in capturing context-specific dependencies (Birge & Louveaux, 2011). Robust optimization (RO) addresses uncertainty by optimizing for the worst-case scenario within predefined uncertainty sets, which may not reflect real-world variability (Ben-Tal, 2009). Similarly, dynamic programming (DP) is powerful for sequential decision-making but requires predefined state transition probabilities or deterministic rules, limiting its applicability to highly dynamic environments (Bellman & Dreyfus, 2015).

In contrast, contextual optimization dynamically incorporates side information (covariates) to model the dependencies between crisis variables and their context (Ban & Rudin, 2019; Bertsekas, 2015). By leveraging contextual data—such as demographic information, environmental conditions, and real-time crisis parameters—decision-makers can adapt strategies dynamically, resulting in more effective and efficient resource allocation, risk mitigation, and response planning (Mišić & Perakis, 2020; Sadana et al., 2024).

The strength of contextual optimization lies in its ability to bridge predictive modeling and prescriptive decision-making. By integrating context-aware insights into optimization frameworks, decision-makers can address uncertainties and complexities that static models fail to capture (Sadana et al., 2024). For instance, in the case of earthquake response, contextual optimization can adaptively allocate resources based on population density, infrastructure vulnerabilities, and predicted aftershock risks, ensuring a more targeted and impactful response.

This paper leverages three foundational frameworks of contextual optimization—Decision Rule Optimization (DRO), Sequential Learning and Optimization (SLO), and Integrated Learning and Optimization (ILO)—to address the complexities inherent in crisis management decision-making. These frameworks are specifically designed to tackle different dimensions of uncertainty. DRO prioritizes robustness by focusing on optimizing decisions for known covariates, ensuring reliable outcomes even under uncertain conditions (Ben-Tal, 2009). SLO enhances decision-making by employing iterative learning, enabling strategies to evolve dynamically based on real-time feedback and updates, as demonstrated in applications such as supply chain management and healthcare operations (Mišić & Perakis, 2020). ILO advances these approaches further by combining learning and optimization into a single integrated process, directly aligning predictions with downstream cost minimization objectives (Donti et al., 2017; Qi et al., 2021).

The comprehensive capabilities of these frameworks allow crisis management systems to adapt dynamically and contextually, addressing the unique demands of evolving scenarios with precision and efficiency. Recent innovations in ILO illustrate its transformative potential, showcasing methods such as task-based end-to-end learning and imitation-based modeling to refine decisions and minimize regret (Elmachtoub & Grigas, 2022; Kong et al., 2022). Using a case study on earthquake response, this paper highlights how contextual optimization effectively integrates predictive analytics and prescriptive decision-making, enabling actionable insights and impactful solutions for real-world crises.

The remainder of this paper explores how contextual optimization transforms crisis management. Section 2 lays the theoretical foundation for this approach, while Section 3 delves into the three decision-making frameworks that drive it—DRO, SLO, and ILO. Section 4 evaluates their effectiveness through a comprehensive earthquake case study, and Section 5 concludes by discussing the broader implications for future crisis response systems.

THEORETICAL FOUNDATIONS OF CONTEXTUAL OPTIMIZATION IN CRISIS MANAGEMENT

In this section, we establish the theoretical underpinnings necessary to understand the integration of contextual optimization into crisis management systems. The analysis begins with a formalized mathematical model for crisis situations, capturing the inherent uncertainty, dynamic evolution, and complexity of these environments. By modeling the decision-making landscape of crises, we create a robust foundation for employing contextual optimization frameworks. These frameworks, as we will demonstrate, enhance decision-making by tailoring strategies to the unique and evolving contexts of crises, enabling more effective, adaptive, and resource-efficient responses.

Mathematical Model for Crisis Situations

A crisis situation, such as a natural disaster or large-scale emergency, is characterized by uncertainty, interdependent variables, and rapidly changing conditions (Altay & Green III, 2006; Barthe-Delanoë et al., 2018; Comfort et al., 2004; Domdouzis et al., 2017). These elements complicate decision-making processes, necessitating models that effectively integrate uncertainty, constraints, and contextual information (Birge & Louveaux, 2011; Friedrich et al., 2000; Sheu, 2007). In mathematical terms, a crisis can be represented by three key components:

- **y**: uncertain parameters (e.g., severity, duration, or affected population) (Fiedrich et al., 2000),
- **z**: decision variables (e.g., resource allocation, deployment of personnel) (Yi & Kumar, 2007),
- **x**: contextual information (e.g., geographic, demographic, and real-time crisis indicators) (Rogova, 2016).

The decision-making process aims to minimize expected costs while accounting for uncertainties and constraints. Formally, this can be expressed as:

$$\min_{\mathbf{z} \in \mathcal{Z}} \mathbb{E}_{\mathbf{y} \sim P(\mathbf{y}|\mathbf{x})} [c(\mathbf{z}, \mathbf{y})] \quad (1)$$

where:

$\mathbf{z} \in \mathcal{Z}$	represents the decision vector constrained to a feasible set $\mathcal{Z}$,
$\mathbf{y} \in \mathcal{Y}$	denotes uncertain crisis parameters with support $\mathcal{Y}$,
$\mathbf{x} \in \mathcal{X}$	comprises contextual covariates providing side information,
$c(\mathbf{z}, \mathbf{y})$	is the cost function, representing the consequences of decisions,
$P(\mathbf{y} \mathbf{x})$	is the conditional probability distribution of $\mathbf{y}$ given the context $\mathbf{x}$.

This formulation aligns with established methodologies in decision-making under uncertainty, providing a robust framework for crisis management. This model highlights the pivotal role of $\mathbf{x}$, or contextual data, in informing decision-making by capturing the dependencies between $\mathbf{y}$ and the surrounding context (Ansell & Boin, 2019; Ban & Rudin, 2019; Kamyabniya et al., 2024; Paulus et al., 2024; Zamanifar & Hartmann, 2020).

Incorporating Contextual Data. Contextual data $\mathbf{x}$ is vital in differentiating between crises of similar types but different natures. Examples include:

- **Demographic Context:** Population density, socioeconomic indicators, age distribution, and access to resources (Paulus et al., 2024; Sheu, 2007).
- **Geographic and Environmental Context:** Infrastructure resilience, geographic spread, weather conditions, and proximity to resources (Galindo & Batta, 2013; Zamanifar & Hartmann, 2020).
- **Crisis-Specific Parameters:** Severity indices (e.g., earthquake magnitude), resource constraints, and historical patterns of crisis progression (Birge & Louveaux, 2011; Kamyabniya et al., 2024).

Optimization Objectives and Constraints. The optimization objective $\mathbb{E}_{\mathbf{y} \sim P(\mathbf{y}|\mathbf{x})} [c(\mathbf{z}, \mathbf{y})]$ minimizes expected costs, subject to operational constraints:

$$g_i(\mathbf{z}, \mathbf{y}) \leq 0, \quad \forall i \in \{1, \dots, m\}. \quad (2)$$

The constraints $g_i(\mathbf{z}, \mathbf{y})$ represent specific restrictions on decision variables ($\mathbf{z}$) and uncertain parameters ($\mathbf{y}$). These constraints are crucial for ensuring the feasibility and effectiveness of decisions in crisis management scenarios. Examples include:

- **Resource Constraints:** Ensure that allocated resources ($\mathbf{z}$) do not exceed availability ($g_i(\mathbf{z}, \mathbf{y}) = \mathbf{z}_i - \text{AvailableResources}_i \leq 0$) (Kamyabniya et al., 2024; Sheu, 2007).
- **Time Constraints:** Guarantee that actions such as evacuation or medical supply delivery are completed within designated deadlines ($g_i(\mathbf{z}, \mathbf{y}) = \text{TimeRequired}_i - \text{Deadline}_i \leq 0$) (Galindo & Batta, 2013; Zamanifar & Hartmann, 2020).
- **Equity Constraints:** Promote fair resource distribution by addressing proportionality across affected populations ($g_i(\mathbf{z}, \mathbf{y}) = \text{FairDistributionIndex}_i - \text{Threshold} \leq 0$) (Ban & Rudin, 2019; Paulus et al., 2024).
- **Dynamic Constraints:** Account for real-time updates to resource availability or hazard impacts ($g_i(\mathbf{z}, \mathbf{y}) = \text{UpdatedHazardImpact}_i - \mathbf{z}_i \leq 0$) (Birge & Louveaux, 2011; Boin et al., 2016).

This detailed formulation ensures that constraints are not only mathematically consistent but also practically aligned with the unique requirements of real-world crisis management systems.

Dynamic Nature of Crises. Crisis situations are inherently dynamic, requiring models that evolve over time. Changes in $\mathbf{x}$ (e.g., weather updates), $\mathbf{y}$ (e.g., aftershocks in an earthquake), or the probability distribution $P(\mathbf{y}|\mathbf{x})$ demand continuous reevaluation of decisions. Adaptive decision-making frameworks are essential to ensure that strategies remain effective as crises unfold (Birge & Louveaux, 2011; Kamyabniya et al., 2024; Sheu, 2007).

Recent approaches integrate real-time data assimilation and sequential decision-making frameworks to address this dynamic nature. For instance:

- **Real-Time Updates:** Models incorporate evolving data inputs, such as updated hazard maps or live resource tracking (Gowda et al., 2024).
- **Feedback Mechanisms:** Iterative adjustments in response to stakeholder input enhance decision robustness (Sediri et al., 2013).
- **Scenario-Based Planning:** Leveraging simulations to evaluate strategies under varying hypothetical developments of a crisis (Xue et al., 2022).

Toward Contextual Optimization. While the mathematical model of crises (1) provides a general structure for decision-making under uncertainty, it lacks mechanisms for dynamically incorporating contextual information. Crisis management research has leveraged techniques such as robust optimization, stochastic programming, and scenario analysis to incorporate contextual factors like geographic and demographic data into static models (Ben-Tal, 2009; Birge & Louveaux, 2011; Galindo & Batta, 2013). However, these approaches often treat context ($\mathbf{x}$) as fixed or exogenous, failing to capture its evolving dependencies with uncertain parameters ($\mathbf{y}$) or to adapt to real-time changes. Contextual optimization overcomes these limitations by embedding $\mathbf{x}$ directly into the decision-making process, enabling adaptive and data-driven strategies that align with the dynamic and evolving nature of crises. This paradigm forms the foundation for the advanced frameworks discussed in subsequent sections, which dynamically tailor decisions to the unique characteristics of each crisis.

CONTEXTUAL OPTIMIZATION FRAMEWORKS IN CRISIS MANAGEMENT

Contextual optimization frameworks enhance decision-making under uncertainty by dynamically linking contextual covariates to uncertain parameters. In crisis management, where traditional models with fixed assumptions lack real-time adaptability, these frameworks provide a critical advantage. Sadana et al. (Sadana et al., 2024) highlight the capability of contextual optimization to incorporate dynamic contextual information, addressing the unpredictable nature of crises. This section focuses on three principal frameworks—Decision Rule Optimization (DRO), Sequential Learning and Optimization (SLO), and Integrated Learning and Optimization (ILO)—starting with a detailed exploration of DRO, its theoretical foundations, and integration into the crisis model from equation 1.

Decision Rule Optimization (DRO) in Crisis Management

Decision Rule Optimization (DRO) employs parameterized mappings, known as decision rules, to connect contextual information ($\mathbf{x}$) directly to actions ($\mathbf{z}$). These mappings, denoted as $\pi_\theta(\mathbf{x})$, are optimized to minimize expected costs, where θ represents the parameters defining the rule. DRO enables the translation of contextual insights into actionable decisions, ensuring robustness and context sensitivity.

Empirical Risk Minimization (ERM) underpins this optimization process by approximating the expected cost using empirical data (Roy et al., 2021; Vapnik, 1991). Specifically, the optimization problem is formulated as:

$$\theta^* = \underset{\theta}{\operatorname{argmin}} \mathbb{E}_{\hat{P}_N} [c(\pi_\theta(\mathbf{x}), \mathbf{y})], \quad (3)$$

where:

- $\hat{P}_N$ is the empirical distribution derived from historical data $\mathcal{D}_N = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^N$, which serves as an approximation of the true conditional distribution $P(\mathbf{y}|\mathbf{x})$,
- $c(\pi_\theta(\mathbf{x}), \mathbf{y})$ represents the cost function, which quantifies the loss or consequences associated with applying the decision rule $\pi_\theta(\mathbf{x})$ given the uncertain parameters $\mathbf{y}$,
- $\pi_\theta(\mathbf{x})$ is the parameterized decision rule that maps contextual data $\mathbf{x}$ to an actionable decision $\mathbf{z}$.

C

ERM provides a practical solution for optimizing π_θ by leveraging finite samples from historical data, allowing DRO to approximate the true expected cost without requiring full knowledge of the underlying distribution. In this framework, $\pi_\theta(\mathbf{x})$ acts as a surrogate for the decision variable $\mathbf{z}$, establishing a direct mapping between contextual data and actions. Effectively, $\mathbf{z} = \pi_\theta^*(\mathbf{x})$ represents the optimal decision that accounts for both the observed context ($\mathbf{x}$) and the calibrated decision rule (θ^*).

DRO integrates seamlessly into the previously introduced crisis model:

$$\min_{\mathbf{z} \in \mathcal{Z}} \mathbb{E}_{\mathbf{y} \sim P(\mathbf{y}|\mathbf{x})} [c(\mathbf{z}, \mathbf{y})], \quad (4)$$

where parameterizing $\mathbf{z}$ as $\pi_\theta(\mathbf{x})$ transforms the problem into equation 3. This connection demonstrates how DRO bridges contextual information and uncertain parameters, enabling tailored decisions for crisis management. For instance, in earthquake scenarios, contextual data ($\mathbf{x}$) such as population density and infrastructure resilience are linked to uncertain parameters ($\mathbf{y}$) like aftershock probabilities and resource constraints. DRO translates this data into actionable strategies, such as resource allocations or evacuation plans, that are precisely aligned with the dynamic and uncertain conditions of the crisis.

While DRO offers robust and scalable solutions by leveraging contextual data, it is inherently limited to static optimization. This means that DRO cannot iteratively refine its decisions based on new, incoming data. In rapidly changing crisis scenarios, such as escalating floods or sudden aftershocks following earthquakes, the inability to update predictions and actions reduces DRO's adaptability. To overcome this limitation, Sequential Learning and Optimization (SLO) introduces a predictive component that continuously learns from both historical and real-time data, enhancing its ability to adapt to evolving crises.

Sequential Learning and Optimization (SLO) in Crisis Management

Sequential Learning and Optimization (SLO), widely recognized as the "predict-then-optimize" framework, combines a predictive model that estimates conditional distributions of uncertain parameters ($\mathbf{y}$) based on contextual covariates ($\mathbf{x}$) with an optimization model that prescribes actions ($\mathbf{z}$) to minimize costs. Recent works have established generalization bounds to ensure robust decision-making, even under uncertain or noisy data, making the framework reliable for high-stakes scenarios (El Balghiti et al., 2019). Additionally, methods such as decision trees improve interpretability and scalability for real-time applications, while decision-driven regularization aligns predictions more closely with optimization goals, reducing potential mismatches between the two stages (Elmachtoub et al., 2020; Loke et al., 2022). By dynamically integrating predictive and prescriptive capabilities, SLO provides a robust and adaptive approach for decision-making under uncertainty, particularly in crisis scenarios.

Building on the crisis model introduced earlier (1), where $\mathbf{y}$ encapsulates uncertain parameters (e.g., crisis severity or resource availability), $\mathbf{x}$ represents contextual information (e.g., geographic, demographic, or real-time indicators), and $\mathbf{z}$ corresponds to decision variables (e.g., resource allocation strategies), SLO employs a two-stage process:

1. **Predictive Component:** The predictive model, $f_\theta(\mathbf{x})$, parameterized by θ , leverages machine learning techniques to estimate the conditional distribution $P(\mathbf{y}|\mathbf{x})$. This involves training on historical data using a supervised learning framework to minimize a loss function:

$$\min_{\theta} \rho(f_\theta, \hat{P}_N) + \Omega(\theta), \quad (5)$$

where:

- $\rho(f_\theta, \hat{P}_N) = \mathbb{E}_{\hat{P}_N} [\mathcal{D}(f_\theta(\mathbf{x}), \mathbf{y})]$ measures prediction error using a divergence metric $\mathcal{D}$, such as negative log-likelihood or mean squared error,
- $\Omega(\theta)$ is a regularization term that mitigates overfitting by penalizing model complexity,
- $\hat{P}_N$ is the empirical distribution derived from the historical dataset $\mathcal{D}_N = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^N$, approximating the true conditional distribution $P(\mathbf{y}|\mathbf{x})$.

Machine learning techniques such as regression models, decision trees, or neural networks can be employed depending on the complexity of the data and the relationships between $\mathbf{x}$ and $\mathbf{y}$. The goal is to learn a model $f_\theta(\mathbf{x})$ that generalizes well to unseen data while capturing meaningful patterns in the historical data. This predictive stage ensures the inputs for the optimization step reflect the latest contextual and historical insights.

2. **Optimization Component:** The optimization component uses the predictive model's output $f_\theta(\mathbf{x})$ to generate context-specific decisions. This step solves:

$$\mathbf{z}^*(\mathbf{x}, f_\theta) = \underset{\mathbf{z} \in \mathcal{Z}}{\operatorname{argmin}} \mathbb{E}_{f_\theta(\mathbf{x})} [c(\mathbf{z}, \mathbf{y})], \quad (6)$$

where:

- $\mathbf{z}^*(\mathbf{x}, f_\theta)$ represents the optimal decision that minimizes the expected cost, conditioned on the predicted distribution $f_\theta(\mathbf{x})$,
- $c(\mathbf{z}, \mathbf{y})$ is the cost function quantifying the consequences of decisions $\mathbf{z}$ under uncertain outcomes $\mathbf{y}$.

This component transforms predictions into actionable decisions, ensuring alignment with the dynamic requirements of crises. For instance, in flood management, the predictive model may estimate flood severity ($\mathbf{y}$) based on real-time rainfall data ($\mathbf{x}$), while the optimization component determines the optimal deployment of resources ($\mathbf{z}$).

SLO integrates seamlessly into the crisis model by coupling predictive and optimization components. The predictive function $f_\theta(\mathbf{x})$ approximates the conditional distribution $P(\mathbf{y}|\mathbf{x})$, while the optimization component uses this distribution to prescribe context-sensitive actions. For instance, during a flood scenario, $\mathbf{x}$ may include real-time rainfall data, and $\mathbf{y}$ could represent flood severity or affected populations. The SLO framework predicts the conditional distribution of flood impacts and optimizes resource allocation ($\mathbf{z}$), such as deploying emergency shelters or allocating sandbags, ensuring adaptive and effective responses.

While SLO bridges prediction and decision-making, its reliance on separate training phases for predictive modeling and optimization introduces potential challenges. Misalignment between the predictive model and optimization objectives can lead to suboptimal decisions, especially when prediction errors propagate to the decision stage (see section 2.2 of (Sadana et al., 2024)). Furthermore, SLO's dual-stage framework depends on high-quality training data and can become computationally expensive in real-time crisis scenarios.

Despite these challenges, SLO remains a powerful framework for adapting to dynamic and uncertain environments, enabling tailored, context-aware strategies in crisis management. To address its limitations, Integrated Learning and Optimization (ILO) aligns prediction and optimization within a unified framework, directly optimizing decisions to minimize costs and improve efficiency in evolving crises.

Integrated Learning and Optimization (ILO) in Crisis Management

Integrated Learning and Optimization (ILO) builds on SLO by unifying the predictive and optimization components into a single iterative framework. Rather than independently training a predictive model and subsequently using its output for decision-making, ILO aligns the learning process directly with the optimization objective. This end-to-end approach ensures that the model focuses on features and patterns most relevant to the downstream task, enabling more effective and context-sensitive decision-making in crisis management scenarios. By addressing the disconnect between prediction and optimization, ILO provides a cohesive framework for tackling complex and dynamic crises.

Unlike SLO, which minimizes estimation error during the predictive stage and separately solves the optimization problem using its output, ILO directly learns a predictive function that minimizes the overall cost of decisions. By integrating prediction and optimization into a single learning process, ILO ensures that the predictive model is optimized with the final objective in mind, eliminating misalignments that can arise when predictions are treated as intermediate outputs. This makes ILO particularly advantageous in crises, where the relationships between contextual covariates ($\mathbf{x}$), uncertain parameters ($\mathbf{y}$), and decision variables ($\mathbf{z}$) are highly complex and require decisions that dynamically adapt to evolving conditions.

Mathematically, the ILO framework adapts the crisis model (1) by learning a predictive function $f_\theta(\mathbf{x})$, parameterized by θ , that minimizes the total cost over the empirical distribution $\hat{P}_N$. The optimization problem is formulated as:

$$\min_{\theta} \mathbb{E}_{\hat{P}_N} [c(\mathbf{z}^*(f_\theta(\mathbf{x})), \mathbf{y})], \quad (7)$$

where:

$$\mathbf{z}^*(f_\theta(\mathbf{x})) = \underset{\mathbf{z} \in \mathcal{Z}}{\operatorname{argmin}} c(\mathbf{z}, f_\theta(\mathbf{x})).$$

This formulation highlights the end-to-end nature of ILO, where the predictive model is iteratively updated to minimize the actual cost of decisions *rather than merely improving the accuracy of predictions*.

To illustrate the advantages of ILO, consider the flood response example introduced in the context of SLO. In that scenario, SLO uses rainfall data ($\mathbf{x}$) to predict flood severity and affected populations ($\mathbf{y}$), which are then used to optimize resource allocation ($\mathbf{z}$). While effective, this approach separates prediction and optimization, potentially leading to suboptimal resource allocation due to inaccuracies in the predicted distribution of $\mathbf{y}$.

In contrast, ILO directly aligns the learning process with the resource allocation objective. By iteratively updating the predictive function $f_{\theta}(\mathbf{x})$ based on the cost of decisions, the framework ensures that the predictive model prioritizes features and patterns most relevant to the optimization task. For instance, instead of merely improving flood severity predictions, ILO emphasizes predicting resource needs in high-risk areas, ensuring decisions minimize response costs and maximize impact.

This integration of learning and optimization offers significant advantages over SLO. By aligning the predictive model with the downstream task, ILO reduces the propagation of errors from prediction to decision-making. Its end-to-end nature eliminates the need for a separate estimation error objective, streamlining the learning process. Furthermore, ILO excels in scenarios where the optimization problem involves complex interactions between contextual and uncertain parameters, as it learns these relationships directly without relying on intermediate approximations.

However, ILO introduces additional computational complexity, as the predictive model must be iteratively updated within the optimization loop. This can pose challenges in real-time crisis management, where rapid and efficient decision-making is critical. Techniques to mitigate this complexity, such as model simplification or parallel computation, are essential to enhance ILO's practical applicability in high-stakes scenarios.

By integrating prediction and optimization into a unified framework, ILO represents a significant advancement in contextual optimization for crisis management. Its ability to directly minimize decision costs ensures that actions are both data-driven and rigorously aligned with the overarching objectives of minimizing harm and maximizing efficiency. This makes ILO particularly well-suited for addressing the intricate and dynamic challenges of modern crises, where traditional models often fall short.

EVALUATING CONTEXTUAL OPTIMIZATION FRAMEWORKS IN EARTHQUAKE RESPONSE

Natural disasters present intricate challenges for resource allocation and decision-making, making them a critical domain for evaluating contextual optimization frameworks. Among these, earthquakes are particularly significant due to their sudden onset, lack of early warning systems, and the urgent need for rapid, informed decisions to mitigate widespread devastation. Unlike floods or hurricanes, which often allow for some degree of preparation, earthquakes demand immediate and adaptive responses under severe uncertainty, emphasizing the importance of robust optimization approaches.

This case study uses earthquake scenarios to illustrate the effectiveness of the proposed frameworks. Earthquakes involve a complex interplay of contextual variables—such as population density, infrastructure quality, proximity to the epicenter, and seismic magnitude—that create a highly dynamic and nonlinear environment for resource allocation. The severe downstream consequences of unmet resource demand, including economic losses and human suffering, further underscore the need for optimization approaches that can adapt to evolving conditions and anticipate future needs.

Earthquakes serve as an ideal case study due to their relevance across diverse geographic regions and their importance to decision-makers in both developed and developing contexts. Effective modeling in these scenarios demands capturing the intricate dependencies between contextual factors and downstream impacts, such as the cost of unmet resource demands. This case study provides a robust testbed to evaluate the strengths and limitations of static decision rules, as well as the adaptive capabilities of integrated frameworks like SLO and ILO, highlighting their advantages over traditional DRO approaches.

Modeling Earthquake Dynamics: Contextual and Outcome Variables

The crisis model introduced in Equation 1 establishes the interplay between contextual variables ($\mathbf{x}$), decision rules ($\mathbf{z}$), and downstream outcomes ($\mathbf{y}$). While this model captures critical dependencies and non-linear dynamics, it is an artificial construct designed specifically for simulation purposes.

The artificial nature of the model enables controlled experimentation, isolating decision-making challenges in disaster response without relying on noisy or limited real-world data. Although this reduces direct real-world

applicability, it provides a robust testbed for evaluating the performance of optimization frameworks such as SLO and ILO. Future research should validate these findings with real-world disaster data to enhance their practical relevance and applicability.

Contextual Variables (x)

The contextual variables represent the key factors influencing resource needs and response strategies. These include:

- **Population Density** (x_1): Reflects the number of individuals potentially affected in each region, measured in individuals per square kilometer (individuals/km²). Higher population densities typically demand greater resource allocations.
- **Infrastructure Quality** (x_2): A normalized measure of a region's resilience, represented as a dimensionless value where higher values indicate stronger infrastructure. Regions with better infrastructure generally require fewer resources, as they are less vulnerable to earthquakes.
- **Distance to Epicenter** (x_3): The geographical proximity of a region to the earthquake's epicenter, measured in kilometers (km). Closer regions experience more severe impacts and higher resource needs.
- **Earthquake Magnitude** (x_4): The Richter scale magnitude of the earthquake, which directly scales the resource needs across all regions. Larger magnitudes correspond to greater devastation and increased demand.

Outcome Variables (y)

The outcome variables (y) capture the stochastic and dynamic results of resource allocation and decision-making under uncertainty. They include:

- **Demand** (y_1): The total resource requirements for a region during an earthquake event, modeled as a non-linear and uncertain function of x . This variable is measured in resource units and reflects the overall needs of the affected population.
- **Allocated Resources** (y_2): The actual resources distributed to each region, also measured in resource units. These are determined by decision rules and constrained by resource availability and priorities.
- **Unmet Demand** (y_3): The gap between y_1 and y_2 , representing shortfalls in resource allocation. This variable is also measured in resource units and highlights decision-making inefficiencies or systemic constraints.
- **Total Cost** (y_4): The economic impact of unmet demand, calculated as a monetary value (e.g., dollars). This accounts for penalties related to shortages and stochastic fluctuations in response costs, reflecting the broader consequences of unmet needs.

Non-Linear Resource Needs and Decision Dynamics

To simulate earthquake response and evaluate the proposed optimization frameworks, we define resource needs (y_1) as a function of contextual variables (x_1, x_2, x_3, x_4). This equation is designed to model the complex and interdependent effects of these variables on resource demands while enabling controlled experimentation. Although the equation is artificial, it is grounded in disaster modeling principles and serves as a foundation for generating data under realistic assumptions. The goal is to approximate meaningful relationships between contextual factors and resource needs, providing a robust testbed for the optimization frameworks. The proposed equation for resource needs is:

$$y_1 = 0.5\sqrt{x_1} \cdot x_4^2 - 2.0e^{-x_2} \cdot x_4 + 0.3 \ln(1 + x_3) \cdot x_4 + \epsilon,$$

where ϵ introduces stochastic noise to reflect real-world variability. The resource needs are clipped to non-negative values to ensure feasibility.

The structure of y_1 is designed to reflect the interplay of key contextual factors influencing resource demand during earthquakes:

- **Population Density** (x_1): y_1 increases with population density because denser areas typically require more resources for evacuation, medical aid, and recovery. The square root scaling ($\sqrt{x_1}$) models the diminishing marginal effect of population density, while the positive coefficient reflects the direct correlation with resource needs. This term is further amplified by the quadratic dependence on earthquake magnitude (x_4^2), which scales the intensity of resource requirements based on the severity of the earthquake (Altay & Green III, 2006; Freire & Aubrecht, 2012).

- *Infrastructure Quality* (x_2): y_1 decreases with better infrastructure quality, as resilient buildings and systems reduce the need for resources. The negative exponential term ($-2.0e^{-x_2} \cdot x_4$) reflects how better infrastructure mitigates resource needs, with diminishing returns at higher resilience levels (Galindo & Batta, 2013; Thompson, 2011).
- *Distance to Epicenter* (x_3): y_1 increases logarithmically with distance to the epicenter. While closer regions ($x_3 \rightarrow 0$) experience the most severe impacts, the logarithmic form ($\ln(1 + x_3)$) models the non-linear attenuation of seismic intensity with distance, consistent with seismic attenuation models (Pitilakis et al., 2024; Wald et al., 1999).
- *Earthquake Magnitude* (x_4): As the common scaling factor in all terms, earthquake magnitude (x_4) directly amplifies the effects of population density, infrastructure quality, and distance to the epicenter. This reflects the fact that more severe earthquakes impact broader areas and necessitate significantly greater resources (Freire & Aubrecht, 2012; Pitilakis et al., 2024).
- *Stochastic Noise* (ϵ): The addition of stochastic noise (ϵ) accounts for variability in unmodeled factors, such as localized terrain effects, aftershock impacts, and human response inconsistencies, ensuring the realism of the model (Birge & Louveaux, 2011; Geng & Hou, 2021).

The combination of positive and negative contributions in this model ensures that resource needs (y_1) dynamically balance the amplifying effects of population density and earthquake severity with the mitigating effects of infrastructure resilience and distance. While the equation is artificial, its components are grounded in established disaster modeling principles, offering a controlled environment for evaluating decision optimization frameworks. Although it is not derived from empirical disaster data, its structure was intentionally designed to simulate plausible demand dynamics in the absence of real-world data. Future research could refine this model with real-world data and empirical validation as such data becomes available.

Summary of Variables and Relationships

Table 1 summarizes the contextual and outcome variables, their relationships, and their significance in the crisis model.

Table 1. Summary of Variables in the Earthquake Model

Variable	Description	Significance
x_1 : Population Density	Number of individuals per km ² .	Drives demand proportionally.
x_2 : Infrastructure Quality	Normalized measure of resilience.	Reduces demand for resources.
x_3 : Distance to Epicenter	Proximity to earthquake epicenter.	Correlates with severity of impact.
x_4 : Earthquake Magnitude	Richter scale value.	Scales overall resource needs.
y_1 : Demand	Resource requirements.	Output of the crisis model.
y_2 : Allocated Resources	Resources distributed.	Affected by decision rules.
y_3 : Unmet Demand	Residual demand after allocation.	Indicates response inefficiency.
y_4 : Total Cost	Economic consequence of unmet demand.	Reflects downstream impacts.

Relevance and Challenges

This model underscores the multi-dimensional interplay of contextual variables in shaping disaster response. By incorporating non-linear interactions and stochastic variability, it captures the complexities faced in decision-making under uncertainty, serving as the foundation for evaluating static (DRO) and dynamic (SLO and ILO) optimization frameworks, it provides a controlled yet insightful testbed for assessing their effectiveness in minimizing unmet demand and downstream costs.

Instance Demonstration and Insights

To illustrate the dynamics of the earthquake model, a specific instance is presented in Figure 1. The figure showcases key relationships between contextual and outcome variables. Unique colors consistently represent regions ($R_1, R_2, \dots, R_8$) across all plots for clarity.

The top-left plot captures the stochastic occurrence of earthquakes across regions over time, with each region labeled distinctly. The top-right plot examines the efficiency of resource distribution by comparing demand and allocated resources, revealing how closely resource allocations align with regional demands.

These visualizations highlight the model's capacity to simulate disaster scenarios and provide insights into critical variables affecting response dynamics. They set the stage for evaluating the performance of various optimization frameworks, such as SLO and ILO, in minimizing unmet demand and associated costs compared to static approaches (DRO).




Figure 1. Instance visualizations of earthquake response dynamics. Unique colors are used consistently across plots to distinguish regions ($R_1, R_2, \dots, R_8$).

Contextual Decision Optimization for Earthquake Response: DRO Framework

The Decision Rule Optimization (DRO) framework addresses the challenge of resource allocation under uncertainty by approximating the conditional relationship $P(y|x)$ from the crisis model (Equation 1). By leveraging decision rules, DRO enables predictions of key outcomes (y) based on contextual variables (x) to optimize resource distribution.

Implementation of DRO

To implement the DRO framework, contextual variables—population density (x_1), infrastructure quality (x_2), distance to the epicenter (x_3), and earthquake magnitude (x_4)—were utilized to model $P(y|x)$, which predicts demand (y_1) and guides resource allocation (y_2) to minimize unmet demand (y_3) and total costs (y_4).

The framework involves the following key steps:

1. Decision Rule Models: Two decision rules were employed:

- **Linear Decision Rule:** Assumes a linear relationship between x and y . A linear regression model was trained to predict demand (y_1) based on the contextual variables (x).
- **Kernel Decision Rule:** Employs a Radial Basis Function (RBF) kernel to capture non-linear dependencies in $P(y|x)$. Hyperparameters (α and γ) were tuned using grid search with cross-validation.

2. Resource Allocation and Cost Computation: Using the predicted $P(y|x)$, resources were allocated as:

$$y_2 = \min(y_1, \text{available resources}),$$

ensuring resource constraints were adhered to. The unmet demand (y_3) was computed as:

$$y_3 = \max(0, y_1 - y_2),$$

and total costs (y_4) were calculated using:

$$y_4 = \kappa \cdot y_3 + \xi,$$

where κ is the penalty for unmet demand, and ξ accounts for stochastic cost variations.

Modeling $P(y|x)$

The conditional relationship $P(y|x)$ was modeled using supervised learning:

$$P(y|x) \sim f(x; \theta),$$

where $f(x; \theta)$ represents the decision rule (linear or kernel-based). The model parameters θ were optimized by minimizing the mean squared error (MSE) between observed and predicted demand (y_1):

$$\text{MSE} = \frac{1}{N} \sum_{i=1}^N (y_{1,i} - f(x_i; \theta))^2.$$

For the kernel-based decision rule, grid search over α (regularization) and γ (kernel width) was performed to identify the optimal parameters:

$$\{\alpha, \gamma\} = \arg \min_{\alpha, \gamma} \text{MSE}_{\text{validation}}.$$

Results and Discussion

The DRO framework was evaluated using earthquake simulation data. The Linear Decision Rule achieved a Demand Prediction MSE of 25562.65 and a Total Cost MSE of 133574025.09. After hyperparameter optimization ($\alpha = 0.01$, $\gamma = 0.001$), the RKHS Decision Rule achieved a Demand Prediction MSE of 23760.79 and a Total Cost MSE of 122300074.46. The kernel-based model (RKHS) provided marginal improvements in prediction accuracy and total cost, likely due to its ability to capture non-linear patterns in $P(y|x)$. These results (see Table 2) demonstrate that the DRO framework establishes a robust baseline for contextual optimization, effectively modeling $P(y|x)$ to guide resource allocation while balancing model complexity and predictive accuracy.

Sequential Learning Optimization (SLO)

Building upon the static nature of DRO framework, SLO enhances adaptability by introducing an iterative feedback mechanism. While DRO employs static decision rules to approximate $P(y|x)$ for resource allocation, SLO dynamically refines these rules using feedback from observed outcomes, enabling a more responsive and evolving approach to disaster scenarios.

Components of SLO

The SLO framework comprises two interdependent components:

Predictive Component. Similar to DRO, SLO models the conditional relationship $P(y|x)$ between contextual variables (x) and demand (y_1). However, instead of relying on static rules, SLO iteratively improves the decision rules through feedback. A Random Forest model is used as the predictive engine, leveraging its capacity to model complex, non-linear relationships between inputs and outputs. The hyperparameters, including the number of trees ($n = 100$) and maximum depth ($d = 10$), are carefully tuned to ensure generalizability while capturing detailed patterns in the data.

At each iteration, the model predicts demand ($\hat{y}_1$) for the current testing set. The predicted values are evaluated against observed outcomes, with performance measured using the mean squared error (MSE). This iterative feedback allows the predictive component to progressively refine its approximation of $P(y|x)$.

Optimization Component. Given the predicted demand ($\hat{y}_1$), the optimization component allocates resources to minimize unmet demand (y_3) and associated costs (y_4). At each iteration, resource allocation is calculated as:

$$y_2 = \min(\hat{y}_1, \text{available resources}),$$

ensuring that resource constraints are respected. Unmet demand is computed as:

$$y_3 = \max(0, \hat{y}_1 - y_2),$$

while total costs are determined by:

$$y_4 = \kappa \cdot y_3 + \xi,$$

where κ represents the penalty for unmet demand, and ξ accounts for random fluctuations in downstream costs.

After each iteration, observed outcomes—contextual variables, demand, and costs—are added to the training dataset. This feedback-driven update allows the predictive model to adapt to new information and improve its predictions for future iterations.

Addressing DRO's Limitations

While DRO operates under static decision rules, SLO introduces an iterative feedback loop that continuously updates these rules based on observed outcomes. This dynamic learning process enables SLO to adapt to the evolving nature of disaster scenarios, capturing patterns and uncertainties that static frameworks cannot address effectively.

Limitations and Transition to ILO

Despite its adaptive nature, SLO treats resource allocation and downstream costs (y_4) as separate components, optimizing them sequentially rather than simultaneously. This segmented approach may result in suboptimal outcomes, as the interdependencies between prediction accuracy and economic impacts are not fully accounted for. To overcome these limitations, Integrated Learning Optimization (ILO) integrates the predictive and optimization components into a unified framework, directly minimizing unmet demand and costs in an end-to-end manner. The next section presents the ILO framework and its implementation.

Integrated Learning Optimization (ILO)

Integrated Learning Optimization (ILO) directly integrates prediction and optimization into a unified, end-to-end framework, aligning predictions with downstream objectives during training.

Components of ILO

The ILO framework consists of two interconnected components:

End-to-End Predictive Optimization. Unlike SLO, where predictions (y_1) are optimized in isolation, ILO incorporates resource allocation and cost calculations directly into the loss function. During training, the predictive model minimizes a custom loss function defined as:

$$\mathcal{L}(\theta) = \frac{1}{N} \sum_{i=1}^N \left[\kappa \cdot \max(0, \hat{y}_{1,i} - y_{2,i}) + \xi \right],$$

where $\hat{y}_{1,i}$ represents the predicted demand, $y_{2,i}$ represents allocated resources, κ penalizes unmet demand, and ξ models random cost fluctuations. This design ensures that predictions account for their downstream impact on total costs (y_4).

Dynamic Feedback Mechanism. ILO employs an iterative feedback loop to continuously refine predictions and allocations. After each iteration, the training data is augmented with observed outcomes, including contextual variables (x), demand predictions (y_1), and total costs (y_4). This iterative process updates the predictive model with smoothed feedback:

$$y'_{\text{feedback}} = \alpha y_{\text{true}} + (1 - \alpha) \hat{y}_{\text{pred}},$$

where α is a smoothing parameter balancing observed and predicted values. This mechanism allows the model to adapt dynamically to changing conditions.

Implementation of the ILO Framework

The implementation of the ILO framework is structured around four key components:

Predictive Component: A Random Forest Regressor is employed to model the relationship between contextual variables (x) and demand (y_1). This choice leverages the ensemble structure of Random Forests to capture non-linear patterns and complex dependencies in the data effectively.

Optimization Component: The resource allocation process minimizes the penalty of unmet demand (y_3) and adjusts total costs (y_4) dynamically. This is achieved through an iterative mechanism where:

$$y_3 = \max(0, y_1 - y_2), \quad y_4 = \alpha \cdot (y_3 \cdot (100 + \beta y_3) + y_{\text{true}} \cdot (1 - \beta)),$$

where α is the learning rate for cost feedback, and β is the weight decay for penalizing unmet demand.

Iterative Feedback Mechanism: The framework integrates iterative feedback by retraining the predictive model on an extended dataset after each iteration. This dataset includes updated contextual variables, demand predictions, and dynamically computed total costs, enabling the framework to adapt to evolving disaster contexts.

Hyperparameter Adjustment: The iterative process incorporates hyperparameters (α and β) to balance predictive accuracy and cost minimization. These parameters are tuned to ensure that the model optimally penalizes unmet demand while smoothing cost feedback loops for improved convergence.

Addressing SLO's Limitations

ILO addresses key limitations of SLO by unifying prediction and optimization. This integration reduces inefficiencies arising from treating these processes sequentially and ensures predictions are directly optimized for downstream cost minimization. By incorporating resource allocation into the predictive process, ILO aligns decision-making with real-world objectives. However, the computational complexity of ILO increases due to the combined training and optimization process, requiring careful tuning of hyperparameters such as α and β . These parameters also play a stabilizing role: α governs the smoothing of feedback between observed and predicted values, while β introduces a decaying penalty on unmet demand, both of which help mitigate the impact of dynamic changes in the feasible region. Additionally, by bounding resource allocations and clipping predictions, the model avoids extreme decisions in response to minor input fluctuations. While the current formulation does not enforce explicit robustness constraints, it exhibits practical resilience through these mechanisms. Formalizing these aspects via sensitivity analysis or robustness-enhancing techniques remains a promising direction for future work.

Conclusion of Framework Evaluation

The comparative analysis highlights the evolution and trade-offs across decision frameworks for contextual optimization in crisis management. Table 2 presents the summarized performance metrics, with the best performance highlighted in blue, and all improvements (column *Gain (Demand, Cost)*) calculated relative to the Linear Decision Rule (DRO) as the baseline. These results are computed over all eight simulated regions (R1 to R8) and represent average performance across regions and simulation runs, as the models operate collectively rather than at the individual regional level. While Figure 2 illustrates the iterative improvement in demand prediction and cost reduction for SLO and ILO, the table underscores the strengths and limitations of each framework.

Table 2. Performance Comparison of Decision Frameworks with Improvements

Framework	Demand Prediction MSE	Total Cost MSE	Gain (Demand, Cost)
Linear Decision Rule (DRO)	25562.65	133574025.09	(0.00%, 0.00%)
RKHS Decision Rule (DRO)	23760.79	122300074.46	(7.05%, 8.44%)
SLO	14667.47	65306149.30	(42.62%, 51.11%)
ILO	20111.87	943498.70	(21.32%, 99.29%)

- **Linear vs. RKHS Decision Rule (DRO):** The RKHS Decision Rule outperforms the Linear Decision Rule in both demand prediction accuracy (MSE_{demand}) and total cost (MSE_{cost}). This improvement reflects RKHS's ability to model non-linear dependencies more effectively, leading to better overall performance in balancing prediction and cost minimization compared to the Linear Decision Rule.

- **SLO:** SLO demonstrates substantial improvements in both demand prediction and total cost reduction over DRO frameworks. It achieves the lowest demand prediction error, showcasing the strength of iterative feedback mechanisms in dynamically refining predictions and resource allocations. However, the gap in prediction accuracy between SLO and ILO is marginal, while ILO significantly outperforms SLO in total cost minimization, highlighting its stronger focus on downstream optimization.

- **ILO:** ILO emerges as the most effective framework for total cost minimization, achieving an order-of-magnitude reduction compared to other frameworks. This performance underscores the power of end-to-end optimization, which directly aligns predictive models with downstream objectives. Although ILO incurs a slight increase in demand prediction error compared to SLO, its holistic design ensures that total costs are prioritized in resource allocation strategies.

These findings reflect a progression in framework sophistication:

- **DRO:** Establishes a baseline by separating prediction and optimization.
- **SLO:** Introduces iterative feedback to balance demand prediction and cost reduction.
- **ILO:** Aligns prediction and optimization through an integrated, end-to-end approach, excelling in total cost minimization.

Figure 2 further demonstrates that while both SLO and ILO iteratively improve prediction and cost metrics, ILO's end-to-end optimization design ensures superior downstream performance. In the left panel, the dashed line represents the ideal alignment between predicted and true demand, highlighting how SLO progressively refines its predictions. The right panel showcases ILO's exceptional ability to minimize total costs through integrated prediction and optimization. This analysis is conducted on one disaster response scenario using Random Forest as the predictive model. Other predictive models (e.g., Support Vector Machines, Gradient Boosting, Neural Networks) could also be explored, but such exploration is outside the scope of this study. Ultimately, framework selection should align with the operational goals, whether prioritizing prediction accuracy, total cost reduction, or computational efficiency.




Figure 2. Iterative improvement in Demand Prediction and Total Cost MSE for SLO and ILO. The plots highlight how iterative feedback in SLO and end-to-end optimization in ILO refine both metrics over successive iterations.

CONCLUSIONS AND FUTURE DIRECTIONS

This study emphasizes the transformative role of contextual optimization in advancing crisis response systems. By integrating machine learning and operations research, it enables data-driven decision-making in dynamic and resource-constrained environments. The following key conclusions and future research directions are presented:

Contextual Optimization for Crisis Management

Contextual optimization demonstrates significant potential in enhancing decision-making agility during crises. Its ability to dynamically incorporate contextual data, such as resource availability and disaster impact, allows for more targeted and adaptive resource allocation strategies. While this study focuses on earthquake response, the proposed frameworks are broadly applicable to other dynamic crises, including climate-induced disasters, pandemics, wildfires, and supply chain disruptions. Adapting the model to these contexts would involve redefining contextual variables—for example, infection rates and mobility data for epidemic response, or supply chain delays and logistics bottlenecks for operational crises—and integrating domain-specific constraints into the optimization

process. This flexibility makes contextual optimization a promising foundation for cross-domain crisis management. Furthermore, validating the framework against historical crisis datasets (e.g., past earthquake or epidemic responses) will be a key step toward ensuring its robustness and operational relevance in real-world deployments.

From Static to Adaptive Frameworks: The Power of End-to-End Integration

The progression from DRO to SLO and ILO underscores the shift from static to adaptive decision-making frameworks in crisis response. While DRO serves as a baseline for resource allocation, SLO incorporates iterative feedback to dynamically enhance accuracy, and ILO achieves superior cost minimization by integrating learning and optimization into a unified process. ILO's end-to-end optimization ensures that resource allocations are directly aligned with total cost minimization objectives, offering a holistic approach to large-scale logistical operations in dynamic crisis environments. Future advancements should focus on incorporating real-time data streams, such as IoT sensor networks and satellite imagery, to enhance responsiveness and scalability in rapidly changing disaster scenarios.

Balancing Prediction and Cost Efficiency

The trade-off between prediction accuracy and cost minimization observed in SLO and ILO suggests opportunities for hybrid frameworks. Future research could explore multi-objective optimization models to balance these priorities, particularly for crises where both prediction fidelity and cost efficiency are critical.

Limitations and Practical Considerations

Despite the demonstrated strengths of contextual optimization, several practical limitations warrant discussion. The ILO framework, while effective in aligning predictions with downstream cost minimization, introduces significant computational complexity due to its end-to-end optimization loop and iterative feedback mechanism. This can pose challenges for real-time deployment in rapidly evolving disaster contexts where decisions must be made quickly. To address this, promising strategies include warm-starting with pre-trained models from similar crisis scenarios, leveraging meta-learning to adapt more quickly across tasks, or employing surrogate models and pruning techniques to approximate costly optimization steps with lower-complexity alternatives (Donti et al., 2017).

Another key limitation is the reliance on high-quality, timely contextual data. In real-world crisis settings, data may be delayed, incomplete, or affected by noise due to sensor failure, infrastructure damage, or reporting lags. While the current study assumes idealized data availability for simulation, future implementations should incorporate techniques such as probabilistic imputation, data assimilation, and ensemble learning to manage uncertainty and recover missing inputs. Frameworks based on robust optimization or decision-making under partial observability (e.g., POMDPs or contextual bandits) offer additional robustness in data-scarce environments. Embedding such resilience-enhancing methods into contextual optimization pipelines is a critical direction for improving their real-world applicability and operational reliability.

While the current study focuses on minimizing total cost over time, future work could incorporate explicit stability metrics or penalty terms for decision volatility. Scenario-based analysis across different crisis intensities would also provide insight into how robust and consistent the learned policies remain under varied conditions.

Although ILO offers strong performance in cost minimization, there are important scenarios where simpler frameworks may be more appropriate. For instance, DRO may be preferable in disaster preparedness or compliance-driven contexts where decision logic must remain transparent and auditable, and where regulatory constraints demand consistent, rule-based planning. Similarly, SLO can serve as a practical middle ground in settings where real-time adaptation is valuable but full end-to-end optimization may be too resource-intensive. As such, the selection of a decision framework should consider not only downstream performance but also practical factors like model interpretability, data availability, urgency of response, and deployment feasibility.

Advancing Predictive Modeling in Crisis Contexts

While Random Forest was employed in this study, future work could evaluate advanced machine learning models, including neural networks and gradient boosting, for improved prediction robustness in complex crises. Additionally, integrating explainable AI (XAI) could enhance model transparency, fostering trust among stakeholders in high-stakes decision-making.

REFERENCES

- Alshoubaki, W. (2020). Understanding of crisis management: A theoretical literature review. *International Journal of Management Concepts and Philosophy*, 13(3), 236–246.
- Altay, N., & Green III, W. G. (2006). Or/ms research in disaster operations management. *European journal of operational research*, 175(1), 475–493.
- Ansell, C., & Boin, A. (2019). Taming deep uncertainty: The potential of pragmatist principles for understanding and improving strategic crisis management. *Administration & Society*, 51(7), 1079–1112.
- Ban, G.-Y., & Rudin, C. (2019). The big data newsvendor: Practical insights from machine learning. *Operations Research*, 67(1), 90–108.
- Barthe-Delanoë, A.-M., Montarnal, A., Truptil, S., Benaben, F., & Pingaud, H. (2018). Towards the agility of collaborative workflows through an event driven approach—application to crisis management. *International journal of disaster risk reduction*, 28, 214–224.
- Bellman, R. E., & Dreyfus, S. E. (2015). *Applied dynamic programming* (Vol. 2050). Princeton university press.
- Ben-Tal, A. (2009). Robust optimization. *Princeton University Press google schola*, 2, 35–53.
- Bertsekas, D. (2015). *Convex optimization algorithms*. Athena Scientific.
- Birge, J. R., & Louveaux, F. (2011). *Introduction to stochastic programming*. Springer Science & Business Media.
- Boin, A., & McConnell, A. (2007). Preparing for critical infrastructure breakdowns: The limits of crisis management and the need for resilience. *Journal of contingencies and crisis management*, 15(1), 50–59.
- Boin, A., Stern, E., Sundelius, B., et al. (2016). *The politics of crisis management: Public leadership under pressure*. Cambridge University Press.
- Comfort, L. K., Ko, K., & Zagorecki, A. (2004). Coordination in rapidly evolving disaster response systems: The role of information. *American behavioral scientist*, 48(3), 295–313.
- Dantzig, G. B. (1998). *Linear programming and extensions*. Princeton university press.
- Domdouzis, K., Akhgar, B., Andrews, S., & Day, T. (2017). Identification of critical factors in large crisis decision making processes using computational tools: The case of athena. *International Journal of Strategic Information Technology and Applications (IJSITA)*, 8(2), 11–28.
- Donti, P., Amos, B., & Kolter, J. Z. (2017). Task-based end-to-end model learning in stochastic optimization. *Advances in neural information processing systems*, 30.
- El Balghiti, O., Elmachtoub, A. N., Grigas, P., & Tewari, A. (2019). Generalization bounds in the predict-then-optimize framework. *Advances in neural information processing systems*, 32.
- Elmachtoub, A. N., & Grigas, P. (2022). Smart “predict, then optimize”. *Management Science*, 68(1), 9–26.
- Elmachtoub, A. N., Liang, J. C. N., & McNellis, R. (2020). Decision trees for decision-making under the predict-then-optimize framework. *International conference on machine learning*, 2858–2867.
- Fiedrich, F., Gehbauer, F., & Rickers, U. (2000). Optimized resource allocation for emergency response after earthquake disasters. *Safety science*, 35(1-3), 41–57.
- Freire, S., & Aubrecht, C. (2012). Integrating population dynamics into mapping human exposure to seismic hazard. *Natural Hazards and Earth System Sciences*, 12(11), 3533–3543.
- Galindo, G., & Batta, R. (2013). Review of recent developments in or/ms research in disaster operations management. *European journal of operational research*, 230(2), 201–211.
- Geng, S., & Hou, H. (2021). Demand stratification and prediction of evacuees after earthquakes. *Sustainability*, 13(16), 8837.
- Gowda, V. D., Sharma, A., Prasad, K., Saxena, R., Barua, T., & Mohiuddin, K. (2024). Dynamic disaster management with real-time iot data analysis and response. *2024 International Conference on Automation and Computation (AUTOCOM)*, 142–147.
- Hale, J. E., Dulek, R. E., & Hale, D. P. (2005). Crisis response communication challenges: Building theory from qualitative data. *The Journal of Business Communication* (1973), 42(2), 112–134.
- Kamyabniya, A., Sauré, A., Salman, F. S., Bénichou, N., & Patrick, J. (2024). Optimization models for disaster response operations: A literature review. *OR Spectrum*, 1–47.

- Kong, L., Cui, J., Zhuang, Y., Feng, R., Prakash, B. A., & Zhang, C. (2022). End-to-end stochastic optimization with energy-based model. *Advances in Neural Information Processing Systems*, 35, 11341–11354.
- Loke, G. G., Tang, Q., & Xiao, Y. (2022). Decision-driven regularization: A blended model for predict-then-optimize. Available at SSRN 3623006.
- McConnell, A. (2011). Success? failure? something in-between? a framework for evaluating crisis management. *Policy and Society*, 30(2), 63–76.
- Mišić, V. V., & Perakis, G. (2020). Data analytics in operations management: A review. *Manufacturing & Service Operations Management*, 22(1), 158–169.
- Paulus, D., Fathi, R., Friedrich, F., de Walle, B. V., & Comes, T. (2024). On the interplay of data and cognitive bias in crisis information management: An exploratory study on epidemic response. *Information Systems Frontiers*, 26(2), 391–415.
- Pitilakis, K., Fotopoulou, S., Manakou, M., Karafagka, S., Petridis, C., Pitilakis, D., & Raptakis, D. (2024). Towards seismic risk reduction of critical facilities combining earthquake early warning and structural monitoring: A demonstration study. *Bulletin of Earthquake Engineering*, 1–35.
- Qi, M., Grigas, P., & Shen, Z.-J. M. (2021). Integrated conditional estimation-optimization. *arXiv preprint arXiv:2110.12351*.
- Rogova, G. L. (2016). Crisis management and context. *Fusion Methodologies in Crisis Management: Higher Level Fusion and Decision Making*, 107–119.
- Roy, A., Balasubramanian, K., & Erdogdu, M. A. (2021). On empirical risk minimization with dependent and heavy-tailed data. *Advances in Neural Information Processing Systems*, 34, 8913–8926.
- Sadana, U., Chenreddy, A., Delage, E., Forel, A., Frejinger, E., & Vidal, T. (2024). A survey of contextual optimization methods for decision-making under uncertainty. *European Journal of Operational Research*.
- Sediri, M., Matta, N., Loriette, S., & Hugerot, A. (2013). Decision support by handling experience feedback of crisis situations. *International Conference on Knowledge Management and Information Sharing*, 2, 351–359.
- Sheu, J.-B. (2007). An emergency logistics distribution approach for quick response to urgent relief demand in disasters. *Transportation research part E: logistics and transportation Review*, 43(6), 687–709.
- Thompson, K. (2011). Earthquake risk reduction in buildings and infrastructure program.
- Vapnik, V. (1991). Principles of risk minimization for learning theory. *Advances in neural information processing systems*, 4.
- Wald, D. J., Quitoriano, V., Heaton, T. H., & Kanamori, H. (1999). Relationships between peak ground acceleration, peak ground velocity, and modified mercalli intensity in california. *Earthquake spectra*, 15(3), 557–564.
- Xue, L., Zhang, Q., & Zhong, K. (2022). Building a modern crisis management system: An analysis of the decision-making process. In *Crisis management in china: Challenges of the transition* (pp. 113–137). Springer.
- Yi, W., & Kumar, A. (2007). Ant colony optimization for disaster relief operations. *Transportation Research Part E: Logistics and Transportation Review*, 43(6), 660–672.
- Zamanifar, M., & Hartmann, T. (2020). Optimization-based decision-making models for disaster recovery and reconstruction planning of transportation networks. *Natural Hazards*, 104(1), 1–25.