Optimization Letters
https://doi.org/10.1007/s11590-024-02168-z

ORIGINAL PAPER

Adaptive opening times for evacuation shelters during disasters

Hanwen Liu[1] · **Qi Luo**[2] · **Yongjia Song**[1]

Received: 1 March 2024 / Accepted: 3 November 2024

Abstract
The increasing frequency and severity of extreme weather events, such as hurricanes and tropical cyclones, address the importance of prompt and effective natural disaster response strategies. Our research aims at developing a learning-based decision-making framework tailored for evacuation shelter opening time (ESOT), with a focus on prioritizing the demands of vulnerable populations. This approach seamlessly integrates various complex supply- and demand-related factors, including evacuation demand forecasting and shelter operations requirements. The shelter opening time is formulated as a multi-class optimal stopping problem, which readily addresses the trade-off between the risks of false alarms and the perilous consequences of delayed responses, accommodating the uncertainties in disaster state evolution. To improve the computational and sample efficiency, we created a hierarchical policy approximation approach, providing provable optimality guarantees. Through a case study of Hurricane Florence in 2018 using historical wind speed data, our findings demonstrate the efficiency and flexibility of the ESOT policy, clearly outperforming standard stochastic optimization methods. For example, the total cost saving using our approach ranges from 6.6 to 28.2%, and the cost saving is more significant when the variance of the predictor is larger. These results highlight the benefits of integrating learning-based disaster management strategies with physics-informed forecasting models for protecting vulnerable populations in the face of disasters.

Keywords Emergency response · Evacuation shelter management · Optimal stopping problem

1 Introduction

The increasing frequency and intensity of extreme weather events, driven by the confluence of physical and societal causes [1], have highlighted the importance of adaptive disaster response strategies. In 2021, the United States was struck by 20

Extended author information available on the last page of the article

Published online: 18 November 2024

severe natural disasters, including wildfires, floods, tornadoes, and tropical cyclones, resulting in damages amounting to over $145 billion [2]. An alarming trend in the intensification of major natural disasters was discovered by [3–5]. Consequently, disaster response management has become a critical research field aimed at optimizing the allocation of resources and mitigating their impacts within local communities. These management strategies consider a wide array of prediction and decision-making problems, including forecasting behavioral responses [6], calculating evacuation routes [7], operating shelters [8], and coordinating humanitarian aid [9, 10].

This paper focuses on the ESOT problem because of its critical role in the transition from pre-disaster preparedness to during-disaster responses. The ESOT policy outlines how authorities dynamically determine evacuation and/or sheltering directives to local communities [11], taking into account the evolving forecasts and damage assessments. This policy can guide disaster management entities to initiate disaster preparedness and communicate with the public via mass media and mobile technologies [12, 13]. Nevertheless, deriving adaptive ESOT policies necessitates a combination of advanced weather forecasting techniques and reliable decision-support tools [14].

Since ESOT orders and the corresponding hurricane preparedness policies are irrevocable in the face of unpredictable hurricane trajectory and intensity scenarios [15, 16], the optimal ESOT policy must balance between *false alarms* and *delayed responses*. Specifically, issuing evacuation and sheltering directives too early may lead to false alarms, squandering preparedness efforts and diminishing the public's trust in authorities. Conversely, delayed directives can directly endanger lives and property [15, 17, 18]. The evacuation order for Hurricane Ian in September 2022, criticized for being over 24 h late in Lee County, led to 61 deaths [19]. The delay was primarily attributed to an unforeseen shift in the projected storm despite prior warnings of severe flooding by meteorologists. Given the unpredictable nature of hurricane events, our paper presents a learning-based method adapting to hurricane predictors modeled by stochastic processes, including shifts in hurricane trajectories characterized by jumps. Our research focuses on determining *when and where* to issue evacuation orders and initiate shelter operations. Nevertheless, the insights obtained from this analysis are also applicable to the broader context of disaster response management [20, 21].

Previous studies have primarily concentrated on two distinct areas: the development of precise statistical models for predicting hurricane severity and forecasting demand [22, 23], and the formulation of the optimal policy for shelter operations [21], with the assumption that such models are readily available. For example, the predict-and-optimize (PAO) framework distinguishes between estimating parameters and determining shelter opening times as two sequential tasks, aiming to strike a balance between them. Recently, there is a notable shift toward employing *end-to-end learning* for leveraging physics-informed forecasting models. For example, [24] developed a deep-learning approach to predict evacuation traffic conditions. [25] extracted useful hurricane-related information from social media text. Yet, these models lack adaptability within constantly evolving hurricane events. An alternative solution is integrating dynamic programming [26, 27] with simulation-based environments to address sequential decision-making. However, this method faces

a nuanced trade-off between the model complexity and computational complexity. Using physics-informed forecasting and simulation models will improve predictive accuracy but may adversely complicate the assessment of ESOT strategies. This is exemplified by the intense computational resource required for running the FEMA's HURREVAC [28] and Hazus [29] models for risk assessment and damage estimation. In conclusion, optimizing ESOT policy in conjunction with simulated scenarios at granular scales necessitates substantial computing resources.

This study proposes a *direct policy approximation* (DPA) approach to address the challenges of interdependent shelter opening time decisions with continuous state spaces. The DPA approach constructs hierarchical neural networks (HNNs) to estimate the optimal opening times for different classes of shelters, which are arranged according to a predetermined priority sequence. When the demand exceeds the capacity of shelters in the current priority class, it is transferred to the next class in the sequence only if they are open. This strategy aims to reduce both the operational costs of the shelters and the penalties associated with failing to meet demand from the most vulnerable populations. The DPA method can automatically balance the total cost induced by early and delayed shelter opening time decisions. We conduct numerical experiments that show notable robustness with respect to prediction and model misspecification errors. The results outperform the PAO benchmark, using the sample average approximation (SAA) method for simulation-based policy evaluations.

Our study enhances disaster response strategies by addressing a key tradeoff in ESOT decisions: the balance between false alarms and delayed responses due to predictive accuracy. Our main contributions include: (1) We apply a deep-learning-based DPA approach to integrate hurricane forecasts with evacuee behavior and decision-making processes; (2) We propose an HNN model to approximate ESOT policies with theoretical guarantees for the optimality of these policies with a small sample complexity. Since the DPA approach and its HNN-based policy representation are tailored to this particular problem, they are anticipated to surpass the effectiveness of generic model-free methods in approximate dynamic programming. Our contributions to optimizing disaster preparedness and shelter management lay the foundation for advanced computational techniques to deploy scientifically reliable and practically feasible disaster response strategies.

The remainder of this paper is organized as follows. Section 2 provides a formal statement of the ESOT problem as multiclass optimal stopping and proposes a deep-learning-based optimal stopping framework. Section 3 conducts comparative numerical experiments based on Hurricane Florence data in 2018. Section 4 concludes and discusses future research directions.

2 Optimal opening times for evacuation shelters

2.1 Problem statement

Since the execution of hurricane evacuation plans is delegated to local government agencies [19], local emergency management authorities must determine shelter

operations plans for their local communities after receiving the hurricane watch. Consider a decision-maker (DM) who determines whether and when to send a sheltering notification to its administered areas over N periods. The DM's decision-making process is modeled as a Markov decision process (MDP) where the ESOT decisions are interdependent optimal stopping time and the state transitions are determined by simulation-based hurricane and evacuation demand forecasts. The beginning of the planning horizon is set at the first time that the National Hurricane Center (NHC) issues a tropical storm or hurricane watch (tropical storms sustained winds of 39 to 73 miles per hour (mph); hurricanes sustained winds of 74 mph or higher), which is typically 48 h in advance of the anticipated onset of tropical storm or hurricane conditions within the specified area [30]. The terminal period of the planning horizon corresponds to the hurricane landfall time or when the hurricane watch is canceled due to weakened tropical storm forces. The DM reassesses the shelter opening decision using 12-hour intervals to match with the NHC's hurricane forecast update frequency. For example, the case study of Hurricane Florence considers a planning horizon of six days, which include twelve intervals.

MDP formulation Hurricane preparedness is organized within a network consisting of a set of residential nodes R and shelters S (see examples in Fig. 3). The hurricane forecast data at time n is denoted by the predicted hurricane state $X_n \in \mathbb{R}^{|R|\times l}$, where l represents the number of attributes of the hurricane, such as wind intensity, wind speed, the stage of the cyclone, and environmental variables such as the sea-surface temperature. For example, $(X_n)_{n\in[N]}$ can be predicted by multiple different forecast models developed by NHC [31]. Let $\mathcal{X}$ denote this high-dimensional hurricane state space. Given hurricane forecasts at time n, the evacuation demands are estimated by $f_D : \mathcal{X} \to \mathbb{R}_+^{|R|}$ [32]. Specifically, the evacuation demand vector for each observed X_n is represented by $\tilde{D}_n(X_n) = f_D(X_n) = \{d_i(n)\}_{i\in R}$, enabling to integrate simulation-based damage prediction and evacuation behavior models.

According to the operational guidelines from various states [33], shelters are organized into K priority levels (classes), where: (1) shelters within the same priority level are activated simultaneously, and (2) a shelter of a certain priority level can only be activated if all shelters of higher priority levels have already been activated. The prioritization of shelters is determined by their capacity and proximity to the most vulnerable households, as defined in [33]. The allocation of individuals from their origins to shelters, denoted as $D_n \in \mathbb{R}_+^K$, can be inferred from $\tilde{D}_n$ through related evacuation behavior models [16, 18]. Since this aspect falls outside the scope of our research, we assume that such allocations are pre-established, meaning that the demand for shelters across the K classes, contingent upon the hurricane condition X_n at each time period n, is predetermined.

Transition process Upon the realization of the current hurricane state, the DM will predict the future demand to guide subsequent decisions. Since the transition function, denoted as $P_n(X_{n+1}|X_n)$, has incorporated complex physics-informed hurricane forecasts, it does not have a closed-form expression for transitioning from state

X_n to X_{n+1}. We approach the hurricane evolution process as a grey-box environment, indicating the limited knowledge about the system evolution through data-driven simulations. Our case study applies a simulation-based approach to reconstruct the hurricane's wind speed forecast model using historical hurricane data, a commonly used method in literature [34]. Therefore, we consciously avoid making specific assumptions about the transition function, acknowledging the inherent uncertainties and complexities in accurately modeling hurricane dynamics.

ESOT policy An ESOT policy maps the hurricane state space $\mathcal{X}$ to multi-class shelter opening decisions, which is denoted by $\pi : \mathcal{X} \to \mathcal{A}$. The decision space $\mathcal{A}$ is defined by $\mathcal{A} = \{a \in \{0,1\}^K : \forall k \in \{1,\dots,K-1\}, a_k \geq a_{k+1}\}$. Given the demand model, the DM's decision at time n reduces to a vector $a_n = (a_{1,n}, \dots, a_{k,n}, \dots, a_{K,n}) \in \mathcal{A}$, where each $a_{k,n}$ denotes whether or not shelters in class k are open at time n. The implementable ESOT policy follows the following assumptions: (1) Since the opening and operations of shelters require substantial resources and maintenance personnel, once a shelter is open, it will remain in operation until the end of the planning horizon N. (2) The DM can commence evacuation procedures prior to the recommended action time, i.e., when $a_n = \mathbf{1}$ is implemented, it is assured that all evacuees have securely arrived at the shelters by the time point n.

Cost function The shelters' cost at time n, denoted as $c_n(X_n, a_n)$, consists of two types of cost, the shelter operational cost and penalty cost. Specifically, an operational shelter will incur a fixed cost of $h = \gamma_0 \cdot \sum_{i \in R} b_i$ where γ_0 is the unit cost and $b_i = (b_{i,1}, \dots, b_{i,K})^\intercal$ is the shelter capacities allocated to residential node i for all classes. This fixed cost h accounts for personnel, material supplies, and maintenance costs depending on the shelters' capacities. The penalty cost incorporates the number of evacuees who are underserved, including the ones who cannot find sheltering spaces due to late notification or insufficient sheltering capacity. This is because only qualified evacuees [33] for each class of shelters will be assigned to facilities following a first-come-first-serve policy. Thus, the cost function $c_n(X_n, a_n)$ can be written as:

$$c_n(X_n, a_n) = \begin{cases} h^\intercal \cdot a_n + \sum_{i \in R} o_{i,n}(X_n) \min\{D_{n,i}(X_n), (\mathbf{1} - a_n)^\intercal b_i\} & n \in [N-1] \\ c_N(X_N) & n = N, \end{cases} \tag{1}$$

where $o_{i,n}(X_n) \in \mathbb{R}$ is the unit penalty for the underserved population in node i at time n.

The objective is to find the optimal ESOT policy π^* that minimizes the expected total cost over the planning horizon. The expected total cost under any policy π is given by

$$J^\pi(X) = \mathbb{E}_\pi \left[\sum_{n=1}^{N-1} c_n(X_n, a_n) + c_N(X_N) \mid X_0 = X \right]. \tag{2}$$

Challenges of solving the optimal ESOT policy The challenges of solving the optimal opening time problem described in (2) include: (i) The transition functions and demand estimation models are based on simulations or other physics-informed disaster forecast models. Thus, the state space is continuous and the transition functions do not have explicit forms; (ii) A feasible multi-class optimal stopping policy must satisfy the irrevocable decisions and the order of priority in opening shelters, i.e., a feasible policy satisfies that $a_{k,n} \geq a_{k+1,n}$ for $n = 1, \ldots, N, k = 1, \ldots, K-1$ and $a_{k,n} \leq a_{k,n+1}$ for $n = 1, \ldots, N-1, k = 1, \ldots, K$. Alternative approximate dynamic programming methods may encounter challenges related to sample complexity stemming from the expense involved in evaluating policies through physics-informed simulations.

2.2 Optimal ESOT policy with priority constraints

Define J_n^* as the optimal total cost at time n. Given the irrevocable nature of opening shelters, the feasible policy must consider the current state of the hurricane, X_n, and the decisions made in the previous time step, a_{n-1}. Thus, the optimal total cost, $J_n^*(X_n, a_{n-1})$ for $n \leq N-1$ is given by the following expression:

$$J_n^*(X_n, a_{n-1}) = \min_{a_n \in \mathcal{A}, a_n \geq a_{n-1}} \left\{ c_n(X_n, a_n) + \mathbb{E}_P[J_{k,n+1}^*(X_{n+1}, a_n)] \right\}. \tag{3}$$

This optimality equation highlights the challenge of dealing with a continuous state space problem constrained by the dependent and irrevocable decisions of shelter opening times. Various discretization and approximation approaches have been proposed to manage the computational and analytical hurdles associated with such constrained decision variables. We consider a more direct approach to address the challenges associated with the dependency among the decisions by segmenting the problem into a sequential HNN structure.

In the following discussion, we elaborate on the methodology for determining the optimal stopping decision $a_{k,n}$ for any single class k. Let $J_{k,n}^*(X_n)$ represent the expected optimal total cost for shelter class k at time n, which is conditional on the optimal actions of classes $k' > k$. The optimal ESOT policy for class k is to seek $a_{k,n}$ that achieves $J_{k,n}^*(X_n)$. In this structure, each class k is treated as an independent subproblem. When disaggregating by class, we first consider the condition of irrevocable decisions. Specifically, once $a_{k,n} = 1$, the shelter class cannot be closed in subsequent periods, eliminating further decision-making for class k and introducing a terminal condition and cost $\psi_{k,n} = (N-n+1)h_k$, where $h_k = \gamma_0 \sum_{i \in R} b_{i,k}$ is the holding cost for class k. Notably, activating $a_{k,n} = 1$ allows decisions for classes greater than k at time n, thereby $J_{k+1,n}^*(X_n)$ is given by

$$J_{k,n}^*(X_n) = \psi_{k,n} + J_{k+1,n}^*(X_n), \text{ for } a_{k,n} = 1.$$

Conversely, if $a_{k,n} = 0$, implying sequential opening, all classes greater than k must remain closed, accruing a penalty cost $c_{k,n}(X_n)$ at time n for classes k to K

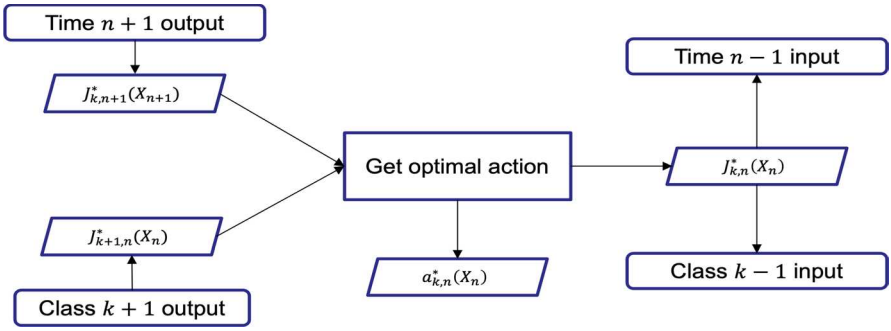

Fig. 1 The two-layer backward induction framework

$$c_{k,n}(X_n) = \sum_{i \in R} o_{i,n}(X_n) \min \left\{ (D_{n,i}(X_n) - \sum_{j=0}^{k-1} b_{i,j})^+, \sum_{j=k}^{K} b_{i,j} \right\},$$

where $b_{i,j}$ is the shelter capacities allocated to residential node i for class j, $b_{i,0} = 0$, $o_{i,n}(X_n) \in \mathbb{R}$ is the unit penalty. Furthermore, $a_{k,n} = 0$ implies a potential opening in future periods, the optimal cost for next period $J^*_{k,n+1}(X_n)$ must be considered:

$$J^*_{k,n}(X_n) = c_{k,n}(X_n) + \mathbb{E}_P[J^*_{k,n+1}(X_{n+1})], \text{ for } a_{k,n} = 0.$$

Combining these conditions, the expected total cost at time $J^*_n(X_n)$, is recursively computed as:

$$\begin{aligned} &J^*_n(X_n) = J^*_{1,n}(X_n), \text{ and for all } k = 1, 2, \ldots, K, \\ &J^*_{k,n}(X_n) = \min_{a \in \{0,1\}} \left\{ a(\psi_{k,n} + J^*_{k+1,n}(X_n)) + (1-a)(c_{k,n}(X_n) + \mathbb{E}_P[J^*_{k,n+1}(X_{n+1})]) \right\}, \end{aligned} \tag{4}$$

where $J^*_{K+1,n}(X_n) = 0$. This approach involves two layers of backward inductions, starting from time N to 1 and from class K down to class 1. Appendix A provides a detailed proof of the validity of this optimality equation. The advantage of this approach lies in its simplicity for the design of DPA methods. The decision $a_{k,n}$ is directly derived from the set $\{0, 1\}$, eliminating the need to account for interdependent decision vectors without loss of optimality. The flowchart of the methodology is provided by Fig. 1.

The translation from a sequence of actions a_t to the opening times for K classes of shelters is provided below. From time n onwards, we define $\tau_n \in \{n, n+1, \ldots, N\}^K$ as a vector of opening times where each entry $\tau_{k,n} \in \{n, n+1, \ldots, N\}$ is the ESOT for shelters in class k. Let τ^*_n denote the optimal ESOT accordingly. Because of the shelter opening priority rules, $\tau_{k,n} \leq \tau_{k-1,n}$. Therefore, $\tau_{k,n}$ ($\tau^*_{k,n}$) can be represented as a function of $a_{k,n}$ ($a^*_{k,n}$). For any shelter class k during time n to N, the optimal opening time, $m \in [n, N]$ should be the first instance following $\tau_{k-1,n}$ which $a^*_{k,m} = 1$. By definition,

$$\tau^*_{k,n} = \sum_{m=\tau_{k-1,n}}^{N} m a^*_{k,m} \prod_{j=\tau_{k-1,n}}^{m-1} (1 - a^*_{k,j}). \tag{5}$$

Apply this calculation across all classes from $k = 1$ to $k = K$ with $\tau_{0,n} = 0$, the DM can calculate $\tau_{k,n}$. The product-form optimal stopping time representation [26] can substitute the direct calculation of optimal policies. In the next section, we will derive a policy approximation framework that outputs the binary action with a continuum approximation.

2.3 Neural network approximation for ESOT policy

Recall that deriving optimal ESOT policies is challenging because: (i) The classic curse of dimensionality in dynamic programming arises due to the compact hurricane state space $\mathcal{X}$, potentially continuous, and the corresponding high dimensional and potentially implicit demand forecast models; (ii) The state transition $P_n(X_{n+1}|X_n)$ is driven by physics-informed models, preventing the application of standard solution methods for solving the Bellman's optimality equation (4). An alternative is to use approximate dynamic programming techniques [35], such as Q-learning, to compute approximate $\tilde{J}$. However, these generic methods may require extensive parameter tuning and are sensitive to input data.

DPA method We propose a DPA method that exploits the structure of the ESOT policy. Let $f^{\theta_{k,n}}(X_n) : \mathbb{R}^{|R|\times l} \to \{0, 1\}$ with parameter $\theta_{k,n} \in \mathbb{R}^q$ denoting the HNNs that approximate the optimal policy for opening times. For simplicity, we assume that the terminal conditions are fixed with decisions $a_N = \mathbf{1}$. The HNN approximation is implemented as follows:

1. In the backward induction of $n \in \{0, 1, \ldots, N-1\}$, let $\theta_{k,n+1}, \theta_{k,n+2}, \ldots, \theta_{K,N} \in \mathbb{R}^q$ denote parameters such that

$$\tau_{k,n+1} = \sum_{m=\tau_{k-1,n+1}}^{N-1} m f^{\theta_{k,m}}(X_m) \prod_{j=\tau_{k-1,n+1}}^{m-1} (1 - f^{\theta_{k,j}}(X_j))$$

   produces an near-optimal value for $J^*_{k,n+1}(X_{n+1})$.
2. For notational convenience, we use a vector of $f^\theta : \mathbb{R}^{|R|\times l} \to \{0, 1\}$ to denote the HNN approximation for the policy function. Compared to generic value or policy approximation methods, this HNN approximation is more sample-efficient. The architecture and operational dynamics of the HNNs for sequential opening time decisions are shown in Fig. 2. The effectiveness of this policy approach is validated in Appendix B, demonstrating its utility in efficiently solving complex ESOT problems in disaster management.

   f^θ is given by

$$f^\theta = \mathbf{1}_{[0,\infty]} \circ g^\theta_I \circ \varphi_{q_{I-1}} \circ g^\theta_{I-1} \circ \ldots \circ \varphi_{q_1} \circ g^\theta_1, \tag{6}$$

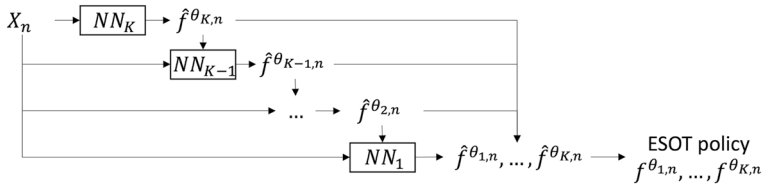

Fig. 2 Sequential hierarchical NNs structure

where (1) '$\circ$' denotes the composition of functions; (2) $I, q_1, q_2, \ldots, q_{I-1}$ are positive integers specifying the depth of the network and the number of nodes in the hidden layers (if there are any); (3) $g_1^\theta : \mathbb{R}^{|R|\times l} \to \mathbb{R}^{q_1}, \ldots, g_{I-1}^\theta : \mathbb{R}^{q_{I-2}} \to \mathbb{R}^{q_{I-1}}$ and $g_I^\theta : \mathbb{R}^{q_{I-1}} \to \mathbb{R}$ are affine functions; (4) for $j \in \mathbb{N}$, $\varphi_j : \mathbb{R}^j \to \mathbb{R}^j$ is the component-wise activation function (e.g., ReLU) given by $\varphi_j(x_1, \ldots, x_j) = (x_1^+, \ldots, x_j^+)$; (5) $\mathbf{1}_{[0,\infty)} : \mathbb{R} \to \{0, 1\}$ is the indicator function of $[0, \infty)$.

3. Since f^θ takes values in $\{0, 1\}$, it does not directly lend itself to a gradient-based optimization method. So, as an intermediate step, we introduce a feed-forward HNN $\hat{f}^\theta : \mathbb{R}^{|R|\times l} \to [0, 1]$ of the form

$$\hat{f}^\theta = \psi \circ g_I^\theta \circ \varphi_{q_{I-1}} \circ g_{I-1}^\theta \circ \ldots \circ \varphi_{q_1} \circ g_1^\theta,$$

where $\psi : \mathbb{R} \to (0, 1)$ is the logistic function $\psi(x) = e^x/(1 + e^x)$, and all other functions are defined the same way with (6).

The parameter $\theta \in \mathbb{R}^q$ of $\hat{f}^\theta$ consists of (1) the entries of the matrices $A_1 \in \mathbb{R}^{q_1\times d}, \ldots, A_{I-1} \in \mathbb{R}^{q_{I-1}\times q_{I-2}}, A_I \in \mathbb{R}^{1\times q_{I-1}}$ and (2) the vectors $b_1 \in \mathbb{R}^{q_1}, \ldots, b_{I-1} \in \mathbb{R}^{q_{I-1}}, b_I \in \mathbb{R}$. The affine functions for $i = 1, \ldots, I$ are given by $g_i^\theta(x) = A_i x + b_i$. So, the dimension of the parameter space is

$$q = \begin{cases} d + 1 & \text{if } I = 1 \\ 1 + q_1 + \cdots + q_{I-1} + dq_1 + \cdots + q_{I-2}q_{I-1} + q_{I-1} & \text{if } I \geq 2 \end{cases}.$$

Based on the HNN approximation of the ESOT policy, the objective is to find $\theta_{k,n} \in \mathbb{R}^q$ and $\hat{f}^{\theta_{k,n}}(X_n) \in [0, 1]$ such that

$$J'_{k,n}(X_n) = \mathbb{E}\Big\{ [\psi_{k,n} + J'_{k+1,n}(X_n)]\hat{f}^{\theta_{k,n}}(X_n) + [c_{k,n}(X_n) + J^*_{k,n+1}(X_{n+1})](1 - \hat{f}^{\theta_{k,n}}(X_n)) \Big\}$$
$$\approx \min_{\theta\in\mathbb{R}^q} \mathbb{E}\Big\{ [\psi_{k,n} + J'_{k+1,n}(X_n)]\hat{f}^{\theta}(X_n) + [c_{k,n}(X_n) + J^*_{k,n+1}(X_{n+1})](1 - \hat{f}^{\theta}(X_n)) \Big\}.$$

4. Set $J'_{K+1,n}(X_n) = 0$. Apply steps 2 and 3 for each class reversely from $k = K$ to 1.
5. The approximate policy calculates the binary action by a threshold rule

$$f^{\theta_{k,n}} = \begin{cases} 1 & \text{if } \hat{f}^{\theta_{k,n}} \geq 1/2 \\ 0 & \text{if } \hat{f}^{\theta_{k,n}} < 1/2 \end{cases}. \tag{7}$$

6. Let $f^{\theta_{k,n}}(X_n) = a^*_{k,n}$, we compute the optimal cost functions $J^*_{k,n}(X_n)$ at time n.

3 Case study: response policies for Hurricane Florence in South Carolina

Our numerical case study focuses on the hurricane response decision-making process related to Hurricane Florence in September 2018. While Hurricane Florence started with a modest tropical storm, a rapid intensification on September 4 upgraded it to a Category 4 hurricane on the Saffir-Simpson scale (SSHWS) [36]. It reached estimated peak winds of 130 miles per hour (mph) and caused catastrophic damage in the Carolinas, including significant destruction of homes and infrastructure and widespread power outages, with some areas experiencing inundation from the storm surge. Hurricane Florence caused an estimated $24 billion in damages across the Carolinas, marking its place as one of the most destructive storms in recent history [36]. The DM in this context is the South Carolina Emergency Management Division (SCEMD), deciding whether to evacuate residents across South Carolina and open shelters for those most vulnerable. The specific elements included in the SCEMD's decision are described below.

3.1 Experiment setup and data description

Hurricane data Our numerical experiments collected Hurricane Florence forecast data from September 8 at 0:00 to September 14 at 0:00 with a 12-hour time interval for each period, resulting in $N = 12$ periods in total. The goal of this numerical experiment is to trace back to the incident's starting time and check the performance of our algorithm in identifying the optimal ESOT policy to reduce potential damages. The hurricane made landfall on September 14, 2018, near Wrightsville Beach, North Carolina, with maximum sustained winds of 90 mph. The NHC's Geographic Information Systems archive provided various hurricane forecast data and products, allowing us to reconstruct the evolution of the hurricane's progression over the planning horizon. To test the impacts of uncertainties in system states, our numerical experiments also examine the impacts of varying volatility levels through simulations.

Shelter data The South Carolina Hurricane Management Plan [37] divides all hurricane-related general population shelters in South Carolina into four regions: Central, North, South, and West. The division of each region is shown in Fig. 3. Within each region, shelters are further divided into four shelter classes (C1 to C4), ranked by their priority in operations. Each class contains a group of shelters (16, 18, 10, 4) for C1 to C4 to serve potential evacuees in each county in disaster events (Fig. 3). Fig. 3 shows an example of the distribution of shelters in SC and the connections between counties and shelter classes in the Central region.

Hurricane state representation Our first set of experiments focuses on the shelter opening plans within the Central region of South Carolina, which includes 48 shelters located throughout seven counties. The uncertainties in hurricane forecasts are simulated by combining the historical Wind Speed Probabilities (WSP) distribution

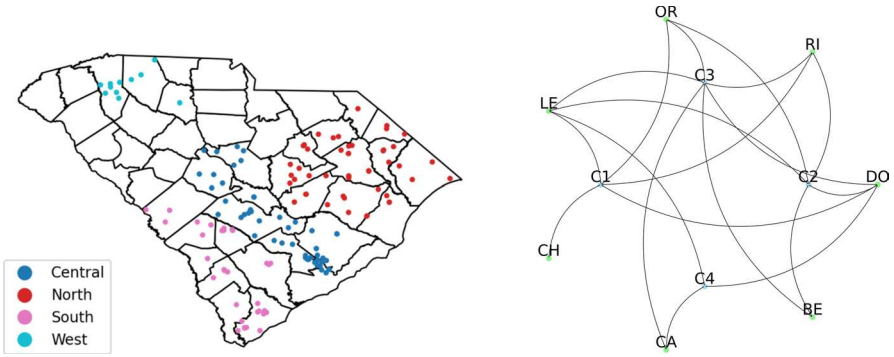

Fig. 3 Shelters in SC are divided into the 4 regions (left); The connections between counties and shelter classes (C1,C2,C3,C4) in the Central region (right) from SCEMD

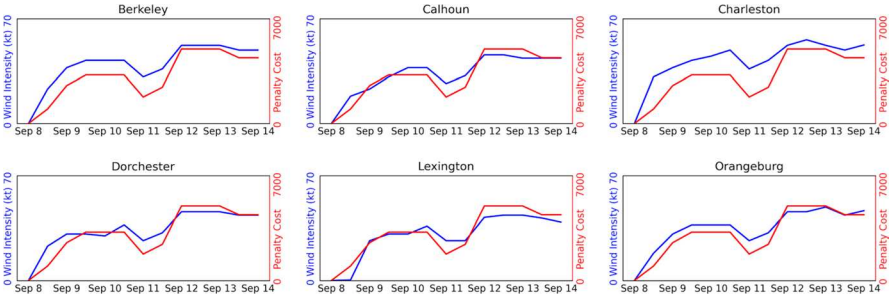

Fig. 4 Mean wind intensity and total penalty cost estimation of Hurricane Florence in Central region, SC. based on NOAA's WSP data

and a continuous-time Brownian motion for spatiotemporal uncertainties. The WSP data provided by the NOAA offers the probability of a specific level of hurricane occurring within a county in the next five days [30]. Specifically, it provides three probabilities: the probability of a hurricane with wind speeds of at least 34 knots (39 mph), at least 50 knots (58 mph), and at least 64 knots (74 mph) hitting the county in the future. The hurricane intensity, denoted by X_t, is represented by Gaussian processes. Given the observation at time t, the tail distributions are combined to estimate the average intensity $\bar{X}_t$ during the interval and its covariance. Figure 4 shows the average hurricane intensity $\bar{X}_t$ used in our experiments. As a result, the wind intensity in X_t can be expressed by:

$$dX_t = \alpha \bar{X}_t dt + \sigma dB_t, \tag{8}$$

where α is a constant parameter, σ represents the infinitesimal variance estimated from the historical data, and B_t is a Brownian motion. This state evolution process can also be extended for the case where the hurricane state X_t includes other attributes.

Table 1 Central coastal hurricane region vulnerable population and capacities

County	Vulnerable population	C1	C2	C3	C4
Berkeley	3507	0	1792	1215	0
Calhoun	389	0	0	265	561
Charleston	18,173	1976	459	243	0
Dorchester	1677	3427	404	459	1906
Lexington	3788	1061	1829	846	500
Orangeburg	2450	3577	2328	174	0
Richland	1478	1196	1672	267	0

The instantaneous demand $D_n(X_n)$ for shelters employs a simplified evacuation behavior model from [38, 39] and the data from the South Carolina hurricane management plan [37]. Initially, the behavioral model determines the number of evacuees by estimating the probability of individuals deciding to evacuate. The evacuation decision follows a probit model $y(X_n) = \beta_0 X_n + \beta_1 I + e$. In this model, y represents the propensity of a household to evacuate. The variable I includes other static personal data, including attributes like gender, race, income, and the nature of the evacuation notice. The coefficient β_0 is tied to hurricane-centric conditions, whereas β_1 concerns personal traits; e is the error term. The derived probability of evacuation is $\Phi(y(X_n))$, with Φ denoting the cumulative normal distribution function. Next, we use the census data in these counties to make a general estimation of I and generate a comprehensive map of $y(X_n)$, representing the percentage of vulnerable individuals in each county who opted to evacuate. Various statistical and discrete choice models have been developed for estimating evacuation demand based on the hurricane forecast and other factors [18, 38, 40]. In our study, the total demand was calculated as $D_n(X_n) = \Phi(y(X_n))p$, where p represents the population in each county. The proportion of individuals choosing public shelters was also provided by the aforementioned study [37]. A summary of shelter capacities and evacuee populations in each county is displayed in Table 1.

Cost functions The holding and penalty costs for each shelter class follow (1), the coefficients of which are specified as follows. Let $h_k = \sum_{i=1}^{|R|} \gamma_0 b_{i,k}$, where $b_{i,k}$ is the total capacity of shelter class k for county i and γ_0 is the unit holding cost. For simplicity, γ_0 has been normalized based on varying operational conditions in different counties. The unit penalty cost function is $o_n(X_n) = \gamma_1 X_n$, where γ_1 is a constant representing the impact of hurricane intensity on unsheltered population. The estimation of demand and the evolution of hurricane states are well-determined by numerous resources with standardized and predictive models. The distinction between the holding cost γ_0 and penalty costs γ_1 has a substantial impact on the optimal policies. Acknowledging that there are no universally accepted standards to define the opening and penalty costs, we propose a separable cost function and test its sensitivity beyond a general examination of hurricane data. This sensitivity analysis also enables us to understand the effects of cost differentials on the associated decisions.

Evaluation and experiment setting The evaluation procedure is as follows: First, using the hurricane state generation model detailed in (8), we choose a relatively

Table 2 Sensitivity analysis for three regions in SC

Region	Parameter	Time (seconds)	Lower bound	Mean	Upper bound	Decisions
Central	$\gamma_0 = 0.5, \gamma_1 = 1/30$	2175	124,491	124,570	124,649	1, 7, 7, 12
Central	$\gamma_0 = 0.5, \gamma_1 = 1/40$	2232	114,048	114,134	114,219	7, 7, 7, 12
Central	$\gamma_0 = 0.5, \gamma_1 = 1/50$	2150	97,969	98,102	98,236	7, 12, 12, 12
Central	$\gamma_0 = 0.5, \gamma_1 = 1/60$	2150	85,668	85,810	85,953	12, 12, 12, 12
North	$\gamma_0 = 0.5, \gamma_1 = 1/30$	2164	83,527	83,669	83,812	7, 12, 12, 12
North	$\gamma_0 = 0.5, \gamma_1 = 1/40$	2195	62,880	63,019	63,158	12, 12, 12, 12
North	$\gamma_0 = 0.5, \gamma_1 = 1/50$	2195	50,499	50,611	50,723	12, 12, 12, 12
North	$\gamma_0 = 0.5, \gamma_1 = 1/60$	2214	42,005	42,097	42,190	12, 12, 12, 12
South	$\gamma_0 = 0.5, \gamma_1 = 1/30$	2174	59,061	59,197	59,333	12, 12, 12, 12
South	$\gamma_0 = 0.5, \gamma_1 = 1/40$	2145	44,298	44,401	44,503	12, 12, 12, 12
South	$\gamma_0 = 0.5, \gamma_1 = 1/50$	2115	35,452	35,532	35,612	12, 12, 12, 12
South	$\gamma_0 = 0.5, \gamma_1 = 1/60$	2115	29,497	29,564	29,631	12, 12, 12, 12

small parameter $\sigma = 1.5$ to ensure that our generated tracks closely align with the hurricane predictions made by NOAA. We simulate 10,000 trajectories and split them into two halves. The first half of these trajectories are used for training the policy f and the second half of them are used for testing. The boundary condition is set by the proportion of $f(X_n)$ values as follows: if more than 50% of trajectories (i.e., over 2500) choose class k shelters to open at time n, then we set $a_{k,n} = 1$, otherwise we set $a_{k,n} = 0$. Then we apply (5) to get the optimal shelter opening time $\tau^*_{k,n}$. The HNN architecture is determined by a distributed hyperparameter tuning tool (RayTune) [41]. Each network comprises two hidden layers between the input and output layers, with each hidden layer containing 40 more neurons than the number of input features. The Adam optimizer [42] was employed for training, with a learning rate set at 0.001.

3.2 Results for central region in South Carolina

To understand the variation in the optimal costs and to verify the efficacy of our stopping time policy, we conducted tests using actual hurricane tracks while varying coefficients γ_0 and γ_1 values. Note that γ_0 and γ_1 have been normalized as relative costs to address the tradeoff between false alarms and delayed responses. Our sensitivity analysis ranges from a high penalty scenario with $\gamma_1 = 1/30$ to moderate penalty scenarios at $\gamma_1 = 1/40$ and $\gamma_1 = 1/50$, and a low penalty scenario at $\gamma_1 = 1/60$. The running time, average cost and optimal decisions are provided in Table 2.

The actual historical hurricane tracks in the central region can be used to check if our solution method can precisely determine the optimal shelter opening time based on hurricane forecasts across various γ_1 situations. In particular, for moderate penalty scenarios, the optimal shelter opening time is pinpointed at time 7 (September 11th, 12:00 PM), which aligns with South Carolina's official evacuation time [43]. In higher penalty scenarios, our model recommends an earlier opening time

of period 1 (September 8th, 12:00 PM), providing a longer duration for evacuation preparations. Conversely, in Southern regions like Beaufort County, evacuation orders were issued but our model withheld an opening decision until the terminal time $n = N$, irrespective of the penalty level, because of the low hurricane intensity predicted for these counties. The considerable resources spent in the preparations were ultimately not needed [44]. The comparison between simulated experiments and historical ground truth data confirms that the implementation of ESOT policy can mitigate adversarial impacts in these areas [43] and improve resource allocations during disaster events. On the other hand, it is important to acknowledge the intricate nature of decision-making in the actual planning. The local governments must consider more nuanced costs and other factor [33] and prioritize the safety of citizens, even if risks are deemed to be low.

3.3 Comparative study and sensitivity analysis

PAO benchmark model The PAO approach is commonly used to solve optimal stopping problems where the state forecasts and system uncertainty are estimated from historical observations. The benchmark model uses an SAA-based MDP model for hurricane evacuation decisions [21]. In this case, the primary challenge in solving MDPs is determining the transition probability matrix after discretizing the state space. This transition model can be estimated using regression on Monte Carlo sampled state evolution trajectories [21, 45]. The general form of the MDP model can be expressed as:

$$J^*_{k,n}(X_n) = \min_{a\in\{0,1\}} \Big\{ a[\psi_{k,n} + J^*_{k+1,n}(X_n)] + (1-a)[c_{k,n}(X_n) + \sum_{X_{n+1}} q_n(X_{n+1}|X_n) J^*_{k,n+1}(X_{n+1})] \Big\}.$$

Different from the DPA approach which calculates the near-optimal policy directly, the PAO approach consists of two consequent steps similar to SAA. In the first step, the transition matrix is estimated based on the simulated training data. In the second step, the optimal shelter opening time decisions are computed by the standard backward induction. The assumption is that geographically proximate counties share identical transition matrices and hurricane intensities due to their similar hurricane intensity evolution trajectories, as depicted in Fig. 4. Given that Hurricane Florence impacted SC was at Category 1, the storm's intensity was capped at 100 knots for our simulation. While the hurricane generation model is consistent with (8), any intensity values surpassing the 100 kt threshold were set to 100 kt. To discretize the continuous state space of hurricane intensity, they were partitioned into 10 non-overlapping intervals (i.e., 0–10, 10–20, ⋯, 90–100). These intervals were denoted by $\omega_1, \omega_2, \ldots, \omega_{10}$. The middle point of each interval (e.g., 15 for the 10–20 range) was employed to represent the hurricane's intensity. The same training sample used in the DPA method was utilized for this analysis. The training sample size is denoted by M, and the transition probability is $q_n(X_{n+1} \in \omega_j | X_n \in \omega_i) = \frac{\sum_{m=1}^{M} 1_{X_n \in \omega_i} 1_{X_{n+1} \in \omega_j}}{\sum_{m=1}^{M} 1_{X_n \in \omega_i}}$.

3.3.1 Parameter sensitivity analysis

Our numerical results include additional sensitivity analyses, aiming to explore the performances of the DPA and the PAO method under various conditions by altering specific parameters related to the forecast model. These analyses provide deeper insights into the robustness characteristics of the DPA method. To enhance the assessment of models' efficacy, we incorporate a quadratic penalty cost function, defined as $o_n(X_n) = (\gamma_1 X_n)^2$. This quadratic formulation is intentionally chosen for its greater sensitivity over linear models, allowing for a more precise calibration of penalty costs in relation to hurricane preparedness. Under this model, the penalty for unsheltering remains relatively low for minor hurricane intensities. However, as the intensity of the hurricane increases, the penalty for unsheltering rises quadratically, reflecting a steep costs increasing with escalating hurricane severity.

Sensitivity analysis for γ_1 In the first experiment, we investigate the effect of varying the penalty parameter γ_1 while keeping other variables fixed. Specifically, we analyze how the DPA method's performance changes relative to the SAA method when the penalty cost sensitivity varies. The results of this experiment are shown in Table 3. As the value of γ_1 increases, the DPA method demonstrates a more pronounced performance. For instance, when γ_1 is 1/60, the improvement proportion is 0.32%, while at 1/30, it increases to 9.78%. This indicates that the DPA method performs more effectively when dealing with sensitive penalty costs. Conversely, when γ_1 is small, the DPA method's cost saving is less significant.

Sensitivity analysis for σ In the second experiment, we examine the impact of changing the prediction variance (σ) while keeping other variables fixed. The results are presented in Table 4. When σ is small (between 1.5-5), the DPA method's saving is relatively modest, ranging from 5.52% to 8.21%. However, as the prediction variance increases (e.g., $\sigma = 10, 20$), the improvement proportion increases to 16.09% and 28.16%. This finding highlights the DPA method's ability to adapt to larger prediction ranges or uncertainties, which can be particularly useful in real-world situations with variable hurricane intensity predictions.

In conclusion, our sensitivity analyses offer insightful revelations about the performance of the DPA method in comparison to the PAO method under a variety of

Table 3 Sensitivity analysis for γ_1 ($\sigma = 1.5$)

Parameter	Mean cost of SAA	Mean cost of DPA	Relative improvement (%)
$\gamma_0 = 0.5, \gamma_1 = 1/60$	62,495	62,297	0.32
$\gamma_0 = 0.5, \gamma_1 = 1/50$	87,260	86,493	0.88
$\gamma_0 = 0.5, \gamma_1 = 1/40$	118,024	110,274	6.57
$\gamma_0 = 0.5, \gamma_1 = 1/30$	143,644	129,595	9.78

Table 4 Sensitivity analysis for σ ($\gamma_1 = 1/40$)

Parameter	Mean cost of SAA	Mean cost of DPA	Relative improvement (%)
$\sigma = 1.5$	117,876	110,044	6.64
$\sigma = 3$	119,603	112,998	5.52
$\sigma = 5$	128,841	118,267	8.21
$\sigma = 10$	155,150	130,185	16.09
$\sigma = 20$	197,184	141,657	28.16

conditions. Based on the same training datasets, the DPA method exhibits a superior performance edge, particularly in scenarios characterized by larger prediction variances and large late-response penalty costs. This enhanced performance can be attributed to the inherent complexities of managing a large-scale continuous state space. The PAO method requires additional assumptions for defining the transition processes and state space. For instance, it discretizes the state space and employs the Monte Carlo method for estimating transition probabilities, which compromises its precision relative to the DPA method. To potentially boost the PAO method's performance, one strategy involves refining the number of grids and employing smaller intervals as in other SAA approaches. However, this improvement in estimation necessitates a significantly larger number of training samples for a large-scale state space, a requirement that is markedly more demanding than for the DPA method. Despite DPA's advantage in facilitating more accurate policy estimations, it requires more computational resources, especially for training HNNs.

3.3.2 Discussion and policy implications

Our numerical results highlight the potential of the DPA method as a robust approach to calculate adaptive operations decisions the context of disaster management. By offering enhanced performance in diverse scenarios, the DPA method can contribute to the development of more effective and efficient disaster management strategies, ultimately helping to save lives and mitigate the societal impacts of severe disaster events.

- The data-driven component in our proposed method balances between model complexity and sample complexity, leading to more accurate estimates than the transition matrix based on state space discretization.
- The structure of HMMs allows for a better representation of constrained opening time decisions, adeptly manages multiple interdependent decisions. This representation of shelter opening time policy can also be used in other hierarchical decision-making that coordinates multi-stage decisions with priority constraints.
- This end-to-end learning method excels in flexibility, not limited by transition function forms or state space continuity. It effectively uses various data sources, from historical events to simulations, enhancing accuracy and robustness in predictions and developing dynamic disaster management strategies.

- Our approach can be adapted to different types of natural disasters by modeling them as Markov processes, defining a suitable cost function for the specific disaster, and ensuring the decision structure can be transformed into binary and irrevocable decisions. This allows for the application of our method across various natural disaster scenarios, maintaining its effectiveness and flexibility.

4 Conclusion

The escalating impacts of extreme weather events necessitate adaptive and prompt data-driven decision-making techniques. This paper tackles the challenge of determining ESOT policies by formulating it as an MDP with a compact state space and sequential opening time decisions with priority constraints. Our novel approach departs from conventional methods by integrating an intricate interplay of estimation and optimization that affects effectiveness. With the aid of HNNs, this learning-based framework can maintain computational feasibility while substantiating the efficacy of the proposed method. With the increasing use of physics-informed models in disaster and evacuee behavior forecasts, our data-driven approach can be expanded as a pivotal tool for data-driven disaster management.

The main limitations of our method include requiring disaster intensity levels to be modeled as Markov processes. Additionally, the irrevocable evacuation decisions must be framed as a set of binary variables with full knowledge about their disaster-related costs. When dealing with more complex stochastic scenarios with multi-dimensional decisions over networks, translating them into this structure can exponentially increase the complexity and lead to the persistent curse of dimensionality in dynamic programming. Future endeavors should focus on extending the methodology to accommodate more complex emergency response decisions, which could pave the way for addressing more complex uncertain situations during natural disasters.

Appendix A: Supplementary Proofs for Sect. 2.2

Proof for optimality equation in (4): According to Bellman's principle of optimality, the optimal expected cost $J^*_{k,n}(X_n)$ for class k at time n is determined by the instantaneous cost, $c_{k,n}(X_n, a)$, based on the action $a \in \{0, 1\}$, as well as the expected future cost contingent on the action a taken at time n. For any class k, if the DM decides to open shelters in this class at time n (i.e., $a = 1$), the expected future cost becomes $J^*_{k+1,n}(X_n)$. This is because the cost function for the next class becomes active immediately after opening the current class, by definition. However, if the DM decides not to open shelters in this class (i.e., $a = 0$), all subsequent shelter classes must be closed. The expected future cost then depends on $J^*_{k,n+1}(X_{n+1})$, which is the cost function for the same class k in time $n + 1$.

To minimize the total cost, we select the action a that minimizes the sum of the immediate cost and the expected future cost. This leads to the following recursive equations for $k = 1, 2, \ldots K-1$:

$$J^*_{k,n}(X_n) = \min_{a\in\{0,1\}} \Big\{ a(\psi_{k,n} + J^*_{k+1,n}(X_n)) + (1-a)(c_{k,n}(X_n) + \mathbb{E}_P[J^*_{k,n+1}(X_{n+1})]) \Big\},$$

where $\mathbb{E}_P[J^*_{k,n+1}(X_{n+1})] = \int_{X_{n+1}} P_n(X_{n+1}|X_n) J^*_{n+1}(X_{n+1}) dX_{n+1}$.

The cost function $J^*_{1,n}(X_n)$ represents the optimal cost for the first class at time n. Given the sequential nature of the multi-class decisions and the dependency of all subsequent classes on the first class, the decision concerning class 1 at time n sets the stage for all other classes. When class 1 is considered, it encapsulates the decisions of all subsequent classes. This expression means that the effect of all further classes on the cost is already considered. As such, the optimal cost of the entire system at time n is simply the optimal cost for class 1:

$$J^*_{1,n}(X_n) = J^*_n(X_n).$$

This completes the proof. □

Appendix B: Supplementary Proofs for Sect. 2.3

Recall that

$$J^*_{K,n}(X_n) = \min_{a\in\{0,1\}} \mathbb{E}\Big\{ \psi_{k,n} a + [c_{K,n}(X_n) + J^*_{K,n+1}(X_{n+1})](1-a) \Big\}$$

Next, we want to show that, using the NN approximation $\hat{f}^{\theta_{k,n}}$ to derive $f^{\theta_{k,n}}$ for all classes k, the corresponding value function will approximate the optimal value with a small error bound.

Theorem 1 *Let $J'_{K,n}(X_n)$ denote the approximate value function using actions $f^{\theta_{k,n}}$ in (7). [26] For any decision k, for every depth $I \geq 2$ and constant $\varepsilon > 0$, there exists a neural network $f^{\theta_{K,n}}(X_n) : \mathbb{R}^{|R|\times l} \to \{0,1\}$ in (6) with parameter $\theta_{K,n} \in \mathbb{R}^q$ that has*:

$$\begin{aligned} J'_{K,n}(X_n) &= \mathbb{E}\Big\{ \psi_{K,n} f^{\theta_{K,n}}(X_n) + [c_{K,n}(X_n) + J^*_{K,n+1}(X_{n+1})](1 - f^{\theta_{K,n}}(X_n)) \Big\} \\ &\leq \min_{a\in\{0,1\}} \mathbb{E}\Big\{ \psi_{K,n} a + [c_{K,n}(X_n) + J^*_{K,n+1}(X_{n+1})](1-a) \Big\} + \frac{\varepsilon}{K} \\ &= J^*_{K,n}(X_n) + \frac{\varepsilon}{K}. \end{aligned}$$

Expanding upon the theorem, we explore its implications for multiple opening time decisions. Initially, we focus on its application for $J'_{K-1,n}(X_n)$.

$$
\begin{aligned}
J'_{K-1,n}(X_n) &= \mathbb{E}\Big\{ [\psi_{K-1,n} + J'_{K,n}(X_n)] f^{\theta_{K-1,n}}(X_n) \\
&\quad + [c_{K-1,n}(X_n) + J^*_{K-1,n+1}(X_{n+1})](1 - f^{\theta_{K-1,n}}(X_n)) \Big\} \\
&\le \mathbb{E}\Big\{ [\psi_{K-1,n} + J^*_{K,n}(X_n) + \frac{\varepsilon}{K}] f^{\theta_{K-1,n}}(X_n) \\
&\quad + [c_{K-1,n}(X_n) + J^*_{K-1,n+1}(X_{n+1})](1 - f^{\theta_{K-1,n}}(X_n)) \Big\}.
\end{aligned}
$$

Given the structure of $f^{\theta_{K-1,n}}(X_n)$ where outputs fall in the set $\{0, 1\}$, it is clear that $f^{\theta_{K-1,n}}(X_n) \le 1$. This property lets us further bind our function:

$$
\begin{aligned}
J'_{K-1,n}(X_n) &\le \mathbb{E}\Big\{ [\psi_{K-1,n} + J^*_{K,n}(X_n) + \frac{\varepsilon}{K}] f^{\theta_{K-1,n}}(X_n) \\
&\quad + [c_{K-1,n}(X_n) + J^*_{K-1,n+1}(X_{n+1})](1 - f^{\theta_{K-1,n}}(X_n)) \Big\} \\
&\le \mathbb{E}\Big\{ [\psi_{K-1,n} + J^*_{K,n}(X_n)] f^{\theta_{K-1,n}}(X_n) \\
&\quad + [c_{K-1,n}(X_n) + J^*_{K-1,n+1}(X_{n+1})](1 - f^{\theta_{K-1,n}}(X_n)) \Big\} + \frac{\varepsilon}{K} \\
&\le \min_{a \in \{0,1\}} \mathbb{E}\Big\{ [\psi_{K-1,n} + J^*_{K,n}(X_n)] a \\
&\quad + [c_{K-1,n}(X_n) + J^*_{K-1,n+1}(X_{n+1})](1 - a) \Big\} + \frac{2\varepsilon}{K} \\
&= J^*_{K-1,n}(X_n) + \frac{2\varepsilon}{K}.
\end{aligned}
$$

By recursively extending this logic, we can extrapolate the results to the optimal total cost. Specifically, when considering the estimation $J^*_{1,n}(X_n) = J^*_n(X_n)$, we observe:

$$
\begin{aligned}
J'_{1,n}(X_n) &= \mathbb{E}\Big\{ [\psi_{1,n} + J'_{2,n}(X_n)] f^{\theta_{1,n}}(X_n) \\
&\quad + [c_{1,n}(X_n) + J^*_{1,n+1}(X_{n+1})](1 - f^{\theta_{1,n}}(X_n)) \Big\} \\
&\le J^*_{1,n}(X_n) + \sum_{i=1}^{K} \frac{\varepsilon}{K} \\
&\le J^*_{1,n}(X_n) + \varepsilon = J^*_n(X_n) + \varepsilon.
\end{aligned}
\tag{B1}
$$

In sum, this theorem and its proof suggest that HNNs, with hyperparameters tuned from backward propagation, can be instrumental in making nearly optimal shelter opening time decisions. These HNNs are structured sequentially in the same order of the multi-class ESOT setup.

Acknowledgements This work was supported by the National Science Foundation: awards CMMI-2308750 and CMMI-2045744.

Data Availability Statement The raw/processed data required to reproduce the above findings cannot be shared at this time as the data also forms part of an ongoing study. We will communicate with administrated entities to seek shareable format.

References

1. NCEI, N.: US billion-dollar weather and climate disasters. https://www.ncdc.noaa.gov/billions (2018). Accessed 30 Sept 2023
2. O'Toole, M., Hasan, S., et al.: Catastrophic events impacting transportation infrastructure: understanding funding and risk management approaches. Transport Res. Internat. Doc. (2022). https://trid.trb.org/View/2051172
3. Bhatia, K.T., Vecchi, G.A., Knutson, T.R., Murakami, H., Kossin, J., Dixon, K.W., Whitlock, C.E.: Recent increases in tropical cyclone intensification rates. Nat. Commun. **10**(1), 635 (2019)
4. Wu, L., Zhao, H., Wang, C., Cao, J., Liang, J.: Understanding of the effect of climate change on tropical cyclone intensity: a review. Adv. Atmos. Sci. **39**(2), 205–221 (2022)
5. Kossin, J.P., Olander, T.L., Knapp, K.R.: Trend analysis with a new global record of tropical cyclone intensity. J. Clim. **26**(24), 9960–9976 (2013)
6. Hasan, S., Mesa-Arango, R., Ukkusuri, S.: A random-parameter hazard-based model to understand household evacuation timing behavior. Transp. Res. Part C: Emerg. Technol. **27**, 108–116 (2013)
7. Campos, V., Bandeira, R., Bandeira, A.: A method for evacuation route planning in disaster situations. Procedia Soc. Behav. Sci. **54**, 503–512 (2012)
8. Bayram, V., Yaman, H.: Shelter location and evacuation route assignment under uncertainty: a benders decomposition approach. Transp. Sci. **52**(2), 416–436 (2018)
9. Chakravarty, A.K.: Humanitarian response to hurricane disasters: coordinating flood-risk mitigation with fundraising and relief operations. Nav. Res. Logist. **65**(3), 275–288 (2018)
10. Zhang, J., Liu, Y., Yu, G., Shen, Z.-J.: Robustifying humanitarian relief systems against travel time uncertainty. Nav. Res. Logist. **68**(7), 871–885 (2021)
11. Sanusi, F., Choi, J., Ulak, M.B., Ozguven, E.E., Abichou, T.: Metadata-based analysis of physical–social-civic systems to develop the knowledge base for hurricane shelter planning. J. Manag. Eng. **36**(5), 04020041 (2020)
12. Dash, N., Gladwin, H.: Evacuation decision making and behavioral responses: individual and household. Nat. Hazard. Rev. **8**(3), 69–77 (2007)
13. Roy, K.C., Hasan, S., Sadri, A.M., Cebrian, M.: Understanding the efficiency of social media based crisis communication during hurricane sandy. Int. J. Inf. Manag. **52**, 102060 (2020)
14. NOAA: Hurricane forecasting. https://www.noaa.gov/explainers/hurricane-forecasting (2016). Accessed 30 Sept 2023
15. Baker, E.J.: Hurricane evacuation behavior. Int. J. Mass Emerg. Disasters **9**(2), 287–310 (1991)
16. Huang, S.-K., Lindell, M.K., Prater, C.S.: Who leaves and who stays? A review and statistical meta-analysis of hurricane evacuation studies. Environ. Behav. **48**(8), 991–1029 (2016)
17. Whitehead, J.C., Edwards, B., Van Willigen, M., Maiolo, J.R., Wilson, K., Smith, K.T.: Heading for higher ground: factors affecting real and hypothetical hurricane evacuation behavior. Global Environ. Change Part B: Environ. Hazards **2**(4), 133–142 (2000)
18. Lindell, M.K., Lu, J.-C., Prater, C.S.: Household decision making and evacuation in response to Hurricane Lili. Nat. Hazard. Rev. **6**(4), 171–179 (2005)
19. NPR: As Ian's death toll rises, questions swirl on why more Floridians didn't evacuate. https://www.npr.org/2022/10/08/1127501943/hurricane-ian-florida-delayed-evacuations-lee-county (2022). Accessed 30 Sept 2023
20. Yin, W., Murray-Tuite, P., Ukkusuri, S.V., Gladwin, H.: An agent-based modeling system for travel demand simulation for hurricane evacuation. Transp. Res. Part C: Emerg. Technol. **42**, 44–59 (2014)
21. Rambha, T., Nozick, L.K., Davidson, R.: Modeling hurricane evacuation behavior using a dynamic discrete choice framework. Transp. Res. Part B: Methodol. **150**, 75–100 (2021)
22. Sarwar, M.T., Anastasopoulos, P.C., Ukkusuri, S.V., Murray-Tuite, P., Mannering, F.L.: A statistical analysis of the dynamics of household hurricane-evacuation decisions. Transportation **45**, 51–70 (2018)

23. Serulle, N.U., Cirillo, C.: The optimal time to evacuate: a behavioral dynamic model on Louisiana resident data. Transp. Res. Part B Methodol. **106**, 447–463 (2017)
24. Rahman, R., Hasan, S.: Short-term traffic speed prediction for freeways during hurricane evacuation: a deep learning approach. In: 2018 21st International Conference on Intelligent Transportation Systems (ITSC), pp. 1291–1296 (2018). IEEE
25. Yu, M., Huang, Q., Qin, H., Scheele, C., Yang, C.: Deep learning for real-time social media text classification for situation awareness-using Hurricanes Sandy, Harvey, and Irma as case studies. Int. J. Digit. Earth **12**(11), 1230–1247 (2019)
26. Becker, S., Cheridito, P., Jentzen, A.: Deep optimal stopping. J. Mach. Learn. Res. **20**(1), 2712–2736 (2019)
27. Becker, S., Cheridito, P., Jentzen, A., Welti, T.: Solving high-dimensional optimal stopping problems using deep learning. Eur. J. Appl. Math. **32**(3), 470–514 (2021)
28. FEMA: Hurricane Decision Support. https://www.hurrevac.com/ (2022). Accessed 30 Sept. 2023
29. FEMA-hazus: Hazus Hurricane Model User Guidance. https://www.fema.gov/sites/default/files/2020-09/fema_hazus_hurricane_user-guidance_4.2.pdf (2018). Accessed 30 Sept. 2023
30. NOAA: GIS data and products. https://www.nhc.noaa.gov/gis/ (2023). Accessed 30 Sept 2023
31. Cangialosi, J.P., Blake, E., DeMaria, M., Penny, A., Latto, A., Rappaport, E., Tallapragada, V.: Recent progress in tropical cyclone intensity forecasting at the National Hurricane Center. Weather Forecast **35**(5), 1913–1922 (2020)
32. Roy, K.C., Hasan, S., Culotta, A., Eluru, N.: Predicting traffic demand during hurricane evacuation using Real-time data from transportation systems and social media. Transp. Res. Part C Emerg. Technol. **131**, 103339 (2021)
33. DHEC: Sheltering During Disasters. https://scdhec.gov/disaster-preparedness/hurricanes-floods/sheltering-during-disasters (2022). Accessed 30 Sept 2023
34. Vickery, P.J., Wadhera, D., Twisdale, L.A., Jr., Lavelle, F.M.: Us hurricane wind speed risk and uncertainty. J. Struct. Eng. **135**(3), 301–320 (2009)
35. Si, J., Barto, A.G., Powell, W.B., Wunsch, D.: Handbook of Learning and Approximate Dynamic Programming, vol. 2. Wiley, Hoboken (2004)
36. NOAA: Hurricane Florence. https://www.nhc.noaa.gov/data/tcr/AL062018_Florence.pdf (2018). Accessed 30 Sept 2023
37. SCEMD: South Carolina Hurricane Plan. https://www.scemd.org/media/1323/annex-h-general-population-shelter-management.pdf (2022). Accessed 30 Sept 2023
38. Whitehead, J.C.: One million dollars per mile? The opportunity costs of hurricane evacuation. Ocean Coast. Manag. **46**(11–12), 1069–1083 (2003)
39. Xu, K., Davidson, R.A., Nozick, L.K., Wachtendorf, T., DeYoung, S.E.: Hurricane evacuation demand models with a focus on use for prediction in future events. Transp. Res. Part A Policy Pract. **87**, 90–101 (2016)
40. Zhu, Y., Xie, K., Ozbay, K., Yang, H.: Hurricane evacuation modeling using behavior models and scenario-driven agent-based simulations. Procedia Comput. Sci. **130**, 836–843 (2018)
41. Liaw, R., Liang, E., Nishihara, R., Moritz, P., Gonzalez, J.E., Stoica, I.: Tune: a research platform for distributed model selection and training. arXiv preprint arXiv:1807.05118 (2018)
42. Kingma, D.P., Ba, J.: Adam: a method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
43. SCEMD: Gov. Henry McMaster Orders Mandatory Evacuations for Coastal Counties Effective Tomorrow, September 11 at Noon. https://tinyurl.com/2p8avsrh (2018). Accessed 30 Sept 2023
44. Islandpacket: Beaufort County residents grabbed thousands of sandbags ahead of Florence. Now what? https://www.islandpacket.com/news/weather/hurricane/article218561720.html (2018). Accessed 30 Sept 2023
45. Ciocan, D.F., Mišić, V.V.: Interpretable optimal stopping. Manag. Sci. **68**(3), 1616–1638 (2022)

Authors and Affiliations

Hanwen Liu[1] · Qi Luo[2] · Yongjia Song[1]

✉ Qi Luo
qi-luo-1@uiowa.edu

Hanwen Liu
hanwen@clemson.edu

Yongjia Song
yongjis@clemson.edu

[1] Department of Industrial Engineering, Clemson University, 211 Fernow Street, Clemson, SC 29634, USA

[2] Department of Business Analytics, University of Iowa, W322 John Pappajohn Business Building, Iowa City, IA 52242, USA