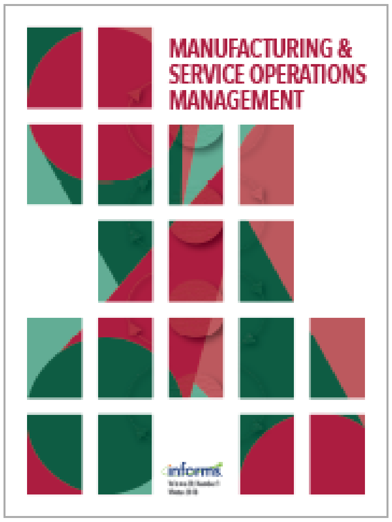

Manufacturing & Service Operations Management

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Shortening Emergency Medical Response Time with Joint Operations of Uncrewed Aerial Vehicles with Ambulances

Xiaoquan Gao, Nan Kong, Paul Griffin

MANUFACTURING & SERVICE OPERATIONS MANAGEMENT
Articles in Advance, pp. 1–18
https://pubsonline.informs.org/journal/msom
ISSN 1523-4614 (print), ISSN 1526-5498 (online)

Shortening Emergency Medical Response Time with Joint Operations of Uncrewed Aerial Vehicles with Ambulances

Xiaoquan Gao,[a] Nan Kong,[b,*] Paul Griffin[c,d]

[a] School of Industrial Engineering, Purdue University, West Lafayette, Indiana 47097; [b] Weldon School of Biomedical Engineering, Purdue University, West Lafayette, Indiana 47907; [c] Department of Industrial and Manufacturing Engineering, Pennsylvania State University, University Park, Pennsylvania 16802; [d] Consortium on Substance Use and Addiction, SSRI, Pennsylvania State University, University Park, Pennsylvania 16802

*Corresponding author

Contact: gao568@purdue.edu, https://orcid.org/0000-0002-8746-0303 (XG); nkong@purdue.edu, https://orcid.org/0000-0002-4047-3414 (NK); pmg14@psu.edu, https://orcid.org/0000-0002-0206-4775 (PG)

Received: April 11, 2022
Revised: January 21, 2023; July 12, 2023; September 23, 2023
Accepted: October 5, 2023
Published Online in Articles in Advance: December 7, 2023

https://doi.org/10.1287/msom.2022.0166

Abstract. ***Problem definition*:** Uncrewed aerial vehicles (UAVs) are transforming emergency service logistics applications across sectors, offering easy deployment and rapid response. In the context of emergency medical services (EMS), UAVs have the potential to augment ambulances by leveraging bystander assistance, thereby reducing response times for delivering urgent medical interventions and improving EMS outcomes. Notably, the use of UAVs for opioid overdose cases is particularly promising as it addresses the challenges faced by ambulances in delivering timely medication. This study aims to optimize the integration of UAVs and bystanders into EMS in order to minimize average response times for overdose interventions. ***Methodology/results*:** We formulate the joint operation of UAVs with ambulances through a Markov decision process that captures random emergency vehicle travel times and bystander availability. We apply an approximate dynamic programming approach to mitigate the solution challenges from high-dimensional state variables and complex decisions through a neural network-based approximation of the value functions (NN-API). To design the approximation, we construct a set of basis functions based on queueing and geographic properties of the UAV-augmented EMS system. ***Managerial implications*:** The simulation results suggest that our NN-API policy tends to outperform several noteworthy rule- and optimization-based benchmark policies in terms of accumulated rewards, particularly for situations that are primarily characterized by high request arrival rates and a limited number of available ambulances and UAVs. The results also demonstrate the benefits of incorporating UAVs into the EMS system and the effectiveness of an intelligent real-time operations strategy in addressing capacity shortages, which are often a problem in rural areas of the United States. Additionally, the results provide insights into specific contributions of each dispatching or redeployment strategy to overall performance improvement.

Funding: This work was supported by the National Science [Grant 1761022].
Supplemental Material: The online appendix is available at https://doi.org/10.1287/msom.2022.0166

Keywords: approximate dynamic programming • large-scale MDP • uncrewed aerial vehicles • dispatching • simulation

1. Introduction

1.1. Motivation

The United States is facing a severe opioid crisis, with over 210 opioid overdose deaths reported each day (Centers for Disease Control and Prevention National Center for Health Statistics 2022). Opioid overdoses can lead to respiratory depression and cardiac arrest, and without timely intervention, the chances of survival decrease by up to 10% per minute (Cao 2005). Brain damage can occur after four minutes, and death can occur within six to eight minutes later (Doe-Simkins et al. 2009). In many cases, trained first responders are unable to reach the patient in time to administer naloxone, typically administered as a nasal spray, and provide ventilation to prevent death. The bystander-enabled uncrewed aerial vehicle (UAV) delivery system is one potential approach to mitigate this problem. In such a UAV-augmented emergency medical service (EMS) system, 911 dispatchers can dispatch a UAV and direct a bystander to render an emergency response to the patient while EMS personnel are en route.

UAV use, as pilotless aircraft, has seen a rapid expansion of applications in recent years. In the United States, this is largely facilitated by the increasingly specified guidelines and relaxed regulations of the U.S. Federal Aviation Administration (FAA) on UAV airspace and operations. For example, Amazon Prime Air, Amazon's special service that delivers packages within 30 minutes, was granted operation by the FAA in mid-2020 to test order delivery via UAVs (Reuters 2020). In addition to parcel delivery, UAV use has a wide range of applications, including land surveillance, wildlife tracking, search and rescue operations, disaster response, and border patrol (Everaerts et al. 2008). UAVs are particularly well suited for these tasks because of their ability to cover large areas quickly and efficiently, as well as their ability to access areas that may be difficult or dangerous for humans to reach. Since 2016, a number of east and central African countries have collaborated with Zipline, the world's largest automated delivery system designer, manufacturer, and operator, to deliver blood supplies, reducing the delivery time from four hours to 15 minutes in some cases (World Health Organization 2019). With state-of-the-art technologies, UAVs designed for on-demand commodity delivery can fly up to an hour and reach distances of up to 45 miles while carrying necessary payloads for emergency responses, such as automated external defibrillators (AEDs) for out-of-hospital cardiac arrests (Boutilier et al. 2017), blood transfusion tool kits for trauma injuries (Ling and Draghic 2019), and naloxone nasal spray for opioid overdoses (Ornato et al. 2020). All of these recent developments, including the improvement of technology, decreases in cost, and changing regulations, will enhance the potential use of UAVs in EMS delivery.

This technology is especially promising given the inherent need for a rapid response to enhance patient outcomes, particularly in remote or hard-to-reach areas. For example, UAVs can respond faster in urban environments that present barriers to emergency services, such as heavy traffic congestion. Additionally, the scarcity of ambulances in many U.S. rural counties can lead to extremely long response times, with 1 in 10 patients waiting nearly 30 minutes for EMS arrival (Mell et al. 2017). Compared with increasing the number of ambulances and corresponding medical personnel, incorporating UAVs for medical delivery is a more realistic and cost-effective approach to optimize medical resources. For emergencies such as opioid overdoses and out-of-hospital cardiac arrests, it is expected that UAV-delivered medical interventions can significantly save critical response times and avert life-threatening conditions. In practice, medical UAVs will be equipped with audio or video assistance devices, such as cameras, to help bystanders quickly assess the situation and follow instructions. For example, Zipline and Intermountain Healthcare have implemented drone deliveries in the Salt Lake Valley to reach patients and deliver medication faster without requiring patients to travel to a clinic or hospital (Gereau 2022). In the realm of EMS delivery, there have been ongoing efforts to encourage bystander intervention, such as educational campaigns (Lockey et al. 2021), training programs (Clark et al. 2014), and liability protection laws (Latimore and Bergstein 2017). More recently, mobile phone applications (e.g., UnityPhilly) (Schwartz et al. 2020) have focused on connecting layperson first responders with people experiencing overdoses. Furthermore, research has demonstrated the benefit of layperson-initiated overdose reversal through the administration of naloxone before the arrival of an ambulance or first responders (Schwartz et al. 2020).

The main objectives of our work are twofold: first, to demonstrate the benefits of using UAVs to deliver lifesaving medication and second, to provide guidance for EMS agencies on how to incorporate UAVs into their operations. By highlighting the potential advantages of using UAVs, we hope to encourage their wider adoption in the EMS field. At the same time, our work aims to provide practical guidance for how to effectively incorporate UAVs into EMS operations, including strategies for dispatching and redeployment.

In the United States, 911 calls are typically received by a public safety answering point (PSAP), which determines the nature of the emergency and next, decides either to dispatch responders immediately or transfer the call to a specialized secondary PSAP. We assume that both UAVs and ambulances are managed in a centralized scheme and aim to improve EMS outcomes through the joint operation of both types of vehicles, including dispatching and redeployment. We develop a Markov decision processes (MDP) framework to capture the interplay between these decisions, spatially distributed stochastic arrivals of requests, and the state of the EMS system involving the concurrent use of UAVs and ambulances. The process flow is shown in Figure 1. When a witness calls 911 to report a case, a request enters the EMS system. During the call, the dispatcher will ask about the presence of bystanders nearby in addition to following the current protocol. Based on the status of the system, a dispatching decision is made, which may include the dispatch of an ambulance or a UAV followed by an ambulance and the specific ambulance (and UAV) to be dispatched. When the UAVs or ambulances have completed their service, they are redeployed to one of their bases to be better prepared for upcoming requests. Note that UAV dispatching is helpful if and only if at least one bystander is willing to help when the UAV arrives at the emergency scene. The probability that at least one bystander is willing to help is estimated based on the number of bystanders acquired during the call and the likelihood that one of them is willing to retrieve and administer the EMS tool kit.

The most straightforward dispatching strategy is to dispatch the UAV closest to the request in order to minimize the response time and maximize the chance of survival (Kim et al. 2009). However, this myopic strategy is typically suboptimal, and a more sophisticated strategy would improve outcomes (Jagtenberg et al. 2015). Because ambulance and UAV resources are limited, minimization of response time for the current requests may

Figure 1. (Color online) An Illustration of a UAV-Based EMS Process Timeline

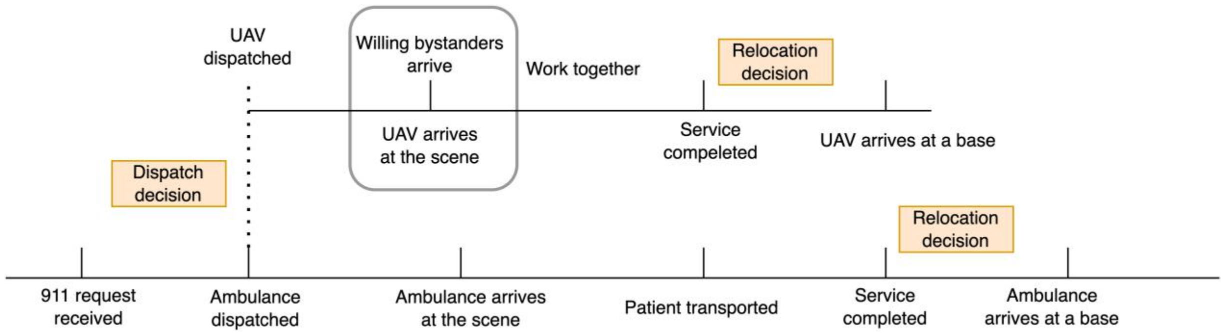

lead to a much longer response time for future requests (i.e., a nuanced balance should be considered between current and future requests). Additionally, we consider redeployment decisions for UAVs and ambulances. Redeploying UAVs and ambulances to one of their bases allows for charging, replenishment, and personnel rest. Proper redeployment decisions aim to balance the distribution of available UAVs and ambulances over the bases for time-varying demand. On top of this delicate tradeoff, uncertainties in travel time, bystander willingness to assist, and emergency request volume and their locations should be modeled explicitly in the MDP framework.

In summary, the integration of UAVs and ambulances brings increased flexibility to EMS delivery but also introduces unique challenges in making intelligent operational decisions in real time. The operational strategy/policy suggested by our model and algorithm addresses the fundamental trade-off between response times for current and future requests while taking into account various uncertainties and the use of heterogeneous servers (ambulances and UAVs).

1.2. Main Contributions and Results

1. We extend the MDP-based analytics framework to consider a range of real-time operational decisions in a dynamic, coordinated logistics system with both conventional (ambulances) and augmenting (UAVs) delivery vehicles. Given the two types of delivery vehicles, we consider two types of requests that can be answered in the logistic system, differing in the use of UAVs. Our approach also accounts for additional sources of uncertainty. In addition to uncertain request arrival times and locations, we use a delayed reward function to reflect the uncertainty of UAV and ambulance travel times. Yet, we preserve the model's high fidelity, including the presence of bystanders.

2. We apply an approximate dynamic programming (ADP) approach and design a tractable approximate policy iteration (API) algorithm for the complex stochastic dynamic optimization problem, which employs value function approximation (VFA) via neural networks. Based on the spatial and temporal characteristics of the EMS system, we design a set of effective basis functions to enhance the algorithm's performance. Our basis functions are novel and distinct from the literature in three aspects. (a) We consider the heterogeneous nature of the ambulance-UAV system and approximate coverage accordingly, especially on the future missed call rate; (b) we approximate the average response time to adapt to the objective of minimizing response times, beyond maximizing the number of requests served; and (c) we consider not only spatial features but also temporal features of the system (i.e., availability reduction). Additionally, the neural network representation of the value function offers a much richer class of nonlinear functions and can be trained iteratively.

3. We acquire important managerial implications about real-time operations for coordinated EMS logistics. Our case studies are based on historical data from the state of Indiana for emergency naloxone administration for opioid overdoses and the National EMS Information System (NEMSIS) database. Simulation results suggest that our policy consistently outperforms several noteworthy rule- and optimization-based benchmark policies in terms of accumulated rewards. This superiority is particularly pronounced when the request arrival rate is high and the availability of ambulances and UAVs is limited. Our results also highlight the benefits of using UAVs in the EMS system and the effectiveness of an intelligent real-time operations strategy in addressing capacity shortages, which are common challenges in rural areas of the United States. Additionally, our results provide valuable insights into the contributions made by each dispatching or redeployment strategy. We also identify scenarios where the use of the neural network-based approximation of the value functions (NN-API) is strongly recommended, as it significantly outperforms benchmark policies, and situations where a simple static policy can perform comparably to the NN-API.

2. Literature Review

To the best of our knowledge, our work is the first that combines the real-time operations of UAVs and ambulances for time-sensitive logistics. Previous research has primarily focused on individual aspects of the problem, such as ambulance dispatching and redeployment using static and dynamic policies, optimization models for vehicle mix in EMS, and the use of UAVs for emergency response. Our research builds upon these efforts by also considering the coordination optimization between UAVs and ambulances.

There is a wealth of literature on ambulance operations management, with many studies focusing on minimizing average response time, as in our research. Early work in this area primarily utilized static policies, including threshold-based policies and policies derived through integer programming (IP) or mixed IP. For example, Daskin (1983) and Marianov and ReVelle (1996) used an IP-based approach and generalized the maximal expected covering location problem (MEXCLP) for public service facility location analysis. Alternatively, the threshold-based policy is also widely used in ambulance dispatching and redeployment. One representation of the threshold policy is the "preparedness" measure proposed by Andersson and Värbrand (2007), which is used to evaluate the ability of an EMS system to serve potential patients.

In the EMS system, patient numbers are highly uncertain, and so, preplanned scheduling or operation solutions may not optimally respond to fluctuating situations. Therefore, real-time decision making is required, which considers systems dynamics, such as time-varying demand (emergency calls), time-varying traffic, and different intervention times required by patients. As a result, several researchers have explored the benefit of dynamic dispatching and redeployment optimization using assumptions such as exponential service time and no-buffer request queue. McLay and Mayorga (2013a, b) and Jagtenberg et al. (2017) built several MDP models and solved them to optimality for small-scale instances. These exactly solved MDPs highlighted the value and suboptimality of the closest idle dispatching policy and how various equity formulations affect the underlying dispatching policies. Recent advances in ADP have improved our ability to solve large-scale problems efficiently. For example, Schmid (2012) and Jenkins et al. (2020) approximated value functions with tabular ADP. Among previous work using ADP, our work is most closely related to Maxwell et al. (2010) and Nasrollahzadeh et al. (2018), which proposed novel basis functions based on the underlying problem structure to approximate value functions. Our work differs in that we consider the joint operation of two delivery modes for EMS logistics. The joint operation requires the consideration of additional novel basis functions for heterogeneous service providers. We refer to two review papers—Aringhieri et al. (2017) and Bélanger et al. (2019)—for comprehensive discussions on optimizing location, redeployment, and dispatching decisions for emergency medical vehicles.

Another stream of literature relates to vehicle mix and response to multiple requests for EMS. Similar to our work, most of these papers differentiate vehicles by their service capability for different types of patients. Previous papers consider multiple responses in the context of deterministic and probabilistic maximal covering ambulance location problems (Schilling et al. 1979, ReVelle and Marianov 1991). McLay (2009) proposed the MEXCLP with two types of servers to efficiently deploy two types of medical units (i.e., advanced life support (ALS) and basic life support (BLS)) to serve multiple types of customers. For UAV-ambulance coordination, Shin et al. (2022) developed a modeling framework to optimize a network of drones, bystanders, and ambulances for cardiac arrest response, taking into account the availability of bystanders. In addition to location problems, researchers have also worked on dispatching and redeployment problems involving multiple types of ambulances with stochastic programming (SP) and MDP. Boujemaa et al. (2020) addressed the ambulance redeployment planning problem in a two-tiered EMS using a two-stage SP model, with the first stage addressing redeployment decisions and the second stage addressing dispatching decisions. Similarly, Yoon et al. (2021) formulated a two-stage SP problem for location and dispatching decisions considering prioritized emergency patients and also extended the model to incorporate nontransport vehicles, similar to UAVs in our model. For real-time operations, Chong et al. (2016) and Yoon and Albert (2020, 2021) constructed MDP models to optimize the dispatching of multiple types of vehicles to (prioritized) patients, demonstrating structural properties. For example, the optimal policy is a control-limit policy, which is more likely to send an ALS unit to calls when more ALS units are available. However, to make the MDP tractable, the authors only decided whether to dispatch an ALS or a BLS but not which specific unit to dispatch. With ADP-based solution techniques, we are able to model and improve upon dispatching and redeployment decisions with greater specificity, including the specific UAV and ambulance to dispatch and the specific base to redeploy the UAV or ambulance to.

Unlike previous work on vehicle mix, we model the information of bystanders for UAV dispatch, which is unique. Additionally, to our best knowledge, among studies of multiple types of ambulances, only Park and Lee (2019) considered the real-time dispatching that makes specific dispatching and redeployment decisions. The authors leveraged ADP with state aggregation and monotonicity-preserving projection operators to solve the complex MDP model. In our work, we maintain a high level of fidelity without using state aggregation and address the curse of dimensionality through neural network-based VFA. We capture differences between UAVs and ambulances by using different transition dynamics and integrating them into the VFA using a queueing model of heterogeneous servers.

Recent studies have explored the feasibility of delivering medical equipment via UAVs, including flotation devices (Claesson et al. 2017b), AEDs (Boutilier et al. 2017, Claesson et al. 2017a), and blood products (Amukele et al. 2017). In addition to strategic decisions, Chu et al. (2021) developed UAV dispatching rules based on the difference between predicted ambulance response time and calculated UAV response time for each out-of-hospital cardiac arrest. There are two categories of existing literature that investigate the joint optimization of UAV planning and operations management. The first category concerns situations where UAVs are carried by truck and dispatched from the truck near the service location. Typical applications include precision agriculture, package delivery, oceanographic sampling, forest fire, or oil spill monitoring (Tokekar et al. 2016, Fawaz et al. 2017, Jia and Zhang 2017). The second category of studies considers cases where UAVs and ground vehicles perform independent tasks, similar to the situation in our work. However, most previous studies in this category only consider facility locations at the strategic level and UAV allocation at the tactical level (Dorling et al. 2016, Agatz et al. 2018). The studies most similar to ours are Ulmer and Thomas (2018) and Chen et al. (2019). Ulmer and Thomas (2018) explored the addition of UAVs to conventional vehicles for the same-day delivery problem. Using an MDP model, the authors presented a dynamic vehicle routing problem with heterogeneous fleets, where the decisions were to reject an order and assign a UAV or a ground vehicle. To address the curse of dimensionality, the authors adopted policy function approximation based on the insight that distant customers should generally be served by UAVs and that closer customers should be served by conventional vehicles. Chen et al. (2019) extended the work to include additional information on resource availability and demand as well and implemented a deep $Q-$learning. However, only the acceptance and general assignment decisions were made in both studies (i.e., the order/request should either be served by UAVs or be served by conventional vehicles). An important difference in our work is that we seek to optimize the dispatching and redeployment decisions of UAVs and ambulances with potentially multiple responses to each EMS request. This difference in the decisions significantly expands our MDP model, making it very difficult to parameterize the policy function directly and use the value function of state-action pairs. Additionally, our application emphasizes the "time criticality" of the service, with emergency response time being the key objective. Therefore, the objective function is different, and different basis functions are required to approximate the value function.

3. MDP Model

This section presents an infinite-horizon average-cost MDP formulation. We adopt event-driven modeling to incorporate on-demand UAV and ambulance dispatching and redeployment decisions in the model and capture the EMS system evolution. Events are triggered by changes in the status of UAVs, ambulances, and requests. Let $\mathcal{N} := \{1,2,\dots,N\}$ be the set of demand nodes, $\mathcal{M}^u := \{1,2,\dots,M^u\}$ be the set of UAV charging stations, and $\mathcal{M}^a := \{1,2,\dots,M^a\}$ be the set of ambulance bases (e.g., hospitals or bases of private EMS agencies). We consider a total of L^u UAVs and L^a ambulances. The home base of UAV l is denoted by $h_l^u \in \{1,2,\dots,M^u\}$, and the home base of ambulance l is denoted by $h_l^a \in \{1,2,\dots,M^a\}$. Let $\mathcal{A}^u(s)$ be the set of available UAVs (i.e., $\mathcal{A}^u(s) := \{l : r_l^u = 0\}$, where r_l^u is the remaining time in its current status of UAV l); similarly, let $\mathcal{A}^a(s)$ be the set of available ambulances (i.e., $\mathcal{A}^a(s) := \{l : r_l^a = 0\}$, where r_l^a is the remaining service time in its current status of ambulance l).

We divide 911 requests into two types based on whether UAVs can be of help. (1) Type 1 includes requests for which UAVs can serve as the first response, such as opioid overdose and out-of-hospital cardiac arrest. (2) Type 2 includes requests that only ambulances can help, such as massive hemorrhage. We assume that arrivals of EMS requests of type 1 and type 2 follow Poisson processes with rates λ^u and λ^a, respectively. We make dispatching decisions for both types of requests. For type 1 requests, we need to choose between a single-ambulance response and a UAV-ambulance sequence response. For type 2 requests, we only consider which ambulance to dispatch. We also make redeployment decisions each time a UAV or ambulance finishes the current emergency response task and returns to a base for replenishment, recharging, and personnel rest. In the following model description, we assume that all the information obtained from the 911 call, such as request location, request type, and bystander information, is accurate.

3.1. State Space

The state space is composed of five parts: vectors $\boldsymbol{B^u} = (b_1^u, b_2^u, \dots, b_{L^u}^u)$, $\boldsymbol{B^a} = (b_1^a, b_2^a, \dots, b_{L^a}^a)$, $\boldsymbol{C} = (c_1, c_2, \dots, c_J)$, e, and τ, where b_l^u, $l = 1,\dots,L^u$ contains information about the state of the lth UAV; b_l^a, $l = 1,\dots,L^a$ contains information about the state of the lth ambulance; c_j, $j = 1,\dots,J$ contains information about the jth request; e denotes the event type; and τ corresponds to the current time. Therefore, the state space of the system is represented by $\mathcal{S} := \{s = (\tau, e, B^u, B^a, C)\}$.

The status of UAV l is given by $\boldsymbol{b_l^u} = (d_l^u, r_l^u, f_l^u)$, $l = 1,\dots,L^u$, where $d_l^u \in \{1,2,\dots,N\}$ is the destination of each UAV. For this work, it is sufficient to consider four possibilities for the status of UAVs (i.e., $f_l^u \in \{0,1,2,3\}$, where 0 indicates that the UAV is available at the base, 1 indicates that the UAV is going to a request location, 2 indicates that the UAV is serving a request on the scene, and 3 indicates that the UAV is returning to a base).

The state of ambulance l is given by $\boldsymbol{b}_l^a = (d_l^a, r_l^a, f_l^a)$, $l = 1, \ldots, L^a$, where $d_l^a \in \{1, 2, \ldots, N\}$ is the destination of each ambulance. Assume that all the opioid-overdosed patients require transport to hospitals and that all other patients require transporting to hospitals with probability p^t. For this work, it is sufficient to consider five possibilities for the status of an ambulance (i.e., $f_l^u \in \{0, 1, 2, 3, 4\}$, where 0 indicates that the ambulance is available at the base, 1 indicates that the ambulance is going to a request location, 2 indicates that the ambulance is serving a request on the scene, 3 indicates that the ambulance is going to the hospital, and 4 indicates that the ambulance is returning to a base).

A request j is represented by $\boldsymbol{c}_j = (g_j, q_j, o_j, \omega_j)$, $j = 1, \ldots, J$, where $g_j \in \{1, 2, \ldots, N\}$ is the request location and q_j is the arrival time of the request. For this work, $o_j \in \{1A, 1B, 2\}$ denotes the type and status of the request, where $o_j = 1A$ ($o_j = 1B$) implies that UAVs are qualified for the first response and the request is waiting for the first (follow-up) response and $o_j = 2$ implies that only ambulances are qualified for serving the request. If request j is finished, it would be marked as "served" and immediately removed from the request set. In addition, ω_j represents bystander helping probability at request j. Specifically, $\omega_j = 1 - (1 - p^b)^{N^b}$ is the probability that at least one bystander is willing to help when the UAV arrives, where N^b is the estimated total number of bystanders and p^b is the probability that each bystander is still willing to help when the UAV arrives. The bystander willingness p^b can be estimated through interviews (Lankenau et al. 2013).

An event is represented by $e, e \in E$, where E is the set of all possible event types. Without loss of generality, we assume that decisions are made at transition times. In our model, transition times are associated with the following events: (1) request j arrives; (2) ambulance l is in transit to request j; (3) ambulance l arrives at the location of request j and starts service; (4) ambulance l finishes serving request j at the scene; (5) ambulance l finishes serving request j at a hospital; (6) ambulance l arrives at a base; (7) UAV l is in transit to request j; (8) UAV l arrives at the location of request j and starts service; (9) UAV l finishes serving request j at the scene and is in transit to a base; and (10) UAV l arrives at a base.

3.2. Action Space

We consider a loss system without request queues (i.e., a request will be outsourced to a nearby EMS agency if there are no available UAVs and ambulances in our system). Another option is to place requests in a queue when all servers are busy. Bandara et al. (2014) showed that the strategy performance relationship remains the same for systems allowing and not allowing request queueing, but the overall system performance with queuing is lower because of increased vehicle utilization. Therefore, we consider outsourcing requests that arrive when no ambulance is available. The action space is described with three event-based cases based on the type of actions required.

Case 1. If request j arrives ($e = 1$), the decision maker has three types of decisions: (1) whether to outsource the request; (2) which ambulance to immediately dispatch to serve the request; and (3) which UAV to immediately dispatch and which ambulance to dispatch as a follow-up.

Define $X_{l,j}^u = 1$ if UAV l is dispatched to request j and $X_{l,j}^u = 0$ otherwise. Also, define $X_{l,j}^a = 1$ if ambulance l is dispatched to request j and $X_{l,j}^a = 0$ otherwise. Therefore, if event $e = 1$, the action space is given by

$$A_1(s) := \left\{ (X_{l,j}^u, X_{l,j}^a) : \sum_{l \in \mathcal{A}^a(s)} X_{l,j}^a \le 1, \sum_{l \in \mathcal{A}^u(s)} X_{l,j}^u \le \sum_{l \in \mathcal{A}^a(s)} X_{l,j}^a, X_{l,j}^u, X_{l,j}^a \in \{0, 1\} \right\},$$

where the first constraint states that for each request, at most one ambulance is dispatched and the second constraint states that the number of UAVs dispatched should not exceed the number of ambulances dispatched. The two constraints together limit the dispatching decisions into three types: (1) dispatching one UAV and one ambulance, (2) dispatching only one ambulance, and (3) dispatching no UAV or ambulance (i.e., outsourcing the request).

Case 2. If event $e \in \{5, 9\}$, the decision is to determine to which base to redeploy the ambulance/UAV.

If event $e = 5$ (i.e., an ambulance finishes service), let $Z_{l,b}^a = 1$ if ambulance l is redeployed to base b and $Z_{l,b}^a = 0$ otherwise. Then, the action space is given by

$$A_2(s) := \left\{ (Z_{l,b}^a) : \sum_{b \in \mathcal{M}^a} Z_{l,b}^a = 1 \right\},$$

which ensures that ambulance l is redeployed to only one base.

If event $e = 9$ (i.e., a UAV finishes service), let $Z_{l,b}^u = 1$ if UAV l is redeployed to base b and $Z_{l,b}^u = 0$ otherwise. Then, the action space is given by

$$A_3(s) := \left\{ (Z_{l,b}^u) : \sum_{n \in \mathcal{M}^u} Z_{l,b}^u = 1 \right\},$$

which ensures that UAV l is redeployed to only one base.

Case 3. If the event $e \in \{2, 3, 4, 6, 7, 8\}$, we set $A(s) = \emptyset$ (i.e., no action will be taken).

3.3. Transitions

Let s_k be the state of the system when the kth event happens. The evolution of state s_k can be characterized by action a_k, random element $\omega(s_k, a_k)$, and a function F (i.e., $s_{k+1} = F(s_k, a_k, \omega(s_k, a_k))$). We assume that the on-scene time follows a lognormal distribution (Ingolfsson et al. 2008).

3.4. One-Step Reward Function

We consider maximizing the average health outcome as the primary objective function in our optimization framework. The health outcome is modeled as a decreasing function of the response time because the likelihood of survival decreases with the time it takes to receive medical treatment (Blackwell and Kaufman 2002, Wilde 2013), especially in cases of opioid overdose where every minute is critical. Let $h(s_k, a_k, s_{k+1})$ denote the cost or reward of a transition from s_k to s_{k+1}, when action a_k is taken. The system only incurs a cost or gains a reward at event $e \in \{1, 3, 8\}$.

When $e = 1$ (i.e., request j arrives), a penalty will be incurred if the request is outsourced, specifically

$$h(s_k, a_k, s_{k+1}) = \begin{cases} C_o & \text{if } \sum_{l \in \mathcal{A}^u(s)} X^u_{l,j} = 0, \ \sum_{l \in \mathcal{A}^a(s)} X^a_{l,j} = 0; \\ 0 & \text{otherwise,} \end{cases} \tag{1}$$

where C_o is the penalty for health outcome decrease from delay in response time caused by outsourcing request j.

When $e = 8$ (i.e., a UAV arrives at the scene),

$$h(s_k, a_k, s_{k+1}) = \begin{cases} g_1(\tau - q_j) \cdot \mathbf{1}_{\{N_j \geq 1\}} & \text{if } o_j = 1A; \\ 0 & \text{otherwise,} \end{cases} \tag{2}$$

where $\tau - q_j$ represents the first response time of request j; $g_1(\cdot)$ denotes the reward function with respect to the response time for opioid overdoses (e.g., $g_1(t) = \frac{[T_1 - t]^+}{T_1}$), where T_1 is the response time threshold (e.g., $T_1 = 8$ minutes); and random variable $N_j \sim Binom(N^b, p^b)$ is the number of willing bystanders when a UAV arrives at the scene. If $o_j \neq 1A$, then the UAV arrives later than the ambulance, or the UAV is dispatched for a request of type 2. In this case, the UAV dispatching becomes redundant.

When $e = 3$ (i.e., an ambulance arrives at the scene),

$$h(s_k, a_k, s_{k+1}) = \begin{cases} g_1(\tau - q_j) & \text{if } o_j = 1A; \\ g_2(\tau - q_j) & \text{if } o_j = 2; \\ 0 & \text{if } o_j = 1B, \end{cases} \tag{3}$$

where $g_2(\cdot)$ denotes the reward function with respect to the response time for general (e.g., $g_2(t) = \frac{[T_2 - t]^+}{T_2}$), with T_2 being the response time threshold (e.g., $T_2 = 12$ minutes). Let $o_j = 1A$ and $o_j = 2$ denote the cases where the ambulance serves as the first response; let $o_j = 1B$ denote the case where the ambulance serves as the follow-up response.

3.5. Optimality Criterion

The *expected* average reward value of a policy π is defined for all $s_0 \in \mathcal{S}$ as

$$v_g(\pi, s_0) := \lim_{T \to \infty} \frac{1}{T} \mathbf{E}_{s_t, a_t, s_{t+1}} \left[\sum_{t=0}^{T-1} h(s_t, a_t, s_{t+1}) | s_0, \pi \right], \tag{4}$$

where $h(s_t, a_t, s_{t+1})$ is the one-step cost/reward, which is defined in Section 3.4. The limit in Equation (4) exists for a stationary policy when the MDP is unichained (Puterman 2014, section 8.3.3). Assuming that the Markov chain under the policy π is unichain, we have $v_g(\pi, s_0) = v_g(\pi), \ \forall s_0 \in \mathcal{S}$. Then, the optimal policy π^* is the policy that satisfies the average-reward Bellman *optimality* equation,

$$v_b(\pi^*, s) + v_g(\pi^*) = \max_{a \in \mathcal{A}} \sum_{s' \in \mathcal{S}} p(s'|s, a)[h(s, a, s') + v_b(\pi^*, s')], \quad \forall s \in \mathcal{S}, \tag{5}$$

where $v_b(\cdot)$ is the optimal relative value function, $h(s, a, s')$ is the one-step cost/reward, and $p(s'|s, a)$ is the transition probability. For notational simplicity, we omit π^* and denote the optimal average reward as v_g in the following sections. We refer to Cavazos-Cadena (1991) and Cavazos-Cadena and Sennott (1992) for summaries of results on existence conditions for discrete-time average cost MDPs with countable state space and finite action sets.

4. Approximate Solutions

A conventional method to solve Equation (5) is through policy iteration (PI). The PI algorithm starts with a random policy, computes the value function of that policy (step 1: policy evaluation), and then, determines a new and improved policy based on the previous value function (step 1: policy improvement). These two steps are repeated iteratively until the policy converges. However, to perform steps 1 and 2, it is first necessary to parameterize the associated transition matrix $\mathcal{S} \times \mathcal{A} \times \mathcal{S} \to \mathcal{R}$ and reward matrix $\mathcal{S} \times \mathcal{A} \to \mathcal{R}$. In our MDP model, the state space $|\mathcal{S}|$ is unbounded as the time variable τ is continuous. Even without τ, the dimension of the state space grows exponentially with the number of UAVs and ambulances. This makes it infeasible to store all $v_b(s), s \in \mathcal{S}$, not to mention enumerating the state space to solve the Bellman Equation (5) to optimality. To tackle the curse of dimensionality resulting from the need to enumerate the state space, we conduct a simulation-based API. The main framework of the API algorithm is described in Section 4.1.

To represent value functions of the high-dimensional state space, we approximate the relative value function (i.e., $v_b(s), s \in \mathcal{S}$) with a neural network model of a finite set of basis functions (or features). That is, for each $s \in \mathcal{S}$,

$$v_b(s) \approx Z(\mathbf{\Phi}(\boldsymbol{s})) = Z(\phi_1(s), \phi_2(s), \ldots, \phi_f(s)). \qquad (6)$$

Here, $\mathbf{\Phi}(\boldsymbol{s}) = (\phi_1(s), \phi_2(s), \ldots, \phi_f(s))$ is the set of f basis functions, and $Z(\cdot)$ is a model mapping basis functions to the relative value function. In our algorithm, $Z(\cdot)$ represents a neural network model. Using Equation (6), the approximate relative value function is determined by the neural network model $Z(\cdot)$ and a set of prespecified basis functions, which are described in detail in Sections 4.1 and 4.2, respectively.

4.1. Approximate Policy Iteration

The NN-API algorithm is extended from the basic PI algorithm and adapted in the following four aspects.

4.1.1. Simulation-Based Policy Evaluation. The core concept of approximate dynamic programming is to follow a sample path rather than enumerating the state space to update value functions. In NN-API, sample paths are generated based on the current policy and predefined distributions of randomness. Only rewards from the states visited on the sample paths would be used to update $Z(\cdot)$.

4.1.2. Postdecision State Variable. A postdecision state variable is the state of the system after we have made a decision but before any new information has arrived (Powell 2007). Rather than estimating the expectation of the value around the next predecision state s_{k+1}, we directly estimate $\overline{V}(s_k^a)$ for the postdecision state s_k^a. That is, we make decisions by optimizing $\hat{v}_k^n = \min_{a_k}(C(s_k^n, a_k) + \gamma \overline{V}_k^{n-1}(S^{M,a}(s_k, a_k)))$ instead of estimating the expectation and optimizing $\tilde{v}_k^n = \min_{a_k}(C(s_k^n, a_k) + \gamma \mathbb{E}\{\overline{V}_{k+1}^{n-1}(S^{M,W}(s_k^n, a_k, W_{k+1}))\})$. Here, M and W represent our model and the uncertainty in the model, respectively; $S^{M,a}$ denotes the system state immediately after decision a, and $S^{M,W}$ denotes the system state after the uncertainty W is realized.

4.1.3. Average Reward Computation. For the average reward v_g computation, instead of solving the Poisson Equation (5), NN-API estimates v_g iteratively with $v_g \leftarrow v_g + \beta_g \Delta_g$, where $\beta_g = \frac{1}{n_u+1}$ (n_u is the number of updates of v_g so far) and $\Delta_g = r(s,a) - v_g$.

4.1.4. Value Function Representation. We use a feed-forward neural network to approximate the value function rather than using a tabular form. The neural network consists of three layers: an input layer, a hidden layer, and an output layer. The information provided to the input layer is a set of $|\boldsymbol{\phi}|$ basis functions associated with a postdecision state s^a. The hidden layer consists of a set of nonlinear activation units, and the size of this layer is a tunable parameter. The output layer produces a single scalar output by applying the activation function, which is the final approximation for the value function with respect to the input. For more information on the design and training of the neural network, see section 1 in the online appendix.

The four solution ideas help alleviate the curse of dimensionality and are incorporated into the framework of the Neural Network-based Approximate Policy Iteration (NN-API) (Algorithm 1) by using temporal difference learning.

Algorithm 1 (NN-API: Neural Network-Based Approximate Policy Iteration)

Result: A trained neural network whose input is basis functions and output is an approximate value function.

Construct the basis functions $\mathbf{\Phi} = (\phi_1, \phi_2, \ldots, \phi_f)'$;
Initialize the neural network $Z_0(\mathbf{\Phi})$ using a myopic policy. Initialize the average reward $\hat{v}_g^{1,1} = 0$.
for $n = 1, 2, \ldots, N$ **do**
  Policy evaluation starts.
  Sample an initial state $S^{n,1}$ and choose a sample path ω^n;
  for $m = 1, 2, \ldots, M$ **do**
    Compute $a^{n,m} = \arg\min_{a \in \mathcal{A}^{n,m}}(C(s^{n,m}, a) + Z_{n-1}(\Phi^{n,m}(S^{M,a}(s^{n,m}, a))))$;
    Compute $S^{n,m+1} = S^M(S^{n,m}, a^{n,m}, W_m(\omega^n))$;
    Compute $\hat{v}_g^{n,m+1} = \hat{v}_g^{n,m} + \beta_g(C(s^{n,m}, a) - C(s^{n,m}, a))$, $\beta_g = 1/((n-1)M + m + 1)$;
  end
  Let $\mathbf{\Psi^n}$ be an $M \times F$ matrix where the (i,k)th entry is given by $\phi_k^{n,i}$;
  Let $\hat{\boldsymbol{V}}^{\boldsymbol{n}}$ be a vector of M dimensions with elements, $\hat{v}^{n,m} = C(s^{n,m}, a^{n,m}) - \hat{v}_g^{n,M} + Z_{n-1}(\Phi^{n,m}(S^{M,a^{n,m}}(S^{n,m}, a^{n,m})))$, $m = 1, \ldots, M$;
  Policy evaluation ends.
  Retrain the neural network model with feature $\mathbf{\Psi^n}$ and label $\hat{\boldsymbol{V}}^{\boldsymbol{n}}$. Denote the updated neural network model with Z_n; Policy improvement.
end

4.2. Basis Functions

A key difficulty in the design of value function approximation is to select a set of basis functions that enable us to approximate the downstream costs. Based on the problem property, we conjecture the following basis functions (Sections 4.2.1–4.2.6). We adopt an iterative process of testing and refining to identify effective basis functions for our ADP implementation. Although simulation-based comparison would need to verify the effectiveness of these basis functions, we ensure that the constructed basis functions have the same monotonicity as the optimal value function at the design phase. Specifically, the optimal value function is monotone with bystander helping probability and availability of ambulances/UAVs, as stated in Proposition 1.

Proposition 1. *The optimal value function for the average-reward MDP described in Section 3 has the following properties.*

1. *It increases with an increase in the bystander helping probability w_j given the assumption that the on-scene time of the UAV is insignificant.*
2. *It increases when the system has one additional available UAV or ambulance.*

The proof of Proposition 1 is based on the coupling arguments between the systems under optimal and suboptimal policies; see section 2 in the online appendix for details. To further refine the design of basis functions, we follow Nasrollahzadeh et al. (2018) by approximating the system dynamics by constructing a queueing model. We use an $M/G/c/c$ queue to approximate our system because we do not consider putting 911 requests in the queue as reasonable. Another adaptation is from the fact that we cannot treat UAVs and ambulances as homogeneous servers with the same service time distribution. Thus, based on the approximations provided by Fakinos (1980), we construct basis functions by dealing with the queueing system with heterogeneous servers. In Sections 4.2.1–4.2.6, when mentioning "server," we refer to both UAVs and ambulances for type 1 requests and ambulances for type 2 requests.

4.2.1. Expected Delayed Rewards. When dispatching decisions are made (i.e., $e = 1$), the immediate reward only includes the penalty for outsourcing requests. The reward for serving the request is not realized until a UAV or ambulance arrives at the scene (i.e., $e = 8$ or $e = 3$). To reflect the direct impact of the dispatching decision on the reward, we include the expected reward as one of the basis functions. We denote the expected travel time for UAVs to travel between locations d_1 and d_2 as $\hat{t}^u(d_1, d_2)$ and the expected travel time for ambulances as $\hat{t}^a(d_1, d_2)$.

If both a UAV and an ambulance are dispatched (i.e., $X^u_{l_1,j} = 1, X^a_{l_2,j} = 1$),

$$\phi_1(s) = \begin{cases} \omega_j \cdot g_1(\hat{t}^u(d^u_{l_1}, g_j)) & \text{if } o_j = 1A, \hat{t}^u(d^u_{l_1}, g_j) < \hat{t}^a(d^a_{l_2}, g_j), \\ g_1(\hat{t}^a(d^a_{l_1}, g_j)) & \text{if } o_j = 1A, \hat{t}^u(d^u_{l_1}, g_j) \geq \hat{t}^a(d^a_{l_2}, g_j), \\ g_2(\hat{t}^a(d^a_{l_2}, g_j)) & \text{otherwise,} \end{cases} \tag{7}$$

where $o_j = 1A, \hat{t}^u(d^u_{l_1}, g_j) < \hat{t}^a(d^a_{l_2}, g_j)$ refers to the scenario where UAVs are eligible for the first response and the dispatched UAV is expected to arrive earlier than the dispatched ambulance. Here, $g(\hat{t}^u(d^u_{l_1}, g_j))$ represents the reward from UAV's first response time if there is more than one willing bystander (with probability ω_j); $o_j = 1A, \hat{t}^u(d^u_{l_1}, g_j) \geq \hat{t}^a(d^a_{l_2}, g_j)$ refers to the scenario where UAVs are eligible for the first response, but the dispatched ambulance is expected to be the first response. Otherwise, $o_j = 2$ (i.e., the dispatched ambulance) would serve as the first response.

If only an ambulance is dispatched (i.e., $\sum_l X^u_{l,j} = 0$, $X^a_{l_2,j} = 1$), a delayed reward from ambulance response is expected: that is,

$$\phi_1(s) = \begin{cases} g_1(\hat{t}^a(d^a_{l_2}, g_j)) & \text{if } o_j = 1A, \\ g_2(\hat{t}^a(d^a_{l_2}, g_j)) & \text{if } o_j = 2. \end{cases} \tag{8}$$

If the request is outsourced (i.e., $\sum_l X^u_{l,j} = 0, \sum_l X^a_{l,j} = 0$), the penalty for outsourcing would be incurred immediately when the dispatching decision is made so that there would not be any delayed rewards: that is,

$$\phi_1(s) = 0. \tag{9}$$

4.2.2. Uncovered Request Rate. This basis function captures the rate of request arrivals that cannot be reached within the response time threshold by any of the available servers. Let $\mathcal{A}^u(s)$ and $\mathcal{A}^a(s)$ be the set of available UAVs and the set of ambulances, respectively, when the system state is s. Specifically, $\mathcal{A}^u(s) = \{l | f^u_l = 0\}$, and $\mathcal{A}^a(s) = \{l | f^a_l = 0\}$. Then, the coverage of demand node i can be written as

$$N_i(s) = \sum_{l \in \mathcal{A}^u(s)} \mathbf{1}_{\{d(d^u_l(s), i) \leq \Delta\}} + \sum_{l \in \mathcal{A}^a(s)} \mathbf{1}_{\{d(d^a_l(s), i) \leq \Delta\}}. \tag{10}$$

We can then compute the rate of request arrivals that are not covered by any available servers with

$$\phi_2(s) = \sum_{i \in \mathcal{N}} \lambda_i \mathbf{1}_{\{N_i(s) = 0\}}. \tag{11}$$

Note that a type 1 request is considered as covered either when (i) an ambulance is expected to reach it within time T_A or when (ii) a UAV is expected to reach it within T_A and an ambulance is expected to reach it within T_B.

4.2.3. Future Uncovered Request Rate. When making redeployment decisions, the state where the redeployed server reaches its new base is more important than the current state. This basis function is parallel to the second basis function, but it replaces the current location of redeployed servers by its destination if we are making redeployment decisions (i.e., $e \in \{5, 9\}$). Denote the future state with the redeployed server arriving at the base by s'. Then, the future uncovered rate can be written as

$$\phi_3(s') = \sum_{i \in \mathcal{N}} \lambda_i \mathbf{1}_{\{N_i(s') = 0\}}, \tag{12}$$

where the coverage $N_i(s')$ is defined in the same way as in (10) except that s is replaced by s'. When making dispatching decisions, we do not perform the replacement because the server will still be unavailable when it arrives at its destination. In other words, with the "future uncovered request rate," we would like to maximize the future coverage when the redeployed server becomes available.

4.2.4. Future Missed Call Rate. This basis function captures the rate that a request is outsourced or delayed because all the servers are busy with other requests, which is represented as

$$\phi_4(s) = \sum_{i=1}^{N} \lambda_i P_i(s), \tag{13}$$

where $P_i(s)$ is the probability that all servers that can reach a request at demand node i are busy with other requests. We estimate $\{P_i(s), i = 1, \ldots, N\}$ by treating the request service processes in different demand areas as *Erlang loss systems*. In an Erlang loss system with arrival rate λ, service rate μ, and n servers, the steady-state probability of losing a request is given by $\psi(\lambda, \mu, n) = \frac{(\lambda/\mu)^n/n!}{\sum_{k=0}^{n}(\lambda/\mu)^k/k!}$. In an EMS system, arrivals follow a Poisson distribution, which satisfies the assumptions of Erlang loss systems. Service times include response time (base to the scene), on-scene time, transport time (scene to hospital), and transition time (hospital to base). However, in our model, service time distributions for ambulances and UAVs are not identical. Fakinos (1980) generalized the Erlang B formula for the case of heterogeneous servers. Specifically, for an $M/G/k/k$ blocking system with heterogeneous servers, denote the arrival rate at demand node i with λ_i, average service time $\beta_{i,j}, j = 1, \ldots, k$, and their product $\rho_{i,j} = \lambda_i \beta_j, j = 1, \ldots, k$. Then, the probability of blocking is

$$P_i(s) = B_{i,k}(\rho_{i,1}, \ldots, \rho_{i,k}) = \frac{\frac{1}{k!}\rho_{i,1}\rho_{i,2}\cdots\rho_{i,k}}{\sum_{v=0}^{k}\frac{(k-v)!}{k!}\sum_{j_1<\cdots<j_n}\rho_{i,j_1}\rho_{i,j_2}\cdots\rho_{i,j_k}},$$

where $j_1 < j_2 < \cdots < j_v$ is a permutation of v servers in $\{1, 2, \ldots, k\}$. To estimate parameters in the generalized Erlang B formula, we let $\mathcal{L}_i$ be the set of available servers that can serve a request in each demand node i within the threshold response time so that $\mathcal{L}_i(s) = \{l \in \mathcal{A}(s) : d(d_l^{a/u}, i) \le \Delta\}$. Then, we use $k = |\mathcal{L}_i(s)|$ as the number of servers in the Erlang loss system for demand area i.

4.2.5. Average Response Time. This basis function captures the average response time of the two closest servers to each request. We only count the number of servers satisfying certain conditions in the previous three basis functions. With "average response time," we emphasize the exact distance, directly affecting the response time. Denote by ρ_{li} the probability that server l is dispatched to the request at node i (Chelst and Jarvis 1979). Then, the average travel time to each demand node is

$$T_i = \frac{\sum_l \rho_{li} t_{li}}{\sum_l \rho_{li}},$$

where t_{li} is the average travel time from l to i. The dispatching probability is estimated by the hypercube queueing model developed by Larson (1974), which characterizes the operations of an EMS system with a multiserver-queuing system comprising distinguishable servers. Additionally, we estimate the dispatching probability following the approximation procedures described in Larson (1975). Thus, the demand-weighted average response time can be written as

$$\phi_5(s) = \sum_{i \in \mathcal{N}} \lambda_i T_i = \sum_{i \in \mathcal{N}} \lambda_i \frac{\sum_l \rho_{li} t_{li}}{\sum_l \rho_{li}}. \tag{14}$$

4.2.6. Availability Reduction. The previous five basis functions capture the spatial features of the system, whereas this one captures the temporal aspect of availability reduction associated with the traveling of UAVs and ambulances. The travel time matters because UAVs and ambulances are unavailable to serve requests during the travel. For example, consider a scenario where an ambulance can be redeployed to two bases, A and B. Redeployment to A results in a slightly higher coverage but requires a much longer travel time. Without this basis function, we will choose to redeploy the ambulance to base A, but it may not always be the optimal choice. Therefore, we represent availability reduction caused by the travel time of UAVs and ambulance redeployment with the following basis function:

$$\phi_6(s) = \begin{cases} t(d_l(s), d_l(s')) & e \in \{5, 9\} \text{ and } \sum_{b \in \mathcal{M}^a} Z_{l,b}^a = 1 \text{ or} \\ & \sum_{b \in \mathcal{M}^u} Z_{l,b}^u = 1, \\ 0 & \text{otherwise.} \end{cases} \tag{15}$$

5. Numerical Experiments

In this section, we present a simulation-based comparative study with realistic UAV design parameters and real EMS data from the NEMSIS database and naloxone administration heat map in the state of Indiana. We start by introducing the experimental setup and benchmark policies. Then, we compare the policies' performance in scenarios with varying numbers of UAVs, service areas, and base locations. When we investigate the influence of one factor, everything else is fixed to ensure a fair comparison. Next, in Sections 5.6 and 5.7, we analyze the underlying mechanism behind the superior performance of NN-API and provide insights into situations where a simpler approach, such as the static ad hoc policy, exhibits comparable performance with the NN-API.

5.1. Experimental Setup

In this section, we present an overview of our experimental setup and the calibration criteria employed for the determination of base locations, bystander willingness, EMS demands, performance metrics, and validation of results. For a comprehensive understanding of the estimation and evidential support for other parameters, such as fleet sizes and time parameters, please refer to section 3 in the online appendix. To ensure statistical confidence in our comparisons, we set the simulation time horizon to one day and conducted 400 replications for each case.

5.1.1. Base Locations and Initial Layouts. To evaluate the efficiency of dynamic operations of the joint EMS system, we consider the home bases and initial layouts of UAVs and ambulances as predetermined. We leverage a status quo layout. Ambulance bases are located at hospitals, police stations, and fire departments; bases of UAVs are located at low-ambulance coverage areas. Initially, numbers of ambulances and UAVs located at each base are proportional to the demand density of the base location. We investigate the impact of optimizing base locations in Section 5.5.

5.1.2. Bystander Willingness. The willingness of bystanders to provide assistance in emergency situations exhibits significant variation, as reported in the literature, with estimates ranging from 27% to 76% (Strang et al. 2000, Kerr et al. 2009, Barbic et al. 2020). This variability can be attributed to several factors, including sociodemographic characteristics, prior witnessing experience, prior overdose experience, perceived risk of arrest, and the specific location of the overdose incident (Tobin et al. 2005, Burn 2017). Considering the wide range of estimates, we model bystander willingness using a uniform distribution within the interval of $[0.2, 0.8]$.

5.1.3. EMS Demands. For our case studies, we extract EMS request distribution from the naloxone administration heat map in the state of Indiana, with the assumption that opioid overdose requests share the same distribution with other EMS requests. The geographic distribution of these request incidences (measured by the centrality of the distribution) is consistent with the rural-urban area classification (Figure 2). Accordingly, we consider three catchment areas based on their geographic delineation: Marion County (urban), Tippecanoe County (semiurban), and Marshall County (rural).

5.1.4. Performance Metrics. We calculate six metrics: accumulated total rewards, average response time for all requests and for opioid overdose requests, fraction of calls served within the response time threshold for all requests and for opioid overdose requests, and fraction of outsourced requests.

5.1.5. Results Validation. The response time of the practical policy is in line with the performance of the current EMS system. Specifically, median EMS arrival times for all call types are between seven and eight minutes (Mell et al. 2017). In rural, remote, geographically challenging, or high-traffic urban areas, this response time can average more than 14 minutes (Hanna 2018).

5.2. Benchmark Policies

In this section, we introduce several policies that are used as benchmarks in our study.

5.2.1. Static Ad Hoc Policy. In current practice, experienced dispatchers make ambulance dispatching decisions in the following ad hoc manner (Schmid 2012). They can view a dashboard with a regional map that shows the position and status of each server. In case of an emergency, the closest available server is usually dispatched. Servers will return to their home bases after serving a request.

5.2.2. Dynamic Heuristic Policy. When an EMS request is received, the closest available server is dispatched. The redeployment policy uses the heuristic developed by Jagtenberg et al. (2015), which is based on solving the

Figure 2. (Color online) Spatial Distribution of Naloxone Administration for Opioid Overdose

Notes. In panels (a)–(c), the dotted areas represent regions with medium-to-high incident arrival rates (In colored version, purplish and yellowish colors represent medium and high incident arrival rates, respectively). Panel (d) shows the probability of an incident occurring in each grid based on the incident arrival probability distribution heat map for Tippecanoe County. For example, the value "0.270" in the center grid indicates that there is a 0.27 probability of the next incident occurring in that grid in Tippecanoe County. (a) Marion. (b) Tippecanoe. (c) Marshall. (d) Extracted distribution.

MEXCLP. This heuristic is easy to implement and has shown good performance for instances considered by the authors. When a server finishes its current task, it will be redeployed to a base that results in the largest marginal contribution to coverage according to the MEXCLP model.

5.2.3. Maxwell ADP. Maxwell et al. (2010) and Nasrollahzadeh et al. (2018) proposed two ADP-based policies for optimal ambulance redeployment decisions and/or dispatching decisions. We regard them as predecessors of our solution methodology and thus, use them as benchmark policies. We employ the basis functions and coefficient training algorithms proposed by Maxwell et al. (2010) and Nasrollahzadeh et al. (2018) while keeping our modeling framework, which includes the system dynamics and the objective function.

5.2.4. L-ADP. In this ADP-based benchmark, we employ the basis functions developed in Section 4.2. However, for feature combination and ADP training, we adopt a linear representation for the value function approximation and choose the coefficient training algorithms in Maxwell et al. (2010) and Nasrollahzadeh et al. (2018).

5.3. Benefit of Introducing UAVs

In this section, we explore the benefit of the joint operation of UAVs with ambulances. The comparison is based on the arrival profile in Tippecanoe County. We consider three settings of the fleet size (i.e., no UAVs, 8 UAVs, and 16 UAVs). Table 1 suggests that the NN-API policy achieves a shorter response time than the benchmark in all the tested scenarios (the complete comparison results can be found in section 4.1 in the online appendix).

Table 1 demonstrates that the incorporation of UAVs into an EMS system can significantly reduce the response time, especially for opioid overdose requests. The response time reduction is more significant in rural areas, where ambulance resources are typically limited. For both counties, the response times can be reduced by over 30% with the use of UAVs, even when using a simple operational strategy. With our proposed strategy, response times could be reduced by as much as 50%. This magnitude of reduction in response time could potentially save more lives, as UAVs would be able to arrive at the scene faster and provide lifesaving treatment in a timely manner. This could be especially beneficial for opioid overdose victims who might not have survived without prompt medical intervention. Moreover, Table 1 also demonstrates comparable performances between the NN-API policy with no UAV and the static policy with eight UAVs in Tippecanoe County (semiurban). This suggests that intelligent real-time operations can compensate for the lack of EMS resources. In general, the addition of UAVs to the EMS system can improve the overall performance and efficiency of emergency response.

Table 1. Average Response Time (Minutes) Comparison (95% Confidence Interval) Among Different Numbers of UAVs in Tippecanoe County (Semiurban) and Marshall County (Rural)

	General	Opioid
Tippecanoe County		
0 UAV		
NN-API	10.8 ± 0.3	8.3 ± 0.4
Static	13.7 ± 0.3	12.8 ± 0.5
8 UAVs		
NN-API	9.1 ± 0.4	5.7 ± 0.3
Static	10.3 ± 0.2	7.7 ± 0.4
Marshall County		
0 UAV		
NN-API	12.6 ± 0.5	10.8 ± 0.5
Static	13.5 ± 0.6	11.5 ± 0.7
8 UAVs		
NN-API	7.8 ± 0.3	5.4 ± 0.3
Static	9.3 ± 0.4	7.9 ± 0.5

5.4. Sensitivity Analysis on Service Areas

In this section, we investigate the sensitivity of the policy performances with respect to the urban-rural delineation of the service areas. As shown in Figure 3, a service area that is more urban (e.g., Marion County) tends to have a higher volume of requests and more pockets of spatially concentrated requests, whereas a service area that is more rural (e.g., Marshall County) tends to be the reverse. Table 2 summarizes our results, suggesting that the NN-API policy performs consistently better than the benchmarks in terms of accumulated rewards. The complete comparison results can be found in section 4.2 in the online appendix.

Table 2 suggests that in areas with high demand, such as Marion County, more requests arrived continuously. Hence, ADP-based policies take advantage of considering future requests and lead to significantly better performance, as indicated by a p-value of less than 0.01 in a paired t-test. In low-demand cases, like Marshall County, the intervals between requests were longer, making the decision process more similar to a one-shot decision. In these cases, there is barely any benefit from sacrificing current rewards to prepare for future demand, resulting in similar performances between the NN-API policy and the static policy, as indicated by a p-value of 0.12 in a paired t-test.

5.5. Performance Improvement with Optimized Base Locations

In this section, we investigate the impact of the optimization of the home base locations on the policy performance. We use a maximal coverage location problem (MCLP) to jointly optimize the locations of ambulance and UAV bases for maximum coverage. More details

Figure 3. (Color online) Spatial Distribution of Naloxone Administration and Settings of Ambulances and UAVs in Tippecanoe, Marion, and Marshall Counties

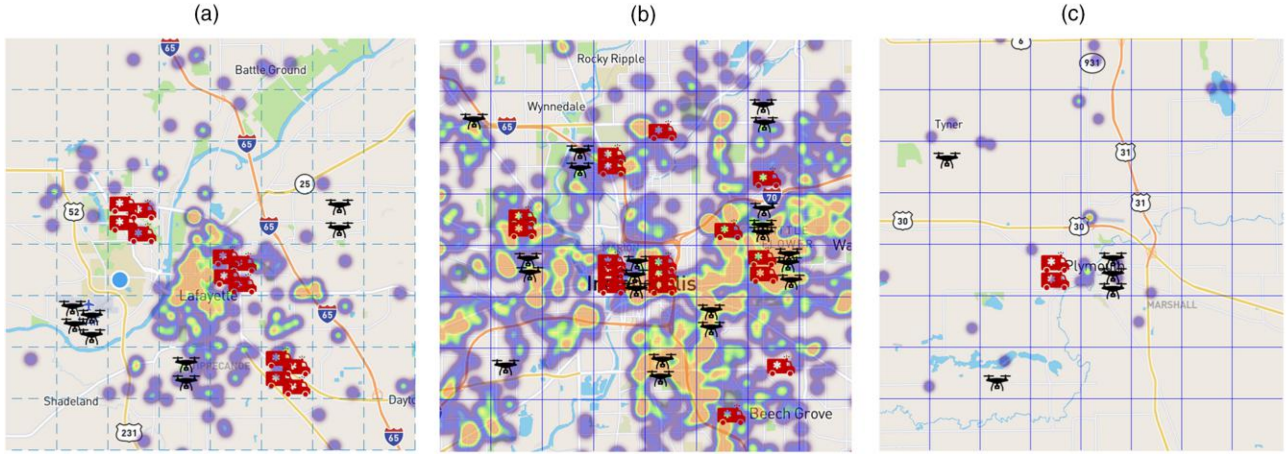

Notes. (a) Tippecanoe. (b) Marion. (c) Marshall.

about this location model can be found in section 4.3.1 in the online appendix. We compare the performance of our NN-API policy with the benchmark policies under four different sets of base locations in Tippecanoe County (Figure 4): (a) both random, where both ambulance and UAV bases are randomly located across the service area; (b) status quo supplementary (the baseline setting used in the previous two sections), where ambulance bases are located at hospitals, fire departments, and police stations, whereas UAV bases are located in low-ambulance coverage areas; (c) status quo optimized, where ambulance bases are fixed at status quo locations and only UAV bases are optimized using an MCLP; and (d) both optimized, where both ambulance and UAV bases are jointly optimized using a larger-scale MCLP.

To simplify the result presentation and clarify the comparison across different settings, we only present the performance of the NN-API and static ad hoc policies in Table 3. Full comparative results can be found in section 4.3.2 in the online appendix.

Table 3 shows that our NN-API policy consistently outperforms the benchmark policies in terms of accumulated rewards by reducing response times, particularly for opioid overdoses. This suggests that even with predetermined base locations, there is always room for improvement in real-time operations, which our NN-API policy can achieve. Furthermore, the table demonstrates a significant improvement from the *Both random* setting to the *status quo-supplementary* setting under both the NN-API and static policies, highlighting the importance of base locations. Base locations determine the feasible choices for dispatching and relocation decisions, and optimizing them could further improve performance and shorten response times, especially for the static policy. Although optimizing base locations can improve coverage, the impact of such optimization for the NN-API policy may be limited to Tippecanoe County because of the centralized pattern of requests. The *status quo-supplementary* setting yields base locations that adequately cover high-demand areas, and optimizing base locations would primarily benefit low-demand areas through enhancing coverage. Therefore, the resulting improvement in average response time may not be significant. Nevertheless, optimizing base locations in regions

Table 2. Performance Comparison (95% Confidence Interval) Among Different Service Areas

		Response time (minutes)		Respond within threshold (%)		
	Reward	General	Opioid	General	Opioid	Outsourced (%)
Tippecanoe						
NN-API	2,132 ± 40	9.1 ± 0.4	5.7 ± 0.3	76.6 ± 0.4	78.8 ± 0.4	13.9 ± 0.4
Static	1,882 ± 39	10.3 ± 0.2	7.7 ± 0.4	74.9 ± 0.5	74.6 ± 0.7	14.8 ± 0.4
Marion						
NN-API	5,420 ± 133	6.8 ± 0.3	4.9 ± 0.2	70.2 ± 0.5	72.6 ± 0.6	21.9 ± 0.4
Static	3,672 ± 104	7.9 ± 0.6	7.4 ± 0.5	64.5 ± 0.7	61.2 ± 0.7	25.9 ± 0.8
Marshall						
NN-API	135 ± 10	12.6 ± 0.6	8.1 ± 0.6	72.6 ± 1.2	74.1 ± 1.7	17.5 ± 1.0
Static	112 ± 9	13.5 ± 1.1	9.2 ± 0.8	67.5 ± 1.1	62.2 ± 1.8	18.1 ± 0.9

Figure 4. (Color online) Various Settings of Base Locations of Ambulances and UAVs in Tippecanoe County

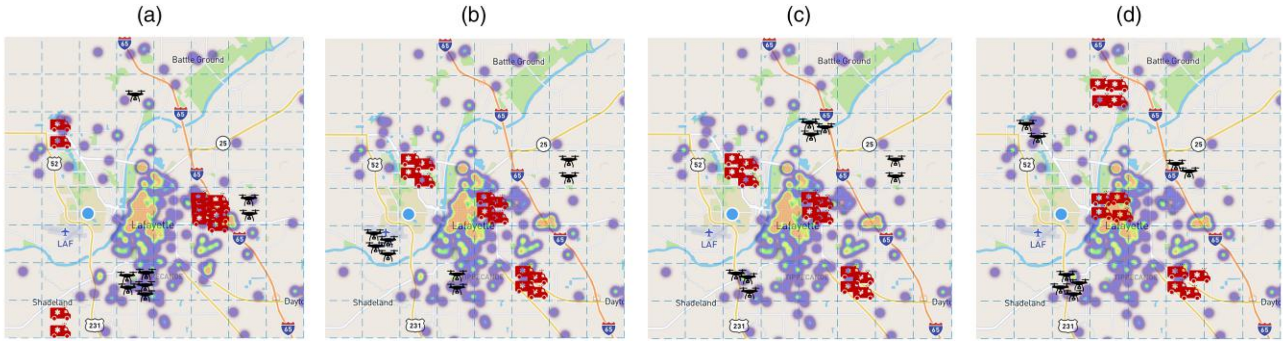

Notes. (a) Both random. (b) Status quo supplementary. (c) Status quo optimized. (d) Both optimized.

with more dispersed demands could lead to more significant reductions in response time.

5.6. Mechanism Behind Performance Superiority

Our experiments in Sections 5.3–5.5 consistently demonstrate that the NN-API policy outperforms the benchmark policies, as shown in Tables 1–3. The NN-API policy is particularly effective at reducing response times for opioid overdose emergency requests. Table 4 provides an example of the performance comparison between the NN-API policy and all the tested benchmarks. The superior performance of the NN-API policy is driven by the consideration of the following components: dynamic redeployment, sequential events and long-term cost, heterogeneity between UAVs and ambulances, and the complex relationship between value functions and basis functions. The benefits of considering these components are verified by the comparison between different benchmarks.

The ad hoc static policy generally performs poorly compared with the other policies. With the same dispatching policy, the heuristic policy can achieve a shorter response time by maximizing coverage in dynamic redeployment. In general, because of the same dispatching policy, the performances of the static policy and the dynamic heuristic policy do not differ a lot from each other. In most of the test scenarios, the three ADP-based policies significantly outperform the other two policies, with a much shorter response time, particularly for opioid overdose emergency requests. This is because these policies consider long-term rewards and are able to learn toward an objective function that gives higher weights to opioid overdose cases. Among the three ADP-based policies, the NN-API and L-ADP policies can lead to higher rewards in most of the test scenarios. Although Maxwell ADP occasionally outperforms L-ADP, the NN-API policy, which uses a well-trained neural network, can consistently improve on the performance of L-ADP and outperform Maxwell ADP. The superior performance of the NN-API policy can be attributed to two factors: the design of the basis functions and the modeling flexibility of the neural network. The basis functions we use take into account the heterogeneity of ambulances and UAVs in terms of service time and incorporate temporal features

Table 3. Performance Comparison (95% Confidence Interval) Under Different Base Locations and Initial Layouts at Tippecanoe County

		Response time (minutes)		Respond within threshold (%)		
	Reward	General	Opioid	General	Opioid	Outsourced (%)
Both random						
NN-API	1,975 ± 44	11.5 ± 0.4	7.4 ± 0.3	76.3 ± 0.4	78.6 ± 0.5	15.2 ± 0.3
Static	1,735 ± 38	14.1 ± 0.6	8.7 ± 0.4	74.0 ± 0.4	69.5 ± 0.6	17.5 ± 0.4
Status quo supplementary						
NN-API	2,132 ± 40	9.1 ± 0.4	5.7 ± 0.3	76.6 ± 0.4	78.8 ± 0.4	13.9 ± 0.4
Static	1,882 ± 39	10.3 ± 0.2	7.7 ± 0.4	74.9 ± 0.5	74.6 ± 0.7	14.8 ± 0.4
Status quo optimized						
NN-API	2,169 ± 41	9.1 ± 0.4	5.6 ± 0.3	77.1 ± 0.3	79.1 ± 0.4	13.9 ± 0.4
Static	1,916 ± 40	10.0 ± 0.2	7.5 ± 0.4	75.3 ± 0.4	75.2 ± 0.6	14.3 ± 0.5
Both optimized						
NN-API	2,192 ± 41	9.0 ± 0.4	5.5 ± 0.3	77.8 ± 0.3	79.4 ± 0.4	13.6 ± 0.5
Static	1,932 ± 39	9.9 ± 0.2	7.4 ± 0.4	75.4 ± 0.4	75.9 ± 0.6	14.1 ± 0.4

Table 4. Performance Comparison (95% Confidence Interval) Among All Benchmark Policies at Tippecanoe County with Realistic Base Locations and a Moderate Number of UAVs

		Response time (minutes)		Respond within threshold (%)		
	Reward	General	Opioid	General	Opioid	Outsourced (%)
Static ad hoc	1,882 ± 39	10.3 ± 0.2	7.7 ± 0.4	74.9 ± 0.5	74.6 ± 0.7	14.8 ± 0.4
Dynamic heuristic	1,943 ± 41	9.8 ± 0.2	7.6 ± 0.3	74.9 ± 0.4	74.6 ± 0.7	14.8 ± 0.4
Maxwell ADP	2,002 ± 39	9.3 ± 0.3	7.1 ± 0.4	75.4 ± 0.4	75.1 ± 0.3	14.8 ± 0.2
L-ADP	2,090 ± 38	9.7 ± 0.3	6.9 ± 0.3	75.9 ± 0.3	76.2 ± 0.4	14.7 ± 0.3
NN-API	2,132 ± 40	9.1 ± 0.4	5.7 ± 0.3	76.6 ± 0.4	78.8 ± 0.4	13.9 ± 0.4

of the system. This allows our solution method to approximate the value function more accurately, even with a linear approximation model.

5.7. Managerial Insights on Policy Selection

To facilitate the integration of UAVs in EMS and provide operational guidance for policy selection, we collect and analyze results from 60 environment scenarios. These scenarios span a range of demand patterns (centralized, scattered), demand rates (low, medium, high), ambulance and UAV quantities (low, medium, high), and base locations (optimized, unoptimized). Our collected results demonstrate that the NN-API policy leads to the best performances among all the tested policies in 59 of 60 scenarios, with no other policies showing statistically better results at a 95% confidence interval. The L-ADP policy and the static ad hoc policy perform comparably with NN-API in 28 and 11 scenarios, respectively, with no significant difference from NN-API at the 95% confidence interval. These findings reaffirm the superiority of the NN-API policy and also suggest that the adoption of a "simpler" policy may yield sufficiently good performance under specific circumstances. Thus, we investigate two questions. (1) Under which circumstances is NN-API statistically significantly better than other policies so that we would recommend adopting it ("NN best")? (2) In which cases can the static ad hoc policy achieve statistically comparable performance with NN-API, indicating that it may be reasonable to use the simpler static policy ("static NN comp")? To answer these questions, we utilize exact logistic regression with regularization to examine the relationship between the scenario settings and the policy performance. Further details about scenario generation and logistic regression can be found in section 5 in the online appendix.

The regression coefficients for "NN best" are shown in Table 5. These results indicate that when demand is not low and ambulance and UAV availability is limited, adopting the NN-API policy is strongly recommended as it significantly outperforms all other policies. These findings align with the observations made in Sections 5.3–5.5, which suggest the use of the NN-API policy in scenarios where dispatching and redeployment decisions are particularly challenging. This typically occurs when consecutive high-demand situations occur and resources are scarce, where each dispatching decision significantly affects the response time of subsequent events.

Panel B in Table 5 presents the regression coefficients for "static NN comp," where the coefficient for "demand_high" is regularized to zero. The regression results suggest that the static ad hoc policy has the potential to achieve performance comparable with the NN-API policy under specific conditions. These conditions include low and scattered demand, sufficient availability of ambulances and UAVs, and optimized base locations. On the contrast to the favorable scenarios for "NN best," the "Static NN comp" requires an adequate number of ambulances and UAVs, along with low demand that allows for less consideration of subsequent requests. Additionally, in the case of scattered demand, redeploying to the closest base becomes less disadvantageous as predicting the location of the next request becomes more challenging compared with centralized demand scenarios. Further, optimizing the base locations enhances the potential of improving performance with the static policy. Under optimized base locations, each base can effectively cover its designated area,

Table 5. Regularized Exact Logistic Regression Coefficients for "NN Best" and "Static NN Comps"

	Coefficient	Standard error	*p*-value
Panel A: NN best			
demand_central	0.037	0.359	0.918
demand_high	−0.119	0.439	0.786
demand_low	−1.229	0.531	**0.021**
fleet_high	0.375	0.430	0.383
fleet_low	1.737	0.468	**0.016**
base_optimized	0.131	0.387	0.735
Panel B: Static NN comp			
demand_central	−2.530	0.121	**0.000**
demand_high	0.000	—	—
demand_low	1.055	0.153	**0.000**
fleet_high	−0.073	0.105	0.485
fleet_low	−0.635	0.110	**0.000**
base_optimized	0.853	0.099	**0.000**

Note. Bold *p*-values indicating significance at the 95% confidence level.

thereby rendering the redeployment to home bases a satisfactory strategy.

6. Conclusions

In this paper, we develop an ADP approach that makes intelligent real-time decisions in UAV-augmented logistic operations, with a focus on the emergency response to opioid overdose incidents. We formulate this problem with an MDP model that captures the complex system evolution and interplay between the high-dimensional state and a variety of actions. In the face of the solution challenges, we adopt an ADP-based solution approach, which is shown to be efficient in solving real-sized instances. We extend the literature on emergency response operations management by incorporating autonomous vehicles into the modeling. Additionally, we extensively investigate the use of basis functions and neural networks in the ADP framework.

With a detailed event-based simulation model, we evaluate the system performance metrics of the NN-API policy together with various benchmark policies proposed in the previous literature. We use historical data on naloxone administration for opioid overdose in Indiana and the NEMSIS database for our case studies. Our comparative results suggest that our NN-API policy significantly improves system performance over the benchmark policies, especially when the base locations are scattered and the incident rate is high. This improvement is mainly because of the following reasons. First, by adopting an MDP framework, we incorporate future rewards into consideration when making both dispatching and redeployment decisions. Second, by explicitly considering the heterogeneity of UAVs and ambulances, we can better approximate the value functions. Additionally, our use of basis functions designed to minimize response times and consider the heterogeneity of the system leads to better performance compared with the benchmarks in the literature that only consider a single type of server. Finally, the use of neural networks helps improve value function approximation in a wide range of scenarios. Notably, when demand is significant and ambulance and UAV availability is limited, we strongly recommend the adoption of the NN-API policy because of its outstanding performance surpassing all other tested policies.

In brief, to implement the NN-API algorithm, we first obtain catchment area details from the partnering EMS agency, including the locations of their anticipated UAV and ambulance bases. Using this information, we calibrate the EMS logistics system simulator and develop prediction models for travel times, 911 request times and locations, and bystander response rates. We then train our recommended policies using the calibrated simulator, our algorithm, and the prediction models, taking into account the level of uncertainty and capacity of the ambulances and UAVs. These policies provide dispatching and redeployment recommendations to the dispatcher based on the status of the ambulances and UAVs (i.e., location and idle/busy status) and the location of any incoming requests.

In this work, we demonstrate the benefits of using UAVs to deliver lifesaving medication and provided guidance on how EMS agencies can incorporate UAVs into their operations. Our findings highlight the potential advantages of using UAVs in the EMS field and offer practical guidance on how to effectively utilize UAVs, including strategies for dispatching and redeployment. We hope that this work will contribute to the wider adoption of UAVs in EMS.

Our future research will be in the following directions. First, we will incorporate additional realistic features of UAVs and the EMS system, such as time-varying weather conditions and nonstationary, request distributions in the future. For example, individuals are likely to be at different places over our considered time horizon (e.g., the workplace during daytime and home during nighttime). Second, we will investigate the issue of EMS access equity across diverse communities within a service area, for which we will formulate constrained MDPs with constraints on the allowable system outcome inequity. We will develop an ADP approach to the resultant constrained MDPs. Finally, it is also worth investigating the optimal locations of ambulances and UAV bases given an RL-based operational strategy instead of assuming a closest dispatch rule in the facility location phase. This will allow us to further improve system efficiency and reduce response times.

References

Agatz N, Bouman P, Schmidt M (2018) Optimization approaches for the traveling salesman problem with drone. *Transportation Sci.* 52(4):965–981.

Amukele T, Ness PM, Tobian AAR, Boyd J, Street J (2017) Drone transportation of blood products. *Transfusion* 57(3):582–588.

Andersson T, Värbrand P (2007) Decision support tools for ambulance dispatch and relocation. *J. Oper. Res. Soc.* 58(2):195–201.

Aringhieri R, Bruni ME, Khodaparasti S, Theresia Van Essen J (2017) Emergency medical services and beyond: Addressing new challenges through a wide literature review. *Comput. Oper. Res.* 78:349–368.

Bandara D, Mayorga ME, McLay LA (2014) Priority dispatching strategies for EMS systems. *J. Oper. Res. Soc.* 65(4):572–587.

Barbic D, Duncan K, Trainor R, Ertel EA, Enos MK, Philips H, Besserer F, et al. (2020) A survey of the public's ability to recognize and willingness to intervene in out-of-hospital cardiac arrest and opioid overdose. *Acad. Emergency Medicine* 27(4):305–308.

Bélanger V, Ruiz A, Soriano P (2019) Recent optimization models and trends in location, relocation, and dispatching of emergency medical vehicles. *Eur. J. Oper. Res.* 272(1):1–23.

Blackwell TH, Kaufman JS (2002) Response time effectiveness: Comparison of response time and survival in an urban emergency medical services system. *Acad. Emergency Medicine* 9(4):288–295.

Boujemaa R, Jebali A, Hammami S, Ruiz A (2020) Multi-period stochastic programming models for two-tiered emergency medical service system. *Comput. Oper. Res.* 123:104974.

Boutilier JJ, Brooks SC, Janmohamed A, Byers A, Buick JE, Zhan C, Schoellig AP, Cheskes S, Morrison LJ, Chan TCY (2017) Optimizing a drone network to deliver automated external defibrillators. *Circulation* 135(25):2454–2465.

Burn SM (2017) Appeal to bystander interventions: A normative approach to health and risk messaging. Nussbaum J, ed. *Oxford Research Encyclopedia of Communication* (Oxford University Press, Oxford, UK), 1–25.

Cao X-R (2005) Basic ideas for event-based optimization of Markov systems. *Discrete Event Dynamic Systems* 15(2):169–197.

Cavazos-Cadena R (1991) Recent results on conditions for the existence of average optimal stationary policies. *Ann. Oper. Res.* 28(1):3–27.

Cavazos-Cadena R, Sennott LI (1992) Comparing recent assumptions for the existence of average optimal stationary policies. *Oper. Res. Lett.* 11(1):33–37.

Centers for Disease Control and Prevention National Center for Health Statistics (2022) Drug overdose deaths in the U.S. top 100,000 annually. Accessed March 30, 2022, https://www.cdc.gov/nchs/pressroom/nchs_press_releases/2021/20211117.htm.

Chelst K, Jarvis JJ (1979) Estimating the probability distribution of travel times for urban emergency service systems. *Oper. Res.* 27(1):199–204.

Chen X, Ulmer MW, Thomas BW (2019) Deep Q-learning for same-day delivery with a heterogeneous fleet of vehicles and drones. Preprint, submitted October 25, https://arxiv.org/abs/1910.11901.

Chong KC, Henderson SG, Lewis ME (2016) The vehicle mix decision in emergency medical service systems. *Manufacturing Service Oper. Management* 18(3):347–360.

Chu J, Leung KHB, Snobelen P, Nevils G, Drennan IR, Cheskes S, Chan TCY (2021) Machine learning-based dispatch of drone-delivered defibrillators for out-of-hospital cardiac arrest. *Resuscitation* 162:120–127.

Claesson A, Bäckman A, Ringh M, Svensson L, Nordberg P, Djärv T, Hollenberg J (2017a) Time to delivery of an automated external defibrillator using a drone for simulated out-of-hospital cardiac arrests vs emergency medical services. *JAMA* 317(22):2332–2334.

Claesson A, Svensson L, Nordberg P, Ringh M, Rosenqvist M, Djarv T, Samuelsson J, et al. (2017b) Drones may be used to save lives in out of hospital cardiac arrest due to drowning. *Resuscitation* 114:152–156.

Clark AK, Wilder CM, Winstanley EL (2014) A systematic review of community opioid overdose prevention and naloxone distribution programs. *J. Addiction Medicine* 8(3):153–163.

Daskin MS (1983) A maximum expected covering location model: Formulation, properties and heuristic solution. *Transportation Sci.* 17(1):48–70.

Doe-Simkins M, Walley AY, Epstein A, Moyer P (2009) Saved by the nose: Bystander-administered intranasal naloxone hydrochloride for opioid overdose. *Amer. J. Public Health* 99(5):788–791.

Dorling K, Heinrichs J, Messier GG, Magierowski S (2016) Vehicle routing problems for drone delivery. *IEEE Trans. Systems Man Cybernetics Systems* 47(1):70–85.

Everaerts J (2008) The use of unmanned aerial vehicles (UAVs) for remote sensing and mapping. *Internat. Arch. Photogrammetry Remote Sensing Spatial Inform. Sci.* 37(2008):1187–1192.

Fakinos D (1980) The *M/G/K* blocking system with heterogeneous servers. *J. Oper. Res. Soc.* 31(10):919–927.

Fawaz W, Atallah R, Assi C, Khabbaz M (2017) Unmanned aerial vehicles as store-carry-forward nodes for vehicular networks. *IEEE Access* 5:23710–23718.

Gereau J (2022) Intermountain healthcare partners with zipline for drone delivery service. *Utah Public Radio* (October 10), https://www.upr.org/utah-news/2022-10-10/intermountain-healthcare-partners-with-zipline-for-drone-delivery-service.

Hanna M (2018) Utility of unmanned aircraft systems (drones) in inner city emergent response during peak rush hour. *Ann. Emergency Medicine* 72(4):S77–S78.

Ingolfsson A, Budge S, Erkut E (2008) Optimal ambulance location with random delays and travel times. *Health Care Management Sci.* 11(3):262–274.

Jagtenberg CJ, Bhulai S, Van der Mei RD (2015) An efficient heuristic for real-time ambulance redeployment. *Oper. Res. Health Care* 4:27–35.

Jagtenberg CJ, Bhulai S, Van der Mei RD (2017) Dynamic ambulance dispatching: Is the closest-idle policy always optimal? *Health Care Management Sci.* 20(4):517–531.

Jenkins PR, Robbins MJ, Lunday BJ (2020) Approximate dynamic programming for military medical evacuation dispatching policies. *INFORMS J. Comput.* 33(1):2–26.

Jia S, Zhang L (2017) Modelling unmanned aerial vehicles base station in ground-to-air cooperative networks. *IET Comm.* 11(8):1187–1194.

Kerr D, Dietze P, Kelly A-M, Jolley D (2009) Improved response by peers after witnessed heroin overdose in Melbourne. *Drug Alcohol Rev.* 28(3):327–330.

Kim D, Irwin KS, Khoshnood K (2009) Expanded access to naloxone: Options for critical response to the epidemic of opioid overdose mortality. *Amer. J. Public Health* 99(3):402–407.

Lankenau SE, Wagner KD, Silva K, Kecojevic A, Iverson E, McNeely M, Kral AH (2013) Injection drug users trained by overdose prevention programs: Responses to witnessed overdoses. *J. Community Health* 38(1):133–141.

Larson RC (1974) A hypercube queuing model for facility location and redistricting in urban emergency services. *Comput. Oper. Res.* 1(1):67–95.

Larson RC (1975) Approximating the performance of urban emergency service systems. *Oper. Res.* 23(5):845–868.

Latimore AD, Bergstein RS (2017) "Caught with a body" yet protected by law? Calling 911 for opioid overdose in the context of the good Samaritan law. *Internat. J. Drug Policy* 50:82–89.

Ling G, Draghic N (2019) Aerial drones for blood delivery. *Transfusion* 59(S2):1608–1611.

Lockey AS, Brown TP, Carlyon JD, Hawkes CA (2021) Impact of community initiatives on non-EMS bystander CPR rates in West Yorkshire between 2014 and 2018. *Resuscitation Plus* 6:100115.

Marianov V, ReVelle C (1996) The queueing maximal availability location problem: A model for the siting of emergency vehicles. *Eur. J. Oper. Res.* 93(1):110–120.

Maxwell MS, Restrepo M, Henderson SG, Topaloglu H (2010) Approximate dynamic programming for ambulance redeployment. *INFORMS J. Comput.* 22(2):266–281.

McLay LA (2009) A maximum expected covering location model with two types of servers. *IIE Trans.* 41(8):730–741.

McLay LA, Mayorga ME (2013a) A dispatching model for server-to-customer systems that balances efficiency and equity. *Manufacturing Service Oper. Management* 15(2):205–220.

McLay LA, Mayorga ME (2013b) A model for optimally dispatching ambulances to emergency calls with classification errors in patient priorities. *IIE Trans.* 45(1):1–24.

Mell HK, Mumma SN, Hiestand B, Carr BG, Holland T, Stopyra J (2017) Emergency medical services response times in rural, suburban, and urban areas. *JAMA Surgery* 152(10):983–984.

Nasrollahzadeh AA, Khademi A, Mayorga ME (2018) Real-time ambulance dispatching and relocation. *Manufacturing Service Oper. Management* 20(3):467–480.

Ornato JP, You AX, McDiarmid G, Keyser-Marcus L, Surrey A, Humble JR, Dukkipati S, et al. (2020) Feasibility of bystander-administered naloxone delivered by drone to opioid overdose victims. *Amer. J. Emergency Medicine* 38(9):1787–1791.

Park SH, Lee YH (2019) Two-tiered ambulance dispatch and redeployment considering patient severity classification errors. *J. Heathcare Engrg.* 2019:6031789.

Powell WB (2007) *Approximate Dynamic Programming: Solving the Curses of Dimensionality*, vol. 703 (John Wiley & Sons, Hoboken, NJ).

Puterman ML (2014) *Markov Decision Processes: Discrete Stochastic Dynamic Programming* (John Wiley & Sons, New York).

Reuters (2020) Amazon receives U.S. regulatory approval to start drone delivery trials. Accessed February 6, 2021, https://www.reuters.com/article/us-amazon-prime-air/amazon-receives-u-s-regulatory-approval-to-start-drone-delivery-trials-idUSKBN25R2NG.

ReVelle C, Marianov V (1991) A probabilistic fleet model with individual vehicle reliability requirements. *Eur. J. Oper. Res.* 53(1):93–105.

Schilling D, Elzinga DJ, Cohon J, Church R, ReVelle C (1979) The team/fleet models for simultaneous facility and equipment siting. *Transportation Sci.* 13(2):163–175.

Schmid V (2012) Solving the dynamic ambulance relocation and dispatching problem using approximate dynamic programming. *Eur. J. Oper. Res.* 219(3):611–621.

Schwartz DG, Ataiants J, Roth A, Marcu G, Yahav I, Cocchiaro B, Khalemsky M, Lankenau S (2020) Layperson reversal of opioid overdose supported by smartphone alert: A prospective observational cohort study. *EClinicalMedicine* 25:100474.

Shin J, Marla L, Boutilier J (2022) Heterogeneous facility location under coordination and collaboration: The case of ambulance-bystander-drone coordination. Preprint, submitted August 22, https://ssrn.com/abstract=3519243.

Strang J, Best D, Man L-H, Noble A, Gossop M (2000) Peer-initiated overdose resuscitation: Fellow drug users could be mobilised to implement resuscitation. *Internat. J. Drug Policy* 11(6):437–445.

Tobin KE, Davey MA, Latkin CA (2005) Calling emergency medical services during drug overdose: An examination of individual, social and setting correlates. *Addiction* 100(3):397–404.

Tokekar P, Vander Hook J, Mulla D, Isler V (2016) Sensor planning for a symbiotic UAV and UGV system for precision agriculture. *IEEE Trans. Robotics* 32(6):1498–1511.

Ulmer MW, Thomas BW (2018) Same-day delivery with heterogeneous fleets of drones and vehicles. *Networks* 72(4):475–505.

Wilde ET (2013) Do emergency medical system response times matter for health outcomes? *Health Econom.* 22(7):790–806.

World Health Organization (2019) Drones take Rwanda's national blood service to new heights. Accessed February 6, 2021, https://www.who.int/news-room/feature-stories/detail/drones-take-rwandas-national-blood-service-to-new-heights.

Yoon S, Albert LA (2020) A dynamic ambulance routing model with multiple response. *Transportation Res. Part E Logist. Transportation Rev.* 133:101807.

Yoon S, Albert LA (2021) Dynamic dispatch policies for emergency response with multiple types of vehicles. *Transportation Res. Part E Logist. Transportation Rev.* 152:102405.

Yoon S, Albert LA, White VM (2021) A stochastic programming approach for locating and dispatching two types of ambulances. *Transportation Sci.* 55(2):275–296.