Proceedings of the 2021 Winter Simulation Conference
S. Kim, B. Feng, K. Smith, S. Masoud, Z. Zheng, C. Szabo, and M. Loper, eds.

ANALYZING THE IMPACT OF TRIAGE CLASSIFICATION ERRORS ON MILITARY MEDICAL EVACUATION DISPATCHING POLICIES

Emily S. Graves
Phillip R. Jenkins
Matthew J. Robbins

Department of Operational Sciences
Air Force Institute of Technology
2950 Hobson Way
Wright-Patterson AFB, OH 45433, USA

ABSTRACT

This paper analyzes how triage classification errors and blood transfusion kits impact military medical evacuation (MEDEVAC) system performance with regards to dispatching policies. A discounted, infinite-horizon Markov decision process (MDP) model is formulated to analyze the MEDEVAC dispatching problem. A notional, representational scenario based in Azerbaijan is utilized to compare the MDP-generated policies to current practices. Results reveal the MDP-generated dispatching policy outperforms the currently practiced dispatching policy by 2.18%. Triage classification errors negatively impact system performance between 0.6% and 2.5% for the scenarios analyzed. Moreover, the inclusion of blood transfusion kits on board aircraft increase MEDEVAC system performance between 2.83% and 4.37%, depending on which units are equipped.

1 INTRODUCTION

In a deployed environment, military emergency medical service (EMS) response system personnel seek to effectively and efficiently evacuate casualties from the battlefield to medical treatment facilities (MTFs). Two primary resources are available to accomplish this task: medical evacuation (MEDEVAC) and casualty evacuation (CASEVAC). MEDEVAC platforms have dedicated medical personnel on board to treat casualties en route to an MTF, whereas CASEVAC platforms do not (Department of the Army 2019). As such, military medical planners rely on MEDEVAC to serve as the primary link among the roles of medical care across combat operations.

Whereas a variety of platforms can be leveraged when performing an evacuation (e.g., ground ambulances, air ambulances, and sea ambulances), this paper focuses specifically on rotary wing air assets (e.g., helicopters) for evacuating casualties via MEDEVAC. Helicopters were first utilized to evacuate casualties during the Korean War and continue to be employed as the primary MEDEVAC platform. The United States (U.S.) Army employs HH-60M Black Hawk helicopters for MEDEVAC missions, which are capable of air crash rescue support; expeditious delivery of whole blood and medical supplies to meet critical requirements; rapid movement of medical personnel and accompanying equipment to address changes in battlefield requirements; and movement of patients between hospitals (Department of the Army 2019).

It is imperative that a MEDEVAC system is effective and efficient, not only to increase survivability and decrease the time between injury and medical care, but also to retain confidence among military personnel conducting combat operations on the battlefield (Jenkins et al. 2020a; Jenkins et al. 2021c). More specifically, an effective and efficient system demonstrates to battlefield personnel that rapid and quality care is available upon request. Many decisions impact the effectiveness and efficiency of a MEDEVAC system, including the location, allocation, relocation, dispatching, and redeployment of units (Bastian 2010; Bastian et al. 2013; Grannan et al. 2015; Keneally et al. 2016; Jenkins et al. 2018; Jenkins et al. 2020b).

This paper focuses on the MEDEVAC dispatching problem, which seeks to determine which (if any) unit to dispatch in response to a request for service. The MEDEVAC dispatching problem is formulated as a discounted, Markov decision process (MDP) model over an infinite horizon. Similar to previous research, this study assumes MEDEVAC asset locations are predetermined and that redeployment does not occur. As an augment to the previous research in this area (e.g., Rettke et al. (2016), Robbins et al. (2020), and Jenkins et al. (2021b)), this work accounts for the possibility of triage classification errors. In a deployed environment, assessing injuries is difficult, which may result in inaccurate triage classification reporting. The priority level (i.e., severity) of a casualty is not truly known to MEDEVAC staff until the injury is assessed by trained medical personnel at the casualty collection point (CCP). In addition to accounting for classification errors, this paper explicitly models blood transfusion kits on board select MEDEVAC aircraft. Malsby III et al. (2013) examine the feasibility of performing blood transfusions on board military MEDEVAC vehicles and found that the vehicle and altitude do not have an adverse effect on blood transfusions. Moreover, Kotwal et al. (2018) show that an early blood transfusion is associated with higher chances of battlefield survival. Blood transfusion kits on board military MEDEVAC aircraft allow for patients to be treated more extensively and in a more timely manner prior to arriving at an MTF. These additional problem features create a more accurate depiction of the military MEDEVAC system.

The remainder of this paper is organized as follows. Section 2 describes the MEDEVAC dispatching problem. Section 3 discusses the MDP formulation developed to determine an optimal MEDEVAC dispatch policy. Section 4 covers an application of the formulated MDP based on a representative scenario in Azerbaijan. Section 5 concludes the paper and proposes several directions for future research.

2 PROBLEM DESCRIPTION

When a MEDEVAC request is submitted, the dispatching authority analyzes the request's information and subsequently decides which (if any) MEDEVAC unit to task to service the request. The information in each request includes, but is not limited to, the location, number, and priority level. Each MEDEVAC request submitted to the system is categorized as one of three priority levels: Priority I, Priority II, and Priority III (Department of the Army 2019). Priority I and II requests are considered life-threatening and need to be serviced within one and four hours, respectively, to maximize survivability and minimize long-term disabilities (e.g., loss of limb or eyesight). Priority III requests are considered non-life-threatening, but still need to be serviced within 24 hours to prevent further deterioration in health.

The incorporation of admission control allows the MEDEVAC dispatching authority to turn away requests to other forms of evacuation. When this occurs, the MEDEVAC request is sent to a secondary evacuation service such as CASEVAC. If the request is admitted into the MEDEVAC system, the dispatching authority tasks a unit to service the request. If every MEDEVAC unit is busy and a service request is submitted to the system, the service request is relayed to other forms of evacuation rather than being placed in a queue. There are multiple forms of evacuation in a deployed environment, and we assume it is more beneficial for a request to be relayed to a secondary evacuation service as opposed to waiting for an aerial MEDEVAC unit to become available. During the flight from the CCP to the MTF, medical treatment is administered to the injured individuals, but medical personnel are limited by what is available on board. If the MEDEVAC aircraft is equipped with a blood transfusion kit, medical personnel are able to start blood transfusions as soon as injured individuals are loaded. Kotwal et al. (2018) show that the timing of a blood transfusion is critical to survival and the earlier blood is given the better chance of survival. Figure 1 illustrates the MEDEVAC mission timeline.

As depicted in Figure 1, the time at which a MEDEVAC request reaches the dispatching authority is denoted as T_1. If the request is admitted into the system, a MEDEVAC unit is tasked at time T_2 and departs the staging area at time T_3. The MEDEVAC unit arrives at the CCP at time T_4, loads the injured individuals onto the aircraft, and then departs the CCP in transit to an MTF at time T_5. The unit arrives

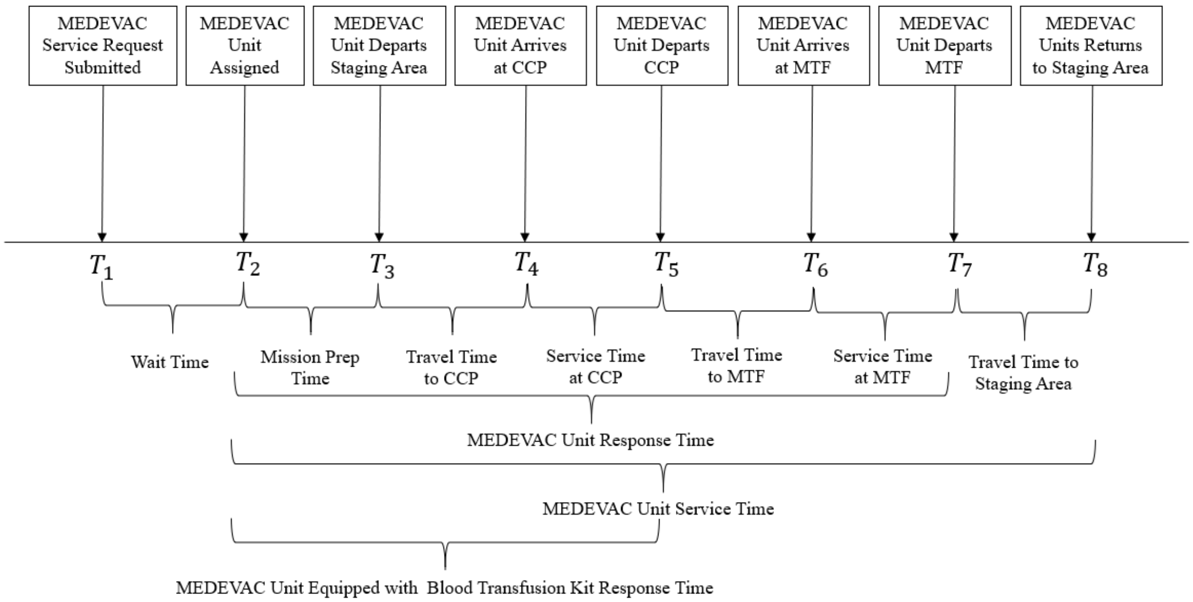

Figure 1: MEDEVAC mission timeline.

at the MTF at time T_6, unloads the injured individuals, and subsequently departs the MTF in transit to the unit's staging area, arriving at time T_8. It is important to note the primary cause of death for battlefield casualties is blood loss (Garrett 2013). As such, response time is defined as the time it takes a MEDEVAC unit to enable blood transfusion for casualties. For MEDEVAC aircraft equipped with blood transfusion kits, this time is $T_5 - T_2$, whereas aircraft without blood transfusion kits have a response time of $T_7 - T_2$. Service time is the total time a MEDEVAC unit is busy, which includes the time it takes to transit back to the original staging facility. Hence, service time is defined as $T_8 - T_2$.

3 METHODOLOGY

The MEDEVAC dispatching problem requires sequential decision making under uncertainty, which makes an MDP modeling approach ideally suitable. The primary components of an MDP model include the decision epochs, state space, action space, transition function, and reward function, which are described in detail below for the MEDEVAC dispatching problem.

The set of decision epochs $\mathscr{T} = \{1,2,...\}$ represents the points in time when a decision is required. These epochs occur when a request is submitted to the system or when a MEDEVAC unit arrives back to its staging facility.

Let $S_t = (M_t, R_t) \in \mathscr{S}$ represent the state of the system at decision epoch $t \in \mathscr{T}$. The state space is comprised of two components: (1) the MEDEVAC status tuple M_t and (2) the request status tuple R_t. The MEDEVAC status tuple is defined as $M_t = (M_{tm})_{m\in\mathscr{M}}$, where $\mathscr{M} = \{1,2,...,|\mathscr{M}|\}$ represents the set of MEDEVAC units in the system. The state variable $M_{tm} = \{0\} \cup \mathscr{Z}$ contains information pertaining to MEDEVAC unit m at epoch t, where $\mathscr{Z} = \{1,2,...,|\mathscr{Z}|\}$ represents the set of zones from which MEDEVAC requests can originate. When $M_{tm} = 0$, unit m is idle, and when $M_{tm} = z$, unit m is busy servicing a request from zone $z \in \mathscr{Z}$.

The request status tuple is defined as $R_t = (Z_t, K_t, C_t)_{Z_t\in\mathscr{Z}, K_t\in\mathscr{K}, C_t\in\mathscr{K}}$, where $\mathscr{K} = \{1,2,...,|\mathscr{K}|\}$ represents the set of possible priority levels. The request status tuple provides the zone Z_t, reported priority level K_t, and true priority level C_t of the request awaiting an admission control decision given there is one in the system at epoch t. If $C_t \neq K_t$, a classification error has occurred. If no requests are awaiting admission control decisions, then $R_t = (0,0,0)$.

A submitted request is either admitted into and serviced by the MEDEVAC system, or the request is rejected from entering the system and is serviced through other evacuation options. The dispatching authority needs to make a decision when a request is submitted to the system and at least one MEDEVAC unit is available to service the request. Once a MEDEVAC unit is tasked to service a request, it is considered unavailable until it returns back to the staging location (i.e., completes the mission). Let $a_t = (a_t^{admit}, a_t^d)$ represent the decision variable tuple at epoch t. The decision variable tuple is comprised of two components: (1) the admission control decision variable and (2) the dispatch decision variable.

The admission control decision variable is given by $a_t^{admit} \in \{0,1\}$. Let $\mathscr{I}(S_t) = \{m : m \in \mathscr{M}, M_{tm} = 0\}$ denote the set of idle MEDEVAC units available to be dispatched and $\mathscr{B}(S_t) = \{m : m \in \mathscr{M}, M_{tm} \neq 0\}$ represent the set of busy MEDEVAC units when the state of the system is S_t at decision epoch t. Of note, $\mathscr{M} = \mathscr{I}(S_t) \cup \mathscr{B}(S_t), \forall\, S_t \in \mathscr{S}$. The decision variable $a_t^{admit} = 1$, if the current request in the request status tuple is admitted into the MEDEVAC system at epoch t, and 0 otherwise. The dispatch decision variable is given by $a_t^d \in \{0\} \cup \mathscr{I}(S_t)$. If the dispatching authority chooses to dispatch MEDEVAC unit $m \in \mathscr{I}(S_t)$ to service the request R_t, then $a_t^d = m$, 0 otherwise.

Let $\mathscr{A}(S_t)$ denote the actions available to the dispatching authority when the system is in state S_t at decision epoch t. The action space is given by

$$\mathscr{A}(S_t) = \begin{cases} (0,0), & \text{if } R_t = (0,0,0) \vee \mathscr{I}(S_t) = \varnothing \\ (\{0,1\} \times \{0 \cup \mathscr{I}(S_t)\}), & \text{if } R_t \neq (0,0,0) \wedge \mathscr{I}(S_t) \neq \varnothing, \end{cases}$$

such that if $a_t^{admit} = 0$ then $a_t^d = 0$, and if $a_t^{admit} = 1$ then $a_t^d \neq 0$. The first case represents the feasible actions when the request status tuple is empty or there are no MEDEVAC units available to be dispatched. In this case, the only action is to reject the request from entering the system. The second case represents the feasible actions when the request status tuple in nonempty and at least one MEDEVAC unit is idle. In this case, the actions available are to either reject the request from entering the system or to admit request into the system and dispatch one of the available MEDEVAC units to service the request.

This system transitions when an event occurs. The first event type is when a request is submitted, whereas the second event type is when a MEDEVAC unit completes service. Requests for MEDEVAC service arrive to the dispatching authority according to a Poisson process with a rate parameter of $\lambda = \sum_{z \in \mathscr{Z}} \sum_{k \in \mathscr{K}} \sum_{c \in \mathscr{K}} \lambda_{zkc}$, where λ_{zkc} represents the MEDEVAC request arrival rate in minutes for zone z, reported priority level k and true priority level c requests. Utilizing a Poisson process to model request arrivals is reasonable within the military MEDEVAC context and is utilized in related efforts; see Jenkins et al. (2021a) for more details. Let μ_{mz} represent the expected service rate in minutes for MEDEVAC unit m servicing a zone z request. When the MEDEVAC system is in state S_t and action a_t is taken, the system immediately transitions to a post decision state denoted by S_t^a. The transition time from this post decision state to the next pre-decision state S_{t+1} (i.e., the state-action sojourn time) is exponentially distributed with parameter

$$\beta(S_t, a_t) = \lambda + \sum_{m \in \mathscr{B}(S_t)} \mu_{m,M_{tm}} + \sum_{m \in \mathscr{I}(S_t)} a_{tm}^d \mu_{m,Z_t}.$$

If $\mathscr{B}(S_t) = \varnothing$ and $a_t^d = \{0\}^{\mathscr{I}(S_t)}$, indicating all MEDEVAC units are idle and there are no requests in the system, then $\beta(S_t, a_t)$ represents the state-action pair sojourn time wherein the next decision epoch occurs upon the arrival of a new MEDEVAC service request. Otherwise, $\mathscr{B}(S_t) \neq \varnothing$ and/or a unit is tasked to service an incoming request (e.g., $a_{tm}^d = 1$ for some $m \in \mathscr{I}(S_t)$), indicating at least one MEDEVAC unit is servicing a request. In this case, $\beta(S_t, a_t)$ represents the state-action pair sojourn time wherein the next decision epoch occurs after the arrival of a new request or a MEDEVAC unit returns from servicing a request. The probabilistic behavior of the MEDEVAC system can be summarized using an infinitesimal generator as follows,

$$G(S_{t+1}|S_t,a_t) = \begin{cases} -[1-p(S_t^a|S_t,a_t)]\beta(S_t,a_t), & \text{if } S_{t+1} = S_t^a \\ p(S_{t+1}|S_t,a_t)\beta(S_t,a_t), & \text{if } S_{t+1} \neq S_t^a. \end{cases}$$

Given that the system is in state S_t and action a_t is taken, the probability that the system transitions to state S_{t+1} is denoted by,

$$p(S_{t+1}|S_t,a_t) = \begin{cases} \frac{\lambda_{zkc}}{\beta(S_t,a_t)}, & \text{if } R_{t+1} = (z,k,c), z \in \mathscr{Z}, k \in \mathscr{K}, c \in \mathscr{K} \\ \frac{\mu_{mz}}{\beta(S_t,a_t)}, & \text{if } R_{t+1} = (0,0,0), M_{t+1,m} = 0, M_{tm}^a = z, m \in \mathscr{M}, z \in \mathscr{Z} \\ 0, & \text{otherwise.} \end{cases}$$

The system transitions to a different state S_{t+1} at the end of a sojourn in state S_t^a. It is important to note that decision epochs may occur at any point due to the continuous nature of the problem, and the time between decisions follows an exponential distribution. Uniformization is applied to transform the continuous process to an equivalent discrete process to make subsequent analysis easier to perform (Puterman 1994; Kulkarni 2017). To uniformize the system, the maximum rate of transition must be determined and is given by $\nu = \lambda + \sum_{m\in\mathscr{M}} \tau_m$, where $\tau_m = \max_{z\in\mathscr{Z}} \mu_{mz}$, $\forall\, m \in \mathscr{M}$. Applying uniformization yields the following transition probabilities:

$$\tilde{p}(S_{t+1}|S_t,a_t) = \begin{cases} 1 - \frac{[1-p(S_t^a|S_t,a_t]\beta(S_t,a_t)}{\nu}, & \text{if } S_{t+1} = S_t^a \\ \frac{[p(S_{t+1}|S_t,a_t]\beta(S_t,a_t)}{\nu}, & \text{if } S_{t+1} \neq S_t^a \\ 0, & \text{otherwise.} \end{cases}$$

The uniformization process relaxes the restriction of no self-transitions and may be viewed as inducing extra transitions from a particular state to itself.

The MEDEVAC systems earns rewards when the dispatching authority tasks units to respond to service requests. More specifically, the immediate expected reward obtained by the system is a function of the tasked MEDEVAC unit and the request being serviced, which includes the location of the MEDEVAC unit as well as the location, reported priority level, and true priority level of the request. The immediate expected reward is defined as $r(S_t,a_t) = \delta_K \phi_{KC} f(\zeta_{a_{tm}^d,Z})$, where δ_K is a trade-off parameter that varies the reward attained base on K, ϕ_{KC} denotes the probability the reported priority level k of a request is in reality the true priority level c of a request, where $\sum_{c\in\mathscr{K}} \phi_{kc} = 1, \forall\, k \in \mathscr{K}$, and $f(\zeta_{a_{tm}^d,Z})$ is a monotonically decreasing utility function with respect to response time $\zeta_{a_{tm}^d,Z}$. If no MEDEVAC unit is dispatched $r(S_t,a_t) = 0$. Uniformization is applied to transform the reward function to an equivalent discrete-time form as follows

$$\tilde{r}(S_t,a_t) = r(S_t,a_t)\frac{\alpha + \beta(S_t,a_t)}{\alpha + \nu},$$

where $\alpha > 0$ represents the continuous-time discounting rate.

Let $A^\pi(S_t) \in \mathscr{A}(S_t)$ represent the decision function that maps the state space to the action space. This function indicates action a_t to be taken given the system is in state S_t according to policy π. The MDP model seeks to determine the optimal policy π^* from all available policies, $(A^\pi(S_t))_{\pi\in\Pi}$. The optimal policy maximizes the expected total discounted reward and is given by

$$\max_{\pi\in\Pi} \mathbb{E}^\pi \left[\sum_{t=1}^{\infty} \gamma^{t-1} \tilde{r}(S_t, A^\pi(S_t)) \right],$$

where $\gamma = \frac{\nu}{\nu+\alpha}$ is the uniformized discount factor. The optimal policy is found via the Bellman Equation

$$V(S_t) = \max_{a_t \in \mathscr{A}(S_t)} \left(\tilde{r}(S_t, a_t) + \gamma \sum_{S_{t+1} \in \mathscr{S}} \tilde{p}(S_{t+1} | S_t, a_t) V(S_{t+1}) \right).$$

The policy iteration algorithm is implemented in MATLAB to solve the Bellman Equation and determine the optimal dispatching policy π^* exactly.

4 TESTING, RESULTS, & ANALYSIS

This section examines a notional military MEDEVAC problem instance in Azerbaijan to demonstrate the applicability of the MDP model. The problem instance includes four MEDEVAC staging areas (i.e., bases) located in Agdzhabedi, Karachala, Goradiz, and Salyany as well as casualty cluster centers (red circles) as depicted in Figure 2. Of the four bases, two are collocated with MTFs (blue squares) and two are not

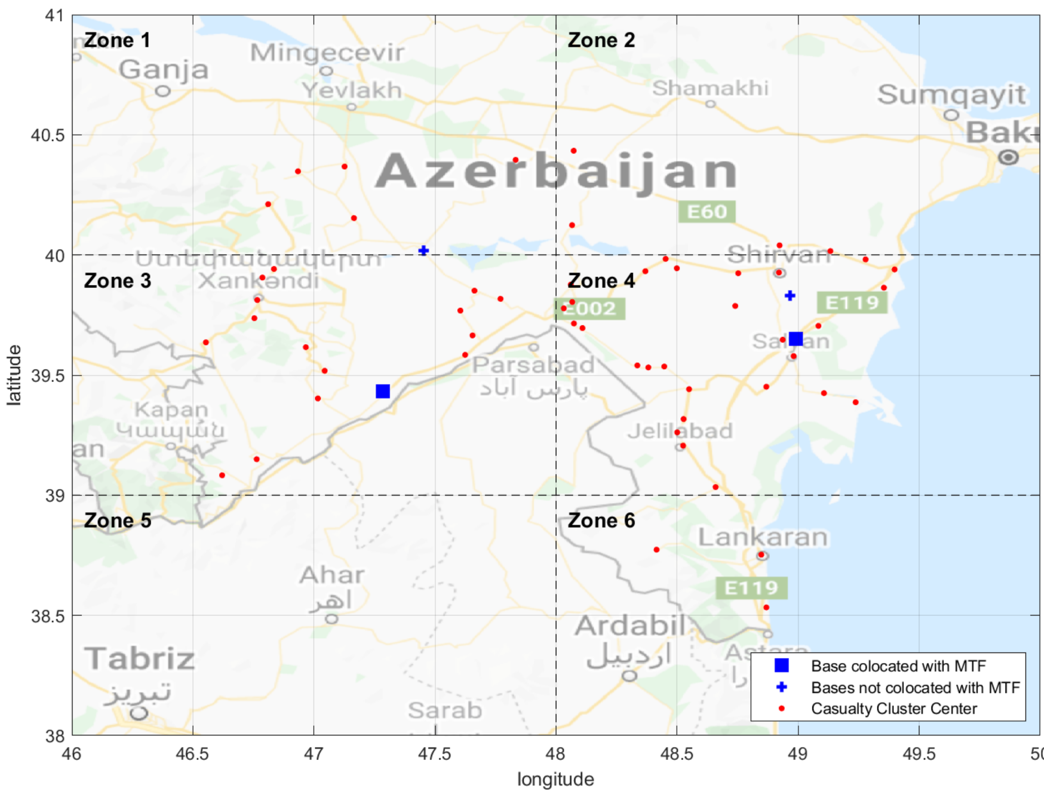

Figure 2: Notional Azerbaijan problem instance.

collocated with MTFs (blue crosses). The locations of the bases and casualty cluster centers are notional representations on where friendly and enemy forces are likely to be located. The problem instance consists of six zones, three priority classes (i.e., Priority I, II, and III), and four MEDEVAC units (one unit at each base). The MEDEVAC units stationed at the bases collocated with MTFs may be equipped with blood transfusion kits, but it is originally assumed that no units are equipped with blood transfusion kits. Subsequent analysis determines the performance trade-offs when one or both MEDEVAC units collocated with the MTFs have blood transfusion kits.

The problem instance assumes an arrival rate $\lambda = \frac{1}{60}$, which indicates an average of one service request per 60 minutes to be submitted to the MEDEVAC system. The likelihood of a request being any one of the three potential priority levels considered is the same (i.e., $\frac{1}{3}$). The speed of each MEDEVAC aircraft is set to 150 knots, which corresponds with the average airspeed of HH-60M Black Hawk aircraft. The expected proportion of MEDEVAC requests originating from each zone, the expected response time by MEDEVAC unit and zone, and the expected service time by MEDEVAC unit and zone are calculated via a Monte Carlo simulation in the same manner described in Jenkins et al. (2018). Each MEDEVAC request is either classified correctly or overestimated. For example, if a MEDEVAC request is reported as Priority I, the request may actually be Priority I or it may have been misclassified and is truly a Priority II or Priority III request. We set the missclassification rates to $\phi_{12} = \phi_{13} = 0.15$ and $\phi_{23} = 0.3$. We let δ_1, δ_2, and δ_3 equal 10, 1, and 0.01, respectively, to rank Priority I requests more important than Priority II requests and Priority II requests more important than Priority III requests. The discount factor is set to $\gamma = 0.99$ to motivate the system to position itself to efficiently respond to future service requests.

The MDP-generated optimal policy for the Azerbaijan problem instance is determined in approximately 32 minutes using MATLAB 2019a on a Intel Zeon E5-2687W workstation with 64GB RAM, 10 cores, and MATLAB Parallel Computing Toolbox. We compare the optimal policy to the currently employed MEDEVAC dispatching policy (i.e., the myopic policy), which immediately dispatches the nearest available MEDEVAC unit when a service request enters the system. When the myopic policy is implemented, the dispatching authority does not take system factors (e.g., priority level or number of units available) into consideration.

The optimal policy outperforms the myopic by 1.71% with regards to the objective function of maximizing expected total discounted reward. Because of the life or death nature of the MEDEVAC system, any improvement over the myopic policy is significant and may increase casualty survivability rates. The optimal policy differed from the myopic in 18,945 of the 132,055 states. That is, 14.3% of the optimal policy decisions are different than the myopic policy. Of the 18,945 differences, 15,969 (or 85.5%) of the differences resulted from the myopic policy recommending to send a unit to service an incoming request and the optimal policy recommending to reject the MEDEVAC request to an alternate form of evacuation. The remaining 14.5% are differences regarding which of the available units to dispatch. Table 1 illustrates the myopic and optimal decisions for five MEDEVAC system state scenarios. Scenario 1 represents a

Table 1: Comparison of MEDEVAC dispatching policies.

Scenario	$S_t = (M_t, R_t)$	Myopic Policy	Optimal Policy
1	$((0,1,1,1),(5,3,3))$	Dispatch Unit 1	Reject
2	$((4,0,0,1),(5,3,3))$	Dispatch Unit 2	Reject
3	$((5,6,0,5),(1,1,3))$	Dispatch Unit 3	Dispatch Unit 3
4	$((3,1,0,0),(3,2,3))$	Dispatch Unit 3	Dispatch Unit 4

MEDEVAC system state wherein 3 of the 4 MEDEVAC units are not available and a request is in the system for Zone 5 with a reported and true priority of Priority III. The myopic policy recommends dispatching the last remaining MEDEVAC unit, whereas the optimal policy recommends reserving MEDEVAC Unit 1 and rerouting the request to other forms of evacuation. The optimal policy reserves the remaining MEDEVAC for potential Priority I requests that may arrive before one of the other three MEDEVAC units completes its service request. Considering the rerouted request is Priority III (i.e., non-life-threatening), CASEVAC or other forms of evacuation can effectively and safely service the request while avoiding the possibility of a Priority I or Priority II (i.e., life-threatening) request being rerouted to other evacuation services. Scenario 2 represents a similar situation, but there are two MEDEVAC units available when the request arrives. The optimal policy recommends to not dispatch either of the units and the myopic recommends dispatching MEDEVAC Unit 2. This scenario emphasizes the optimal policy's tendency to reserve units when compared to the myopic policy of always dispatching a unit. Scenario 3 represents a similar situation to Scenario 1 in which there is one MEDEVAC available and a request arrives in the system. The request arrives from

Zone 1 and is reported as Priority I, but the true priority is Priority III. Because the request priority level is overestimated, the system falsely expects a bigger benefit to dispatching a MEDEVAC unit to service the request. Had the true priority level of routine been presented, the optimal policy may be to reserve the final MEDEVAC unit for a true Priority I request. This scenario represents the detriment to the system that triage classification errors cause. Finally, Scenario 4 represents the difference between the myopic and optimal policies regarding which unit to send to service an incoming request when two or more units are available. For this scenario, MEDEVAC Unit 3 and MEDEVAC Unit 4 are available and a request to service Zone 3 arrives in the system. The myopic policy recommends dispatching the MEDEVAC Unit 3 (i.e., the closest MEDEVAC unit), whereas the optimal policy recommends reserving MEDEVAC Unit 3 and instead dispatching MEDEVAC Unit 4 because the optimal policy takes into account possible future states and dispatching the nearest MEDEVAC unit does not always maximize performance.

In the MEDEVAC system, prompt service leads to better outcomes for service requests of all three priority levels, but prompt service is imperative for Priority I casualties due to the nature of their classification. If at least one MEDEVAC unit is available when a Priority I request is submitted to the system, both the myopic and optimal policies admit and service the request. Let the availability rate represent the expected proportion of time a specific number of MEDEVAC units is available. For example, if the availability rate for four units is 0.44, we expect that 44% of the time there are four units available to service an incoming request in the MEDEVAC system. Table 2 shows the proportion of time the system expects to have any number of units available for the Azerbaijan problem instance. A difference between the optimal

Table 2: MEDEVAC unit availability rates.

Policy	0	1	2	3	4
Myopic	0.023	0.092	0.224	0.364	0.298
Optimal	0.017	0.090	0.236	0.387	0.271

and myopic policies is that the optimal policy reserves a MEDEVAC unit for possible incoming Priority I requests. When implementing the optimal policy, the system has zero MEDEVAC units available 1.7% of the time. This indicates if a Priority I request is reported, only 1.7% of the time would we expect the MEDEVAC system to reject the request and redirect it to another form of evacuation. The myopic policy's availability rate for zero units is 2.3%. For both the myopic and optimal policy, three MEDEVAC units are available the largest amount of time. An availability rate with a right skew would represent a system that struggles to meet the demand in incoming service requests, whereas a rate with a significant left skew would represent a MEDEVAC system that does not efficiently use MEDEVAC units (i.e., they are sitting idle the majority of the time). Table 2 shows a good balance between the two extremes for both the myopic and optimal policies.

Four MEDEVAC units are considered in the Azerbaijan problem instance. Let the utilization rate for each MEDEVAC unit represent the expected proportion of time a MEDEVAC unit spends actively servicing requests. For example, a utilization rate of 0.4 for MEDEVAC Unit 3 indicates that we expect MEDEVAC Unit 3 to be servicing requests 40% of the time in the long run. A high utilization rate for MEDEVAC units may lead to increased wear and tear on equipment and overworking of MEDEVAC personnel. A low utilization rate may indicate that there are possibly too many MEDEVAC units assigned. Table 3 provides the utilization rate of each unit for both the myopic and optimal policies. Comparing the myopic and the

Table 3: MEDEVAC unit utilization rates.

Policy	1	2	3	4
Myopic	0.36	0.26	0.26	0.31
Optimal	0.31	0.26	0.29	0.34

optimal policies, the myopic policy has an unbalanced use of the MEDEVAC units. When following the myopic policy, MEDEVAC Unit 1 has the highest utilization rate and we expect it to be used at least 5%

more than the other three units. Alternatively, when following the optimal policy, MEDEVAC Unit 4 has the highest utilization rate we expect it to be used at least 3% more than the other three units. We expect each of the four MEDEVAC units to be be busy less than 40% of the time for the both myopic and optimal policies. The utilization rate for MEDEVAC Unit 1 is disproportionately higher than the other three units for the myopic policy because more casualties occur near MEDEVAC Unit 1 than any other unit. The optimal policy balances the burden of responding to MEDEVAC service requests more evenly across the four available units.

Misclassification rates introduce more uncertainty into the model. As the priority level of an incoming request is less certain, the MDP model adapts, taking into consideration the additional uncertainty. We explore four scenarios with different misclassification rates to gain insights as to how model performance is impacted. We expect the optimal policy will perform more similarly to the myopic policy as misclassification rates increase due to the incorrect information received. Table 4 compares the expected total discounted reward for each scenario implementing both the optimal and myopic policies as well as the percent improvement over the myopic policy. As the misclassification rates increase, the expected total discounted

Table 4: Comparison of expected total discounted reward for misclassification rate scenarios.

Scenario	Myopic Policy	Optimal Policy	% Improvement over Myopic
$\phi_{12} = \phi_{13} = \phi_{23} = 0$	8.25	8.43	2.18
$\phi_{12} = \phi_{13} = 0.15$ and $\phi_{23} = 0.3$	8.20	8.34	1.71
$\phi_{12} = \phi_{13} = 0.25$ and $\phi_{23} = 0.5$	8.17	8.29	1.47
$\phi_{12} = \phi_{13} = 0.40$ and $\phi_{23} = 0.8$	8.13	8.22	1.11

reward for both policies decreases, and the optimal policy percent improvement over the myopic policy for each scenario decreases. A higher misclassification rate means the system has to make less informed recommendations. When the system cannot differentiate between service request priority levels accurately, it acts similarly to the myopic policy. The expected total discounted reward does not decrease significantly for the increased misclassification rates because the system is not under high stress. The priority level of the incoming request is considered more important when there is only one or two MEDEVAC units available. If there are more than two units available, the optimal policy is more likely to recommend dispatching a MEDEVAC unit and performs more similarly to the myopic policy.

The presence of blood transfusion kits on board MEDEVAC units decreases the time until an injured individual receives necessary medical care. As more MEDEVAC units are equipped with blood transfusion kits, more individuals receive critical care earlier, and the overall efficiency of the MEDEVAC system increases. MEDEVAC Units 1 and 4 are collocated with MTFs, indicating that they may be equipped with blood transfusion kits and deliver life saving care earlier. We explore four scenarios, a scenario where neither MEDEVAC unit is equipped with blood transfusion kits, only MEDEVAC Unit 1 is equipped, only MEDEVAC Unit 4 is equipped, and both MEDEVAC Units 1 and 4 are equipped with blood transfusion kits. Table 5 compares the expected total discounted reward for each scenario implementing both the optimal and myopic policies as well as the percent improvement over the myopic policy. In an ideal environment,

Table 5: Comparison of expected total discounted reward for blood transfusion kit scenarios.

Scenario	Myopic Policy	Optimal Policy	% Improvement over Myopic
None Equipped	8.20	8.34	1.71
Unit 1 Equipped	8.66	9.02	4.23
Unit 4 Equipped	9.05	9.31	2.83
Units 1 & 4 Equipped	9.50	9.92	4.37

both MEDEVAC Units 1 and 4 would be equipped with blood transfusion kits as indicated in Table 5. More specifically, when both MEDEVAC Units 1 and 4 are equipped with blood transfusion kits, the model receives the highest expected total discounted reward and the highest percent improvement over the myopic policy. If resources are limited and only one blood transfusion kit is available, MEDEVAC Unit 4 should be equipped rather than MEDEVAC Unit 1. Although the percent improvement over the myopic policy is higher when MEDEVAC Unit 1 is equipped and MEDEVAC Unit 4 is not, the expected total discounted reward is more important and is higher when MEDEVAC Unit 4 is equipped and MEDEVAC Unit 1 is not.

5 CONCLUSIONS

The objective of this paper is to determine the optimal dispatching policy of MEDEVAC units to complete MEDEVAC service requests by examining the MEDEVAC dispatching problem. Improving the performance of the MEDEVAC system leads to higher efficiency and ultimately improves battlefield survivability rates. We develop a discounted, infinite-horizon MDP model to examine military medical planning scenarios. As an augmentation to previous research, this paper incorporates the possibility of triage classification errors and the placement of blood transfusion kits on board select MEDEVAC aircraft. Triage classification errors occur when the reported priority level of a casualty event is different from the true priority level, which is assessed when the MEDEVAC unit arrives at the casualty site. Blood transfusion kits allow for life saving care to be administered as soon as an injured individual is loaded onto the MEDEVAC aircraft as opposed to waiting until the MEDEVAC unit arrives back at the MTF and the injured individual(s) are offloaded. When a request arrives to the MEDEVAC system, the dispatching authority can accept or reject the request, and if it accepts the request, the dispatcher decides which of the available MEDEVAC units to dispatch. The dispatching authority considers the priority level and zone of the incoming request to determine the dispatching decision. Requests that are rejected from the system are serviced by other means of evacuation such as CASEVAC. This paper measures system performance based off of the true priority level of the incoming request and the response time of the MEDEVAC unit servicing the request. To explore the MDP model a notional scenario in Azerbaijan is created utilizing simulation of historical data.

When the system receives a service request the MDP model considers the reported priority and the probability that the priority classification level is reported incorrectly. The MEDEVAC system is not informed of the true priority level until a MEDEVAC unit is servicing a request and the immediate expected reward is calculated accordingly. The decisions from the MDP model are made when a service request arrives in the system or when a MEDEVAC unit completes a service request. The entire state of the system is considered when a decision is made by the dispatching authority.

The current policy in place is a myopic policy that always dispatches the nearest available MEDEVAC unit without considering priority level classification. Results indicate that this myopic policy is not optimal. The optimal policy, which considers the entire state of the MEDEVAC system (i.e., the MEDEVAC units' status, location of incoming request, priority level of incoming request, and probability of a triage classification error), increases the expected total discounted reward. The baseline scenario is 1.71% more efficient when the optimal policy is implemented when compared to the myopic policy. The improvement gaps between the myopic and optimal policies range from 1.11% to 4.37% via the scenarios examined. In the long run, the optimal policies will substantially increase the survivability rates of battlefield casualties and should be considered by MEDEVAC system planners.

As the misclassification rates of the MEDEVAC system increased, the difference between the myopic and optimal decreased. Additionally, the expected total discounted reward decreased but only slightly. The optimal policy's expected total discounted reward for when the total misclassification rate is zero is 8.43, and when the total misclassification rate increases to 0.8 the expected total discounted reward is 8.22, a decrease of only 0.19.

Including blood transfusion kits on board MEDEVAC Units 1 and/or 4 increases the expected total discounted reward. If only one unit may be equipped with a blood transfusion kit, we recommend MEDEVAC Unit 4, which leads to the greatest increase in expected total discounted reward. The best case scenario is to equip both aircraft with blood transfusion kits for an expected total discounted reward of 9.92.

This research is limited by the computational constraints of the MDP model. Whereas the Azerbaijan problem instance has 132,055 states and is tractable, any additional MEDEVAC aircraft and/or zones included in the model increase the computational complexity, and the MDP model may not be able to solve to completion in a tractable amount of time. The application of approximate dynamic programming (ADP), which has been implemented in similar research (e.g., Jenkins et al. (2021a)), may lead to near optimal solutions and solve the issue of computational complexity. Further research should focus on applying ADP to the MEDEVAC dispatching problem with triage classification errors and blood transfusion kits.

DISCLAIMER

The views expressed in this article are those of the authors and do not reflect the official policy or position of the United States Air Force, United States Department of Defense, or United States Government.

ACKNOWLEDGEMENTS

This research is partially supported by the Air Force Office of Scientific Research (AFOSR) under the Dynamic Data and Information Processing (DDIP) portfolio. Likewise, the authors thank the editorial board and anonymous reviewers for their constructive comments that helped improve the content and presentation of this paper.

REFERENCES

Bastian, N. D. 2010. "A robust, multi-criteria modeling approach for optimizing aeromedical evacuation asset emplacement". *The Journal of Defense Modeling and Simulation* 7(1):5–23.

Bastian, N. D., D. Brown, L. V. Fulton, R. Mitchell, W. Pollard, M. Robinson, and R. Wilson. 2013. "Analyzing the future of army aeromedical evacuation units and equipment: a mixed methods, requirements-based approach". *Military Medicine* 178(3):321–329.

Department of the Army, U. S. 2019, July. *Army Techniques Publication 4-02.2 Medical Evacuation*. Department of Defense.

Garrett, M. X. 2013. "USCENTCOM review of MEDEVAC procedures in Afghanistan". Technical report, United States Central Command.

Grannan, B. C., N. D. Bastian, and L. A. McLay. 2015. "A maximum expected covering problem for locating and dispatching two classes of military medical evacuation air assets". *Optimization Letters* 9(8):1511–1531.

Jenkins, P. R., B. J. Lunday, and M. J. Robbins. 2020a. "Aerial MEDEVAC Operations Decision-making under Uncertainty to Save Lives". *Phalanx* 53(1):63–66.

Jenkins, P. R., B. J. Lunday, and M. J. Robbins. 2020b. "Robust, Multi-Objective Optimization for the Military Medical Evacuation Location-Allocation Problem". *Omega* 97(2020):1–12.

Jenkins, P. R., M. J. Robbins, and B. J. Lunday. 2018. "Examining military medical evacuation dispatching policies utilizing a Markov decision process model of a controlled queueing system". *Annals of Operations Research* 271(2):641–678.

Jenkins, P. R., M. J. Robbins, and B. J. Lunday. 2021a. "Approximate Dynamic Programming for Military Medical Evacuation Dispatching Policies". *INFORMS Journal on Computing* 33(1):2–26.

Jenkins, P. R., M. J. Robbins, and B. J. Lunday. 2021b. "Approximate dynamic programming for the military aeromedical evacuation dispatching, preemption-rerouting, and redeployment problem". *European Journal of Operational Research* 290(1):132–143.

Jenkins, P. R., M. J. Robbins, and B. J. Lunday. 2021c. "Optimising Aerial Military Medical Evacuation Dispatching Decisions via Operations Research Techniques". *BMJ Military Health*:1–3.

Keneally, S. K., M. J. Robbins, and B. J. Lunday. 2016. "A Markov decision process model for the optimal dispatch of military medical evacuation assets". *Health Care Management Science* 19(2):111–129.

Kotwal, R. S., L. L. Scott, J. C. Janak, B. W. Tarpey, J. T. Howard, E. L. Mazuchowski, F. K. Butler, S. A. Shackelford, J. M. Gurney, and Z. T. Stockinger. 2018. "The effect of prehospital transport time, injury severity, and blood transfusion on survival of US military casualties in Iraq". *Journal of trauma and acute care surgery* 85(1S):S112–S121.

Kulkarni, V. 2017. *Modeling and Analysis of Stochastic Systems*. 3rd ed. North Carolina: Chapman and Hall.

Malsby III, R. F., J. Quesada, N. Powell-Dunford, R. Kinoshita, J. Kurtz, W. Gehlen, C. Adams, D. Martin, and S. Shackelford. 2013. "Prehospital blood product transfusion by US Army MEDEVAC during combat operations in Afghanistan: a process improvement initiative". *Military Medicine* 178(7):785–791.

Puterman, M. 1994. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. 2nd ed. New Jersey: Wiley-Interscience.

Rettke, A. J., M. J. Robbins, and B. J. Lunday. 2016. "Approximate dynamic programming for the dispatch of military medical evacuation assets". *European Journal of Operational Research* 254(3):824–839.

Robbins, M. J., P. R. Jenkins, N. D. Bastian, and B. J. Lunday. 2020. "Approximate dynamic programming for the aeromedical evacuation dispatching problem: Value function approximation utilizing multiple level aggregation". *Omega* 91:102020.

AUTHOR BIOGRAPHIES

EMILY S. GRAVES is an operations research analyst in the United States Air Force at Shaw Air Force Base, South Carolina. She earned her M.Sc. in Operations Research at the Air Force Institute of Technology (AFIT). Her e-mail address is emily.graves@afcent.af.mil.

PHILLIP R. JENKINS is an assistant professor of operations research in the Department of Operational Sciences at AFIT. He earned his Ph.D. in Operations Research at AFIT. His research interests include dynamic programming, approximate dynamic programming, machine learning, Markov decision processes, simulation, and optimization. His e-mail address is phillip.jenkins@afit.edu. His website is https://www.afit.edu/BIOS/bio.cfm?facID=613.

MATTHEW J. ROBBINS is an associate professor of operations research in the Department of Operational Sciences at AFIT. He earned his Ph.D. in Industrial Engineering from the University of Illinois. He currently serves as an associate editor for the *Military Operations Research* and *Naval Research Logistics* journals. His academic research interests include the application of computational stochastic optimization methods to the dispatch of medical evacuation assets, the routing of military distribution assets under uncertainty, the control of autonomous mobile sensors and weapons systems, and the management of military workforce systems. His e-mail address is matthew.robbins@afit.edu. His website is https://www.afit.edu/BIOS/bio.cfm?facID=618.