

Examining military medical evacuation dispatching policies utilizing a Markov decision process model of a controlled queueing system

Phillip R. Jenkins¹ · Matthew J. Robbins¹  · Brian J. Lunday¹

Published online: 1 February 2018
© US Government 2018

Abstract Military medical planners must develop dispatching policies that dictate how aerial medical evacuation (MEDEVAC) units are utilized during major combat operations. The objective of this research is to determine how to optimally dispatch MEDEVAC units in response to 9-line MEDEVAC requests to maximize MEDEVAC system performance. A discounted, infinite horizon Markov decision process (MDP) model is developed to examine the MEDEVAC dispatching problem. The MDP model allows the dispatching authority to accept, reject, or queue incoming requests based on a request's classification (i.e., zone and precedence level) and the state of the MEDEVAC system. A representative planning scenario based on contingency operations in southern Afghanistan is utilized to investigate the differences between the optimal dispatching policy and three practitioner-friendly myopic policies. Two computational experiments are conducted to examine the impact of selected MEDEVAC problem features on the optimal policy and the system performance measure. Several excursions are examined to identify how the 9-line MEDEVAC request arrival rate and the MEDEVAC flight speeds impact the optimal dispatching policy. Results indicate that dispatching MEDEVAC units considering the precedence level of requests and the locations of busy MEDEVAC units increases the performance of the MEDEVAC system. These results inform the development and implementation of MEDEVAC tactics, techniques, and procedures by military medical planners. Moreover, an analysis of solution approaches for the MEDEVAC dispatching problem reveals that the policy iteration algorithm substantially outperforms the linear programming algorithms executed by CPLEX 12.6 with regard to computational effort. This result supports the claim that policy iteration remains the superlative solution algorithm for exactly solving computationally tractable Markov decision problems.

Keywords Markov decision processes · Military medical evacuation (MEDEVAC) · Admission control · Queueing · Emergency medical service (EMS)

✉ Matthew J. Robbins
matthew.robbins@afit.edu

¹ Department of Operational Sciences, Air Force Institute of Technology, 2950 Hobson Way, Wright-Patterson AFB, OH 45433, USA

1 Introduction

The primary objective of a deployed military emergency medical services (EMS) system is to successfully evacuate casualties from the battlefield in a timely manner. Casualty evacuation (CASEVAC) and medical evacuation (MEDEVAC) are the two main options available for transporting combat casualties to a medical treatment facility (MTF). CASEVAC refers to the transport of casualties to an MTF via non-medical vehicles or aircraft *without* en route medical care by onboard medical professionals. Casualties transported via CASEVAC may not receive the necessary medical care nor be transported to an appropriate MTF. MEDEVAC refers to the transport of casualties to an appropriate MTF via standardized medical evacuation platforms with onboard medical professionals who are equipped to provide en route medical care and emergency medical intervention (Department of the Army 2016). As such, MEDEVAC is the preferred and primary method of transporting combat casualties.

Whereas MEDEVAC operations utilize several different types of evacuation platforms, this paper focuses on the aerial aspect of MEDEVAC operations (i.e., aeromedical helicopter operations). Helicopters have the capability and flexibility to fly directly to a predetermined casualty collection point (CCP), meet battlefield casualties when they are at their most vulnerable and critical stages, and either land in an area where no other platform (e.g., ground vehicle or fixed-wing aircraft) can or utilize a rescue hoist to lift casualties to the helicopter. After securing the casualties, helicopters can fly directly to dedicated trauma centers or hospitals, unencumbered by roads, at speeds often exceeding 150 miles per hour, all while providing definitive en route care via well trained and highly skilled medics (O'Shea 2011). These helicopter capabilities greatly contribute to recent increases in casualty survivability rates.

Helicopter ambulances were first introduced in the military during the Korean conflict and immediately became a high visibility asset of the MEDEVAC system. By the end of the Vietnam War, the capabilities of helicopters (i.e., speed and versatility) in austere conditions far exceeded the capabilities of ground platforms. The ability to travel across terrain in remote areas not accessible to ground vehicles makes helicopters well suited for MEDEVAC operations (De Lorenzo 2003; Clarke and Davis 2012). The United States Army operates HH-60M helicopters specifically designed for the MEDEVAC mission. HH-60M helicopters are equipped with the necessary resources (e.g., oxygen generator, integrated EKG machine, electronically controlled litters, built-in external hoist, and an infrared system that can locate patients by their body heat) to provide medical personnel the ability to simultaneously treat and transport casualties from a CCP to an appropriate MTF. The urgency of the MEDEVAC mission is critical to the survivability of battlefield casualties and the HH-60M helicopter has proved to be advantageous to the Army with its ability to lift-off on a mission within 7 min of notification (O'Shea 2011). The United States (U.S.) military recognizes the unique capabilities of MEDEVAC helicopters and utilizes them as the primary evacuation platform for battlefield casualties. For example, during the Afghanistan conflict between September 11, 2001 and March 31, 2014, the U.S. military incurred 21,089 casualties, of which 19,148 were transported via MEDEVAC helicopter (Kotwal et al. 2016). Eastridge et al. (2012) report that the survivability of combat casualties has continued to increase over time since World War II (WWII). Approximately 80% of casualties occurring on the battlefield survived in WWII, whereas 84% survived during the Vietnam War. An increase to 90% casualty survivability was observed in the continuous decade of United States' conflicts between 2001 and 2011. The improved casualty rates are attributed to improvements in the versatility and speed of

MEDEVAC helicopters and the resulting decrease in the time required for casualties to receive proper medical care (De Lorenzo 2003).

Military medical planners are responsible for designing deployed MEDEVAC systems. An effective and efficient MEDEVAC system boosts the *esprit de corps* of deployed military personnel, who understand that rapid and quality care will be provided if they are injured in combat (Department of the Army 2016). Important decisions include determining where to locate MEDEVAC units and MTFs, identifying a MEDEVAC dispatching policy, and recognizing when redeployment of aeromedical helicopters is necessary and possible. The location of MEDEVAC units is usually determined while considering two objectives: maximizing coverage and minimizing response time subject to logistical, resource, and force protection constraints. Deciding which MEDEVAC unit to dispatch to a given service request is a vital aspect of any EMS, including a MEDEVAC system, and is the primary focus of this paper. The military often defaults to a myopic dispatching policy wherein the closest available MEDEVAC unit is dispatched to retrieve combat casualties from a CCP regardless of the request's evacuation precedence category (e.g., Priority I—*Urgent*, Priority II—*Priority*, and Priority III—*Routine*). Redeployment of MEDEVAC units prior to returning to their originating base is possible but poses challenges due to the numerous resource and availability requirements (e.g., refueling, resupply, and armed escort). These reasons also render temporary relocation of idle MEDEVAC units uncommon within a theater of operations (Rettke et al. 2016).

This paper examines the MEDEVAC dispatching problem wherein a dispatching authority must decide which MEDEVAC unit to dispatch to a particular 9-line MEDEVAC request. The location of MTFs and MEDEVAC assets are known and all MEDEVAC helicopters are assumed to have the capability to meet the mission requirements of any 9-line MEDEVAC request. Redeployment is not considered. The reported dispatch policy is based on the location and status of MEDEVAC units, the location of the casualty event, and the evacuation precedence category of the casualty event.

An infinite horizon, discounted Markov decision process (MDP) model is formulated to determine how to optimally dispatch MEDEVAC helicopters to casualty events occurring in combat to maximize the expected total discounted reward attained by the system. A computational example is applied to a MEDEVAC system in Afghanistan in support of combat operations. Comparisons are made between myopic policies that are typically utilized in practice and the optimal policy derived from the formulated MDP model. Herein, we consider three specific myopic policies that, while adopting the rule of dispatching the closest available MEDEVAC unit to service a request, if all MEDEVAC units are busy, respectively queue any MEDEVAC requests, queue only urgent MEDEVAC requests, or reject (i.e., do not queue) any MEDEVAC requests.

An important difference between this paper and other papers in this research area is the incorporation of admission control and queueing. Consideration of admission control and queueing can greatly improve the performance of existing and proposed systems in many contexts, including manufacturing, distributed computing, and communications (Stidham 1985; Shenker and Weinrib 1989; Stidham and Weber 1993; Stidham 2002) and has yet to be examined in the context of MEDEVAC dispatching. Admission control allows the dispatching authority to observe the current state of the MEDEVAC system before making the decision to accept or reject an incoming request. This provides the dispatching authority the power to reject incoming requests, thereby reserving MEDEVAC units for higher precedence requests instead of satisfying all requests for service. The rejected requests are not simply discarded; rather, they are redirected to another supporting agency to be serviced (i.e., CASEVAC). If the dispatch authority allows a request to enter the MEDEVAC system but all MEDEVAC

units are currently servicing other requests, the entering request will be allocated to a queue based on its precedence level and location as categorized by geographic zone. Once a request has entered the system, it will be serviced; however, the dispatching authority dictates which available MEDEVAC unit will service each request in the system regardless of when the request entered the system. For example, an *urgent* request will be serviced before a *routine* request regardless of the order in which they entered the system. It is important to note that MEDEVAC units will not interrupt service to a request in the case of a higher precedence request arriving. Once a MEDEVAC unit is assigned a specific request, it will be considered unavailable until it completes the service of that request.

The remainder of this paper is organized as follows. Section 2 provides a review of research relating to MEDEVAC systems. Section 3 presents a description of the MEDEVAC dispatching problem. Section 4 describes the MDP formulation developed to determine an optimal MEDEVAC dispatch policy. Section 5 examines an application of the formulated MDP model based on a representative scenario in southern Afghanistan. Section 6 concludes the paper and proposes several directions for future research.

2 Literature review

For nearly half a century, research has been conducted on optimizing both civilian and military emergency medical services (EMS) response systems. The main features of this research include determining the location of servers; dictating the number of servers per location, the server dispatch policy, and the size and number of response zones (if a partitioning strategy for the service area is implemented); identifying which performance measure to focus on as the objective: response time thresholds (RTTs) or patient survivability rates; and recognizing if and when server relocation is necessary due to either a service completion or an incoming service request. Another complicating feature concerns the location of hospitals. In research examining civilian EMS systems, the locations of hospitals are usually given as fixed; however, in some military planning contexts the medical treatment facility (MTF) locations are not given. Military medical planners must decide where to best place MTF locations when designing a military medical evacuation (MEDEVAC) system (Rettke et al. 2016). Operations research (OR) methods provide rigorous, defensible, and quantitative insights to researchers examining EMS systems. Applied OR methods include stochastic modeling, queueing, discrete optimization, and simulation modeling (Green and Kolesar 2004).

The research presented in this paper examines the optimal dispatch of military EMS vehicles (i.e., HH-60M MEDEVAC helicopters) to prioritized requests for service. Consideration of the precedence category (e.g., Priority I—*Urgent*, Priority II—*Priority*, and Priority III—*Routine*) is important. A substantial amount of research seeks to improve the overall performance of EMS systems, but most research endeavors do not account for the precedence of the call (Bandara et al. 2014). When the precedence of the call is not considered, the default dispatching rule sends the closest available emergency response vehicle to satisfy required service requests with no regard as to how that specific vehicle's absence impacts the overall EMS system. Sending the closest available vehicle to a service request regardless of other factors (e.g., precedence, or severity) is commonly referred to as a myopic policy. Many researchers (Carter et al. 1972; Nicholl et al. 1999; Kuisma et al. 2004) show that myopic policies tend to be suboptimal. Incorporating precedence categories into the construction of dispatching policies can ultimately lead to more lives being saved on the battlefield.

EMS research exists that focuses specifically on military MEDEVAC systems. Zeto et al. (2006) develop a goal programming model that seeks to maximize the aggregate expected demands covered and minimize the spare capacities of air ambulances. The authors leverage the work of Alsalloum and Rand (2006) to examine both the problems of resource allocation and coverage in a three-phased approach. In the first phase, they characterize the demand for MEDEVAC missions using a multivariate hierarchical cluster analysis. In the second phase, they estimate the parameters of the model via a Monte Carlo simulation. In the third phase, they utilize a bi-criteria model to emplace the minimum number of required aircraft at each location to maximize the probability of meeting the MEDEVAC demand in the Afghanistan theater. Bastian et al. (2012) investigate the capabilities required for MEDEVAC aircraft platforms to successfully perform the necessary duties and provide coverage within a brigade operating space. The authors develop a decision support tool that military medical planners can utilize to analyze the risk associated with different MEDEVAC strategies. Fulton et al. (2009) evaluate the planning factors and rules of allocation associated with Army air ambulance companies. Military medical planners typically use the rules of allocation, which are based on strategic planning documents, to estimate the number of MEDEVAC units required for tactical and operational scenarios. The authors quantitatively analyze different rules through a Monte Carlo simulation and record the impact that they respectively have on major combat operations. The results indicate that 0.4 aircraft per admission would be a reasonable planning factor. Sundstrom et al. (1996) incorporate linear programming techniques to develop a model based on the probabilistic location set-covering problem that provides the required numbers of MEDEVAC assets needed as well as the optimal positioning of those assets to ensure orderly transport of battlefield casualties to an appropriate medical facility.

The allocation of MEDEVAC units during steady-state combat operations is studied by Fulton et al. (2010) and Bastian (2010). Fulton et al. (2010) formulate a stochastic optimization model that manages the locations of deployable military hospitals, hospital beds, and both aerial and ground MEDEVAC units prior to the reception of a 9-line MEDEVAC request. Their model uses an objective of minimizing the total travel time, weighted by the urgency level of the casualty, from the POI to an appropriate MTF. The weights associated with the urgency levels of casualties are derived from historical data of patient injury severity scores collected from Operation Iraqi Freedom (OIF) combat operations. Bastian (2010) formulates a stochastic optimization goal programming model to meet three separate objectives: maximize the coverage of theater-wide casualty demand in Afghanistan, minimize the spare capacity of MEDEVAC units, and minimize the maximal MTF evacuation site vulnerability to enemy attack. The aforementioned research endeavors alternatively focus on optimizing the location, allocation, or reallocation of MEDEVAC assets. Although such problems are important to consider, this research assumes the locations of MEDEVAC staging areas and MTFs to be fixed and the allocation of MEDEVAC helicopters to be both known and fixed throughout steady-state combat operations. These assumptions are reasonable and are adopted in other works focusing solely on MEDEVAC dispatching policies (e.g., see Keneally et al. (2016) and Rettke et al. (2016)).

Keneally et al. (2016) examine MEDEVAC dispatch policies in the Afghanistan theater via a Markov decision process (MDP) model. The authors assume that each service call arrives sequentially and the locations of each service center are predetermined. Their work classifies each service call into one of three evacuation precedence categories: *urgent*, *priority*, and *routine*. Moreover, they consider the possibility that an armed escort may be required to accompany the MEDEVAC unit. The authors utilize a reward function based on an RTT and conduct computational experiments wherein MEDEVAC units operate in support of Operation Enduring Freedom (OEF). The results highlight that the myopic policy (i.e., the

default policy in practice) is not always the optimal dispatching policy. This work herein extends research by Keneally et al. (2016) via the consideration of admission control and queueing. Moreover, this research measures performance via a survivability function rather than an RTT since survival probability more accurately represents casualty outcomes (Bandara et al. 2014). Grannan et al. (2015) develop a binary linear programming (BLP) model to determine where to locate and how to dispatch multiple types of military MEDEVAC air assets. A spatial queueing approximation model provides inputs to the BLP model. The BLP model incorporates the precedence of each service call to maintain a high likelihood of survival for the most urgent casualties. The overall objective is to maximize the proportion of high-precedence calls responded to within a pre-determined RTT.

Rettke et al. (2016) formulate an MDP model to examine the MEDEVAC dispatching problem. The problem instance size in their study is too large for an exact dynamic programming solution approach, so the authors employ approximate dynamic programming (ADP) techniques to determine a high-quality dispatch policy. The computational experiments in this study indicate that the authors' ADP-generated policy is nearly 31% better than the myopic policy. Military medical planners can use these results to improve existing MEDEVAC tactics and techniques. The problem instances in this research can be solved via exact dynamic programming methods and, therefore, do not utilize ADP techniques to generate MEDEVAC dispatching policies. Moreover, Rettke et al. (2016) assume that all incoming requests must be serviced if there are any MEDEVAC units available and queue incoming requests if all MEDEVAC units are busy. This paper relaxes the assumption that all incoming requests must be serviced and gives the dispatching authority the option to reject incoming requests based on the MEDEVAC system state. Lejeune and Margot (2016) propose a MEDEVAC model that considers endogenous uncertainty in the delivery times of casualties. The objective of their model is to provide prompt medical treatment and evacuation to soldiers injured in combat. The model determines where to locate MEDEVAC units and MTFs. Moreover, it helps the dispatch authority to determine which helicopters to dispatch and to which MTF each call should be transported to. Results indicate a reduction in battlefield deaths due to an increase in timely treatment of combat casualties when compared to a myopic policy. The dispatching policies generated by Lejeune and Margot (2016) assign response districts to each MEDEVAC staging area regardless of the system state. In contrast, the research herein utilizes the benefits of MDP models to determine the optimal dispatching decision for every feasible state in which the MEDEVAC system can be, considering all MEDEVAC assets in the enterprise.

3 Problem description

One of the primary missions of the Army Health System is to provide medical evacuation (MEDEVAC) across a wide range of military operations. The dedicated Army helicopters (i.e., rotary-wing aircraft or air ambulances) utilized in MEDEVAC missions are under the command of the general support aviation battalion (GSAB). Any use of air ambulances must first be coordinated with the supporting GSAB to synchronize evacuation procedures. The GSAB manages all activities related to the execution of aerial operations and serves as the primary decision-making authority for the military MEDEVAC system (Department of the Army 2016). An Army aeromedical evacuation officer (AEO) that works within the GSAB acts as the MEDEVAC dispatching authority in a deployed military emergency medical service (EMS) system (Fish 2014). AEOs direct the use of medical aircraft, personnel, and equipment in support of operational and strategic medical evacuations within a theater of operations.

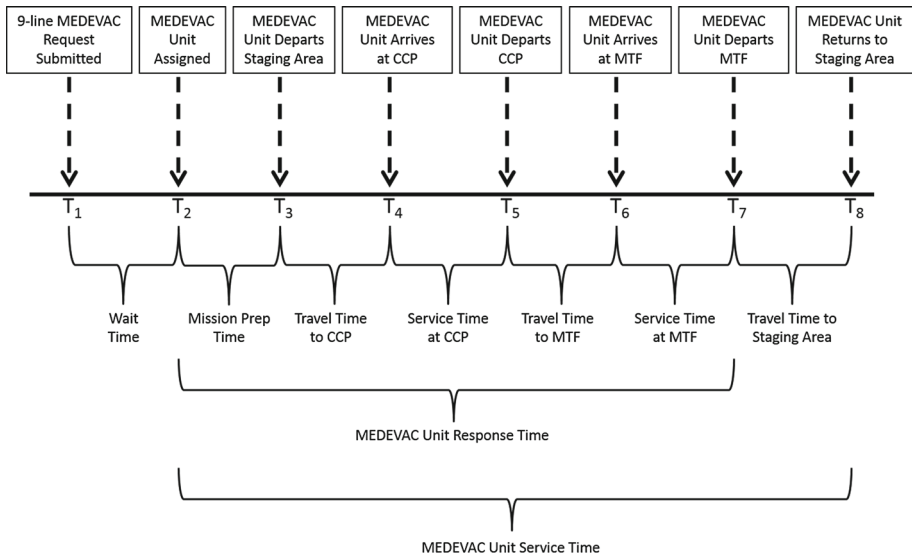


Fig. 1 MEDEVAC mission timeline

When a casualty event occurs and a 9-line MEDEVAC request is submitted, the AEO must decide quickly which MEDEVAC unit (if any) to dispatch. The casualty survivability rate will decrease if there are delays in decision making. To complicate matters further, there are many situations in which MEDEVAC units require a team of armed helicopters to escort them to the casualty site due to high threat-level conditions (e.g., enemy troops in the area). Armed escort requirements can potentially increase the overall response time, which ultimately decreases the chances of casualties surviving. Therefore, it is vital that the GSAB implements a dispatching policy resulting in rapid and high-quality transport of life-threatening battlefield casualties from a pre-determined casualty collection point (CCP) to the nearest, most appropriate medical treatment facility (MTF). The procedures outlined in the Army's Medical Evacuation Field Manual (Department of the Army 2016) and the graphical representations that Keneally et al. (2016) and Rettke et al. (2016) offer in their problem descriptions are utilized as a basis for the MEDEVAC mission timeline depicted in Fig. 1.

A 9-line MEDEVAC request is transmitted in a standardized message format with a prescribed amount of information that helps expedite the process of transporting casualties. When a 9-line MEDEVAC request is determined to be necessary, it should be transmitted over a secure communication system via a dedicated frequency. However, a 9-line MEDEVAC request can still be transmitted without such precautions if no secure communication systems are available. In wartime conditions, the information required in a 9-line MEDEVAC request is reported in the following order: the location of the pickup site (i.e., CCP), radio frequency and call sign, number of casualties by precedence, special equipment required, number of casualties by type, security of pickup site, method of marking pickup site, casualty nationality and status, and any chemical, biological, radiological, and nuclear contamination. The United States Army utilizes a three-category casualty triage rubric that governs the evacuation precedence of 9-line MEDEVAC requests. Priority I (i.e., *urgent*) and Priority II (i.e., *priority*) requests are life-threatening and must be serviced within 60 min and 4 h, respectively. Priority III (i.e., *routine*) requests are not life-threatening but still must be serviced within 24 h

(Department of the Army 2016). Either the senior military member or the senior medical person (if available) at the scene identifies the evacuation precedence category of each casualty and determines whether a 9-line MEDEVAC request is necessary. The tactical situation and the condition of each casualty are taken into consideration when making this decision. The overall precedence of a 9-line MEDEVAC request is based on the most time sensitive precedence among the casualties. Correct casualty event category identification is vital and cannot be overemphasized because mistakes may burden the evacuation system. Aerial ambulances are a low-asset, high-demand resource that must be managed accordingly.

In a combat situation, requests for MEDEVAC units are typically made at the point-of-injury (POI) once enemy fire has been suppressed. MEDEVAC requests are transmitted through several layers of command before reaching an AEO working within the GSAB headquarters. The specific information flow depends on the communication infrastructure within the command, the communication equipment available to the requesting unit, and the command and control organization of the MEDEVAC system (Rettke et al. 2016). Once the request has been made, casualties are transported to a CCP, which is a predesignated point along the evacuation route for collecting the wounded (Department of the Army 2000). The time at which the MEDEVAC request reaches the AEO is denoted by T_1 .

Once the GSAB receives the 9-line MEDEVAC request, the AEO must then decide whether to immediately assign a MEDEVAC unit to the request, depending on any pre-existing requests in the MEDEVAC system, the location of the pick-up site, the number and precedence of the casualties, and the status of the MEDEVAC units. If the MEDEVAC system is burdened with a high number of requests, the AEO may reject the incoming request from entering the system and redirect the request to be handled by casualty evacuation (CASEVAC). Assuming the request enters the system, the AEO will wait for a suitable MEDEVAC unit to become available. At time T_2 , the AEO assigns the MEDEVAC unit to service the request with an armed escort, if required.

The amount of time between an AEO receiving the 9-line MEDEVAC request, T_1 , and the assignment of the MEDEVAC unit, T_2 , is the total wait time for the request in the MEDEVAC system. As stated earlier, once a 9-line MEDEVAC request is received by the GSAB, the AEO must decide whether the request should enter the MEDEVAC system or the request should be serviced by another organization (i.e., CASEVAC). If the AEO allows the request to enter the MEDEVAC system and at least one suitable MEDEVAC unit is available to service the request, another decision must be made regarding whether the request should be assigned immediately or the request should be placed in a queue based on the its precedence category and location (i.e., zone). If the AEO allows a request to enter the MEDEVAC system and no suitable MEDEVAC units are available to service the request, then the request is placed in its respective zone-precedence queue. Figure 2 depicts the multiple-server, multiple-buffer queueing model employed in this paper. The MEDEVAC queueing system represented in Fig. 2 visually depicts the wait time between points T_1 and T_2 in Fig. 1.

Decision epochs occur when a 9-line MEDEVAC request is received by the GSAB or when a MEDEVAC unit completes a service request and becomes available. When a 9-line request is submitted and received by the GSAB, the AEO's decision consists of sending the just-arrived 9-line MEDEVAC request to its respective zone-precedence queue (if the queue is not full), immediately assigning an available MEDEVAC unit to service the request, or rejecting the request from ever entering the system. Once a MEDEVAC unit reaches service completion and at least one of the zone-precedence queues is not empty, the AEO must make a decision. The AEO's decision consists of either assigning a queued 9-line MEDEVAC request to one of the idle MEDEVAC units or waiting for either another (possibly higher precedence) request to enter the system or another MEDEVAC unit to reach service completion.

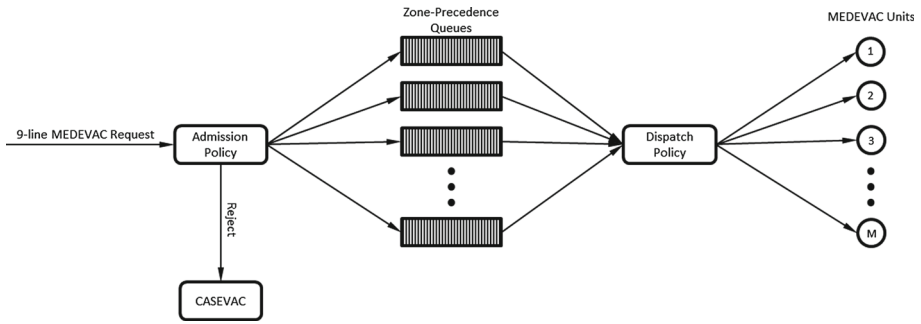


Fig. 2 MEDEVAC queuing system

The information from the 9-line MEDEVAC request is transmitted to the assigned MEDEVAC unit through the command’s communication system. T_3 denotes the time at which the assigned MEDEVAC unit departs its station for the CCP. The amount of time between the MEDEVAC unit being assigned the 9-line MEDEVAC request, T_2 , and the MEDEVAC unit departure, T_3 , is the total mission preparation time, which includes preparing the medical equipment, medical personnel, and helicopters for the MEDEVAC mission. Typically, if an armed escort is required, it will take off with the MEDEVAC unit at the staging area, but there are situations in which the MEDEVAC unit must meet an armed escort at a predetermined rally point en route to the CCP. The MEDEVAC unit cannot land at a high-threat level CCP site without an armed escort, and the additional coordination for an armed escort can increase total response time.

T_4 denotes the time at which the MEDEVAC unit lands at the CCP site. Upon arrival to the CCP site, the MEDEVAC unit immediately loads casualties and begins initial medical treatment. T_5 denotes the time at which the MEDEVAC unit departs the CCP site and proceeds towards an MTF. The destination MTF is selected in a deterministic manner based on sufficiency of medical capability to treat casualties and proximity to the CCP site. The sufficiently capable MTF that is located closest to the CCP site is the one that the MEDEVAC unit departs to at time T_5 .

The MEDEVAC unit arrives at the MTF site at time T_6 . After arriving, the MEDEVAC unit immediately begins to unload casualties and transfers the responsibility of subsequent care of the casualties to the medical staff at the MTF. After all casualties have been unloaded, the MEDEVAC unit departs the MTF and travels back to its own staging area. Once a MEDEVAC unit has finished unloading and transferring the subsequent care of casualties to the MTF medical staff, it must return to its own staging area before being tasked to service another 9-line MEDEVAC request. This requirement comes from concerns about low fuel levels, crew bed down limitations, on-board equipment configurations, and other logistical issues (Rettke et al. 2016). Typically, MEDEVAC units must return to their home staging areas to refuel before being dispatched for another mission. T_7 denotes the time at which the MEDEVAC unit departs the MTF.

The MEDEVAC unit arrives back at its staging area at time T_8 . Once the MEDEVAC unit arrives back at its staging area, the mission is considered complete. The MEDEVAC unit then becomes available for dispatch to another 9-line MEDEVAC request.

It is important to note that battlefield conditions (e.g., enemy disposition, required equipment being transported, weather conditions, and the air density due to flight altitude) are expected to affect the travel times from the MEDEVAC staging area to the CCP site, from the

CCP site to the selected MTF location, and from the MTF location back to the MEDEVAC staging area.

Military medical planners must consider the measurement of MEDEVAC system performance when examining dispatch policies. In civilian operations, the efficacy of EMS systems has been a difficult area to evaluate due to the multitude of variables present (MacFarlane and Benn 2003). The search for a reliable measure of performance remains a topic of interest in the EMS field (e.g., see McLay and Mayorga (2010)). Practitioners and researchers employ various means of assessment. The most common method for evaluating EMS systems utilizes ambulance response times. EMS systems commonly define the response time as the time required to reach the patient after receiving the emergency call. Since EMS systems are evaluated on response time, one of their primary focuses is the rapid response to cardiac arrest situations. This emphasis exists because the ability to provide effective treatment to patients undergoing cardiac arrest is time-sensitive. Another reason behind this rationale is as follows. If the EMS system has the capability to respond quickly to cardiac arrest patients, then it is more likely to be able to service similar life-threatening medical situations. Therefore, defining the response time for a civilian EMS system to be the time between receiving the emergency call and the time the first emergency response vehicle arrives on scene is quite intuitive.

Nonetheless, MEDEVAC system performance cannot be measured using the same evaluation criteria as the civilian EMS system. Several additional factors complicate the medical evacuation of a casualty from a battlefield. The travel times, load times, and unload times can be much greater and vary more in military EMS systems when compared to a civilian EMS system. Moreover, the primary cause of death for battlefield casualties is blood loss, not cardiac arrest. Garrett (2013) indicates that blood loss is the primary cause of death for nearly 85% of soldiers killed in action. Due to this issue, some MEDEVAC units have been recently equipped with in-flight blood transfusion capabilities; however, the majority are not, and there is a lack of data to confirm whether this addition improves the ability to handle casualties with severe blood losses (Malsby III et al. 2013). Without sufficient data to determine the effectiveness of in-flight transfusion, there has not been a change in the MEDEVAC system's evaluation measure. Therefore, unlike civilian EMS systems, it is vital to stabilize and transport battlefield casualties to an appropriate MTF (e.g., one that has the capability and resources to perform necessary care such as blood transfusions) and into surgery rather than simply providing medical aid at the CCP. So, while civilian EMS systems measure performance by response time (i.e., the time it takes to reach the patient after receiving the emergency call), military EMS systems are evaluated in terms of how long it takes to transport the casualties from the CCP to an MTF. Therefore, it is appropriate to define the response time for a MEDEVAC unit as $T_7 - T_2$. Moreover, the service time for a MEDEVAC unit is defined as $T_8 - T_2$, which is commonly associated as the time expended to service a request.

The primary objective of the MEDEVAC system presented in this paper is to dispatch MEDEVAC units in a manner that maximizes the expected total discounted reward attained by the system. The dispatch authority (i.e., AEO) must make sequential decisions under uncertainty regarding which available MEDEVAC unit to dispatch to service a 9-line MEDEVAC request. The system earns rewards based on the response times associated with servicing 9-line MEDEVAC requests. It is impossible to know exactly when and where casualty events will occur, which prevents the dispatch authority from having a priori information on subsequent 9-line MEDEVAC requests. The knowledge and details of any 9-line MEDEVAC request only become known to the MEDEVAC system upon receipt of the request. Once the GSAB receives the request and the AEO selects a MEDEVAC unit to dispatch, the assigned MEDEVAC unit must initiate mission protocols immediately. The mission protocols of a

MEDEVAC unit include preparing medical personnel and equipment prior to departure, traveling to the CCP to pick up casualties, providing appropriate en route medical care, and transporting casualties to the nearest MTF in a rapid and efficient manner. Delaying any mission tasks negatively impacts the total response time and ultimately decreases the survivability rates of casualties awaiting service.

Both a dynamic and stochastic approach are needed when analyzing the dispatch of either civilian or military emergency response vehicles. The stochastic aspect of this problem derives from the uncertainty concerning the manifestation of casualty events. Moreover, the dispatch, travel, and service times vary for each request and cannot be predicted precisely. When examining civilian EMS systems, the data relating to dispatch, travel, and service times are easily accessible and can be leveraged to parameterize decision models. Unfortunately, as noted earlier, one of the implicit challenges for military medical planners is having to develop and identify a dispatching policy prior to commencement of combat operations. No casualty event data exists for such a situation. Therefore, this paper utilizes a rubric that emulates the judgment and expertise of military planners with regard to the future interactions of enemy and friendly forces to identify the locations and arrivals of casualty events.

4 Methodology

This section presents the Markov decision process (MDP) model of the military's medical evacuation (MEDEVAC) dispatching problem. One of the key benefits of formulating an MDP model is that it provides a framework in which dynamic programming algorithms can be utilized to compute exact optimal policies. In most cases, MDP formulations have clear definitions for the state space, action space, rewards, transition probabilities, and optimality equations.

The objective of the MDP model formulated in this paper is to determine which available MEDEVAC unit to dispatch in response to a 9-line MEDEVAC request with the purpose of maximizing the expected total discounted reward over an infinite horizon.

The MDP model assumes that 9-line MEDEVAC requests arrive according to a Poisson process with parameter λ that is denoted by $PP(\lambda)$. Recall that a Poisson process possesses independent and stationary increments. The assumption of independent increments is reasonable in the context of MEDEVAC request arrivals because a large number of small, widely dispersed units perform combat operations that result in localized casualty events that are unrelated to one another, and therefore the numbers of arrivals that occur in disjoint time intervals are independent. The assumption of stationary increments is reasonable due to the underlying presumption that the implicit sizes, locations, and dispositions of friendly and adversary forces remain fixed with respect to time, and therefore the number of arrivals that occur in any interval of time depends only on the length of the time interval. Military medical planners must ensure the MEDEVAC system is tailored to effectively support friendly forces within an assigned area of operations (AO) (Department of the Army 2016). In large-scale combat operations, military medical planners should examine the expected conditions of the operation and carefully select an appropriate λ -value based on these conditions to investigate system performance during the peak hours of operation. Each casualty event that leads to a 9-line MEDEVAC request submission is categorized by its precedence level, which is determined by the senior military member and/or medical personnel present at the injury site.

The Army utilizes three casualty event precedence categories (i.e., *urgent*, *priority*, and *routine*) when submitting a 9-line MEDEVAC request (Department of the Army 2016). A routine evacuation precedence level is assigned to casualties that are triaged as minimally

injured (i.e., non-life-threatening), and typically results in standard ground or waterborne assets responding within 24 h of the initial event (De Lorenzo 2003). Since the focus of this paper is on the aerial aspect of MEDEVAC operations and routine 9-line requests typically do not utilize dedicated air evacuation assets, this paper only considers 9-line MEDEVAC requests that have a precedence level of either urgent or priority.

The arrival of urgent and priority 9-line MEDEVAC requests from different zones is modeled utilizing a *splitting* technique. Splitting consists of generating two or more counting processes out of a single Poisson process (Kulkarni 2009). Let the original counting process $\{N(t') : t' \geq 0\}$ denote the $PP(\lambda)$ that counts the number of 9-line MEDEVAC request arrivals to the general support aviation battalion (GSAB) that have occurred during the time interval $(0, t']$. The original counting process can be split into counting processes that are categorized by both the zone $z \in \mathcal{Z} = \{1, 2, \dots, |\mathcal{Z}|\}$ and the precedence level $k \in \mathcal{K} = \{1, 2, \dots, |\mathcal{K}|\}$ of the request. The sets \mathcal{Z} and \mathcal{K} represent the set of zones and the set of precedence levels in the system, respectively. Let $\mathcal{R} = \{(z, k) : (z, k) \in \mathcal{Z} \times \mathcal{K}\}$ be the set of request categories. There is a total of $|\mathcal{R}| = |\mathcal{Z}||\mathcal{K}|$ request categories. The original process $\{N(t') : t' \geq 0\}$ is split into $|\mathcal{R}|$ independent processes $\{N_{zk}(t') : t' \geq 0\}, \forall (z, k) \in \mathcal{R}$. It is clear that

$$N(t') = \sum_{(z,k) \in \mathcal{R}} N_{zk}(t') \quad (1)$$

since each request belongs to one and only one category. The nature of the split processes $\{N_{zk}(t') : t' \geq 0\}, \forall (z, k) \in \mathcal{R}$ depends on how the requests are categorized. The process of categorizing each request is called the splitting mechanism. The Bernoulli splitting mechanism generates the split processes $\{N_{zk}(t') : t' \geq 0\}, \forall (z, k) \in \mathcal{R}$, given parameters $p_{zk} > 0, \forall (z, k) \in \mathcal{R}$ such that $\sum_{(z,k) \in \mathcal{R}} p_{zk} = 1$. Each request is independently categorized by its zone z and precedence level k combination with probability p_{zk} independent of any other considerations. The splitting mechanism allows the characterization of each split process $\{N_{zk}(t') : t' \geq 0\}, (z, k) \in \mathcal{R}$ as a Poisson process with parameter λp_{zk} , which is denoted by $PP(\lambda p_{zk})$.

There may be times when a 9-line MEDEVAC request is admitted into the system, but all MEDEVAC units are currently servicing other requests. When this occurs, the submitted 9-line MEDEVAC request is placed in its respective zone-precedence queue to be serviced at a later time. Moreover, there may be system states wherein an idle MEDEVAC is available for assignment, but placing the submitted request in its respective zone-precedence queue rather than assigning the idle MEDEVAC to the request could prove more advantageous in the long run. For example, the decision not to assign an available MEDEVAC unit immediately could prove beneficial if a lower precedence request enters the system while many MEDEVAC units are busy. In such a situation, waiting for another MEDEVAC unit to become available before servicing the lower precedence request allows the idle MEDEVAC unit to remain available for a possibly higher precedence request, yet to arrive.

The service time for a MEDEVAC unit consists of the time between the initial assignment notification and the return to the staging area. This paper assumes that the service times of the MEDEVAC units are exponentially distributed. Kotwal et al. (2016) report real-world summary statistics concerning MEDEVAC service times that support this assumption. Moreover, the exponential distribution is commonly used to represent random, real-world phenomena because it provides a reasonable, simplifying approximation of the actual empirical distribution and enables the construction of a tractable mathematical model due to its ease of use and favorable properties (e.g., the memoryless property). Indeed, this simplifying

assumption is often utilized (and investigated) in related literature. For example, Jarvis (1985) performs several computational experiments, and the results suggest that the shape of the service-time distribution has little impact on the overall behavior of the system. Similarly, research by Gross and Harris (1998) also indicates the insensitivity of service time distributions to system performance. McLay and Mayorga (2013) perform simulation analyses utilizing different types of service time distributions to study the impact of modeling the system with exponential service times versus more realistic service times. Results indicate that the assumption of exponential service times does not significantly impact the optimal policies. This suggests that the optimal policies determined utilizing the MDP model from this paper provide military medical planners relevant insight regarding how to dispatch MEDEVAC units despite the simplifying assumption of exponentially distributed service times.

Having introduced the characteristics of the arrival process and the nature of the service times, formulation of the MDP model can now proceed. The development of the MDP model components leverage Maxwell et al. (2010), Keneally et al. (2016), and Rettke et al. (2016). The decision epochs, state space, action space, transition probabilities, rewards, objective, and optimality equation are described in detail below.

The decision epochs of the MEDEVAC system are the points in time that require a decision. The set of decision epochs is denoted as $\mathcal{T} = \{1, 2, \dots\}$. Two event types in the MEDEVAC system constitute all decision epochs. The first event type is the submission of a 9-line MEDEVAC request. The second event type is the change in the status of a MEDEVAC unit from busy to available upon completing a mission.

The MEDEVAC system MDP model follows the properties of semi-Markov decision processes (SMDPs). SMDPs generalize MDPs by requiring the decision-maker to select a feasible action whenever the system changes, allowing the time spent in a specific state to follow an arbitrary probability distribution, and modeling the system evolution in continuous time (Puterman 1994). The MEDEVAC system MDP model is viewed as a continuous time MDP (CTMDP), which is a special case of an SMDP wherein the inter-transition times are exponentially distributed and decisions are made at every transition. There are several different ways that CTMDPs can be analyzed, but the primary method utilized in this paper is *uniformization*. Uniformization is applied to the CTMDP model to obtain an equivalent discrete-time discounted model with constant transition rates (Puterman 1994). The transformation allows the results and algorithms for discrete-time MDP models to be applied directly.

The state $S_t \in \mathcal{S}$ describes the status of the entire MEDEVAC system at decision epoch $t \in \mathcal{T}$. The MEDEVAC system state is represented by the tuple $S_t = (M_t, Q_t, \hat{R}_t)$ wherein M_t represents the MEDEVAC status tuple at epoch t , Q_t represents the queue status tuple at epoch t , and \hat{R}_t represents the request arrival status tuple at epoch t .

The MEDEVAC status tuple M_t describes the status of every MEDEVAC unit in the system at epoch t . The tuple M_t can be written as

$$M_t = (M_{tm})_{m \in \mathcal{M}}, \quad (2)$$

where $\mathcal{M} = \{1, 2, \dots, |\mathcal{M}|\}$ represents the set of MEDEVAC units in the system. The state variable $M_{tm} \in \{0\} \cup \mathcal{Z}$ contains the information pertaining to MEDEVAC unit $m \in \mathcal{M}$ at epoch t . Each MEDEVAC unit can either be idle or servicing a request in one of the zones in the system. When $M_{tm} = 0$, MEDEVAC unit m is idle. When $M_{tm} = z$, MEDEVAC unit m is servicing a request from zone $z \in \mathcal{Z}$.

The queue status tuple Q_t describes the status of every zone-precedence queue in the system at epoch t . The tuple Q_t can be written as

$$Q_t = (Q_{tzk})_{z \in \mathcal{Z}, k \in \mathcal{K}} \tag{3}$$

The state variable $Q_{tzk} \in \{0, 1, \dots, q^{max}\}$ contains the information pertaining to the $(z, k) \in \mathcal{R}$ zone-precedence queue at epoch t . Each zone-precedence queue can hold no more than q^{max} requests at any point in time.

The request arrival status tuple \hat{R}_t indicates whether there is a request arrival awaiting an admission decision at epoch t ; it also provides the zone and precedence level of the request arrival, if one is present at epoch t . Let $\hat{R}_t = (0, 0)$ when there is not a request arrival at the GSAB at epoch t . Otherwise, let

$$\hat{R}_t = (\hat{Z}_t, \hat{K}_t)_{\hat{Z}_t \in \mathcal{Z}, \hat{K}_t \in \mathcal{K}} \tag{4}$$

The random variable \hat{Z}_t represents the zone of the request arrival at epoch t , and the random variable \hat{K}_t represents the precedence level of the request arrival at epoch t . At epoch t , the information in \hat{Z}_t and \hat{K}_t has just been realized and is no longer uncertain. However, \hat{Z}_t and \hat{K}_t are random variables at epochs $1, 2, \dots, t - 1$ because the information they contain is still uncertain at those epochs.

The size of the state space \mathcal{S} depends on $|\mathcal{M}|, |\mathcal{Z}|, |\mathcal{K}|$, and q^{max} . The following expression indicates the cardinality of the state space for the MEDEVAC system:

$$|\mathcal{S}| = (1 + |\mathcal{Z}|)^{|\mathcal{M}|} (1 + q^{max})^{|\mathcal{Z}||\mathcal{K}|} (1 + |\mathcal{Z}||\mathcal{K}|) \tag{5}$$

The size of the state space grows exponentially with respect to the number of state variables. This is commonly referred to as the curse of dimensionality and renders dynamic programming intractable for analyzing practical scenarios (i.e., large-scale problem instances). The purpose of formulating and analyzing small-scale problem instances is to examine the general efficacy of currently practiced (myopic) policies, identify possible structural properties of high-quality solutions, and inform the subsequent development of approximate solution approaches for application to the analysis of large-scale problems.

Events are triggered when a 9-line MEDEVAC request is submitted to the system or if a busy MEDEVAC unit completes a service request and becomes available. An admission control decision only occurs when a 9-line MEDEVAC request is submitted to the system. A dispatching decision may be necessary when either of these two event types occur.

The MEDEVAC system employs an inter-zone policy regarding airspace access that allows any MEDEVAC unit to service any 9-line MEDEVAC request, regardless of the zone from which the request originated. Once a MEDEVAC unit is tasked, it will be considered unavailable until the task is completed and the MEDEVAC unit has returned to its own staging area. Although rerouting a MEDEVAC unit during mid-flight can be accomplished, potential delays and communication difficulties can create issues in the MEDEVAC system that may ultimately cost casualties their lives. Furthermore, most military operations do not utilize a MEDEVAC unit rerouting strategy during combat operations (Rettke et al. 2016). Due to these reasons, rerouting MEDEVAC units mid-flight is not incorporated in this MDP model.

When a 9-line MEDEVAC request is submitted, the AEO must take into account the current state of the system and make an admission control and possibly a dispatching decision. There are three possible alternatives: allowing the request to enter its respective zone-precedence queue; assigning an available MEDEVAC unit to service the request immediately; or rejecting the request from entering the system, which forces the request to be serviced by an outside

agency (i.e., CASEVAC). If a request arrival is present at epoch t and its queue is not full, i.e., $\hat{R}_t = (\hat{Z}_t, \hat{K}_t)$ and $Q_{t\hat{Z}_t\hat{K}_t} < q^{max}$, $\hat{Z}_t \in \mathcal{Z}$, $\hat{K}_t \in \mathcal{K}$, then the AEO can either accept or reject the request from entering the system. If the request is accepted, it can either be placed in its respective zone-precedence queue or an available MEDEVAC unit can be tasked to service the request immediately. Moreover, if a request arrival is present at epoch t and its queue is full, i.e., $\hat{R}_t = (\hat{Z}_t, \hat{K}_t)$ and $Q_{t\hat{Z}_t\hat{K}_t} = q^{max}$, $\hat{Z}_t \in \mathcal{Z}$, $\hat{K}_t \in \mathcal{K}$, then the AEO must reject the request from entering the system. Practically speaking, q^{max} should be set high enough so that requests are not routinely rejected due to a full queue.

Let the decision variable $x_t^{reject} \in \{\Delta, 0, 1\}$ denote the admission control decision at epoch t . If an arrival request is not present at epoch t , i.e., $\hat{R}_t = (0, 0)$, the only available decision is $x_t^{reject} = \Delta$, which indicates the system will continue to transition without any impact from x_t^{reject} . When $x_t^{reject} = 0$, the arrival request at epoch t is admitted to the MEDEVAC system, whereas when $x_t^{reject} = 1$, the arrival request at epoch t is rejected from entering the MEDEVAC system.

Dispatching decisions may be required when either a 9-line request is submitted or a busy MEDEVAC unit completes a service request and becomes available. Let $\mathcal{I}(S_t) = \{m : m \in \mathcal{M}, M_{tm} = 0\}$ denote the set of idle MEDEVAC units available for dispatching when the state of the system is S_t at epoch t . Let $\mathcal{W}(S_t) = \{(z, k) : (z, k) \in \mathcal{R}, Q_{tzk} > 0\}$ denote the set of zone-precedence queues that have at least one casualty event awaiting service when the state of the system is S_t at epoch t . The dispatching decision is represented by the tuple $x_t^d = (x_t^{ar}, x_t^{qr})$ wherein x_t^{ar} represents the arrival request dispatch decision tuple and x_t^{qr} represents the queued requests dispatch decision tuple at epoch t .

The arrival request dispatch decision tuple x_t^{ar} describes the AEO’s dispatching decision with regard to arrival requests at epoch t . The tuple x_t^{ar} can be written as

$$x_t^{ar} = (x_{tm}^{ar})_{m \in \mathcal{I}(S_t)}. \tag{6}$$

The decision variable $x_{tm}^{ar} = 1$ if MEDEVAC unit $m \in \mathcal{I}(S_t)$ is dispatched to service the arrival request $\hat{R}_t = (\hat{Z}_t, \hat{K}_t)$, where $\hat{Z}_t \in \mathcal{Z}$ and $\hat{K}_t \in \mathcal{K}$, at epoch t , and 0 otherwise.

The queued requests dispatch decision tuple, x_t^{qr} , describes the AEO’s dispatching decision with regard to queued requests at epoch t . The tuple x_t^{qr} can be written as

$$x_t^{qr} = (x_{tmzk}^{qr})_{m \in \mathcal{I}(S_t), (z,k) \in \mathcal{W}(S_t)}. \tag{7}$$

The decision variable $x_{tmzk}^{qr} = 1$ if MEDEVAC unit $m \in \mathcal{I}(S_t)$ is dispatched to service a queued request from the (z, k) zone-precedence queue, where $(z, k) \in \mathcal{W}(S_t)$, at epoch t , and 0 otherwise.

Let $x_t = (x_t^{reject}, x_t^d)$ denote a compact representation of the decision variables at epoch t . Several constraints bound the decisions being made at epoch t . The first constraint,

$$I_{\{\hat{R}_t \neq (0,0)\}} \sum_{m \in \mathcal{I}(S_t)} x_{tm}^{ar} + \sum_{m \in \mathcal{I}(S_t)} \sum_{(z,k) \in \mathcal{W}(S_t)} x_{tmzk}^{qr} \leq 1, \tag{8}$$

requires that there is at most one MEDEVAC unit dispatched at epoch t . The next constraint,

$$x_t^{reject} \leq 1 - \sum_{m \in \mathcal{I}(S_t)} x_{tm}^{ar}, \tag{9}$$

indicates that, if an arrival request is present at epoch t and a MEDEVAC unit is tasked to service the arrival request at epoch t , as indicated by $x_{tm}^{ar} = 1$ for some $m \in \mathcal{I}(S_t)$, then the

arrival request must enter the system, as indicated by $x_t^{reject} = 0$. Otherwise, $x_{tm}^{ar} = 0$ for all $m \in \mathcal{I}(S_t)$, and the arrival request is either queued (i.e., $x_t^{reject} = 0$) or rejected (i.e., $x_t^{reject} = 1$) from the system at epoch t . The set of available actions when a decision is required is denoted as follows

$$\mathcal{X}(S_t) = \begin{cases} (\Delta, (\{0\}^{|\mathcal{I}(S_t)|}, \{0, 1\}^{|\mathcal{I}(S_t)| \times |\mathcal{W}(S_t)|})), & \text{if } \hat{R}_t = (0, 0), \mathcal{I}(S_t) \neq \emptyset, \mathcal{W}(S_t) \neq \emptyset \\ (\Delta, (\{0\}^{|\mathcal{I}(S_t)|}, \{0\}^{|\mathcal{I}(S_t)| \times |\mathcal{W}(S_t)|})), & \text{if } \hat{R}_t = (0, 0), \mathcal{I}(S_t) \neq \emptyset, \mathcal{W}(S_t) = \emptyset \\ (1, (\{0\}^{|\mathcal{I}(S_t)|}, \{0, 1\}^{|\mathcal{I}(S_t)| \times |\mathcal{W}(S_t)|})), & \text{if } \hat{R}_t = (\hat{Z}_t, \hat{K}_t), Q_{t, \hat{Z}_t, \hat{K}_t} = q^{max}, \mathcal{I}(S_t) \neq \emptyset \\ (1, (\{0\}^{|\mathcal{I}(S_t)|}, \{0\}^{|\mathcal{I}(S_t)| \times |\mathcal{W}(S_t)|})), & \text{if } \hat{R}_t = (\hat{Z}_t, \hat{K}_t), Q_{t, \hat{Z}_t, \hat{K}_t} = q^{max}, \mathcal{I}(S_t) = \emptyset \\ (\{0, 1\}, (\{0, 1\}^{|\mathcal{I}(S_t)|}, \{0, 1\}^{|\mathcal{I}(S_t)| \times |\mathcal{W}(S_t)|})), & \text{if } \hat{R}_t = (\hat{Z}_t, \hat{K}_t), Q_{t, \hat{Z}_t, \hat{K}_t} < q^{max}, \mathcal{I}(S_t) \neq \emptyset, \mathcal{W}(S_t) \neq \emptyset \\ (\{0, 1\}, (\{0\}^{|\mathcal{I}(S_t)|}, \{0\}^{|\mathcal{I}(S_t)| \times |\mathcal{W}(S_t)|})), & \text{if } \hat{R}_t = (\hat{Z}_t, \hat{K}_t), Q_{t, \hat{Z}_t, \hat{K}_t} < q^{max}, \mathcal{I}(S_t) = \emptyset \\ (\{0, 1\}, (\{0, 1\}^{|\mathcal{I}(S_t)|}, \{0\}^{|\mathcal{I}(S_t)| \times |\mathcal{W}(S_t)|})), & \text{if } \hat{R}_t = (\hat{Z}_t, \hat{K}_t), \mathcal{I}(S_t) \neq \emptyset, \mathcal{W}(S_t) = \emptyset \end{cases} \tag{10}$$

where Constraints (8) and (9) must be satisfied. The first two cases in Eq. (10) represent all feasible actions when the decision epoch occurs due to a MEDEVAC unit completing a service request and becoming available, whereas the last five cases represent all feasible actions when the decision epoch occurs due to a 9-line MEDEVAC request submission.

State transitions are Markovian with two possible events dictating the transition. The first event type is the submission of a 9-line MEDEVAC request. Recall that 9-line MEDEVAC requests arrive according to a $PP(\lambda)$. The second event type is the change in the status of a MEDEVAC unit from busy to available upon completing a mission. Let μ_{mz} denote the service rate of MEDEVAC unit $m \in \mathcal{M}$ when servicing a 9-line MEDEVAC request in zone $z \in \mathcal{Z}$. Let $\mathcal{B}(S_t) = \{m : m \in \mathcal{M}, M_{tm} \neq 0\}$ denote the set of busy MEDEVAC units when the state of the system is S_t at epoch t . If the MEDEVAC system is in pre-decision state S_t and action x_t is taken, the system will immediately transition to a post-decision state S_t^x . The sojourn time in S_t^x (i.e., the time the system remains in post decision state S_t^x before transitioning to the next pre-decision state S_{t+1}) follows an exponential distribution with parameter $\beta(S_t, x_t)$. Simple calculations reveal that

$$\beta(S_t, x_t) = \lambda + \sum_{m \in \mathcal{B}(S_t)} \mu_{m, M_{tm}} + \sum_{m \in \mathcal{I}(S_t)} \mu_{m, \hat{Z}_t} x_{tm}^{ar} + \sum_{m \in \mathcal{I}(S_t)} \sum_{(z, k) \in \mathcal{W}(S_t)} \mu_{mz} x_{tmzk}^{qr} \tag{11}$$

If $\mathcal{B}(S_t) = \emptyset$, $x_{tm}^{ar} = 0 \forall m \in \mathcal{I}(S_t)$, and $x_{tmzk}^{qr} = 0 \forall m \in \mathcal{I}(S_t), (z, k) \in \mathcal{W}(S_t)$, then $\beta(S_t, x_t)$ represents the sojourn time for the state-action pairs for which the next decision epoch occurs upon the arrival of a 9-line MEDEVAC request. Otherwise, $\beta(S_t, x_t)$ represents the sojourn time for the state-action pairs for which the next decision epoch occurs after either a 9-line MEDEVAC request arrives to the GSAB or one of the busy MEDEVAC units completes a service request and becomes available. Let T_a denote the time until the next 9-line MEDEVAC request arrival. Let T_s denote the time until the next service completion. The time until the next decision epoch T_e satisfies $T_e = \min\{T_a, T_s\}$. Since both T_a and T_s follow an exponential distribution, standard calculations show that T_e follows an exponential distribution with parameter $\beta(S_t, x_t)$.

The probabilistic behavior of the process is summarized in terms of its infinitesimal generator. The infinitesimal generator is an $|\mathcal{S}| \times |\mathcal{S}|$ matrix G with components:

$$G(S_{t+1}|S_t, x_t) = \begin{cases} -[1 - p(S_t|S_t, x_t)]\beta(S_t, x_t), & \text{if } S_{t+1} = S_t \\ p(S_{t+1}|S_t, x_t)\beta(S_t, x_t), & \text{if } S_{t+1} \neq S_t \end{cases} \tag{12}$$

wherein

$$p(S_{t+1}|S_t, x_t) = \begin{cases} \frac{\lambda_{zk}}{\beta(S_t, x_t)}, & \text{if } \hat{R}_{t+1} = (z, k), z \in \mathcal{Z}, k \in \mathcal{K} \\ \frac{\mu_{mz}}{\beta(S_t, x_t)}, & \text{if } \hat{R}_{t+1} = (0, 0), M_{t+1,m} = 0, M_{tm}^x = z, m \in \mathcal{M}, z \in \mathcal{Z} \\ 0, & \text{otherwise} \end{cases} \quad (13)$$

denotes the probability that the system transitions to state S_{t+1} given that it is currently in state S_t and decision x_t is made. The post-decision state variable $M_{tm}^x \in \{0\} \cup \mathcal{Z}$ contains the information pertaining to MEDEVAC unit $m \in \mathcal{M}$ when decision x_t is made at epoch t . Note that $p(S_t|S_t, x_t) = 0$, which means that the system will transition to a different state at the end of a sojourn in state S_t^x .

Puterman (1994) argues that converting CTMDPs to equivalent discrete-time MDPs via the uniformization approach makes subsequent analysis easier to perform. To uniformize the system, the maximum rate of transition must be determined and is calculated by

$$v = \lambda + \sum_{m \in \mathcal{M}} \tau_m, \quad (14)$$

wherein

$$\tau_m = \max_{z \in \mathcal{Z}} \mu_{mz}, \quad \forall m \in \mathcal{M}. \quad (15)$$

The restriction that there are no self-transitions from a state to itself is removed when uniformization is applied to the process. Applying uniformization yields the following transition probabilities:

$$\tilde{p}(S_{t+1}|S_t, x_t) = \begin{cases} 1 - \frac{[1-p(S_t|S_t, x_t)]\beta(S_t, x_t)}{p(S_{t+1}|S_t, x_t)\beta(S_t, x_t)}, & \text{if } S_{t+1} = S_t \\ \frac{p(S_{t+1}|S_t, x_t)\beta(S_t, x_t)}{v}, & \text{if } S_{t+1} \neq S_t. \end{cases} \quad (16)$$

This transformation may be viewed as inducing extra (i.e., “notional”) transition opportunities from a state to itself. This modified process has the same probabilistic structure as the CTMDP.

The decision epochs in CTMDPs follow each state transition, and the times between decision epochs are exponentially distributed. Several factors impact the amount of reward gained from making a decision to service a 9-line MEDEVAC request. These factors include the zone and precedence level of the 9-line MEDEVAC request as well as the staging area of the servicing MEDEVAC unit. Let $c(S_t, x_t) = \psi_{mzk}$ denote the immediate expected reward (i.e., contribution) if MEDEVAC unit $m \in \mathcal{M}$ is dispatched to service a zone $z \in \mathcal{Z}$, precedence level $k \in \mathcal{K}$ 9-line MEDEVAC request (i.e., $x_{tm}^{ar} = 1$ or $x_{tmzk}^{qr} = 1$). The immediate expected reward is computed as follows:

$$\psi_{mzk} = \begin{cases} \delta e^{-\frac{\zeta_{mz}}{60}}, & \text{if } k = 1 \text{ (i.e., urgent)} \\ e^{-\frac{\zeta_{mz}}{240}}, & \text{if } k = 2 \text{ (i.e., priority)} \\ 0, & \text{otherwise,} \end{cases} \quad (17)$$

wherein ζ_{mz} is the expected response time when MEDEVAC $m \in \mathcal{M}$ is dispatched to service a request in zone $z \in \mathcal{Z}$, and $\delta \geq 1$ is a tradeoff parameter utilized to vary the urgent-to-priority immediate expected reward ratio. If a MEDEVAC unit is not dispatched to service a 9-line MEDEVAC request at epoch t , then $c(S_t, x_t) = 0$.

Let $h(S_t, x_t)$ denote the holding cost accumulated when the MEDEVAC system is in state S_t and decision x_t is selected. The MEDEVAC system incurs a holding cost for queued 9-line MEDEVAC requests based on the time requirements outlined in the Army’s Medical Evacuation Field Manual (Department of the Army 2016). The MEDEVAC system seeks to

service urgent and priority 9-line MEDEVAC requests within 60 and 240 min from notification, respectively. Let ϕ_k denote the holding cost rate for holding a single precedence- k request in its queue between decision epochs. The holding cost rate ϕ_k is defined as

$$\phi_k = \xi \frac{\sum_{m \in \mathcal{M}} \sum_{z \in \mathcal{Z}} \psi_{mzk}}{|\mathcal{M}| |\mathcal{Z}|}, \forall k \in \mathcal{K}, \tag{18}$$

where $\xi \in [0, 1]$ is a parameter that scales the holding cost rate for a precedence- k request based on the average immediate expected reward over all possible MEDEVAC-zone combinations. Summing the holding costs over all zone-precedence queues yields the following expression

$$h(S_t, x_t) = \sum_{z \in \mathcal{Z}} \sum_{k \in \mathcal{K}} \phi_k Q_{tzk}. \tag{19}$$

Simple calculations show that, if $\mathcal{W}(S_t) = \emptyset$, then $h(S_t, x_t) = 0$. That is, if no requests are queued, then no holding cost is incurred. Since the system does not change in the time between decision epochs, the expected discounted reward is

$$r(S_t, x_t) = c(S_t, x_t) - \frac{h(S_t, x_t)}{\alpha + \beta(S_t, x_t)}, \tag{20}$$

where $\alpha > 0$ denotes the continuous time discounting rate. Applying uniformization gives

$$\tilde{r}(S_t, x_t) \equiv r(S_t, x_t) \frac{\alpha + \beta(S_t, x_t)}{\alpha + \nu}. \tag{21}$$

Note that the uniformized rewards agree with the rewards in the CTMDP.

Let $X^\pi(S_t)$ be a policy (i.e., decision function) that prescribes AEO dispatch decisions for each state $S_t \in \mathcal{S}$. That is, $x = X^\pi(S_t)$ is the dispatching decision returned when utilizing policy π . The optimal policy π^* is sought from the class of policies $(X^\pi(S_t))_{\pi \in \Pi}$ to maximize the expected total discounted reward earned by the MEDEVAC system. The objective is expressed as

$$\max_{\pi \in \Pi} \mathbb{E}^\pi \left\{ \sum_{t=1}^{\infty} \gamma^{t-1} \tilde{r}(S_t, X^\pi(S_t)) \right\}, \tag{22}$$

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

$$J(S_t) = \max_{x_t \in \mathcal{X}(S_t)} \left\{ \tilde{r}(S_t, x_t) + \gamma \sum_{S_{t+1} \in \mathcal{S}} \tilde{p}(S_{t+1} | S_t, x_t) J(S_{t+1}) \right\}. \tag{23}$$

The policy iteration algorithm is implemented in MATLAB to solve Eq. (23) exactly. Policy iteration starts with an initial policy and then iteratively performs two steps: *policy evaluation*, which computes the expected total discounted reward of each state given the current policy, and *policy improvement*, which updates the current policy if any improvements are available (Puterman 1994). The policy iteration algorithm terminates after the policy converges.

For comparison purposes, a linear programming (LP) model of the Markov decision problem is also constructed. Constructing an LP model of a Markov decision problem is beneficial because it eases the inclusion of constraints and provides a better mechanism with which to conduct sensitivity analyses. However, Puterman (1994) notes that LP has not proven to be an efficient method for solving large discounted Markov decision problems. Yet, recent advancements in LP algorithms have increased the computational efficiency of

LP approaches (e.g., as indicated by the performance testing of CPLEX and Gurobi in Bixby (2012)) and make LP a more viable solution method for solving MDPs.

5 Testing, analysis, and results

This section presents a representative military medical evacuation (MEDEVAC) planning scenario utilized both to demonstrate the applicability of the Markov decision process (MDP) model and to examine the behavior of the optimal dispatching policy. A series of sensitivity analyses and computational excursions identify the model parameters that significantly impact the optimal dispatching policy. Military medical planners should focus on these parameters when developing MEDEVAC dispatching policies. Moreover, this section compares the computational efficiency of policy iteration via MATLAB versus linear programming via CPLEX 12.6. The paper utilizes a dual Intel Xeon E5-2650v2 workstation having 128 GB of RAM and MATLAB's Parallel Computing Toolbox to conduct the computational experiments and analyses presented herein.

5.1 Representative scenario

As of 2017, the United States (U.S.) continues to conduct military operations in Afghanistan. The launch of U.S. military operations in Afghanistan began with the initiation of Operation Enduring Freedom (OEF) on October 7, 2001 in response to the terrorist attacks on New York's World Trade Center and the Pentagon on September 11, 2001. OEF lasted a little over 13 years and officially ended when U.S. combat operations in Afghanistan were terminated on December 31, 2014. However, as part of Operation Freedom's Sentinel, U.S. military forces still remain in Afghanistan to participate in a coalition mission to train and assist the Afghan military and to conduct counter-terrorism operations against Al Qaeda (Department of Defense 2016). While official U.S. combat operations are currently not being conducted in Afghanistan, military medical planners still prepare and plan for potential combat scenarios in the event that a sudden change requires U.S. combat operations.

The computational examples in Bandara et al. (2012), Keneally et al. (2016), and Rettke et al. (2016) inform the development of the representative scenario examined herein. This paper considers a notional planning scenario in which a coalition of allied countries executes combat operations in response to an increase in insurgency operations by remnants of Al-Qaeda militants in southern Afghanistan. For simplicity, this notional scenario (hereafter referred to as the 2×2 case) assumes a MEDEVAC system with two demand zones (i.e., the zones at which 9-line MEDEVAC requests originate) and two MEDEVAC unit staging areas (i.e., the locations in which the MEDEVAC units are stationed) with one medical treatment facility (MTF) co-located at each staging area. Both MTFs are equally capable of treating any casualty and each MTF has an unlimited capacity to treat incoming casualties (i.e., no queuing at the MTF), so only the proximity of an MTF to a casualty collection point (CCP) is utilized to determine where a MEDEVAC will transport casualties.

The 2×2 case assumes that southern Afghanistan is the area of operations (AO) and is divided into two separate demand zones: Helmand province (Zone 1) and Kandahar province (Zone 2). Two MEDEVAC units are considered with one being staged in Zone 1 (i.e., MEDEVAC 1) and the other being staged in Zone 2 (i.e., MEDEVAC 2). The placement of the staging areas and co-located MTFs represents a general realism based on the historical trends in enemy activity in southern Afghanistan. Helmand and Kandahar are the two provinces that have produced the most war-related fatalities in Afghanistan since the start of OEF with 956

and 558 coalition service members killed in action, respectively (White 2016). While these numbers do not account for every type of casualty (e.g., military wounded in action and civilian casualties), they do provide a representative sample that is utilized as an approximation of the threat level present in each zone. Moreover, these numbers are utilized to determine the proportion of 9-line MEDEVAC requests from each zone; the proportion of requests coming from Zone 1 is $p_{z_1} = 0.6314$ and the proportion of requests coming from Zone 2 is $p_{z_2} = 1 - p_{z_1} = 0.3686$.

Each 9-line MEDEVAC request is independently categorized by its zone z (e.g., Helmand and Kandahar) and precedence level k (e.g., urgent, priority, and routine) combination. Fulton et al. (2010) report that the probability of a casualty event being classified with a precedence level of urgent, priority, or routine is 11, 12, and 77%, respectively based on historical MEDEVAC data from U.S. operations in Iraq. Recall that routine requests are assumed to be serviced by non-MEDEVAC units (i.e., casualty evacuation (CASEVAC)). The 2×2 case assumes that the proportion of requests classified with an urgent precedence level is approximately $p_{k_1} = 0.5$ and the proportion of requests classified with a priority precedence level is $p_{k_2} = 1 - p_{k_1} = 0.5$. The proportion of each request categorization p_{zk} is found by multiplying the zone proportion with the precedence level proportion (e.g., $p_{11} = p_{z_1} p_{k_1}$).

Military medical planners estimate the arrival rate of 9-line MEDEVAC requests by estimating when and where future tactical level engagements will occur, along with the likelihood and severity of corresponding casualty events. The reward obtained for servicing a 9-line MEDEVAC request depends on the location of the request, the servicing MEDEVAC unit, and the closest MTF. The response and service times described in Sect. 3 are generated by leveraging the procedure set forth by Keneally et al. (2016).

The procedure utilized to model future 9-line MEDEVAC requests avoids using current data from southern Afghanistan to maintain operational security. Indeed, actual data for current MEDEVAC unit, casualty event, and MTF locations are restricted. Instead, the spatial distribution of future 9-line MEDEVAC requests are modeled with a Monte Carlo simulation via a Poisson cluster process. Casualty cluster centers are selected by leveraging data from the International Council on Security and Development (ICOS) (2008) pertaining to insurgent attacks in southern Afghanistan resulting in death in 2007. It is assumed that all casualty events generated from the casualty cluster centers result in 9-line MEDEVAC requests. Moreover, the distribution of 9-line MEDEVAC request locations from a given casualty cluster center is generated on a uniform distribution with respect to the distance of the request to the casualty cluster center. Military medical planners must keep in mind that data will certainly change with respect to each conflict. Furthermore, the dispatching policy generated depends on the input data and therefore must relate to the scenario being modeled to obtain meaningful results.

Figure 3 depicts the two zones (i.e., Helmand and Kandahar) in southern Afghanistan utilized to generate the data, as well as the MEDEVAC and MTF locations. Recall that the MEDEVAC and MTF locations are collocated for the 2×2 case. The collocated MEDEVAC and MTF locations in each zone are represented by blue stars. The casualty cluster centers in each zone are represented by red diamonds.

The data generated for the MEDEVAC mission task times that comprise the response time vary with each mission and, therefore, are represented as random variables. The response time variables representing mission preparation time, travel time to CCP, service time at CCP, travel time to MTF, and service time at the MTF are defined in Sect. 3 and described in detail in the following four paragraphs.

The mission preparation time is exponentially distributed with a mean of 10 min. The 2008 MEDEVAC after action report (AAR) estimates mission prep time to be 20 min (Bastian

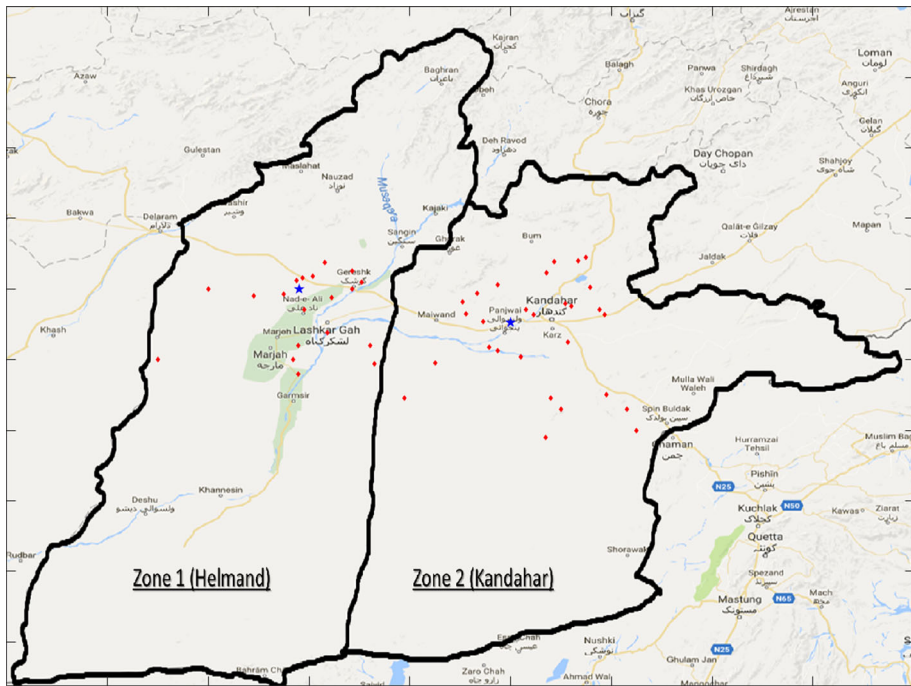


Fig. 3 MEDEVAC and MTF locations with casualty cluster centers

2010). This AAR, along with personal experiences, influences Bastian (2010) to model mission preparation time with a mean of 20 min and standard deviation of 5 min. However, a more recent interview with a MEDEVAC pilot in O’Shea (2011) reports that with proper pre-planning procedures the mission preparation time is often less than 10 min.

The armed escort delay is exponentially distributed with a mean of 10 min. Garrett (2013) reports that there is a 31% chance that a MEDEVAC mission requires an armed escort. Moreover, among the missions requiring an armed escort, approximately 4% are delayed due to issues caused primarily by the escort aircraft. These percentages are included in the computation of the expected response times and the corresponding expected rewards. The delay induced by armed escorts is an important feature of the MEDEVAC problem. This paper applies the same armed escort assumptions found in Keneally et al. (2016), to which we refer a more interested reader for a more in depth description on how armed escorts impact this MDP model.

The flight speed, which accounts for the travel time to the CCP and the travel time to the MTF, is uniformly distributed between 120 and 193 knots with a mean of 156.5 knots. This flight speed is based on currently fielded MEDEVAC helicopters (i.e., HH-60Ms) and on subject matter expertise (Bastian 2010).

The service time at the CCP and the service time at the MTF are exponentially distributed with a mean of 10 and 5 min, respectively. These times are determined by leveraging the data provided by in-theater MEDEVAC pilots and other subject matter experts described in Bastian (2010) and Keneally et al. (2016).

The just-described response time random variables, casualty cluster centers, and MEDEVAC staging areas are utilized in a Monte Carlo simulation to obtain a synthetic, but realistic,

Table 1 Expected response times (min)

Zone, z	MEDEVAC 1	MEDEVAC 2
1 (Helmand)	34.25	48.18
2 (Kandahar)	52.98	36.89

Table 2 Expected service times (min)

Zone, z	MEDEVAC 1	MEDEVAC 2
1 (Helmand)	34.25	67.28
2 (Kandahar)	72.13	36.89

Table 3 Immediate expected rewards

Zone, z	Precedence, k	MEDEVAC, m	
		1	2
1 (Helmand)	1 (Urgent)	5.65	4.48
	2 (Priority)	0.87	0.82
2 (Kandahar)	1 (Urgent)	4.14	5.41
	2 (Priority)	0.80	0.86

spatial distribution of future 9-line MEDEVAC requests and response time data. The means of the response times are computed and presented in Table 1.

After the expected response times are computed, the expected service times can be computed by simply adding the appropriate expected response time to the MEDEVAC unit’s travel time back to its staging area. This travel time is defined in Sect. 3 and is based on the flight speed of the MEDEVAC helicopter. The distribution for the flight speed for the travel time to the staging area is the same as the flight speed distributions for the travel times to the CCP and MTF. The expected service times for the 2×2 case are provided in Table 2.

Recall from Sect. 4 that the MEDEVAC system employs an inter-zone policy regarding airspace access. This means that any MEDEVAC unit can service any 9-line MEDEVAC request, regardless of the zone from which the request originated. For example, the MEDEVAC unit staged in Helmand for the 2×2 case can service requests from both Helmand and Kandahar.

The paper applies a survivability function that is monotonically decreasing in response time to compute the reward obtained from servicing a 9-line MEDEVAC request. The immediate expected reward for servicing a 9-line MEDEVAC request is determined by the precedence level and the response time of the request as indicated in Eq. (17). For the 2×2 case, the immediate expected reward function utilizes $\delta = 10$, which rewards the servicing of *urgent* (i.e., $k = 1$) 9-line MEDEVAC requests much more than *priority* (i.e., $k = 2$) 9-line MEDEVAC requests. Table 3 summarizes the computed immediate expected rewards, ψ_{mzk} .

The continuous expected holding cost is computed based on the number of urgent and priority 9-line MEDEVAC requests that are in the queue between decision epochs. The 2×2 case utilizes $\xi = 0.20$, which scales the holding cost rate for a precedence- k request to be 20% of the average immediate expected reward over all possible MEDEVAC-zone combinations.

The 2×2 case assumes a high operational tempo (i.e., combined frequency and intensity of conflict) with a baseline request arrival rate of $\lambda = \frac{1}{60}$, representing an average 9-line MEDEVAC request rate of one request per 60 min. The military intelligence community, operational planners, and medical planners should work together to determine a reasonable

Table 4 2×2 case parameters

Parameter	Description	Setting
λ	9-line MEDEVAC request arrival rate	$\frac{1}{60}$
$ \mathcal{M} $	Total # of MEDEVAC units	2
$ \mathcal{Z} $	Total # of zones	2
$ \mathcal{K} $	Total # of precedence levels	2
q^{max}	Max (z, k) queue length	5
γ	Uniformized discount factor	0.99
δ	Weight for urgent request	10
ξ	Scale for holding cost rate	0.2
p_{z_1}	Zone 1 proportion of requests	0.6314
p_{z_2}	Zone 2 proportion of requests	0.3686
p_{k_1}	Urgent proportion of requests	0.5
p_{k_2}	Priority proportion of requests	0.5

estimate of the request arrival rate prior to a planned combat operation based on the equipment, size, and disposition of friendly and adversary forces.

5.2 Representative scenario results

A list of parameters associated with the 2×2 case are displayed in Table 4. Utilizing the parameter settings in Table 4 and the expected response times, expected service times, and immediate expected rewards computed in the previous section, the optimal policy for the 2×2 case is determined via policy iteration. Applying Eq. (5) indicates that the size of the state space for the 2×2 case is 58,320. This result shows that even for this relatively simple scenario, the size of the state space is quite large.

For comparison purposes, three myopic dispatching policies are considered. The three myopic policies are all based on a classic inter-zone myopic policy. Recall that an inter-zone myopic policy sends the closest available MEDEVAC unit to service an incoming 9-line MEDEVAC request, regardless of the request’s zone or precedence level. All three myopic policies adopt this strategy when at least one MEDEVAC unit is available. The differences between the three myopic policies are found when both MEDEVAC units are busy. The first myopic policy (i.e., Myopic 1) will queue 9-line MEDEVAC requests if there are no available MEDEVAC units to service the request, regardless of the request’s zone or precedence level. The second myopic policy (i.e., Myopic 2) will queue only urgent 9-line MEDEVAC requests if there are no available MEDEVAC units to service the request, regardless of the urgent request’s zone. The third myopic policy (i.e., Myopic 3) will not queue any 9-line MEDEVAC requests. If there are queued requests, the Myopic 1 and Myopic 2 dispatching policies service requests with a prioritized first-come-first-serve basis. The optimal policy’s dispatching decisions, queue lengths, and MEDEVAC utilization rates are compared against the three myopic policies to obtain a better understanding of where similarities and differences exist. Moreover, the optimality gap for each myopic policy is computed to demonstrate whether a myopic policy is appropriate for the given 2×2 case.

The dispatching decisions for the optimal policy and three myopic policies are compared over three separate scenarios. Each scenario (i.e., Scenarios 1–3) considers a set of MEDEVAC system states with empty zone-precedence queues. The first scenario (i.e., Scenario 1) considers a system state wherein both MEDEVAC units are idle, which can be represented as $S_t \in ((0, 0), (0, 0, 0, 0), \hat{R}_t)$. Regardless of the zone or precedence level of the incoming

Table 5 Dispatching policies for Scenario 2

Policy	\hat{R}_t	MEDEVAC 2 servicing zone 1 Queue\dispatch\reject	MEDEVAC 2 servicing zone 2 Queue\dispatch\reject
Optimal	(1, 1)	Dispatch MEDEVAC 1	Dispatch MEDEVAC 1
	(1, 2)	Dispatch MEDEVAC 1	Dispatch MEDEVAC 1
	(2, 1)	Dispatch MEDEVAC 1	Dispatch MEDEVAC 1
	(2, 2)*	Reject	Queue
Myopic 1	(1, 1)	Dispatch MEDEVAC 1	Dispatch MEDEVAC 1
Myopic 2	(1, 2)	Dispatch MEDEVAC 1	Dispatch MEDEVAC 1
Myopic 3	(2, 1)	Dispatch MEDEVAC 1	Dispatch MEDEVAC 1
	(2, 2)	Dispatch MEDEVAC 1	Dispatch MEDEVAC 1

* Incoming requests for which the optimal policy does not correspond to a myopic policy

Table 6 Dispatching policies for Scenario 3

Policy	\hat{R}_t	MEDEVAC 1 servicing zone 1 Queue\dispatch\reject	MEDEVAC 1 servicing zone 2 Queue\dispatch\reject
Optimal	(1, 1)	Dispatch MEDEVAC 2	Dispatch MEDEVAC 2
	(1, 2)*	Queue	Reject
	(2, 1)	Dispatch MEDEVAC 2	Dispatch MEDEVAC 2
	(2, 2)	Dispatch MEDEVAC 2	Dispatch MEDEVAC 2
Myopic 1	(1, 1)	Dispatch MEDEVAC 2	Dispatch MEDEVAC 2
Myopic 2	(1, 2)	Dispatch MEDEVAC 2	Dispatch MEDEVAC 2
Myopic 3	(2, 1)	Dispatch MEDEVAC 2	Dispatch MEDEVAC 2
	(2, 2)	Dispatch MEDEVAC 2	Dispatch MEDEVAC 2

* Incoming requests for which the optimal policy does not correspond to a myopic policy

9-line MEDEVAC request, \hat{R}_t , all four policies react in a myopic fashion when the system is in state $S_t \in ((0, 0), (0, 0, 0, 0), \hat{R}_t)$, sending the closest MEDEVAC unit to service the request.

The second scenario (i.e., Scenario 2) considers a set of MEDEVAC system states wherein MEDEVAC 1 is idle and MEDEVAC 2 is busy, which can be represented as $S_t \in ((0, z), (0, 0, 0, 0), \hat{R}_t)$ where $z \in \{1, 2\}$. The third scenario (i.e., Scenario 3) considers a set of MEDEVAC system states wherein MEDEVAC 1 is busy and MEDEVAC 2 is idle, which can be represented as $S_t \in ((z, 0), (0, 0, 0, 0), \hat{R}_t)$ where $z \in \{1, 2\}$. The dispatching policies for Scenarios 2 and 3 are displayed in Tables 5 and 6, respectively. Contrary to the findings of Keneally et al. (2016) in their computational example, the best MEDEVAC unit to dispatch to service a 9-line MEDEVAC request does depend on the zone of the request that the busy MEDEVAC is currently servicing. Note that this is an observed result based on the parameter settings for the 2×2 case and that location-independent policies are a possibility, as seen in Keneally et al. (2016). In Tables 5 and 6 an asterisk (*) is placed next to the incoming requests, \hat{R}_t , for which the optimal policy does not correspond to a myopic policy. It is expected that a myopic policy will apply to all urgent 9-line MEDEVAC requests due to the life threatening nature of these requests and the accompanying high rewards for servicing them.

Consider the Scenario 2 results displayed in Table 5. The MEDEVAC system is in a state $S_t \in ((0, z), (0, 0, 0, 0), \hat{R}_t)$ where $z \in \{1, 2\}$. The optimal dispatching policy for $\hat{R}_t = (2, 2)$ (i.e., a Zone 2, priority request) depends on z , the zone where MEDEVAC 2 is currently servicing a request. If MEDEVAC 2 is servicing a Zone 1 (i.e., $z = 1$) request and $\hat{R}_t = (2, 2)$, the optimal decision is to reject the request from entering the system and send the request to be serviced by CASEVAC. The MEDEVAC system is not well positioned to service the (2, 2) request in a timely manner since MEDEVAC 2 is currently servicing a request outside its own zone. If MEDEVAC 2 is servicing Zone 2 (i.e., $z = 2$) and $\hat{R}_t = (2, 2)$, the optimal decision is to accept and queue the request. The MEDEVAC system is in a better position to service the (2, 2) request here because MEDEVAC 2 is currently servicing a request in its own zone. Both of these decisions differ from the myopic decision (i.e., dispatch MEDEVAC 1 to service the request). This result illustrates that, if the system is in a Scenario 2 state and $\hat{R}_t = (2, 2)$, then the optimal policy will reserve MEDEVAC 1 for either an urgent 9-line MEDEVAC request or a Zone 1 request. The difference between rejecting or queueing the request is driven by the holding costs as impacted by the difference in expected service times. Recall that there is large difference in expected service times for MEDEVAC 2 to Zone 1 and MEDEVAC 2 to Zone 2: 67.28 and 36.28 min, respectively.

Consider the Scenario 3 results displayed in Table 6; they mirror those observed for Scenario 2. The MEDEVAC system is in a state $S_t \in ((z, 0), (0, 0, 0, 0), \hat{R}_t)$ where $z \in \{1, 2\}$. The optimal dispatching policy for $\hat{R}_t = (1, 2)$ depends on z , the zone where MEDEVAC 1 is currently servicing a request. If MEDEVAC 1 is servicing Zone 2 (i.e., $z = 2$) and $\hat{R}_t = (1, 2)$, the optimal decision is to reject the request from entering the system and send the request to be serviced by CASEVAC. The MEDEVAC system is not well positioned to service the (1, 2) request in a timely manner since MEDEVAC 1 is currently servicing a request outside its own zone. If MEDEVAC 1 is servicing Zone 1 (i.e., $z = 1$) and $\hat{R}_t = (1, 2)$, the optimal decision is to accept and queue the request. The MEDEVAC system is in a better position to service the (1, 2) request here because MEDEVAC 1 is currently servicing a request in its own zone. Both of these decisions differ from the myopic decision (i.e., dispatch MEDEVAC 2 to service the request). This result illustrates that, if the system is in a Scenario 3 state and $\hat{R}_t = (1, 2)$, then the optimal policy will reserve MEDEVAC 2 for either an urgent 9-line MEDEVAC request or a Zone 2 request. The difference between rejecting or queueing the request is driven by the holding costs as impacted by the difference in expected service times. Recall that there is large difference in expected service times for MEDEVAC 1 to Zone 1 and MEDEVAC 1 to Zone 2: 34.25 and 72.13 min, respectively.

The optimality gaps between the myopic policies and the optimal policy are examined. The expected total discounted reward for the optimal policy and myopic policies when the system is in an *empty and idle* state $S^0 = ((0, 0), (0, 0, 0, 0), (0, 0))$ (i.e., both MEDEVAC units are idle, every zone-precedence queue is empty, and there are no 9-line MEDEVAC requests in the system) are displayed in Table 7, along with the optimality gaps associated with each myopic policy. The results indicate that the superlative myopic policy is Myopic 2, which has the smallest optimality gap of 0.74%. Without having the ability to queue any requests, the Myopic 3 policy performs worse than every other policy and has the largest optimality gap of 5.73%. While these optimality gaps may not seem large, over a long time period the optimal policy will save more lives.

5.3 Computational experiments

Since there are many parameters associated with the MEDEVAC system, a screening experiment is developed to reveal the parameters that significantly impact the value of the optimal

Table 7 Expected total discounted rewards and optimality gaps

Policy, π	$J^\pi(S^0)$	Optimality gap
Optimal	63.50	N/A
Myopic 1	62.09	2.21%
Myopic 2	63.02	0.74%
Myopic 3	59.86	5.73%

Table 8 2^5 Full factorial screening experimental factor levels

Factor	Low level	High level
λ	$\frac{1}{75}$	$\frac{1}{45}$
δ	5	15
ξ	0.1	0.3
p_{z_1}	0.25	0.75
p_{k_1}	0.25	0.75

Table 9 Multiple linear regression analysis

Factors	Coef	SE	T	P
Intercept	- 76.00	17.18	- 4.42	<0.00
λ	2201.80	673.03	3.27	<0.00
δ	5.62	0.60	9.39	<0.00
ξ	- 18.31	29.91	- 0.61	0.55
p_{z_1}	4.57	11.97	0.38	0.71
p_{k_1}	92.51	11.97	7.73	<0.00

dispatching policy. Leveraging the results found from the 2×2 case, a 2^5 full factorial screening experiment is designed to determine the relative significance of factors λ , δ , ξ , p_{z_1} , and p_{k_1} . All five of these factors represent important MEDEVAC problem features of interest. The initial screening design includes all five factors, each specified at two discrete parameter-levels (i.e., low and high). For example, the rate at which 9-line MEDEVAC requests arrive to the system, λ , is designed with low and high factor levels of $\frac{1}{75}$ and $\frac{1}{45}$, respectively, to determine whether λ has a significant impact on the value of the optimal dispatching policy.

The 2^5 full factorial screening experimental factors and the levels associated with each factor are reported in Table 8. Once the results from the 2^5 full factorial screening experiment are examined, the factors that have a statistically significant impact on the value of the optimal dispatching policy are analyzed via a three-level experiment with low, intermediate, and high factor levels.

Multiple linear regression analysis is conducted to examine the relationship between the independent factors λ , δ , ξ , p_{z_1} , and p_{k_1} and the dependent variable $J^{\pi^*}(S^0)$. The results from the multiple linear regression analysis are displayed in Table 9. Starting from the left, the first column lists the experimental factors. The second through fifth columns list the estimated coefficients (Coef), standard errors (SE), test statistics (T), and p -values (P) associated with the experimental factors, respectively.

The results from the multiple linear regression analysis in Table 9 report that the p -values associated with factors λ , δ , and p_{k_1} are all less than 0.01, which indicates that these factors are statistically significant in predicting $J^{\pi^*}(S^0)$. Intuitively, these results make sense. The rate

Table 10 3^3 Full factorial experimental factor levels

Factor	Low level	Intermediate level	High level
λ	$\frac{1}{75}$	$\frac{1}{60}$	$\frac{1}{45}$
δ	5	10	15
p_{k_1}	0.25	0.50	0.75

at which 9-line MEDEVAC requests arrive directly impacts the number of requests that can be serviced, resulting in more or less opportunities to earn rewards. Increasing or decreasing the weight and proportion of urgent requests also directly impacts the amount of reward earned by the system. Moreover, Table 9 reports the p -values associated with ξ and p_{z_1} are both greater than 0.05 and, therefore, do not provide enough evidence to claim that the factors ξ and p_{z_1} are statistically significant in predicting $J^{\pi^*}(S^0)$. The reason that these factors are not significant could be due to the selected experimental design factor levels. Selecting a wider range in factor levels for ξ and p_{z_1} could result in them becoming significant.

Utilizing the results from Table 9, a 3^3 full factorial experiment is generated to examine the differences between the optimal and myopic dispatching policies at different levels for factors λ , δ , and p_{k_1} . The goal of the 3^3 full factorial experiment is to gain insight regarding when medical planners should avoid implementing myopic dispatching policies (e.g., Myopic 1, Myopic 2, and Myopic 3) and to understand how the changes in the factor levels for λ , δ , and p_{k_1} impact the optimal dispatching policy. The 3^3 full factorial experimental factors and the levels associated with each factor are displayed in Table 10.

Table 11 reports the results from the 3^3 full factorial experiment. Starting from the left, the first column indicates the run number. The next three columns indicate the factor levels. The fifth column indicates the dependent variable $J^{\pi^*}(S^0)$, where $J^{\pi^*}(S^0)$ is the value of the optimal policy π^* when the system is empty and idle. The next three columns indicate the optimality gaps for the Myopic 1, Myopic 2, and Myopic 3 policies, respectively. The following four columns indicate the MEDEVAC busy probabilities when the system is operating under the optimal dispatching policy. The four rightmost columns indicate the average zone-precedence queue lengths when the system is operating under the optimal dispatching policy.

The results from Table 11 indicate that the Myopic 2 policy (i.e., only queue urgent) strictly outperforms the Myopic 3 policy (i.e., do not queue). Moreover, the Myopic 2 policy strictly outperforms the Myopic 1 policy (i.e., queue both urgent and priority) when $\frac{1}{\lambda} \in \{45, 60\}$, but not when $\frac{1}{\lambda} = 75$. These results suggest that medical planners should not employ the Myopic 3 policy. More generally, these results suggest that queuing requests is advisable. Additionally, the Myopic 1 policy outperforms the Myopic 2 policy in several instances when $\frac{1}{\lambda} = 75$, because as the inter-arrival time of 9-line MEDEVAC requests increases, it becomes more beneficial to queue all requests versus just queuing urgent requests. Together, these results indicate that, when the arrival rate of requests is relatively low, it is advisable for the MEDEVAC system to queue all requests. However, as the arrival rate of requests increases, there is a point at which it is no longer advisable to queue all requests and it is more beneficial to queue only urgent requests.

The MEDEVAC unit busy probabilities associated with each run in Table 11 also provide interesting results. MEDEVAC 1 is busy servicing Zone 1 requests substantially more often than servicing Zone 2 requests for all 27 runs. MEDEVAC 2 is busy servicing each zone with approximately the same proportion. This result aligns with intuition because the proportion of requests arriving from Zone 1 ($p_{z_1} = 0.6314$) is greater than the proportion of requests arriving from Zone 2 ($p_{z_2} = 0.3686$).

Table 11 3³ Full factorial experimental results

Run #	$\frac{1}{\lambda}$	δ	ρ_{k_1}	$J^{*+} (S^1)$	Optimality gaps		Myopic 3 (%)		MEDEVAC 1 busy		MEDEVAC 2 busy		Average queue lengths			
					Myopic 1 (%)	Myopic 2 (%)	Myopic 3 (%)	Zone 1	Zone 2	Zone 1	Zone 2	Q_{r11}	Q_{r12}	Q_{r21}	Q_{r22}	
1	45	5	0.25	30.63	6.59	1.05	5.27	0.21	0.04	0.17	0.14	0.02	0.03	0.01	0.04	
2	45	5	0.5	41.38	7.11	1.28	5.77	0.20	0.07	0.17	0.14	0.04	0.01	0.02	0.02	
3	45	5	0.75	51.74	7.01	2.93	5.34	0.21	0.04	0.14	0.15	0.05	0.00	0.06	0.00	
4	45	10	0.25	47.25	6.55	1.08	7.39	0.20	0.04	0.15	0.15	0.02	0.01	0.01	0.04	
5	45	10	0.5	74.69	7.87	2.56	8.32	0.20	0.06	0.11	0.14	0.03	0.00	0.02	0.00	
6	45	10	0.75	99.92	7.40	3.74	6.73	0.20	0.09	0.15	0.13	0.05	0.00	0.03	0.00	
7	45	15	0.25	64.20	7.01	1.61	8.89	0.20	0.03	0.05	0.16	0.01	0.00	0.01	0.03	
8	45	15	0.5	108.30	8.41	3.33	9.55	0.19	0.06	0.11	0.14	0.03	0.00	0.02	0.00	
9	45	15	0.75	149.11	8.17	4.68	7.84	0.17	0.08	0.13	0.11	0.04	0.00	0.02	0.00	
10	60	5	0.25	26.49	1.41	1.27	4.47	0.16	0.07	0.13	0.10	0.01	0.03	0.01	0.01	
11	60	5	0.5	35.73	1.64	0.58	4.66	0.16	0.04	0.12	0.11	0.02	0.01	0.01	0.02	
12	60	5	0.75	44.85	1.64	0.68	4.50	0.16	0.06	0.12	0.11	0.03	0.01	0.02	0.01	
13	60	10	0.25	40.36	1.63	0.71	5.34	0.16	0.02	0.11	0.11	0.01	0.02	0.01	0.03	
14	60	10	0.5	63.50	2.21	0.74	5.73	0.17	0.04	0.07	0.12	0.01	0.04	0.01	0.02	
15	60	10	0.75	86.10	2.21	1.11	5.33	0.16	0.06	0.10	0.11	0.03	0.00	0.02	0.00	
16	60	15	0.25	54.38	2.02	0.71	6.03	0.16	0.02	0.11	0.12	0.01	0.02	0.00	0.03	
17	60	15	0.5	91.63	2.83	1.21	6.53	0.16	0.04	0.07	0.11	0.01	0.00	0.01	0.00	
18	60	15	0.75	127.54	2.54	1.40	5.75	0.16	0.06	0.10	0.11	0.02	0.00	0.02	0.00	
19	75	5	0.25	23.05	0.41	1.40	3.83	0.13	0.05	0.09	0.09	0.01	0.02	0.00	0.01	
20	75	5	0.5	31.04	0.41	0.49	3.79	0.13	0.05	0.09	0.09	0.01	0.01	0.01	0.00	
21	75	5	0.75	39.03	0.47	0.28	3.74	0.13	0.04	0.09	0.09	0.02	0.00	0.01	0.01	
22	75	10	0.25	34.90	0.38	0.60	4.05	0.13	0.03	0.09	0.10	0.01	0.01	0.00	0.02	
23	75	10	0.5	54.87	0.74	0.35	4.30	0.13	0.03	0.08	0.09	0.01	0.01	0.01	0.01	
24	75	10	0.75	74.62	0.77	0.39	4.14	0.14	0.04	0.07	0.09	0.02	0.01	0.01	0.01	
25	75	15	0.25	46.88	0.64	0.48	4.43	0.13	0.02	0.08	0.09	0.00	0.01	0.00	0.02	
26	75	15	0.5	78.87	1.09	0.52	4.72	0.14	0.02	0.05	0.10	0.01	0.03	0.00	0.01	
27	75	15	0.75	110.37	1.02	0.58	4.42	0.13	0.04	0.07	0.09	0.01	0.00	0.01	0.00	

Table 12 Dispatching policies for Scenario 1

Policy	\hat{R}_t	Queue\dispatch\reject
Optimal	(1, 1)	Dispatch MEDEVAC 1
	(1, 2)*	Reject
	(2, 1)	Dispatch MEDEVAC 2
	(2, 2)*	Reject
Myopic 1	(1, 1)	Dispatch MEDEVAC 1
Myopic 2	(1, 2)	Dispatch MEDEVAC 1
Myopic 3	(2, 1)	Dispatch MEDEVAC 2
	(2, 2)	Dispatch MEDEVAC 2
Settings: $\lambda = \frac{1}{45}$, $\delta = 15$, and $p_{k_1} = 0.75$		

* Incoming requests for which the optimal policy does not correspond to a myopic policy

An interesting observation from the 3³ full factorial experiment is that the optimal dispatching policy aligns with the myopic policy when the MEDEVAC system is in a Scenario 1 state for 26 out of the 27 runs. Table 12 reports the optimal and myopic dispatching policies for the single run (i.e., Run 9) for which the optimal dispatching policy does not align with myopic policies. The optimal dispatching policy will reject precedence level two requests (i.e., priority requests) when the system is in a Scenario 1 state and $\lambda = \frac{1}{45}$, $\delta = 15$, and $p_{k_1} = 0.75$. This result is intuitive because the inter-arrival times of the requests have increased from one every 60 min to one every 45 min, the immediate expected reward for servicing urgent requests is substantially higher than servicing priority requests, and there is a much higher rate of urgent requests arriving to the system compared to priority requests.

5.4 Excursion 1—request arrival rate

The section considers the impact of the arrival rate λ on the optimal policy when the MEDEVAC system is in a Scenario 1 state $S_t \in ((0, 0), (0, 0, 0, 0), \hat{R}_t)$. The same parameter settings from the 2 × 2 case are utilized for the request arrival rate excursion except for λ ; see Table 4 for a descriptive list of the parameters and their attendant values. The computational results indicate that the optimal policy dispatches the closest MEDEVAC unit when the system is in a Scenario 1 state with an urgent 9-line MEDEVAC request arrival (i.e., $S_t \in ((0, 0), (0, 0, 0, 0), (z, 1))$ where $z \in \{1, 2\}$), regardless of the request arrival rate λ . However, this same result does not hold when the system is in a Scenario 1 state with a priority 9-line MEDEVAC request arrival (i.e., $S_t \in ((0, 0), (0, 0, 0, 0), (z, 2))$ where $z \in \{1, 2\}$). The dispatching policies for when the system is in a Scenario 1 state with a priority 9-line request are displayed in Table 13.

The results from Table 13 indicate that, when $\frac{1}{\lambda} \leq 25$, the optimal policy is to reject priority 9-line MEDEVAC requests regardless of the zone from which the request originated. For $\frac{1}{\lambda} \in \{26, 27, 28\}$ the optimal policy is to reject Zone 1 priority 9-line MEDEVAC requests and to dispatch MEDEVAC 2 to Zone 2 priority requests. Lastly, when $\frac{1}{\lambda} \geq 29$ the optimal policy dispatches MEDEVAC units in a myopic manner. These results indicate that the optimal policy reserves MEDEVAC units for urgent requests as the inter-arrival time of 9-line MEDEVAC requests decreases (i.e., more frequent arrivals).

5.5 Excursion 2—MEDEVAC helicopter flight speed

This section considers the impact of replacing the currently fielded HH-60M MEDEVAC helicopter with a more effective (i.e., faster flight speed) aeromedical aircraft. The

Table 13 Comparison of MEDEVAC dispatching policies for priority requests

$\frac{1}{\lambda}$	Optimal policy		Myopic policy	
	$\hat{R}_t = (1, 2)$	$\hat{R}_t = (2, 2)$	$\hat{R}_t = (1, 2)$	$\hat{R}_t = (2, 2)$
21	Reject	Reject	MEDEVAC 1	MEDEVAC 2
22	Reject	Reject	MEDEVAC 1	MEDEVAC 2
23	Reject	Reject	MEDEVAC 1	MEDEVAC 2
24	Reject	Reject	MEDEVAC 1	MEDEVAC 2
25	Reject	Reject	MEDEVAC 1	MEDEVAC 2
26	Reject	MEDEVAC 2	MEDEVAC 1	MEDEVAC 2
27	Reject	MEDEVAC 2	MEDEVAC 1	MEDEVAC 2
28	Reject	MEDEVAC 2	MEDEVAC 1	MEDEVAC 2
29	MEDEVAC 1	MEDEVAC 2	MEDEVAC 1	MEDEVAC 2
30	MEDEVAC 1	MEDEVAC 2	MEDEVAC 1	MEDEVAC 2

Table 14 MEDEVAC helicopter flight speed results

Flight speed (%)	Optimality gaps		
	Myopic 1 (%)	Myopic 2 (%)	Myopic 3 (%)
- 50	17.07	7.95	10.93
- 25	5.02	1.80	6.56
0 (i.e., current)	2.21	0.74	5.73
25	1.12	0.37	5.30
50	0.63	0.24	5.06

same parameter settings from the 2×2 case are utilized for the MEDEVAC flight speed excursion; see Table 4 for a descriptive list of the parameters and their attendant values. The HH-60M MEDEVAC helicopter still utilizes a power plant that was designed prior to 1989 (Leoni 2007). Significantly faster experimental tiltrotor aircraft could potentially be put into service to replace the HH-60M (Cox 2016). Moreover, programs exist to improve the turbine engines of the current platform (Hoffman 2015). It is reasonable to assume that new aircraft designs or improved helicopter engines will result in a 25–50% increase in average flight speed when compared to the currently fielded HH-60M MEDEVAC helicopter.

To examine the impact of employing faster aircraft, the mean of the flight speed random variable is adjusted while all other random variables modeling the MEDEVAC process remain the same. Incorporating this adjustment leads to immediate changes to response and service times, along with the immediate expected reward. It is expected that, as the mean flight speed increases, the optimal dispatching policy will deploy MEDEVAC units in a more consistently myopic fashion. Moreover, another interesting scenario examined is when the mean flight speed decreases, which can occur due to potential maintenance issues or environmental issues within the area of operations. With limited resources, it is reasonable to assume that slower HH-60M MEDEVAC helicopters would still be utilized in a high intensity conflict.

Table 14 reports the results obtained by increasing and decreasing the mean flight speed, where flight speed is indicated as a percentage increase over the flight speed of the currently employed HH-60M MEDEVAC helicopter.

As expected, the results from Table 14 indicate that as the mean flight speed of the MEDEVAC helicopter increases, the optimality gaps for the Myopic 1, Myopic 2, and Myopic 3 policies all decrease. This shows that, if a new rotary wing aircraft is fielded for MEDEVAC purposes, the optimal dispatching policy will deploy MEDEVAC units in a more myopic fashion. Moreover, the results indicate that as the mean flight speed of the MEDEVAC helicopter decreases, the optimality gaps for the Myopic 1, Myopic 2, and Myopic 3 policies all increase. This is an important observation. Military medical planners must take flight speed issues into consideration when developing dispatching policies. These results should also persuade military medical planners to consider changing dispatching policies during steady state combat operations if the mean flight speed of the MEDEVAC helicopters being utilized decreases due to atmospheric, environmental, or mechanical issues.

5.6 Excursion 3—*intra-zone* policies

This section considers the impact of replacing the MEDEVAC system's *inter-zone* policy with an *intra-zone* policy with regard to airspace access. The same parameter settings from the 2×2 case and the MEDEVAC flight speed excursion are utilized for the *intra-zone* policies excursion; see Table 4 for a descriptive list of the parameters and their attendant values. An *intra-zone* policy prevents MEDEVAC units from operating outside of the zone in which they are staged. Military situations may arise that force strict adherence to an *intra-zone* policy. For example, an execution of a specific, short-duration combat operation may enforce an *intra-zone* policy to reduce the risk of collisions and fratricide (Keneally et al. 2016). Moreover, when separate branches of the U.S. military (e.g., Army and Air Force) and/or allied countries are working together in a combat environment, an *intra-zone* policy restricting MEDEVAC units to serve their own zone may be enforced due to chain of command restrictions, communication limitations, and/or political realities (Keneally et al. 2016).

To examine the impact of enforcing an *intra-zone* policy, each MEDEVAC unit is restricted to operate in their own zone while all other random variables modeling the MEDEVAC process remain the same. The queueing strategies associated with each myopic policy remain the same. Recall that, when both MEDEVAC units are busy, the Myopic 1 policy queues all incoming requests, the Myopic 2 policy only queues incoming urgent requests, and the Myopic 3 policy does not queue any incoming requests. Regardless of the zone or precedence level of the incoming 9-line MEDEVAC request, \hat{R}_t , all four policies react in a myopic fashion when the system is in a Scenario 1 state, sending the closest idle MEDEVAC unit to service the request. Tables 15 and 16 report the dispatching policies associated with the system being in a Scenario 2 and Scenario 3 state, respectively.

These results indicate that the *intra-zone* optimal dispatching policy and the *intra-zone* Myopic 2 dispatching policy (i.e., queue urgent) dispatch MEDEVAC units in the same manner for Scenarios 1-3. Moreover, it is observed that, when a MEDEVAC unit is busy and a request from the MEDEVAC unit's zone arrives to the system, the *intra-zone* optimal dispatching policy always rejects priority requests from entering the system. The difference between the *intra-zone* optimal dispatching policy and the *intra-zone* Myopic 2 dispatching policy is observed when there is at least one urgent 9-line MEDEVAC request in the queue, the MEDEVAC unit able to service the urgent queued request is busy, and there is an incoming request associated with that zone. Many states satisfy this description. Denote such states as *Scenario 4 states*. Table 17 reports the dispatching policies associated with being in a Scenario 4 state when either: MEDEVAC 1 is busy, there is an urgent Zone 1 MEDEVAC request in the queue (i.e., $Q_{111} = 1$), and a Zone 1 MEDEVAC request is submitted; or MEDEVAC 2

Table 15 Intra-zone dispatching policies for Scenario 2

Policy	\hat{R}_t	Queue\dispatch\reject
Optimal	(1, 1)	Dispatch MEDEVAC 1
Myopic 2	(1, 2)	Dispatch MEDEVAC 1
	(2, 1)	Queue
	(2, 2)	Reject
	(1, 1)	Dispatch MEDEVAC 1
Myopic 1	(1, 2)	Dispatch MEDEVAC 1
	(2, 1)	Queue
	(2, 2)	Queue
	(1, 1)	Dispatch MEDEVAC 1
Myopic 3	(1, 2)	Dispatch MEDEVAC 1
	(2, 1)	Reject
	(2, 2)	Reject
	(1, 1)	Reject

Table 16 Intra-zone dispatching policies for Scenario 3

Policy	\hat{R}_t	Queue\dispatch\reject
Optimal	(1, 1)	Queue
Myopic 2	(1, 2)	Reject
	(2, 1)	Dispatch MEDEVAC 2
	(2, 2)	Dispatch MEDEVAC 2
	(1, 1)	Queue
Myopic 1	(1, 2)	Queue
	(2, 1)	Dispatch MEDEVAC 2
	(2, 2)	Dispatch MEDEVAC 2
	(1, 1)	Reject
Myopic 3	(1, 2)	Reject
	(2, 1)	Dispatch MEDEVAC 2
	(2, 2)	Dispatch MEDEVAC 2
	(1, 1)	Dispatch MEDEVAC 2

is busy, there is an urgent Zone 2 MEDEVAC request in the queue (i.e., $Q_{t21} = 1$), and a Zone 2 MEDEVAC request is submitted.

Table 17 indicates that if the MEDEVAC system is in a Scenario 4 state, the optimal policy will reject all incoming requests from the zone with the busy MEDEVAC and the queued urgent request. Conversely, the intra-zone Myopic 2 policy will queue all incoming urgent requests. While rejecting an urgent request may not align with expectations, holding more than one request in the queue is detrimental due to the MEDEVAC units being restricted to service only their own zones. If such a decision is not desired by command authorities, the holding cost rate for urgent requests should be updated to be less detrimental to system performance or the value of servicing urgent requests should be increased to discourage rejecting urgent requests from entering the system. Otherwise, this result reflects the intuition that, if the MEDEVAC system is being overwhelmed with requests, it will divert requests to other command authorities (i.e., CASEVAC).

The optimality gap between the intra-zone optimal policy and the intra-zone myopic policies is examined. The expected total discounted reward for the intra-zone optimal policy

Table 17 Intra-zone dispatching policies for Scenario 4

Policy	\hat{R}_t	MEDEVAC 1 busy and $Q_{t11} = 1$ Queue\dispatch\reject	MEDEVAC 2 busy and $Q_{t21} = 1$ Queue\dispatch\reject
Optimal	(1, 1)	Reject	N/A
	(1, 2)	Reject	N/A
	(2, 1)	N/A	Reject
	(2, 2)	N/A	Reject
Myopic 1	(1, 1)	Queue	N/A
	(1, 2)	Reject	N/A
	(2, 1)	N/A	Queue
	(2, 2)	N/A	Reject
Myopic 2	(1, 1)	Queue	N/A
	(1, 2)	Reject	N/A
	(2, 1)	N/A	Queue
	(2, 2)	N/A	Reject
Myopic 3	(1, 1)	N/A	N/A
	(1, 2)	N/A	N/A
	(2, 1)	N/A	N/A
	(2, 2)	N/A	N/A

Table 18 Expected total discounted rewards and optimality gaps for intra-zone policies

Policy, π	$J^{\pi^*}(S^0)$	Optimality gap
Optimal	60.23	N/A
Myopic 1	45.99	23.64%
Myopic 2	55.74	7.45%
Myopic 3	54.69	9.25%

and intra-zone myopic policies when the MEDEVAC system is in State S^0 is displayed in Table 18, along with the optimality gaps associated with each intra-zone myopic policy. The results indicate that the best intra-zone myopic policy is Myopic 2, which has the smallest optimality gap of 7.45%. The intra-zone Myopic 1 policy performs worse than every other policy with the largest optimality gap of 23.64%. These results indicate that, when intra-zone policy restrictions are enforced, the myopic dispatching policies substantially under-perform compared to the optimal policy. The 2×2 case optimality gaps (for the inter-zone policies) displayed in Table 7 are substantially less than the optimality gaps for the intra-zone policies displayed in Table 18. The Myopic 2 policy has the best optimality gaps for both the 2×2 case and the intra-zone policy excursion. However, the optimality gap for the Myopic 2 policy in the 2×2 case is 0.74% whereas the Myopic 2 optimality gap in the intra-zone policy excursion is 7.45%. Moreover, there is an even larger difference between the Myopic 1 policies (2.21 vs. 23.64%). These results show that intra-zone policies perform substantially worse than inter-zone policies. Moreover, these results inform military medical planners considering the cost associated with an intra-zone dispatching policy.

Table 19 Policy iteration versus linear programming computational efficiency (s)

$ \mathcal{S} $	Policy iteration	CPLEX (dual)	CPLEX (primal)
720	0.03	0.07	1.17
3645	0.10	0.47	4.26
11,520	0.35	3.13	38.15
28,125	1.04	13.32	196.50
58,320	2.23	55.74	656.16
108,045	5.12	134.03	1981.55
184,320	10.91	216.66	4782.90
295,245	17.37	309.83	9754.49
450,000	47.86	412.66	17037.00
933,120	148.24	791.14	33212.01
1,285,245	264.54	967.03	42817.76

5.7 Policy iteration versus linear programming

This section compares the computational efficiency between policy iteration via MATLAB and linear programming (LP) via CPLEX 12.6 for the MEDEVAC dispatching problem. Since each solution algorithm determines the optimal dispatching policy, the focus of the analysis is on how long it takes each algorithm to identify the optimal policy. Comparisons are made on the same computer and on the same problem instances after they have been loaded into memory. The problem instances are generated by adjusting the q^{max} parameter in the 2×2 case. Table 19 reports the total time in seconds required to find the optimal policy for each algorithm.

The results from Table 19 indicate that the computational efficiency in solving the MEDEVAC dispatching problem utilizing CPLEX 12.6 (with either its primal or dual Simplex optimizer) is substantially worse than utilizing policy iteration. Moreover, the gaps between each algorithm increase as $|\mathcal{S}|$ increases, indicating that larger, small-scale problems (i.e., ones that can still be solved to optimality) should be solved via policy iteration. These results comport with the findings of Puterman (1994).

LP problems can be stated in primal or dual form. Moreover, the optimal solution (if one exists) of the dual has a direct relationship to an optimal solution of the primal LP model. The dual Simplex optimizer in CPLEX takes advantage of this relationship, identifying a dual basic feasible solution and iteratively improving it while maintaining complementary slackness until primal feasibility is attained, yielding a solution to the original (primal) formulation. For the primal LP model of the MDP, the number of rows (i.e., inequality constraints) is equal to $|\mathcal{S}| \times \prod_{S \in \mathcal{S}} |\mathcal{X}(S)|$ (i.e., the number of state-action combinations). The number of columns (i.e., the number of variables) is equal to $|\mathcal{S}|$. Modern LP solvers can handle problems with tens of thousands of constraints without difficulty (Powell 2011). Based on the sizes of the state and action space, it may be more efficient to solve the problems in the dual space, and hence via the dual Simplex method, resulting in $|\mathcal{S}|$ rows and $|\mathcal{S}| \times \prod_{S \in \mathcal{S}} |\mathcal{X}(S)|$ columns in the dual formulation's constraint matrix. Despite the greatly increased computational efficiency in LP algorithms reported in Bixby (2012), the results from this analysis indicate that policy iteration substantially outperforms LP via CPLEX (for both primal and dual Simplex methods) for the MEDEVAC dispatching problem.

6 Conclusions

This paper examines the medical evacuation (MEDEVAC) dispatching problem. The objective of this research is to determine how to optimally dispatch MEDEVAC units to 9-line MEDEVAC requests to improve the performance of a deployed medical service system and ultimately maximize battlefield casualty survivability rates. A discounted, infinite horizon Markov decision process (MDP) is developed to enable examination of many different military medical planning scenarios. The MDP model incorporates admission control and queueing, which allows the dispatching authority to accept, reject, or queue incoming 9-line MEDEVAC requests based on the request's classification (i.e., zone and precedence level) and the state of the MEDEVAC system. Rejected requests are not simply discarded; rather, they are redirected to another servicing agency, such as casualty evacuation, to be serviced. The MDP model also accounts for the severity of each call (i.e., urgent and priority) and applies a survivability function that is monotonically decreasing in response time to model the outcome of casualties. While response time thresholds are typically utilized to measure system performance for emergency medical systems, this paper measures performance in terms of casualty survivability since survival probability more accurately represents casualty outcomes. To demonstrate the applicability of the MDP model and to examine the behavior of the optimal dispatching policy, a notional military planning scenario based on contingency operations in southern Afghanistan is developed. A series of sensitivity analyses and computational excursions identifies the model parameters that significantly impact the optimal dispatching policy. Moreover, this paper compares the computational efficiency of policy iteration via MATLAB versus linear programming via CPLEX, utilizing either of two embedded simplex implementation methodologies.

The immediate expected reward obtained from servicing a specific 9-line MEDEVAC request depends on the locations of the request and the servicing MEDEVAC unit's staging area, along with the precedence level of the request. The total holding cost that the MEDEVAC system incurs during each state transition depends on the total number of queued requests and the precedence level of each queued request in the MEDEVAC system. Decisions are made when either a 9-line MEDEVAC request is submitted to the system or when a MEDEVAC unit finishes servicing a request. The dispatching authority examines the entire state of the MEDEVAC system when a decision is required.

Results indicate that dispatching the closest available MEDEVAC unit (i.e., a myopic policy) is not always optimal. Instead, dispatching MEDEVAC units considering the entire MEDEVAC system state (i.e., the MEDEVAC units' status, number and precedence level of queued requests, and location and precedence of the incoming request) increases the casualty survivability. The optimality gaps between the myopic policies examined and the optimal policy range between 0.74 and 5.73% when inter-zone policies are allowed and 7.45 and 23.64% when intra-zone policies are enforced. Over a protracted conflict, these policies will substantially decrease the survivability rates of battlefield casualties, and, therefore, implementation of optimal policies should be considered by medical planners. Myopic policies are often utilized in military practice because they are relatively easy to implement and they perform well as long as 9-line MEDEVAC requests arrive less frequently. Of the myopic policies tested in the 2×2 case, the Myopic 2 policy (i.e., only queue urgent requests) performs best with an optimality gap of 0.74%.

Moreover, results confirm the criticality of the MEDEVAC helicopter's flight speed. Current flight speeds can decrease due to atmospheric, environmental, or mechanical issues. If these problems arise during combat operations and degrade the flight speed of the MEDEVAC helicopters, myopic policies perform even worse compared to the optimal policy. For exam-

ple, if the current flight speeds of MEDEVAC helicopters decrease by 50%, a myopic policy that queues all requests when no MEDEVAC units are available has a 17.07% optimality gap, substantially more than the baseline optimality gap of 2.21%. These results suggest that medical planners should consider changing dispatching policies during combat operations if one or more of these problems arise and negatively impact the flight speed of the MEDEVAC helicopters being utilized. Conversely, current flight speeds can increase if new rotary wing aircraft are employed in combat operations. Were this to occur, initial results indicate that, as the flight speed increases, the performance gap between myopic policies and the optimal policy decreases. For example, if the current flight speed of MEDEVAC helicopters increases by 50%, a myopic policy that queues only urgent requests when no MEDEVAC units are available only has a 0.24% optimality gap, which is less than the baseline optimality gap of 0.74%. This comparison informs current MEDEVAC helicopter designs and development, and it provides promising results for saving lives with a faster MEDEVAC helicopter.

The research presented in this paper is of interest to both military and civilian medical planners and dispatch authorities. Medical planners can apply the MDP model developed to compare different dispatching policies for a variety of planning scenarios with fixed medical treatment facility (MTF) and MEDEVAC staging locations (i.e., hospital and ambulance locations for the civilian sector). Moreover, medical planners can evaluate different location schemes for the medical assets (e.g., MTFs, hospitals, MEDEVAC stations, and ambulances) to maximize the overall performance of the medical system.

One limiting assumption associated with the MDP model developed is that MEDEVAC units are required to return to their own staging areas to refuel and replenish medical supplies after unloading casualties at an MTF prior to servicing a queued request. During combat operations, there are typically bases that have collocated MEDEVAC units and MTFs. It is reasonable to assume that MEDEVAC units staged in different zones can refuel and replenish medical supplies at these locations and immediately proceed to service a queued request instead of first returning to their own staging areas. The MDP model restricts MEDEVAC units from refueling at different locations as a simplifying assumption. Modifying the problem formulation and the corresponding MDP model to allow for refueling, replenishing of supplies, and the ability to immediately service queued requests after casualty delivery at an MTF with a collocated MEDEVAC unit would certainly reduce the response time for many 9-line MEDEVAC requests. This modification is a planned extension for future research.

The computational effort required to solve the MEDEVAC dispatching problem increases substantially as the size of the state space grows. The computational efficiency of policy iteration via MATLAB is compared to linear programming (LP) via CPLEX. The results reveal that, although great improvements have been made concerning the performance of LP algorithms (Bixby 2012), policy iteration still outperforms LP algorithms by a substantial amount. Nevertheless, as the size of the state space grows exponentially, the use of exact dynamic programming techniques becomes intractable. This makes more realistic, large-scale problem instances impossible to analyze via exact algorithms. A planned extension to this work involves incorporating several approximate dynamic programming algorithms to address the issue known as the curse of dimensionality. Although the representative scenario analyzed is not a large-scale scenario, important insights are still drawn concerning the differences between the optimal policy and standard myopic policies utilized today. These insights should be taken into consideration by military medical planners and utilized when planning for major combat operations.

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

the United States Government. The authors would like to thank the United States Army Medical Evacuation Proponency Directorate for its support of this research.

References

- Alsalloum, O. I., & Rand, G. K. (2006). Extensions to emergency vehicle location models. *Computers & Operations Research*, 33(9), 2725–2743.
- Bandara, D., Mayorga, M. E., & McLay, L. A. (2012). Optimal dispatching strategies for emergency vehicles to increase patient survivability. *International Journal of Operational Research*, 15(2), 195–214.
- Bandara, D., Mayorga, M. E., & McLay, L. A. (2014). Priority dispatching strategies for EMS systems. *Journal of the Operational Research Society*, 65(4), 572–587.
- Bastian, N. D. (2010). A robust, multi-criteria modeling approach for optimizing aeromedical evacuation asset emplacement. *The Journal of Defense Modeling and Simulation: Applications, Methodology, Technology*, 7(1), 5–23.
- Bastian, N. D., Fulton, L. V., Mitchell, R., Pollard, W., Wierschem, D., & Wilson, R. (2012). The future of vertical lift: Initial insights for aircraft capability and medical planning. *Military Medicine*, 177(7), 863–869.
- Bixby, R. E. (2012). A brief history of linear and mixed-integer programming computation. *Documenta Mathematica, Extra Volume: Optimization Stories*, 107–121.
- Carter, G. M., Chaiken, J. M., & Ignall, E. (1972). Response areas for two emergency units. *Operations Research*, 20(3), 571–594.
- Clarke, J. E., & Davis, P. R. (2012). Medical evacuation and triage of combat casualties in Helmand Province, Afghanistan: October 2010–April 2011. *Military Medicine*, 177(11), 1261–1266.
- Cox, M. (2016). Bell touts future army helicopter design: 'V-280 is not a V-22'. <http://www.military.com/daily-news/2016/01/15/bell-touts-future-army-helicopter-design-v280-is-not-a-v22.html>. Accessed 7 September 2016.
- De Lorenzo, R. A. (2003). *Military casualty evacuation: MEDEVAC* (pp. 45–59)., Aeromedical evacuation: Management of acute and stabilized patients New York: Springer.
- Department of Defense. (2016). Defense Casualties Analysis System (DCAS) Operation FREEDOM'S SENTINEL (OFS). https://www.dmdc.osd.mil/dcas/pages/casualties_ofs.xhtml. Accessed 22 December 2016.
- Department of the Army. (2000). Field manual 8-10-6, medical evacuation in a theater of operations.
- Department of the Army. (2016). Army Techniques Publication 4-02.2, medical evacuation. Change 1.
- Eastridge, B. J., Mabry, R. L., Seguin, P., Cantrell, J., Tops, T., Uribe, P., et al. (2012). Death on the battlefield (2001–2011): Implications for the future of combat casualty care. *Journal of Trauma and Acute Care Surgery*, 73(6), S431–S437.
- Fish, P. N. (2014). *Army medical officer's guide*. Mechanicsburg: Stackpole Books.
- Fulton, L. V., Lasdon, L. S., McDaniel, R. R., & Coppola, M. N. (2010). Two-stage stochastic optimization for the allocation of medical assets in steady-state combat operations. *The Journal of Defense Modeling and Simulation: Applications, Methodology, Technology*, 7(2), 89–102.
- Fulton, L., McMurry, P., & Kerr, B. (2009). A Monte Carlo simulation of air ambulance requirements during major combat operations. *Military Medicine*, 174(6), 610–614.
- Garrett, M. X. (2013). USCENTCOM review of MEDEVAC procedures in Afghanistan. Technical report, United States Central Command.
- Grannan, B. C., Bastian, N. D., & McLay, L. A. (2015). A maximum expected covering problem for locating and dispatching two classes of military medical evacuation air assets. *Optimization Letters*, 9(8), 1511–1531.
- Green, L. V., & Kolesar, P. J. (2004). Anniversary article: Improving emergency responsiveness with management science. *Management Science*, 50(8), 1001–1014.
- Gross, D., & Harris, C. M. (1998). *Fundamentals of queueing theory* (4th ed.). Hoboken: Wiley.
- Hoffman, M. (2015). Army wants more adaptive HH-60 medical evacuation systems. <http://www.military.com/daily-news/2015/04/03/army-wants-more-adaptive-hh60-medical-evacuation-systems.html>. Accessed 7 September 2016.
- International Council on Security and Development (ICOS). (2008). Afghanistan–Pakistan Insurgent activities in Afghanistan and Pakistan (2007). http://www.icosgroup.net/wp-content/gallery/taliban-presence/016_map.png. Accessed 4 January 2017.
- Jarvis, J. P. (1985). Approximating the equilibrium behavior of multi-server loss systems. *Management Science*, 31(2), 235–239.

- Keneally, S. K., Robbins, M. J., & Lunday, B. J. (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., Howard, J. T., Orman, J. A., Tarpey, B. W., Bailey, J. A., Champion, H. R., et al. (2016). The effect of a golden hour policy on the morbidity and mortality of combat casualties. *JAMA Surgery*, 151(1), 15–24.
- Kuisma, M., Holmström, P., Repo, J., Määttä, T., Nousila-Wiik, M., & Boyd, J. (2004). Prehospital mortality in an EMS system using medical priority dispatching: A community based cohort study. *Resuscitation*, 61(3), 297–302.
- Kulkarni, V. G. (2009). *Modeling and analysis of stochastic systems* (2nd ed.). Boca Raton: CRC Press.
- Lejeune, M. A., & Margot, F. (2016). *Aeromedical battlefield evacuation under endogenous uncertainty*. Technical report, Carnegie Mellon University, Pittsburgh, PA.
- Leoni, R. D. (2007). *Black hawk: The story of a world class helicopter*. Reston: American Institute of Aeronautics.
- MacFarlane, C., & Benn, C. (2003). Evaluation of emergency medical services systems: A classification to assist in determination of indicators. *Emergency Medicine Journal*, 20(2), 188–191.
- Malsby, R. F. I. I., Quesada, J., Powell-Dunford, N., Kinoshita, R., Kurtz, J., Gehlen, W., et al. (2013). Prehospital blood product transfusion by US Army MEDEVAC during combat operations in Afghanistan: A process improvement initiative. *Military Medicine*, 178(7), 785–791.
- Maxwell, M. S., Restrepo, M., Henderson, S. G., & Topaloglu, H. (2010). Approximate dynamic programming for ambulance redeployment. *INFORMS Journal on Computing*, 22(2), 266–281.
- McLay, L. A., & Mayorga, M. E. (2010). Evaluating emergency medical service performance measures. *Health Care Management Science*, 13(2), 124–136.
- McLay, L. A., & Mayorga, M. E. (2013). A model for optimally dispatching ambulances to emergency calls with classification errors in patient priorities. *IIE Transactions*, 45(1), 1–24.
- Nicholl, J., Coleman, P., Parry, G., Turner, J., & Dixon, S. (1999). Emergency priority dispatch systems a new era in the provision of ambulance services in the UK. *Pre-Hospital Immediate Care*, 3(2), 71–75.
- O’Shea, B. (2011). Saving lives on the battlefield. *Military Medical/CBRN Technology*, 15(6), 8–13.
- Powell, W. B. (2011). *Approximate dynamic programming: Solving the curses of dimensionality* (2nd ed.). Princeton: Wiley.
- Puterman, M. L. (1994). *Markov decision processes: Discrete stochastic dynamic programming*. Hoboken: Wiley.
- Rettke, A. J., Robbins, M. J., & Lunday, B. J. (2016). Approximate dynamic programming for the dispatch of military medical evacuation assets. *European Journal of Operational Research*, 254(3), 824–839.
- Shenker, S., & Weinrib, A. (1989). The optimal control of heterogeneous queueing systems: A paradigm for load-sharing and routing. *IEEE Transactions on Computers*, 38(12), 1724–1735.
- Stidham, S. (1985). Optimal control of admission to a queueing system. *IEEE Transactions on Automatic Control*, 30(8), 705–713.
- Stidham, S., Jr. (2002). Analysis, design, and control of queueing systems. *Operations Research*, 50(1), 197–216.
- Stidham, S., Jr., & Weber, R. (1993). A survey of Markov decision models for control of networks of queues. *Queueing Systems*, 13(1–3), 291–314.
- Sundstrom, S. C., Blood, C. G., & Matheny, S. A. (1996). The optimal placement of casualty evacuation assets: A linear programming model. In *Proceedings of the 28th conference on winter simulation* (pp. 907–911). IEEE Computer Society.
- White, M. (2016). Operation ENDURING FREEDOM (OEF), Fatalities by Provinces. <http://icasualties.org/OEF/ByProvince.aspx>. Accessed 22 December 2016.
- Zeto, J., Yamada, W., & Collins, G. (2006). Optimizing the emplacement of scarce resources to maximize the expected coverage of a geographically variant demand function. In *Proceedings of technical report*, US Center for Army Analysis, Ft Belvoir.