Towards an Optimal Policy of Mass Casualty Trauma Triage

Alexandrea K. Ramnarine, MS[1], Nathaniel D. Bastian, PhD, MEng, MS[1], Anne M. Stey, MD, MSc[2]

[1]Northwestern University School of Professional Studies, Chicago, Illinois, USA; [2]Northwestern Medicine, Chicago, Illinois, USA

Abstract

Mass Casualty Incidents (MCIs) have devastating consequences throughout the U.S., highlighting a failure of primary and secondary triage policy in urban healthcare systems. It is imperative to provide optimal care for trauma victims in the shortest amount of time while avoiding burden on level I trauma centers. Reinforcement learning can leverage patient and hospital data to program a computational agent acting as incidence control. By providing the agent with rewards and goals in a simulated emergency environment, it is possible to generate optimal policies for various MCIs to assess and develop a standardized, accepted trauma triage practice.

Background

Emergency departments in the United States generate massive quantities of data annually, captured by a network of hospital trauma centers contributing to a standardized database curated by the American College of Surgeons Committee on Trauma, the National Trauma Data Bank (NTDB). NTDB data can be leveraged to model patient distributions and prescribe triage prioritization of emergency centers responsive to simulated trauma case influx over a time period. The proposed prescriptive research addresses major concerns of retriage and optimizations to avoid burdening bed availability of the highest-level trauma centers . Additionally, there is no "gold-standard" trauma triage protocol adapted throughout the U.S., nor is there clear policy indicating that trauma triage protocols are necessarily distinct given the heterogenous features and outcomes of emergency situations.

Methods

A Q-learning reinforcement learning algorithm was developed to triage one patient or a batch of *n* patients at each decision epoch for a given MCI patient load using queueing theory, discrete event simulation, and Markov processes. Five environments were modeled based on real urban U.S. MCIs using fuzzy logic.

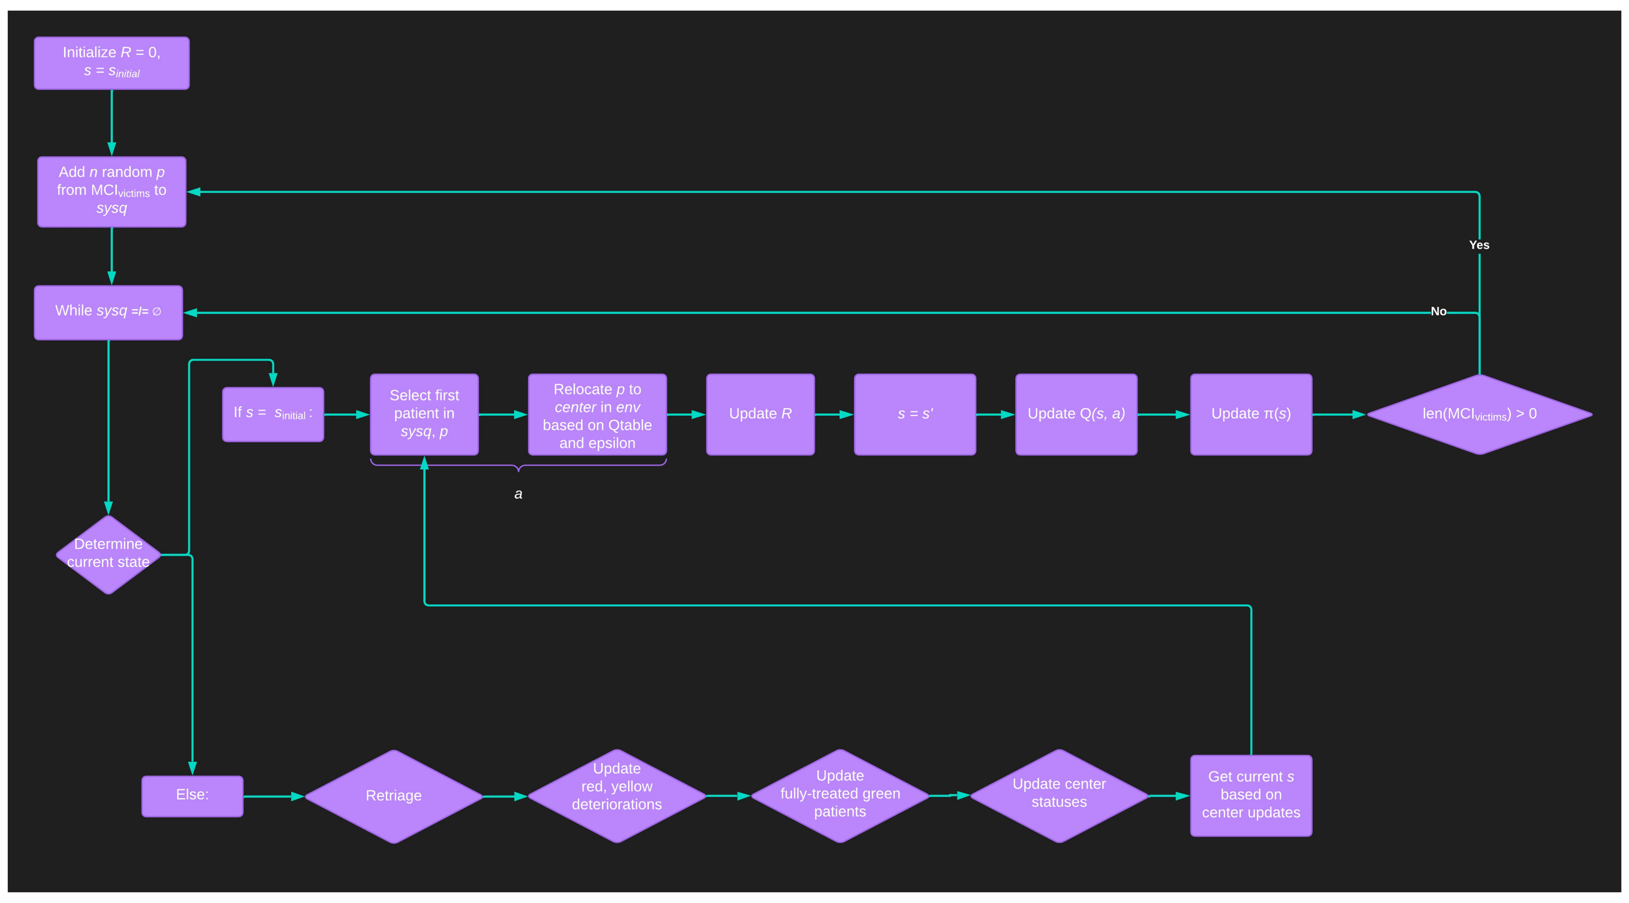

Patient data were generated for each MCI from NTDB using fuzzy logic. The natural trimodal distribution of patient labels from NTDB years 2002-2016 was used to set patient labels: "red" (most severe, subject to deterioration) from (0.0, 0.2] survival probability, "yellow" (moderately injured, subject to deterioration) from (0.2, 0.8), and "green" (minimally injured) from [0.8, 1].

MCI	Number of Patients Needed [2]	Number of Patients from NTDB	Inclusion Criteria	Exclusion Criteria	Patient Labels Present
Continental Flight Crash	46	12166	• Year of Admit = 2008 • Injury Type = Blunt • Mechanism = Transport/other	Intent = Assault	red, yellow, green
Station Nightclub Fire	200	155	• Year of Admit = 2003 • Injury Type = Burn • Mechanism = Fire/flame, Fire/burn, Hot object/substance	Intent = Assault	red, yellow, green
Pulse Nightclub Shooting	271	1084	• Year of Admit = 2016 • Intent = Assault • Injury Type = Penetrating • Mechanism = Firearm	N/A	red, green
Hurricane Katrina	7600	9272	• Year of Admit = 2005 • Mechanism = Natural/environmental, Natural/environmental/Other, Suffocation, Struck by/against, Fall, Other specified and classifiable	Intent = Assault	red, yellow, green
9/11	6000	6015	• Year of Admit = 2002 • Mechanism = Fall, Struck by/against, Cut/pierce, Fire/burn, Suffocation, Fire/flame	Intent = Unintentional	red, yellow, green

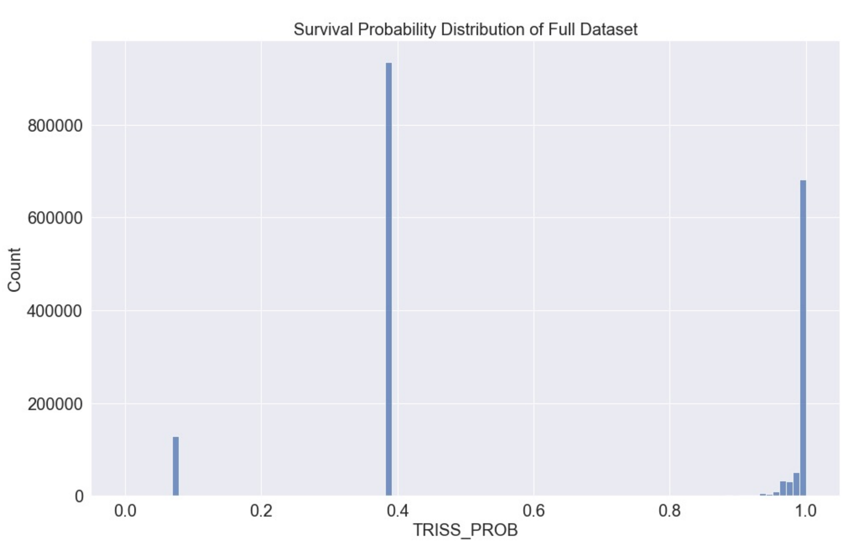

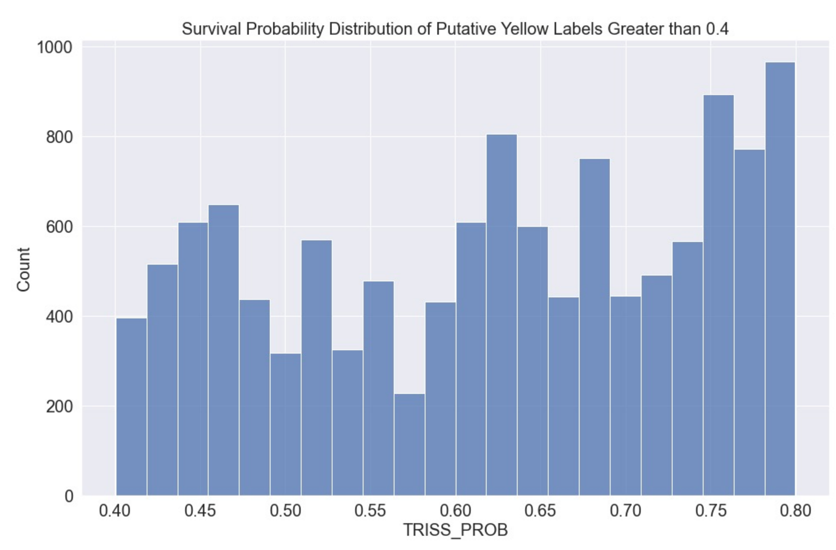

Results

Computation time increased exponentially as patient load increases, where the largest load takes ~3.33 minutes to run and the smallest runs in less than 0.5 seconds.

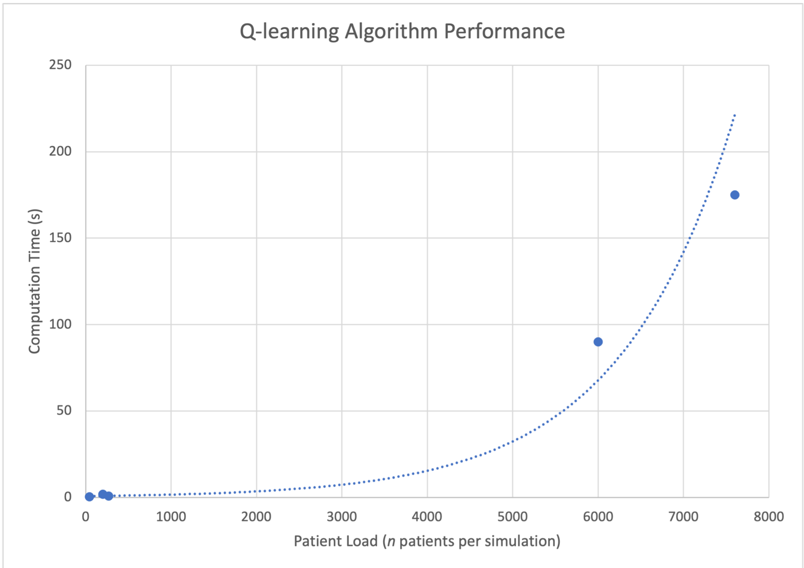

By 200 training episodes, the agent successfully learns to maximize future discounted reward and minimize both under- and over-triage rates at 100% bed availability for each center in the environment using an epsilon-greedy strategy to update the Bellman Equation.

$$Q(s_t, a_t) = (1-\alpha)Q(s_t, a_t) + \alpha[r_{t+1} + \gamma \max_a Q(s_{t+1}, a) - Q(s_t, a_t)]$$

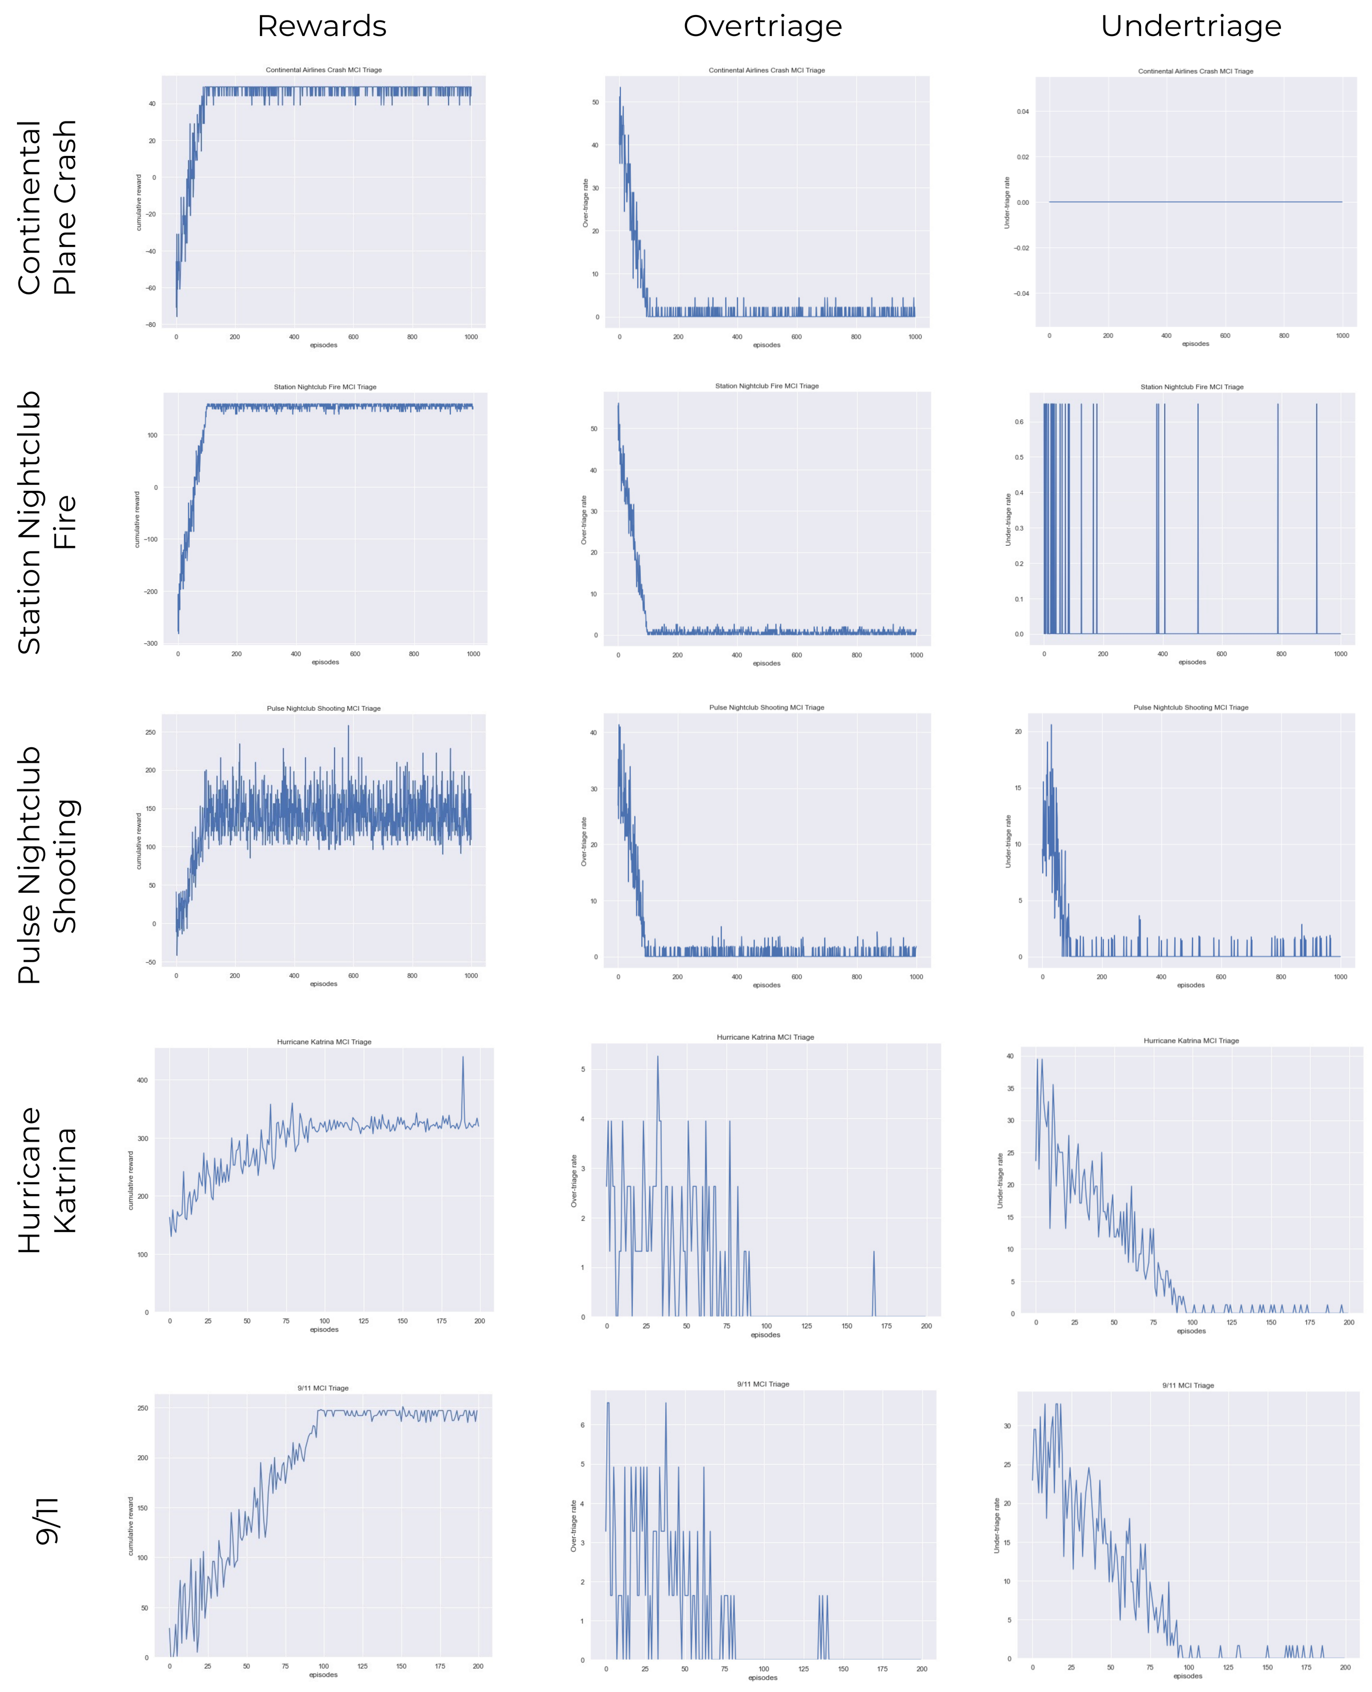

Conclusion

As patient load increases, the agent has more decision epochs to traverse; however, the epsilon-greedy strategy implemented still generates an optimal policy within 200 episodes. For the smaller patient load MCIs, undertriage rate is almost always zero with an optimal policy. The preliminary triage algorithms, based on domain-expertise from CDC Field Triage guidelines and the 2019 ASPR TRACIE "Mass Casualty Trauma Triage Paradigms and Pitfalls" is thus a reliable tool for MCI triage decision-making for incidence response, EMS, and first responders when implemented as a Markov Decision Process for Q-learning.

Future Work

Bed availability and patient influx sensitivity analyses were conducted for each MCI to assess how the Q-learning agent performs under system-wide stress. To increase computational efficiency for larger patient loads, a deep Q-Network with multi-action capabilities will be implemented to update the Bellman Equation. To address the NTDB-based categorical patient labeling, unsupervised natural language processing of injury descriptions will generate an embedding space for a neural network to predict injury label based on injury-specific word tokens. Ultimately, the algorithm and label predictor could be used in mobile applications designed for EMS-incidence response communication at MCI scenes for rapid, accurate triage under duress during catastrophic disasters.