A Data-Based Ambulance Diversion Decision-Support Framework Tailored to Washington, D.C.'s Hospital System

Donia Drias[1], Kaitlyn Frost[1], Collin Schwabicello[1], and Eric Dano[1,*]
[*] *Corresponding Author: ericdano@gwu.edu*

[1] *Department of Engineering Management and Systems Engineering,*
The George Washgington University.
Washington D.C., United States of America

***Abstract*—Emergency Departments (EDs) in Washington, D.C. experience persistent operational strain, including long patient wait times, extended ambulance offload delays, and uneven distribution of emergency patients across the hospital network. Addressing the District's current lack of standardized diversion framework is essential to improving patient safety, strengthening emergency response capacity, and promoting equitable access to emergency care.**

This work focuses on the development of a data-based ambulance diversion decision-support framework tailored to Washington, D.C. 's hospital system. The goal is to identify key indicators of ED strain that are most strongly linked to poor patient outcomes and to use these indicators to guide EMS (Emergency Medical Services) transport decisions across the regional hospital network.

The methodology integrates a structured literature review and statistical and machine learning techniques. Indicators of ED strain are identified from the literature. The most influential indicators are then used to train a random forest model that estimates diversion risk and supports EMS routing decisions. Model performance is compared against diversion strategies documented in other large metropolitan areas. Results include improved prediction of adverse patient outcomes, and demonstrated reductions in ambulance offload delays and patient wait times under simulated diversion scenarios. The resulting framework provides DC Health with a scalable, data-driven tool to improve EMS coordination, enhanced system resilience, and support more equitable emergency care delivery.

Keywords—emergency department, diversion, random forest, machine learning

I. Introduction

A. Problem Statement

Current data show that average ED wait times in the District regularly exceed five hours, and ambulance offload delays often approach one hour before patients can be transferred into care. These extended delays signal sustained pressure on ED operations and contribute to downstream effects such as prolonged throughput, treatment bottlenecks, and reduced staffing flexibility. As a result, D.C. hospitals experience inconsistent patient flow, with some facilities routinely operating at or near capacity while others face fluctuations that are difficult to predict or manage.

This operational strain directly affects the broader emergency response system. When ambulances remain at overcrowded EDs, EMS availability across the city declines, increasing response times for new emergencies and reducing the system's resilience during peak demand or high-acuity events.

B. Purpose

The purpose of this project is to develop an algorithm that improves Emergency Department flow in Washington, D.C. by enhancing coordination between EMS and the hospital system. Specifically, the project seeks to identify the operational indicators most strongly associated with ED strain and negative patient outcomes, and to use these indicators to build a data-driven model that can support real-time EMS decision making. By tailoring the analysis to D.C.'s unique hospital network, this effort aims to provide DC Health with a tool that helps distribute EMS transports more efficiently, reduce bottlenecks in the system, and improve overall patient safety.

C. Significance

Improving Emergency Department flow in Washington, D.C. carries substantial clinical and operational importance. Prolonged wait times, extended ambulance offload delays, and inconsistent patient distribution across hospitals not only compromise individual patient care but also reduce the overall effectiveness of the city's emergency response system. By developing an algorithm grounded in indicators tied to ED strain and patient outcomes, this project aims to strengthen the sequence of emergency care. The framework will support faster time to treatment, increase EMS unit availability, and help stabilize hospital operations during periods of high demand. Importantly, this work also addresses long-standing equity concerns, particularly in areas of the city where access to timely emergency care has historically been limited.

D. Scope

The resulting framework will support EMS coordination, but it cannot account for all unpredictable factors influencing emergency care such as sudden mass-casualty incidents, simultaneous surges, or operational decisions made at the individual hospital level. Additionally, the model does not enforce hospital and/or DC Health's actions. It provides evidence-based guidance intended to support decision-making within existing operational structures. All data used is informed via literature review, not collected from D.C. Emergency Departments.

II. Literature Review

The literature review was divided into four sections: emergency department strain, ED operations, ambulance diversion strategies, and ED mortality rates. Together these four topics establish the technical foundation for the proposed strain equation and decision support framework.

Emergency department overcrowding in urban healthcare systems has been widely documented as a systemic issue driven by rising patient demand, limited bed capacity, and staffing shortages. Sartini et al. identified a broad range of contributing factors and consequences, including increased patient wait times, patients leaving without being seen, and elevated mortality and morbidity [1]. These findings helped the research team identify candidate variables for the strain equation and confirmed that overcrowding is both measurable and consequential.

On ED operations, Kolker demonstrated that ambulance diversion can be meaningfully reduced by managing patient length of stay, establishing that diversion rates are negligible when Length of Stay (LOS) does not exceed six hours and increase substantially when more than eleven patients are waiting [2]. This provided a quantitative foundation for linking operational variables to diversion risk in the context of Washington D.C.

For ambulance diversion strategies, the research team examined both academic and policy sources. Kao, Yang, and Lin found that ambulance-only diversion has little impact on overall ED crowding, as ambulances represent only about 20% of total patient arrivals, and emphasized that regional coordination is essential for any effective diversion strategy [3]. The Los Angeles County EMS Agency's diversion policy framework demonstrated how a large metropolitan region can operationalize diversion standards, offering a practical governance model relevant to the nation's District [4].

Regarding ED mortality rates, Jones et al. found that delays to inpatient admission exceeding five hours are associated with increased 30-day mortality, with the greatest change occurring in patients who waited six to eight hours from arrival. This directly informed the research team's decision to include wait time, ED length of stay, and boarding time as the core variables in the prototype strain equation [5].

The synthesized dataset was informed by several additional sources. National Hospital Ambulatory Medical Care Survey (NHAMCS) provided national benchmarks for rates of patients leaving without being seen, left before treatment complete, left against medical advice, and mode of arrival [6]. Peters et al. informed the Emergency Severity Index (ESI) level distributions for ambulance versus walk-in patients [7]. Horwitz et al. and Lauks et al. provided wait time distributions by ESI level [8][9]. Pitts et al. informed boarding time intervals by ESI level [10]. McCarthy et al. provided hourly and day-of-week arrival rate patterns [11]. Singer et al. linked boarding duration to mortality probability, and Chen et al. informed the relationship between patient-to-provider ratio and adverse outcomes [12][13].

III. Methodologies

The research team employed eight methods organized around four research questions, all in service of a single project goal: developing and testing an evidence-based diversion strain formula for D.C. hospital emergency departments linked to patient mortality. Fig. I sort methodologies by research question.

Methods 4.1.1 and 4.1.2 established the foundation through literature review. The first identified three candidate criteria for inclusion in a prototype strain equation; the second examined how other cities and states implement ambulance diversion programs to inform a later policy comparison. Method 4.1.3 used those candidate criteria to develop a prototype strain equation, which became the backbone of the project. Method 4.1.4 then generated a synthesized patient dataset in R, simulating 25 patient-level variables, including ESI level, wait time, boarding time, and multiple mortality indicators, drawn from published NHAMCS data and peer-reviewed studies. This data was aggregated into hourly summaries for use in modeling. Concurrently, method 4.1.5 built a discrete-event ED simulation in Simio to test how diversion decisions affect patient flow, queue buildup, and system strain under controlled conditions.

Method 4.1.6 used the synthesized dataset to build and optimize a random forest classifier in R, predicting whether any patient death would occur in a given hour. Key hyperparameters were tuned to balance recall, precision, F1, and accuracy. The resulting model was integrated into a Shiny dashboard (4.1.7), giving hospital administrators a real-time decision support tool where they can input current ED metrics and receive a diversion recommendation relative to a dynamic baseline. Finally, method 4.1.8 compared the random forest tool to existing diversion policies, categorizing those policies as rule-based, resource-based, judgment-based, or restricted, and evaluating tradeoffs in flexibility, standardization, and predictive capability.

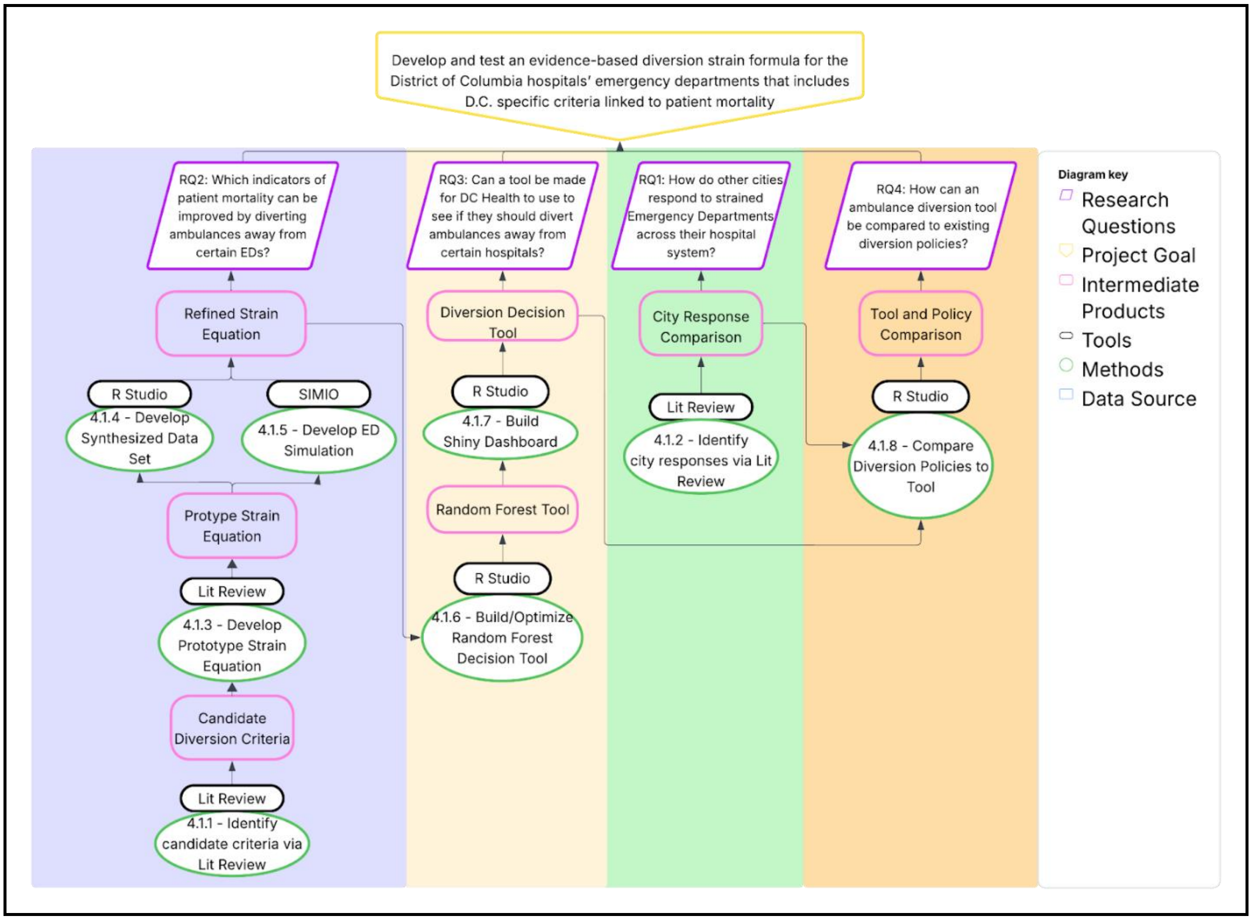

Fig 1. Methodology map describing the research methodology per research question.

IV. Results

This research developed a data-driven decision support tool for emergency department strain management, built on a random forest classification model trained on a synthesized dataset of over 228,000 patient records. The team began by identifying 28 candidate variables from a literature review, ultimately focusing on patient wait time, ED length of stay, and boarding time, which were shown to have statistically significant relationships with patient mortality. From these 28 candidate variables, we narrowed it down to an input, throughput, and output variable: patient wait time, ED length of stay, and ED boarding time. The wait time is how long it takes to see a provider, the ED length of stay is how long a patient is cared for, and the boarding time is how long a patient was boarded in the ED after disposition was assigned.

Initially, the data set consisted of only four columns: patient wait time, ED length of stay, boarding time, and negative patient outcomes. Those criteria were not enough to predict negative patient outcomes, and so the team worked to expand the model. The team also decided to restrict the predicted outcome to patient deaths. Eventually, the team was able to run a logistic regression on the data that showed it was behaving appropriately. The results of that regression are found in Table 1.

TABLE I. Logistic Regression Model Verification

Criteria	Verification Values				
	Estimate	*Std. Error*	*z Value*	*Pr(>|z|)*	*Significant*
Intercept	-1.183	8.251e-02	-14.340	<2e-16	Yes
ESI	-1.933	2.436e-02	-79.345	<2e-16	Yes
WaitTime	8.010e-03	1.216e-04	65.856	<2e-16	Yes
ED LOS	5.438e-03	9.501e-05	57.235	<2e-16	Yes
Boarding	2.168e-03	1.805e-04	12.007	<2e-16	Yes

While the values for the wait time, length of stay, and boarding are very small, they do move in the correct direction and are significant. The inverse relationship between ESI and death also makes sense since lower ESI indicates a more severe case. The research team would also expect that as wait time, LOS, and boarding increase so do the probability of a patient death. Other variables, like mode of arrival, do not have a direct impact on death, so they are not included in the logistic regression above.

The research team then worked to decide what kind of random forest model to develop. This would change depending on whether we try to predict a continuous or discrete outcome. Regressor models were tested for how many patients died in an hour and the rate of patient death over the course of an hour. A classification model was tested for the probability of any patients dying in an hour. Results for all three models can be seen in Table II.

There are different performance measurements for regressor and classifier models. Regressor models were evaluated using the R^2 and root mean squared error (RMSE), while classifier models were evaluated using accuracy, recall, precision, and F1 scores. R^2 shows the amount of variance explained by the model. In this case, using the number of dead patients in an hour as the predicted value results in a slightly more effective model. RMSE reports the average of the model's prediction errors. In this case, both RMSE are nearly the same as the standard deviation. This means that for each prediction, one may as well just guess the average number of deaths and death rate per hour calculated from the data. Neither regressor model is particularly effective.

The recall, precision, and F1 scores of the classifier model are all very low. These scores mean that the model only predicts 56.5% of any hours containing a patient's death and sends up false flags ~86% of the time. However, the accuracy is relatively high, meaning that the model correctly predicted 83.5% of all outcomes. This led the team to conclude that the classifier model is the best choice to move forward with. Hyperparameter tuning further refined the model, with a final configuration of ntree = 550, mtry = 4, nodesize = 6, and a 30% decision threshold optimized for F1 score balance between recall and precision. The optimized recall, precision, and F1 score were 61.1%, 33.6%, and 43.4% respectively.

The random forest tool was deployed through a Shiny dashboard, shown in Fig. II, that allows administrators to enter values into the dashboard regarding the number of patients that arrived at the hospital in the observed hour and how long they waited at each stage of the emergency department. Specifically, they will enter in the number of patients who arrived as walk-ins and via ambulance, the number of patients at each ESI level, the average wait time for a patient to be seen by a provider, the average length of stay in the ED, and the average boarding time. Administrators will then run an assessment at their chosen risk threshold. The research team has found that the threshold of 30% has the highest F1 score, so that threshold is recommended for our model.

For each assessment that is run, the tool will report the predicted probability of any patient death in the next hour alongside a comparison to a historical baseline as well as recall, precision, F1, and accuracy scores. Recall is the number of deaths that were correctly predicted divided by the total number of patient deaths. A high recall means that more dangerous hours were predicted correctly. Precision is the number of deaths that were correctly predicted divided by the total number of predicted deaths. A high precision means a low rate of false flags. F1 is a balance between recall and precision and is useful for comparing performance across levels. Accuracy is the percentage of all predictions that were correct. Lower thresholds will result with higher recalls, while higher thresholds result in higher precision. This is because if you lower the threshold to 15%, the model will constantly predict a patient's death, as the threshold is relatively low. This will catch most dangerous hours, but will also raise a lot of false flags.

When compared against existing diversion policies, the random forest tool was categorized simultaneously as a rule-based, resource-based, and judgment-based system, distinguishing it from the more rigid, single-category systems used elsewhere. These categories, along with "restricted", were developed from research into other diversion policies. Rule-based systems are those that follow fixed thresholds or criteria applied consistently regardless of specific hospital conditions. A good example of this is San Diego County, who recommends diversion for ED's with an ambulance patient offload time over 100 minutes. Resource based systems decisions depend on the current availability or status of critical hospital resources, such as beds, staff, or specialized equipment. A good example of this is Delaware, where one of their diversion categories (cardiac catheterization divert) is triggered when a hospital does not have the resources to provide care for patients who would need cardiac catheterization. Judgment-based systems decisions rely on the discretion and expertise of hospital administrators or staff rather than predetermined rules or measured resources. New York is a good example of this state, as diversion requests are made by hospitals based on administrative judgement, but are accepted or denied by EMS, depending on the health of the system.

TABLE II. MODEL TYPE SELECTION RESULTS

Predicted	**Model Type**	**Results**						
		R^2	RMSE	St. Dev	Accuracy	Recall	Precision	F1
# Patients Dead	Regressor	0.148	0.410	0.444				
Rate of Patient Death	Regressor	0.119	0.057	0.061				
Any Patient Dead in Hour	Classifier				0.835	0.565	0.144	0.229

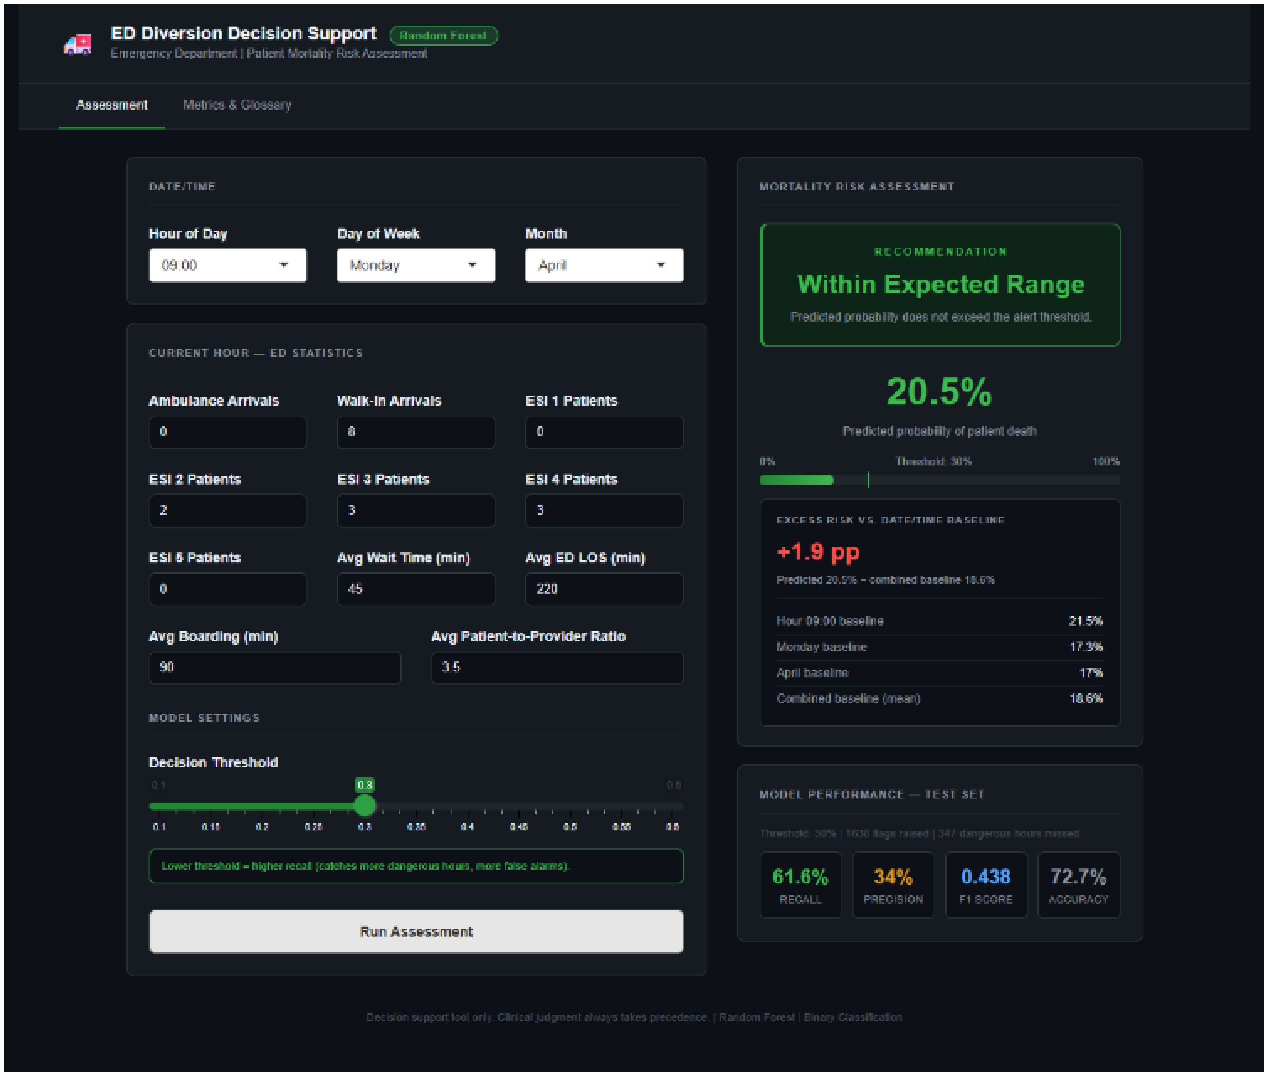

Fig 2. The Ambulance Diversion Decision Support Dashboard utilizing Random Forest Model

The research team has found that traditional policies excel in ease of use and consistency, but the predictive tool offers a more dynamic picture of ED strain that better reflects real operational conditions, while still leaving final diversion decisions to hospital administrators.

V. Discussion

The goal of this project was to develop an ambulance diversion tool for Washington D.C.'s emergency department including linking criteria to patient mortality. Taken as a whole, the results demonstrate that a data-driven approach to ambulance diversion is both technically feasible and operationally meaningful within the context of Washington D.C.'s emergency care network.

The comparison of this tool against existing diversion policies highlighted several tradeoffs between traditional diversion policies and the random forest decision tool. Traditional policies tend to provide clear, uniform rules or guidelines that are easy for staff to understand and implement. However, this simplicity comes at the cost of flexibility and sensitivity. These systems are often unable to account for the rapidly changing conditions that can occur in an emergency department. In contrast, the random forest tool emphasizes flexibility and predictive statistics. By integrating multiple real-time inputs, it can anticipate high risk situations and provide guidance that reflects the combined influence of patient arrivals, acuity levels, and staffing ratios. This allows for a more customized tool that adapts to the unique circumstances of a hospital at any given hour.

Using synthesized data introduces several potential sources of error due to its simulation. The synthesized dataset was generated based off of our literature review, which fails to capture the full picture of emergency department operations. Specifically, the modeled situation lacks the complexity and variability expected from data collected from emergency departments. This synthesized dataset is also not specific to Washington D.C. While the tool was created with D.C. in mind, the synthesized dataset is based on reviewed sources addressing emergency departments outside of Washington D.C.

This research is an important step in developing a path towards a data driven ambulance diversion policy in Washington D.C. The limits on this research are due to the simulated nature of the synthesized data and random forest model, which may not fully communicate the variability of emergency department operations and strain. To continue to iterate on this research would be to collect data from Washington D.C.'s emergency departments and create a validated data set from this information. This could then inform the developed diversion tool, and further policy actions. This continuation of research would strengthen the tool's reliability overall and efficacy in Washington D.C.

VI. Conclusions

This research addresses the challenge of emergency department overcrowding through ambulance diversion policy in Washington D.C. Through a combination of literature review, data synthesis, machine learning, and simulation, the research team developed a prototype framework. This recommendation quantifies emergency department strain and offers decision support for ambulance diversion decisions at an emergency department level.

This project demonstrates that a data-driven approach to ambulance diversion is both possible and promising in Washington D.C. By incorporating multiple strain indicators, this project offers a foundation for more responsive, evidence-based decision making with the potential to reduce emergency department overcrowding and improve patient outcomes.

Acknowledgment

The research team would like to thank Dr. Kristin Raphel of DC Health and Dr. Nathan Danneman of Logistics Management Institute (LMI). Their insights and support were critical to the success of this work.

References

[1] M. Sartini *et al.*, "Overcrowding in emergency department: Causes, consequences, and solutions—A narrative review," *Healthcare*, vol. 10, no. 9, p. 1625, Aug. 2022, doi: 10.3390/healthcare10091625.

[2] A. Kolker, "Process modeling of emergency department patient flow: Effect of patient length of stay on ED diversion," Journal of Medical Systems, vol. 32, no. 5, pp. 389–401, Apr. 2008, doi: 10.1007/s10916-008-9144-x.

[3] C.-Y. Kao, J.-C. Yang, and C.-H. Lin, "The impact of ambulance and patient diversion on crowdedness of multiple emergency departments in a region," *PLOS ONE*, vol. 10, no. 12, p. e0144227, Dec. 2015, doi: 10.1371/journal.pone.0144227

[4] County of Los Angeles, Department of Health Services, EMS Agency, "Diversion request requirements for emergency department saturation (Reference No. 503.1)," revised Jul. 1, 2025, pp. 1–5.

[5] S. Jones *et al*., "Association between delays to patient admission from the emergency department and all-cause 30-day mortality," *Emergency Medicine Journal*, vol. 39, no. 3, pp. 168–173, Jan. 2022, doi: 10.1136/emermed-2021-211234

[6] National Center for Health Statistics, "National Hospital Ambulatory Medical Care Survey: 2022," Centers for Disease Control and Prevention, Hyattsville, MD, USA. [Online]. Available: https://www.cdc.gov/nchs/data/nhamcs/web_tables/2022-nhamcs-ed-web-tables.pdf.

[7] Peters, Gregory A, et al. "Patients Who Use Emergency Medical Services Have Greater Severity of Illness or Injury Compared to Those Who Present to the Emergency Department via Other Means: A Retrospective Cohort Study." *Journal of the American College of Emergency Physicians Open*, vol. 4, no. 4, 31 July 2023, https://pubmed.ncbi.nlm.nih.gov/37529486/

[8] L. I. Horwitz, J. Green, and E. H. Bradley, "US Emergency Department Performance on Wait Time and Length of Visit," *Annals of Emergency Medicine*, vol. 55, no. 2, pp. 133–141, Feb. 2010, doi: https://doi.org/10.1016/j.annemergmed.2009.07.023. Available: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC2830619/

[9] J. Lauks, B. Mramor, K. Baumgartl, H. Maier, C. H. Nickel, and R. Bingisser, "Medical Team Evaluation: Effect on Emergency Department Waiting Time and Length of Stay," *PLOS ONE*, vol. 11, no. 4, p. e0154372, Apr. 2016, doi: https://doi.org/10.1371/journal.pone.0154372

[10] S. R. Pitts, F. L. Vaughns, M. A. Gautreau, M. W. Cogdell, and Z. Meisel, "A cross-sectional study of emergency department boarding practices in the United States," *Academic Emergency Medicine: Official Journal of the Society for Academic Emergency Medicine*, vol. 21, no. 5, pp. 497–503, May 2014, doi: https://doi.org/10.1111/acem.12375. Available: https://pubmed.ncbi.nlm.nih.gov/24842499/

[11] M. L. McCarthy, S. L. Zeger, R. Ding, D. Aronsky, N. R. Hoot, and G. D. Kelen, "The Challenge of Predicting Demand for Emergency Department Services," *Academic Emergency Medicine*, vol. 15, no. 4, pp. 337–346, Apr. 2008, doi: https://doi.org/10.1111/j.1553-2712.2008.00083.x

[12] A. J. Singer, H. C. Thode Jr, P. Viccellio, and J. M. Pines, "The Association Between Length of Emergency Department Boarding and Mortality," Academic Emergency Medicine, vol. 18, no. 12, pp. 1324–1329, Dec. 2011, doi: https://doi.org/10.1111/j.1553-2712.2011.01236.x

[13] Chen, Yi-Ying, et al. "Higher Patient-To-Physician Ratios Associated with Worse Outcomes in the Emergency Department." *Journal of the Formosan Medical Association*, vol. 125, no. 1, 4 Dec. 2024, www.sciencedirect.com/science/article/pii/S0929664624005576, https://doi.org/10.1016/j.jfma.2024.11.020.