ABSTRACT

LEAVITT, PATRICK NICHOLS. Human-Centered Machine Learning for Interpretable Disaster Response Strategies. (Under the direction of Benjamin A. Rachunok and Brandon M. McConnell).

Natural disasters, especially hurricanes, continue to pose significant threats to human life, property, and infrastructure in vulnerable coastal regions. Despite the critical role of evacuation in minimizing disaster impacts, many residents choose to shelter in place—a decision that can lead to heightened distress and strain on emergency services. The devastating effects of Hurricane Harvey in 2017, with its record rainfall and widespread flooding, underscore the complex nature of evacuation decisions, which are influenced by factors such as limited transportation, economic hardship, and intricate household dynamics.

This study investigates the underlying reasons for non-evacuation by integrating distress call data from the Cajun Navy with socio-demographic and economic variables from the American Community Survey, aggregated at the census block group level. By developing a predictive framework using robust cross-validation and machine learning techniques, we aim to identify areas with a higher likelihood of sheltering in place and needing assistance post disaster. This model not only accounts for the diverse risk factors at play but also provides valuable insights that can help emergency planners and policymakers target resources more effectively, especially within the first twenty-four hours following a disaster.

Ultimately, the insights gained from this research are intended to enhance our understanding of why certain populations opt not to evacuate, thereby informing the development of improved disaster preparedness and response strategies. By identifying key vulnerabilities and forecasting areas at risk, this work contributes to more efficient resource allocation and a stronger, more resilient approach to managing future disasters.

Human-Centered Machine Learning for Interpretable Disaster Response Strategies

by
Patrick Nichols Leavitt

A thesis submitted to the Graduate Faculty of
North Carolina State University
in partial fulfillment of the
requirements for the Degree of
Master of Science

Operations Research

Raleigh, North Carolina
2025

APPROVED BY:

Benjamin A. Rachunok Co-chair of Advisory Committee	Brandon M. McConnell Co-chair of Advisory Committee

Frederick Livingston

DEDICATION

To my beloved wife, whose unwavering support, patience, and encouragement have been my greatest source of strength throughout this journey. Your belief in me has been a constant inspiration.

To my three wonderful boys, who show me the true meaning of curiosity, resilience, and joy. You remind me every day why perseverance and hard work matter.

To my parents, whose guidance, sacrifices, and love have shaped the person I am today. Your lessons and values have been the foundation of my determination and ambition.

This work is a testament to the love and support of all of my and my wife's family. I could not have done this without you.

BIOGRAPHY

During this research, I had the honor to be sent to graduate school to pursue my Master in Operations Research and represent the United States Coast Guard, where I was dedicated to search and rescue coordination, ensuring the safety of those in distress at sea. My time in the Coast Guard provided me with first-hand experience in emergency response and disaster management, reinforcing my passion for data-driven decision making in crisis situations.

I also had the privilege of serving on both a Fast Response Cutter and a Medium Endurance Cutter, where I gained extensive operational experience in maritime security, law enforcement, and humanitarian missions. These experiences instilled in me a deep appreciation for rapid response coordination and the importance of leveraging technology and analytics to enhance mission effectiveness.

The skills and insights I gained from my time in the Coast Guard have greatly influenced my academic pursuits, particularly in the field of disaster response modeling. In many ways, this thesis is an extension of my commitment to improving emergency response strategies and ensuring that resources are allocated efficiently to those in need.

ACKNOWLEDGEMENTS

I would like to express my deepest gratitude to Dr. McConnell for believing in me and serving as my mentor throughout this journey. Your guidance, encouragement, and support have been invaluable, and I am truly grateful for the knowledge and insight you have shared with me.

I am also profoundly thankful to Dr. Rachunok for believing in my vision for this thesis and for his invaluable advice and mentorship. Your trust in my ideas and your guidance have played a crucial role in shaping this work.

Additionally, I would like to sincerely thank LCDR Matthew Edes, CAPT Kalie Benson, CAPT Eliza Garrity, and CAPT Thomas Crane for their unwavering support and mentorship throughout my twelve years in the Coast Guard. Your leadership, guidance, and belief in my potential have been instrumental in shaping both my career and my personal growth. I am deeply grateful for the lessons you have imparted and the opportunities you have provided.

I extend my sincerest appreciation to all of you for your time, patience, and belief in my abilities. This thesis would not have been possible without your support.

Disclosures: The views expressed in this paper are those of the author and do not reflect the official policy or position of the United States Coast Guard, the Department of Homeland Security, or the United States Government. The author used Grammarly for spelling, phrasing and grammar assistance. The author used ChatGPT and GitHub CoPilot for Python Coding, designing and debugging functions, and for fixing LaTeX code and making LaTeX tables. The author validated all outputs to ensure correctness.

TABLE OF CONTENTS

List of Tables . . . . . **vi**

List of Figures . . . . . **vii**

Chapter 1 Introduction **1**
1.1 Motivations 1
1.2 Literature Review 2
1.2.1 Who Doesn't Evacuate 2
1.2.2 Vulnerability and Disproportionate Risk 3
1.2.3 The Role of Social Capital in Evacuation and Recovery 3
1.3 Description 4

Chapter 2 Data **5**
2.1 Hurricane Harvey: Impacts and Lessons 5
2.2 Rescue Data 6
2.3 Covariates 6

Chapter 3 Methods **14**
3.1 Description of Models 14
3.2 Model Selection and Evaluation 17

Chapter 4 Results **19**

Chapter 5 Discussion **25**

Chapter 6 Conclusions **27**
6.1 Summary of Results 27
6.2 Future Work 28

References **29**

LIST OF TABLES

Table 2.1 Description of machine learning models used in the analysis. 11

Table 3.1 Description of machine learning models used in the analysis. 15

LIST OF FIGURES

Figure 2.1 ArcGIS map displaying rescue requests 7

Figure 3.1 Performance metrics for the random forest model before and after hyperparameter tuning. 18

Figure 4.1 The left map shows aggregated rescue requests by census tract, while the four right subplots illustrate key socioeconomic covariates in Houston: (a) Unemployed Individuals, (b) Married Couples with Children Under 18, (c) 4+ Person Households with Only One Vehicle, and (d) Population 65+ in Family Households. These factors are analyzed in relation to rescue requests from Hurricane Harvey to identify communities that may be more vulnerable to future disasters due to social, economic, and logistical constraints. This data is acquired with the R `tidycensus` package and taken from the 2013–2017 American Community Survey 20

Figure 4.2 Feature Importance: The analysis is divided into four categories. Household and Family Presence dominates with the top four features (Individual Living Difficulty, Single Parents, Extended Households, and Married with Children), followed by Poverty and Poverty and Transportation Access (Unemployed and Household with one vehicle and 4+ people),and Community Dynamics & Household Stability (moving within the same state and moving within the same county or staying in the same home). 21

Figure 4.3 This figure presents Partial Dependence Plots (PDP) illustrating the marginal effect of various socioeconomic factors on the probability of sheltering in place during Hurricane Harvey. PDPs provide insight into how a single predictor influences the model's predictions while averaging out the effects of other variables.(a) Households in poverty, (b) Households with insufficient vehicles, (c) Individuals with living difficulties, (d) Single-parent households, (e) Households with relatives and non-relatives, (f) Married couples with children under 18, (g) Unemployed individuals, (h) Individuals who moved within the same state, and (i) Individuals aged 65+ living alone. These PDPs highlight the importance of demographic and socioeconomic characteristics in predicting disaster response behaviors. 23

Figure 4.4 (color online) Predicted rescue locations in the Wilmington area using the tuned random forest model. The model, trained on Hurricane Harvey data, forecasts potential rescue sites for a hypothetical hurricane impacting North Carolina's coast. 24

CHAPTER

1

INTRODUCTION

1.1 Motivations

Natural disasters, particularly hurricanes, threaten human life, property, and infrastructure, especially in vulnerable coastal areas. When these storms hit, the situation is often worsened by poor evacuation planning, limited resources, and inadequate communication. For example, Hurricane Harvey, which struck Texas in August 2017, stands as one of the most costly and deadliest storms in U.S. history (National Hurricane Center 2018). Its record-breaking rainfall, widespread flooding, and destructive winds devastated communities, and many people were either unable or unwilling to evacuate. Although evacuations are critical for minimizing loss of life and human suffering, many residents end up staying in their homes even when severe danger is imminent.

Evacuation during disasters is a critical, life-altering decision that depends on a multitude of factors. Access to transportation, available economic resources, family structure, trust in evacuation orders, and past experiences with hurricanes all influence whether people choose to evacuate or shelter in place (Morss et al. 2016). Identifying areas that have higher populations that are more likely to shelter in place and need help after a disaster can help authorities better allocate life-saving goods, like food and water, and rescue and assistance resources, including boats and helicopters, reducing fatalities and long-term displacement. During Hurricane

Harvey, for instance, despite evacuation orders, some residents still choose not to evacuate—whether due to perceived low risk, misinformation, or the prohibitive costs associated with leaving. As Cutter et al. (2003) illustrate, social vulnerability and limited access to resources further exacerbate the challenges faced by these communities.

Given these complexities, there is a clear need to identify areas with a higher likelihood of sheltering in place and needing assistance after the disaster. By analyzing variables such as floodplain locations, income levels, vehicle access, household composition, and social connections, this research aims to develop a framework that policymakers and emergency planners can use to improve evacuation protocols, better target at-risk populations, and allocate resources more efficiently. In this thesis, I build a predictive model that aims to predict the census tracts where there is a higher likelihood that the population will shelter in place, rather than evacuate, during an emergency and require assistance after the disaster with real-time emergency data during Hurricane Harvey. Understanding why certain groups stay behind is critical for mitigating risks in future disasters and enhancing overall emergency preparedness and response.

1.2 Literature Review

Existing work analyzing the impact of major crisis events is focused on who doesn't evacuate, vulnerability, and disproportionate risk, and how these affect recovery in both the short and long term. Evacuation during disasters is a critical decision that can have life-or-death consequences. Various factors, including individual risk perceptions, available resources, and social networks, influence the motivation to evacuate or stay. Studies have shown that people with greater access to financial resources, transportation, and shelter are more likely to evacuate to longer distances and for more time from both inside and outside mandatory evacuation areas, while those without these resources face significant barriers (Yabe and Ukkusuri 2020). Additionally, social networks and connection to the community play a pivotal role in evacuation decisions, as individuals often make evacuation choices based on the actions and advice of family members, friends, and neighbors (Metaxa-Kakavouli et al. 2018; Sankar et al. 2020). However, not all evacuate, and those who stay may do so due to perceived low risk, lack of information, or an inability to afford the costs associated with evacuation (Sankar et al. 2020).

1.2.1 Who Doesn't Evacuate

The more difficult it is to evacuate, the more likely people will shelter in place. The question of *who evacuates* has been studied across various contexts, with research suggesting that

vulnerable populations—such as the elderly, low-income individuals, and those with limited mobility—are less likely to evacuate. This trend is particularly pronounced in low-income communities, where residents may lack the financial means to evacuate, own a vehicle, or access appropriate shelter options (Flanagan et al. 2011). As a result, these individuals often face increased exposure to the disaster's impacts and are less likely to receive timely assistance or information. Understanding why certain groups fail to evacuate and the barriers they face is crucial for developing more effective disaster preparedness strategies.

1.2.2 Vulnerability and Disproportionate Risk

The demographic makeup of disaster-affected populations significantly influences the severity of outcomes. Low-income and vulnerable populations experience disasters differently compared to wealthier groups. These groups tend to be more geographically concentrated in areas with higher exposure to natural hazards, such as floodplains, and they often reside in housing that is more vulnerable to damage (Jonkman et al. 2009). Furthermore, these populations may have less access to disaster preparedness resources, such as insurance, evacuation plans, and social support networks, exacerbating their vulnerability during and after a disaster (Flanagan et al. 2011).

The literature consistently points to increased disaster-related hardships for low-income households, including delayed or inadequate access to emergency services, longer recovery times, and more severe economic impacts (Nakagawa and Shaw 2004). Vulnerable populations, including racial minorities and the elderly, also experience worse outcomes due to limited mobility, language barriers, and greater dependence on public systems. The impacts of these disasters often extend beyond immediate damage, affecting long-term health, economic stability, and social cohesion (Flanagan et al. 2011). Therefore, it is essential to consider how demographic factors intersect with evacuation and recovery outcomes to ensure an equitable national response framework.

1.2.3 The Role of Social Capital in Evacuation and Recovery

Better-connected communities tend to recover faster with less impact on the population following a crisis. Social capital refers to the networks of relationships, trust, and mutual assistance that bind individuals within a community. Social capital can significantly affect evacuation decisions, resource sharing, and recovery efforts in disaster settings. Strong social ties allow individuals to access informal networks for assistance, such as transportation, shelter, or emotional support. Conversely, weaker social networks may limit the ability of individuals to evacuate or recover effectively (Metaxa-Kakavouli et al. 2018). Research has shown that social

capital can serve as a buffer against disaster impacts, enabling individuals to rely on neighbors and community organizations for help (Nakagawa and Shaw 2004).

Social capital can also influence the speed and effectiveness of disaster recovery. Communities with high social cohesion recover more quickly due to shared efforts to rebuild and provide resources to those in need (Nakagawa and Shaw 2004). This is especially important in low-income areas, where residents may need more formal resources but can rely on strong community bonds to address immediate needs. Understanding the role of social capital in evacuation and recovery processes can inform policies that strengthen community networks before and after a disaster.

1.3 Description

This study introduces a novel approach to predicting shelter-in-place decisions during natural disasters by leveraging distress call data—specifically, locations of individuals who chose to shelter in place and later required assistance. Unlike previous research, which focuses primarily on evacuation behaviors or generalized surveys of potential evacuation or sheltering decisions (Maas et al. 2020; Yabe and Ukkusuri 2020), our work utilizes actual post-disaster requests for assistance to model the likelihood of individuals choosing to shelter in place despite warnings or evacuation orders.

This approach fills a gap in the existing literature by moving beyond general evacuation patterns to focus on the actual decisions of people who chose not to evacuate and later needed help. By analyzing these real-world decisions, we can improve pre- and post-disaster response strategies by targeting areas with high shelter-in-place likelihoods and identifying populations most at risk of needing assistance. Our work provides a data-driven framework for more effective resource allocation, helping authorities better anticipate where aid will be required in future disasters.

CHAPTER

2

DATA

2.1 Hurricane Harvey: Impacts and Lessons

Hurricane Harvey, a Category 4 storm that struck Texas on August 25, 2017, caused unprecedented rainfall and catastrophic flooding, particularly in the Houston metropolitan area (National Hurricane Center 2018). The disaster exposed significant vulnerabilities in urban infrastructure, including insufficient zoning regulations and extensive impervious surfaces that exacerbated runoff and flooding.

“More than 19 trillion gallons of rainwater fell on parts of Texas, causing widespread, catastrophic flooding. Nearly 80,000 homes had at least 18 inches of floodwater, 23,000 of those with more than 5 feet. The Houston area experienced the largest amount of rainwater ever recorded in the continental United States from a single storm (51.88 inches). Twenty-four hospitals were evacuated, 61 communities lost drinking water capability, 23 ports were closed and 781 roads were impassable. Nearly 780,000 Texans evacuated their homes. In the days after the storm, more than 42,000 Texans were housed temporarily in 692 shelters. Local, state and federal first responders rescued 122,331 people and 5,234 pets"(FEMA 2017).

Outdated and overwhelmed drainage systems further contributed to the devastation, while many areas outside designated floodplains—where residents were less likely to have flood insurance—suffered severe flooding. Emergency response systems were also tested to their

limits, as 911 call centers were overwhelmed by the sheer volume of rescue requests, leading officials to direct residents to social media platforms for help.

2.2 Rescue Data

We used rescue data from the Cajun Navy. The Cajun Navy is a volunteer organization of private citizens, primarily from Louisiana, who use their boats to assist in search and rescue efforts during natural disasters. Originating after Hurricane Katrina in 2005, the group played a critical role in Hurricane Harvey in 2017, rescuing thousands of stranded residents across flooded areas in Texas, particularly the Houston metropolitan area. Using personal boats, canoes, and airboats, these volunteers braved treacherous, debris-filled floodwaters to reach people trapped in homes, rooftops, or vehicles. Coordinating via social media and smartphone apps, they mapped routes and prioritized critical rescues, working tirelessly alongside official emergency responders. Their efforts saved countless lives and demonstrated the power of community-driven disaster response during times of overwhelming need.

The primary dataset used in this study consists of rescue requests collected by the Cajun Navy during a major disaster event. As show in Figure 2.1 each rescue request was recorded with latitude and longitude coordinates, enabling us to map these incidents using ArcGIS. The spatial distribution of these requests provided valuable insights into the geographic areas where rescue operations were most frequently needed. In order to compare the spatial data with socio-demographic variables, the individual requests were aggregated into census tract groups.

To facilitate predictive modeling, we employed a binary classification framework where each census tract was assigned a value of 0 or 1. A value of 0 indicates that no rescue requests were recorded in the tract, while a value of 1 signifies that at least one rescue request was made. In our dataset, out of the total census tracts analyzed, 652 tracts were classified as 0 (indicating no rescue requests), and 513 tracts were classified as 1 (indicating the presence of one or more rescue requests). This binary approach not only simplified the modeling process but also provided a clear delineation between areas with and without reported distress.

2.3 Covariates

Description of Census Variables

The analysis incorporates a range of census variables drawn from the American Community Survey, each providing insights into various socio-demographic and economic aspects of the

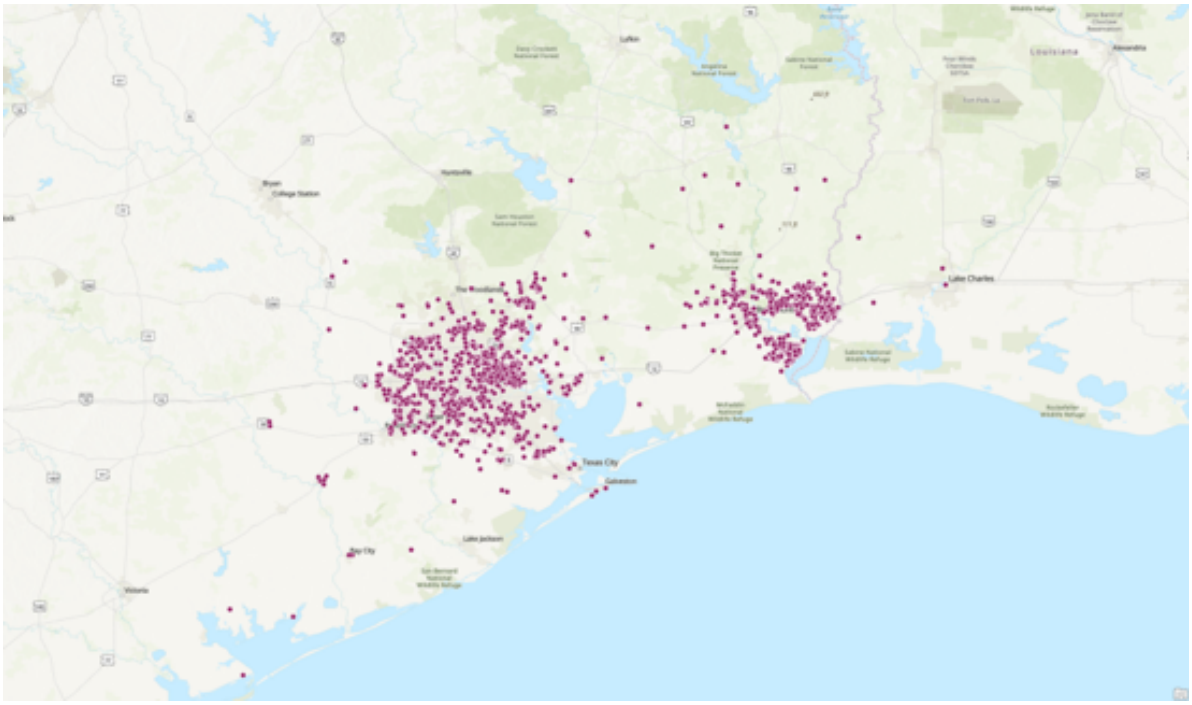

Figure 2.1: ArcGIS map displaying rescue requests

population. These variables serve as key indicators in our study of disaster vulnerability and evacuation behavior.

Poverty and Transportation Access: The variable `hhpovertyE` represents the number of individuals whose income falls below the federally defined poverty line, serving as a measure of economic disadvantage. Additionally, `ANY_HELPE` indicates the number of individuals receiving public assistance or other forms of help, while `unemployedE` provides a measure of the unemployment level in the area. The metric `severe_poverty` captures the proportion of the population experiencing extreme financial hardship, and `poverty_rent` reflects household income expressed as a percentage of gross rent. Employment stability is further assessed by `hours_worked`, which measures the average number of hours worked in the past 12 months. Transportation challenges are addressed by the variables `hhveh_notenE` and `hhveh_noneE`, which identify households—particularly those with four or more residents—that have only one vehicle or no vehicle at all. Together, these measures provide a comprehensive overview of the economic and transportation-related barriers that can impede evacuation efforts.

Household Composition and Family Presence: Family structure and household composition are critical determinants of social support and evacuation capacity. The variable `ft_MAR_wchildE` captures the prevalence of married couples with children, while `ft_MAR_wchild_U6E` specifically denotes married couples with children under the age of six, reflecting the

additional challenges faced by families with very young dependents. The indicator `Elderly_Living_With_Spouse` measures the number of seniors who live with their spouse, suggesting the availability of in-home support for older residents. Conversely, `Single_Parents` identifies households led by a single parent, highlighting potential vulnerabilities due to limited support. Furthermore, `Elderly_Living_Alone` records the number of elderly residents (both male and female, 65+ years) living alone, while `Extended_household` measures the prevalence of extended family living arrangements. Additionally, `Ind_liv_DE` reflects the number of individuals experiencing independent living difficulties. Collectively, these variables offer a detailed view of household composition, which is vital for assessing both the capacity for mutual support and the specific challenges faced during evacuations.

Community Dynamics & Household Stability: This group of variables captures broader community-level factors and the stability of living arrangements, which are essential for understanding overall community resilience. The variable `Spanish_speakE` measures the number of Spanish-speaking residents, providing insights into potential language barriers during emergency communications. Homeownership is indicated by `OwnerE`, which serves as a proxy for residential stability and community investment. `Nearby_Help` reflects the availability of local support networks, a key component in effective community response during crises, while `GroupQE` denotes the population residing in group quarters, such as dormitories, nursing homes, or correctional facilities. Lastly, `SameStateE` and `DifStateE`, as well as `ANY_HELP`, capture aspects of residential mobility by recording the number of residents who have moved within the same state or from a different state, and those residing in the same house or having moved within the county, respectively. These indicators help to assess the cohesion and stability of communities, which are critical for planning and implementing effective disaster response strategies.

Hypotheses

In our analysis, we examined a broad set of variables hypothesized to influence evacuation behavior and the need for assistance during disasters. The conceptual framework driving our study rests on the idea that risk is a function not only of hazard exposure but also of the interplay between vulnerability and available resources (Logan and Guikema 2020). More specifically, we propose that factors such as the ability to comply with evacuation orders, financial constraints (Brodie et al. 2006), access to transportation, household composition, social connectivity, and natural hazards significantly shape an individual's or household's decision to evacuate during a disaster and their subsequent need for assistance.

To capture these dynamics, we gathered data from various sources and measures. For instance, *Access to Vehicle* is measured using the `B08201 Household Size by Vehicles Available` dataset, which underscores that lower rates of vehicle ownership—especially among urban populations (Pucher and Renne 2004; Morrow 2002)—can severely hamper timely evacuation. Similarly, *Generational Living*, captured by the `B11002 Household Type by Relatives and Nonrelatives for Population in Households` dataset, reflects the impact of long-term residence and the stronger community ties on evacuation decisions (Smith and McCarty 2009; Mesa-Arango et al. 2012). Our measure of *Sense of Help*, derived from geographical mobility data (e.g., `B07001` and `B07410`), uses geographical mobility within the same county and remaining in the same residence for longer than a year; while larger immediate networks might reduce the likelihood of evacuating, a broader extended network (or bridging capital), moving from another state or from a county outside the state can significantly increase that liklehood (Collins et al. 2017).

In addition, family dynamics play a crucial role(measured via `B11003 Family Type by Presence of Age of Own Children`) provide insight into the logistical challenges of evacuation. Employment factors, drawn from datasets like `B24124` and `B19052`, further illuminate the socioeconomic pressures that may restrict evacuation options. In this regard, we also consider the `B23020 Mean Usual Hours Worked in Past 12 Months` dataset, as work obligations have been shown to reduce the likelihood of evacuation (Baker 1991). However, if the damage or conditions are hazardous enough, workers may evacuate at the last minute to closer destinations (see Mesa-Arango et al. 2012).

We also pay close attention to the vulnerability of demographic groups using data from `B090202 Relationship by Household Type 65+`, highlighting that older and disabled individuals with limited support are particularly at risk during disasters (Schmidlin and King 1995; Morrow 2002; Peek-Asa et al. 2003; Rosenkoetter et al. 2007).

Moreover, we consider the unique challenges faced by populations residing in group quarters, a category we refer to as "No Place to Go." Using data on group quarters information, we examine the circumstances of residents in college dormitories, military barracks, psychiatric institutions, and prisons (Vogt 1990; Quarantelli 1980). These settings present special concerns during evacuation. In particular, due to their specialized and time-sensitive needs and chronic understaffing during emergencies, nursing homes and long-term care facilities require careful scrutiny. As Flanagan et al. (2011) note, while census data includes numbers from nursing homes, these figures can be obscured by the surrounding population, as was tragically illustrated by the loss of over 30 residents at St. Rita's nursing home in St. Bernard Parish during flooding.

Environmental hazards also factor into our framework through the variable *Location Danger*,

which is informed by NOAA surge data. A recent study of Katrina-related flood deaths (see Flanagan et al. 2011; Jonkman et al. 2009) emphasized the relationship between physical hazards—such as the downflow from levee breaches, water velocity, and water depth—and the vulnerability of nearby populations. In this context, a 2-foot threshold was selected from data classified in two-foot increments because evidence suggests that initial flood mortality tends to occur above approximately 1 meter (Jonkman et al. 2009). Finally, we consider the receipt of evacuation notices as an important component of our multidimensional risk framework, recognizing that effective communication can significantly influence evacuation behavior (Vieira 2022).

Taken together, our hypothesis is that the decision to evacuate—and the subsequent need for disaster assistance—is driven by a confluence of resource availability, social connectivity, housing and employment characteristics, demographic vulnerabilities, and environmental hazards. By examining these interrelated factors, we aim to develop a nuanced understanding of disaster response behavior and inform more effective emergency management strategies.

Table 2.1 below summarizes the key features used in our analysis, organized by category. This comprehensive approach not only enriches our understanding of the determinants of evacuation behavior but also informs the development of more targeted emergency response strategies.

Table 2.1: Description of machine learning models used in the analysis.

Model	Description
Household Composition and Family Presence	• `ft_MAR_wchildE`: Prevalence of married couples with children. • `ft_MAR_wchild_U6E`: Married couples with children under the age of 6. • `Elderly_Living_With_Spouse`: Seniors living with their spouse. • `Single_Parents`: Households headed by a single parent with children. • `Eld_M_aloneE`: Elderly males living alone. • `Extended_household`: Households where extended family members reside together. • `Elderly_Living_Alone`: Elderly individuals living alone. • `Ind_liv_DE`: Individuals experiencing independent living difficulties.

Table 2.1 (continued)

Model	Description
Poverty and Transportation Access	• `hhpovertyE`: Number of individuals with incomes below the federally defined poverty line. • `ANY_HELPE`: Number of individuals receiving public assistance or other forms of help. • `unemployedE`: Unemployment rate among the population aged 16 years and over. • `severe_poverty`: Proportion of the population experiencing extreme poverty (e.g., earning under $35,000 annually and spending 50% or more of income on rent). • `poverty_rent`: Household income expressed as a percentage of gross rent over the past 12 months. • `hours_worked`: Mean usual hours worked in the past 12 months. • `hhveh_notenE`: Households, particularly with four or more residents, that have only one vehicle. • `hhveh_noneE`: Households without any available vehicles.

Table 2.1 (continued)

Model	Description
Community Dynamics & Household Stability	• `Spanish_speakE`: Number of Spanish-speaking residents, indicating potential language barriers. • `OwnerE`: Number of owner-occupied housing units, serving as a proxy for residential stability. • `Nearby_Help`: Availability of local support networks within the community. • `GroupQE`: Population residing in group quarters, such as dormitories, nursing homes, or correctional facilities. • `SameStateE`: Median age of residents who have moved from a different county within the same state. • `DifStateE`: Median age of residents who have moved from a different state.

CHAPTER

3

METHODS

3.1 Description of Models

Our research centers on developing a predictive framework that leverages diverse datasets —ranging from census statistics to real-time rescue requests—to inform emergency response planning. In summary, we have constructed a machine learning model that is designed to predict the vulnerability of different geographic regions during disasters. The goal is to provide actionable insights for decision-makers, helping pinpoint areas where assistance is most likely needed and to optimize the allocation of limited resources during critical times in particular the first 24–48 hours after a disaster when information is sparse.

To identify the best-performing model for predicting shelter-in-place likelihood and the need for assistance post disaster as a binary class in affected census tracts, we evaluated twelve different machine learning algorithms a summary of each is provided in Table 3.1.

A key component of our model development is the use of cross validation, a robust method for evaluating predictive performance. Rather than relying on a single split of the data into training and testing sets, we divide the data into several segments and iteratively train and test the model on different subsets and taking the mean accuracy of all subsets. This process ensures that our evaluation is not overly dependent on any particular partition of the data, thereby giving us a more reliable sense of how well the model will perform on new, unseen data.

By doing so, we can identify and address potential overfitting issues and adjust our approach accordingly.

Based on their accuracy, receiver operating characteristic (ROC) area under the curve (AUC) score, with particular attention to minimizing type II error, interpretability, and ability to handle complex interactions between variables. Cross-validation was employed to ensure robustness and to avoid over-fitting, ensuring that all parts of the dataset contributed to training and testing.

Table 3.1: Description of machine learning models used in the analysis.

Model	**Description**
Random Forest (RF)	We employed Random Forest as a robust ensemble learning method that combines multiple decision trees to enhance predictive accuracy. Its effectiveness in capturing nonlinear relationships and interactions among variables made it a strong candidate for this analysis (Rachka and Liu 2022).
Gradient Boosting Machines (GBM)	Gradient Boosting Machines were utilized to build sequential decision trees, with each tree correcting the errors of its predecessor. While highly effective for handling complex datasets, we noted that its interpretability posed challenges compared to simpler models (Rachka and Liu 2022).
Support Vector Machines (SVM)	SVMs were implemented to handle high-dimensional data by identifying a hyperplane that optimally separates the classes (shelter-in-place versus evacuate). This approach was particularly beneficial for datasets that were not linearly separable (Rachka and Liu 2022).
Neural Networks	We explored neural networks for their capacity to model highly nonlinear interactions among variables. However, their reliance on larger datasets, extended training times, and reduced interpretability presented challenges for this application (Rachka and Liu 2022).
K-Nearest Neighbor (KNN)	The K-Nearest Neighbor algorithm was considered for its simplicity and ability to predict outcomes based on the proximity of data points to their neighbors. It was particularly useful in identifying localized patterns where individuals might require assistance(Rachka and Liu 2022).

Table 3.1 (continued)

Model	**Description**
Stacking Classifier (SCLF)	Stacking classifiers were utilized to combine predictions from multiple base models through a meta-model, improving overall accuracy. This ensemble method leveraged diverse model strengths and minimized overfitting, making it a robust solution.
Logistic Regression (LR)	Logistic Regression was employed to provide probabilistic predictions, allowing for straightforward interpretability. It was used to estimate the likelihood of a land tract containing at least one person who would shelter in place (Rachka and Liu 2022).
Voting Classifier (VC)	The Voting Classifier aggregated predictions from multiple models using both hard and soft voting approaches. This method balanced accuracy and interpretability by integrating the strengths of individual models (Rachka and Liu 2022).
Decision Tree (DT)	Decision Trees were implemented for their intuitive structure, splitting data based on criteria to make predictions. Despite their simplicity, they effectively captured key decision-making processes (Rachka and Liu 2022).
Extreme Gradient Boosting (XGBoost)	XGBoost, an advanced version of Gradient Boosting, was employed for its regularization capabilities to prevent overfitting and its optimizations that sped up training. These features made it a powerful and efficient choice for complex datasets (Rachka and Liu 2022).
Bagging Classifiers (BC)	Bagging was applied using different base models, including Decision Trees (BC(DT)), Logistic Regression (BC(LR)), and KNN (BC(KNN)). These classifiers reduced variance by training multiple models on bootstrapped subsets of data, improving stability and accuracy (Rachka and Liu 2022).
AdaBoost Classifier (AB)	The AdaBoost algorithm was implemented to enhance accuracy by iteratively focusing on misclassified instances and updating their weights. This method effectively boosted the performance of weaker classifiers (Rachka and Liu 2022).

3.2 Model Selection and Evaluation

To determine the models best suited for the problem set, all models were initially run against the dataset to narrow the selection for further experimentation. Mean accuracy scores and their variability during 10-fold cross-validation were the primary criteria, with preference given to models showing reduced variability, higher mean accuracy, and minimal Type II errors (false negatives). Based on these measures, the top models included Random Forest, Logistic Regression, Gradient Boosting, SVM, and all Bagging Classifiers, which were retained for further experimentation. This initial testing phase occurred before Hyperparameter Tuning.

In Phase II, Hyperparameter Tuning was performed via grid search, and results were compared to pre-tuning performance. Most models exhibited performance improvements post-hyperparameter tuning. Random Forest, Logistic Regression, and Bagging Classifiers (using Decision Trees) showed notable increases in average accuracy. Ultimately, the Random Forest model emerged as the most suitable for this dataset due to its high accuracy, minimal outliers, low comparative variance, and reduced Type II errors. The hyperparameters tuned were the number of trees in the forest, which was found to be 105 trees, the maximum depth of the tree, found to be 15. This is not a large change from the original ensemble method, which is to be expected as the out of the box method is quite robust to noise due to averaging across many trees (Rachka and Liu 2022). This combination of characteristics made Random Forest the optimal choice for predicting shelter-in-place behaviors.

Our experiments indicate that combining multiple decision trees using the random forest algorithm results in a mean out-of-sample accuracy of 0.627 when evaluated with 10-fold cross-validation. After fine-tuning the hyper-parameters of the random forest, the mean accuracy improved to 0.652. Notably, this tuning led to several important enhancements: the ROC/AUC increased from 0.671 in the untuned model to 0.695 in the tuned version, and the Type I error rate dropped from 0.285 to 0.231, while the Type II error remained steady. In comparison to alternative models, the random forest algorithm achieved the highest ROC/AUC score and the lowest Type II error rates.

For additional context, if our model predicted a zero for every case, we would see an accuracy of 0.55. However, using 0.55 as a baseline for improvement is not operationally viable, as it would imply not searching at all. A more realistic baseline is 0.45, which approximates the current scenario in disaster response—where help is assumed to be needed universally across the affected area.

After maximizing accuracy, we chose to focus on further enhancing the ROC/AUC and reducing the Type II error. Lowering the Type II error is crucial because it minimizes the risk of false negatives—situations in which our model fails to identify areas where assistance is

needed. This focus is especially important during the first 24–48 hours after a disaster when information is extremely limited.

As shown in Figure 3.1, which presents the performance metrics and ROC/AUC scores across all the different machine learning models used in this study, the improvements achieved by tuning the random forest algorithm are particularly notable. The figure provides a comprehensive visual summary of how each model performed, underscoring the superior performance of the tuned random forest.

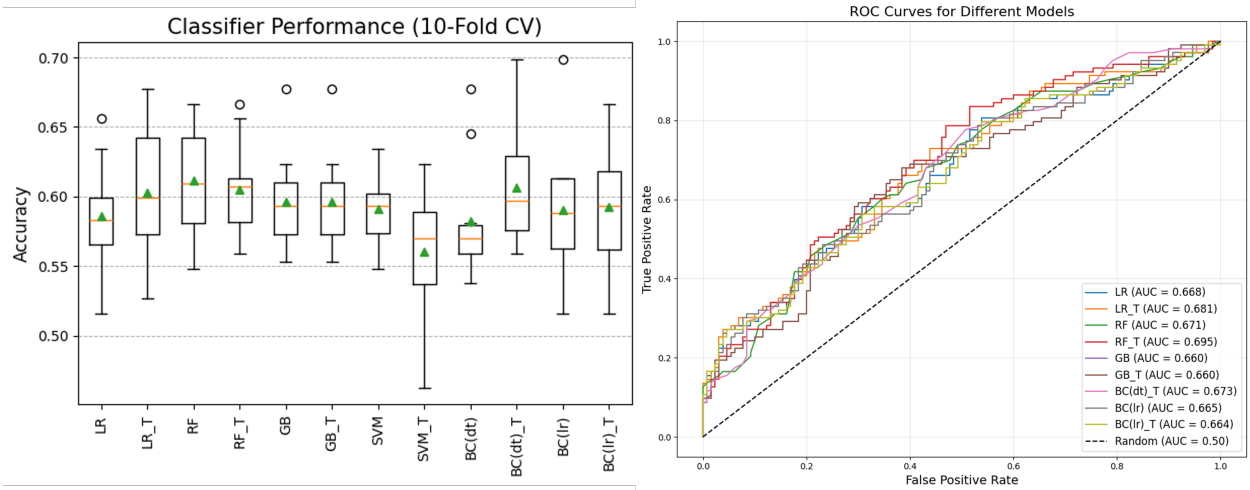

Figure 3.1: Performance metrics for the random forest model before and after hyperparameter tuning.

CHAPTER

4

RESULTS

Figure 4.1 presents a detailed map of rescue locations overlaid with census data that highlights several important covariates. The map shows spatial distributions of unemployment rates, the prevalence of married couples with children, households composed of four or more individuals with only one vehicle, and households with residents aged 65 and older. These layers of information provide a clear visual context of how socio-demographic factors interact with rescue operations, offering insights into community vulnerabilities and capacities during emergencies.

The map suggests that areas marked by higher unemployment rates tend to coincide with clusters of rescue locations as seen in Figure 4.1(a), perhaps reflecting the challenges faced by economically disadvantaged communities during crises. Additionally, regions with a higher concentration of married couples with children also appear to align with certain rescue hotspots shown by Figure 4.1(b). This pattern may indicate that family units, particularly those with children, find it difficult to evacuate and with additional people have lower resilience after the disaster, thereby necessitating prompt support.

Moreover, the census data suggests there is a correlation that households with four or more residents yet only one vehicle are particularly at risk; see Figure 4.1(c). Limited access to transportation in these areas can delay evacuations and complicate rescue efforts. Similarly Figure 4.1(d) shows there is a correlation with communities with a significant proportion

of households containing individuals aged 65 and older are especially vulnerable, as older residents often have a more difficult time evacuating and require additional support during emergencies, due to a higher likelihood of limited mobility and health complications. Together, these trends underscore the importance of incorporating demographic and economic variables into disaster response planning, ensuring that resources are efficiently allocated to those who need them the most.

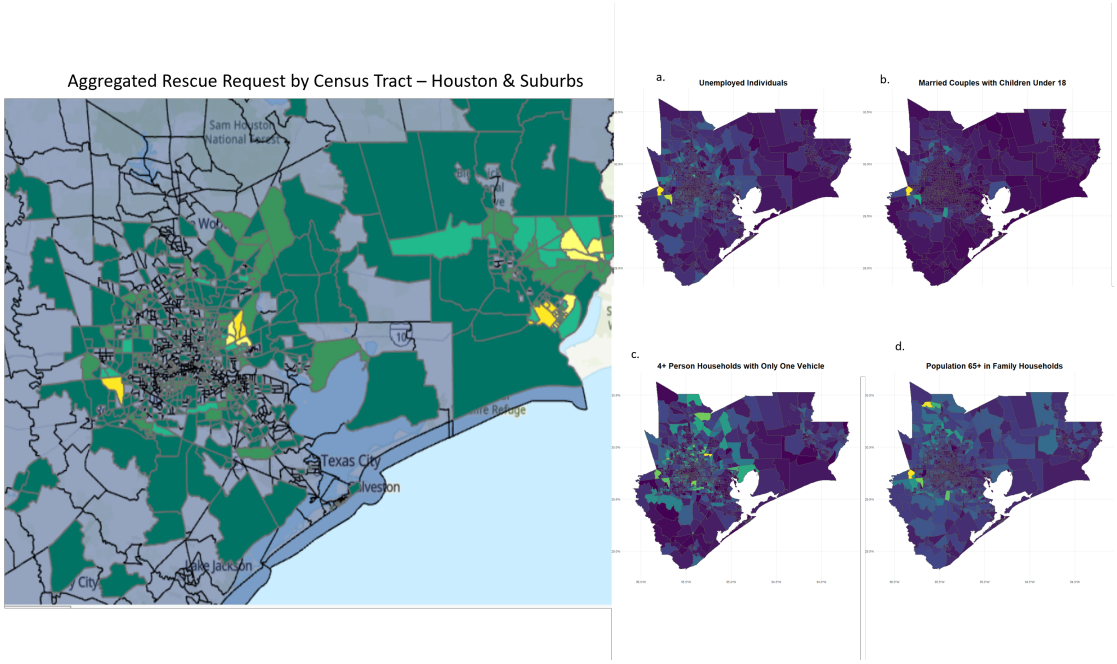

Figure 4.1: The left map shows aggregated rescue requests by census tract, while the four right subplots illustrate key socioeconomic covariates in Houston: (a) Unemployed Individuals, (b) Married Couples with Children Under 18, (c) 4+ Person Households with Only One Vehicle, and (d) Population 65+ in Family Households. These factors are analyzed in relation to rescue requests from Hurricane Harvey to identify communities that may be more vulnerable to future disasters due to social, economic, and logistical constraints. This data is acquired with the R `tidycensus` package and taken from the 2013–2017 American Community Survey

Figure 4.2 presents the feature importance analysis. We generate feature importance in order to visualize how much each variable contributes to the models predictive power, and which features have the most influence on the target variable, in this case the binary variable if someone needs assistance in a census tract after a disaster. We used the mean impurity decrease computed from all the decision trees in the forest, randomly shuffling the values of one feature at a time, features that cause a larger drop in accuracy are considered the most important (Rachka and Liu 2022)

We then organized the features into five distinct categories: Poverty and Financial Conditions, Transportation Issues, Household and Family Presence, Demographic and Household Stability, and Natural Hazard. The analysis reveals that the features related to Household and Family Presence are the most influential, with the top four most important features being Individual Living Difficulty, Single Parents, Extended Households, and Married with Children. This suggests that the structure and composition of households play a critical role in the context of our study.

Following these, the next significant feature falls under the category of Poverty and Financial Conditions, represented by the Unemployed variable. Transportation Issues are then highlighted, particularly the challenge posed by households with only one vehicle and four or more people. Finally, two key aspects of Demographic and Household Stability—namely, moving within the same state and moving within the same county or residing in the same home over the past year—also emerge as influential factors.

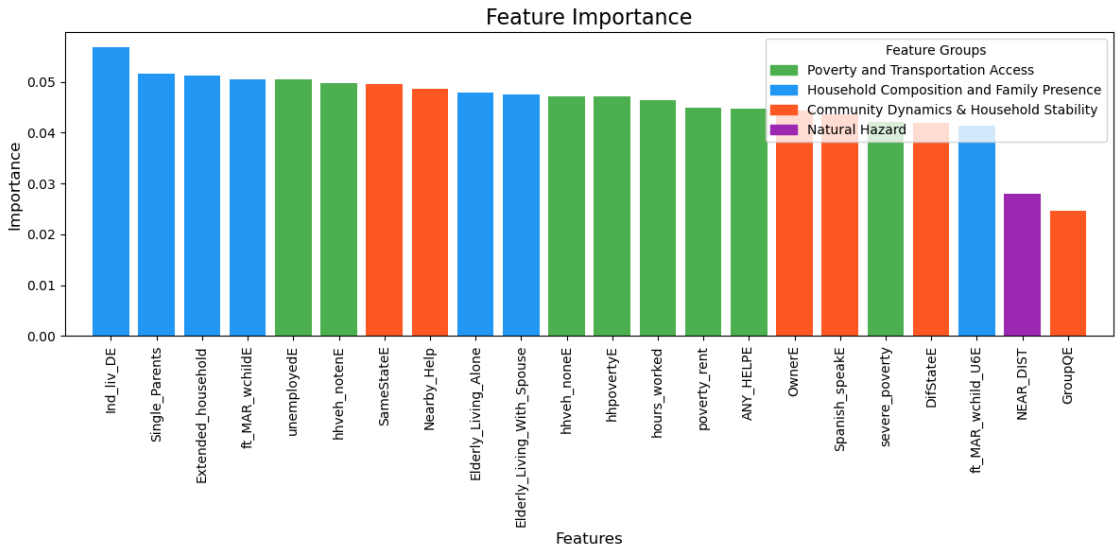

Figure 4.2: Feature Importance: The analysis is divided into four categories. Household and Family Presence dominates with the top four features (Individual Living Difficulty, Single Parents, Extended Households, and Married with Children), followed by Poverty and Poverty and Transportation Access (Unemployed and Household with one vehicle and 4+ people),and Community Dynamics & Household Stability (moving within the same state and moving within the same county or staying in the same home).

We generate Partial Dependence Plots (PDPs) to understand how a specific feature in a machine-learning model influences predictions while accounting for other variables. PDPs

help interpret complex models by showing the marginal effect of a feature on the predicted outcome, averaging over all other variables (Friedman 2001).

The partial dependency graphs, as shown in Figure 4.3, show the marginal effect of a socio-economic feature on the predicted outcome of the model in predicting the need for assistance after a disaster. In particular, our analysis highlights several important trends at the census tract level.

First we find that, shown in Figure 4.3(a) the likelihood of not evacuating and subsequently needing assistance increases as the number of people living under the poverty line in a census tract grows. This suggests that economically disadvantaged communities are particularly vulnerable in disaster scenarios. Similarly, Figure 4.3(b) shows the probability of households without sufficient vehicles rising sharply correlates with an increased need for help. However, this probability falls off as the number of such households nears 100, which may be explained by the urban context—areas with higher counts might also have better access to public transportation, thereby reducing dependency on private vehicles.

In addition, we find that a higher concentration of individuals experiencing difficulties living alone—see Figure 4.3(c)—whether due to ambulatory, vision, or hearing challenges—correlates strongly with an increased need for assistance. A notable threshold emerges for single parents in Figure 4.3(d): once a census tract exceeds 200 single-parent households, the predicted need for support rises sharply. As can be seen in Figures 4.3(i),(e),(f) the need for assistance among people aged 65 or older living alone, households with relatives and non-relatives, as well as couples with children, remains steady initially, it begins to increase significantly at different thresholds such individuals are present in a tract. This could indicate that, beyond this point, the capacity for mutual neighborhood support diminishes. A similar strong increase can be seen in unemployed graph in Figure 4.3(g), possibly showing the change over from more affluent communities where one parent is the caregiver while the other works to less affluent areas where neither adult in the household has employment. Finally in Figure 4.3(h), looking at mobility, people who have moved within the same state dramatically decreases, indicating the possible availability of easier evacuation back to the county where they recently moved from.

To conclude, a case study was conducted to assess the model's predictive power in a different geographic region. We applied our tuned random forest model—originally trained on Hurricane Harvey data—to predict potential rescue locations around Wilmington and the surrounding coastal areas of North Carolina for a hypothetical hurricane scenario (Figure 4.4). This test demonstrated the model's ability to generalize beyond the original training data, offering a practical application for identifying regions where emergency assistance would likely be required.

This case study highlights how predictive mapping can be rapidly deployed in post-disaster

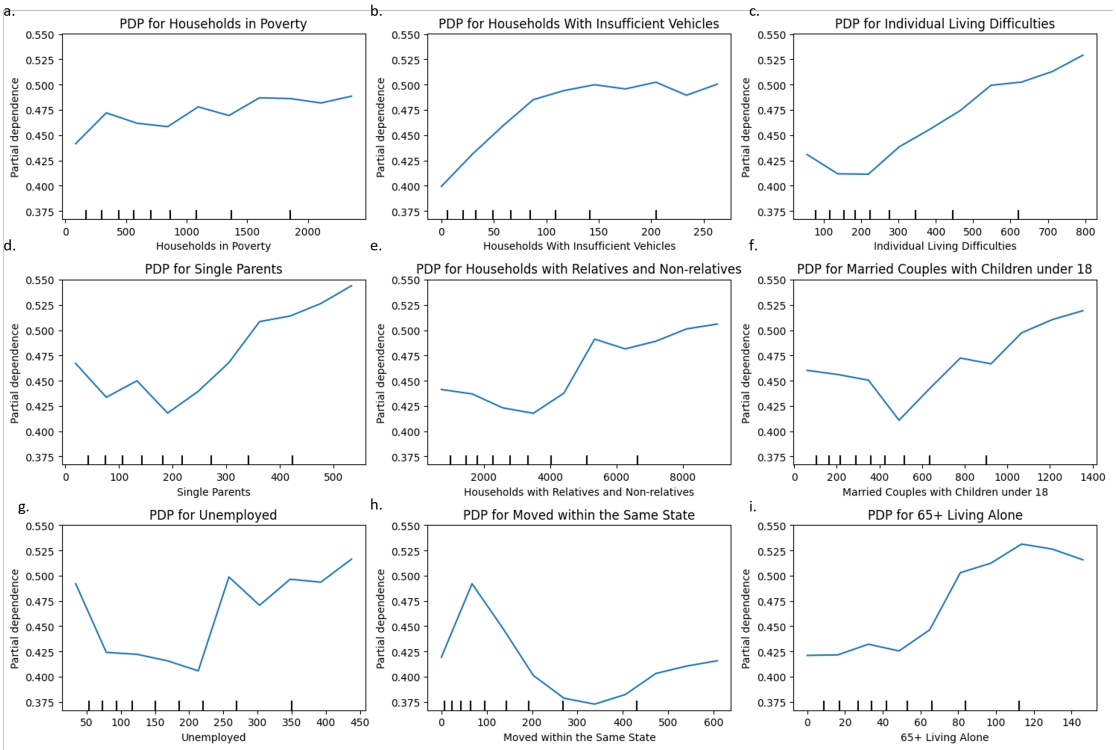

Figure 4.3: This figure presents Partial Dependence Plots (PDP) illustrating the marginal effect of various socioeconomic factors on the probability of sheltering in place during Hurricane Harvey. PDPs provide insight into how a single predictor influences the model's predictions while averaging out the effects of other variables.(a) Households in poverty, (b) Households with insufficient vehicles, (c) Individuals with living difficulties, (d) Single-parent households, (e) Households with relatives and non-relatives, (f) Married couples with children under 18, (g) Unemployed individuals, (h) Individuals who moved within the same state, and (i) Individuals aged 65+ living alone. These PDPs highlight the importance of demographic and socioeconomic characteristics in predicting disaster response behaviors.

scenarios. By analyzing the storm's projected path and determining which counties are most affected, relevant census data can be collected, processed, and incorporated into a model to generate a real-time map. This automated process enables decision-makers to visualize areas with a higher likelihood of residents needing assistance, even without prior local knowledge. Such an approach reduces reliance on potentially outdated or incomplete situational reports, streamlining the allocation of emergency resources.

Furthermore, this method supports both immediate disaster response and long-term planning. Emergency management agencies can use these maps to optimize resource distribution, plan rescue operations, and anticipate areas of heightened vulnerability before a storm makes landfall. By leveraging data-driven insights, decision-makers can make more informed and effective choices in crises, potentially saving lives and mitigating damage.

Figure 4.4 illustrates the predicted rescue locations in the Wilmington area. The map underscores how the tuned random forest model synthesizes multiple variables to offer a clear, actionable visual of potential rescue hotspots that leaders can use in an incident command post to assist with search planning.

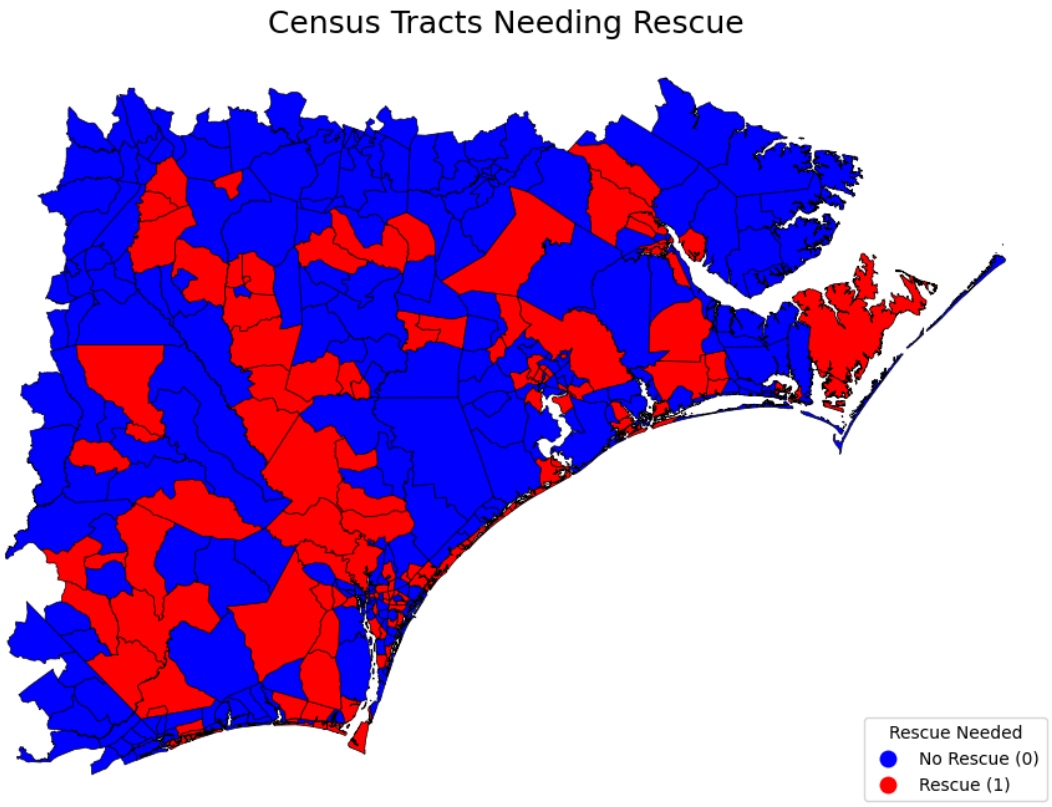

Figure 4.4: (color online) Predicted rescue locations in the Wilmington area using the tuned random forest model. The model, trained on Hurricane Harvey data, forecasts potential rescue sites for a hypothetical hurricane impacting North Carolina's coast.

CHAPTER

5

DISCUSSION

Our machine learning models, tested using tenfold cross-validation, achieved an average out-of-sample accuracy of 65% and a ROC-AUC score of 69%. When selecting the final model, we prioritized accuracy alongside minimizing Type I error. This decision was guided by the operational context of disaster response, where it is preferable to overestimate the presence of individuals needing assistance (false positives) than to risk overlooking someone in need (false negatives). This aligns with the principle that in early rescue coordination, allocating resources to a potential site, even if unnecessary, is less risky than failing to send aid where it is critically needed.

We recognize, however, that in environments with extreme resource constraints, it may also be essential to consider Type II errors, as these represent wasted effort and resources. Balancing these trade-offs is a vital consideration in optimizing disaster response strategies.

A significant limitation of this project was the dataset's quality, which introduced challenges to model effectiveness. The data was sourced from a self-reporting tool, which relies on individuals being aware of and accessing the tool. While Hurricane Harvey presented a unique scenario where 911 call centers were inoperable, our dataset does not likely represent the entire population of those in need. Additionally, certain records aggregated multiple individuals under a single data point based on comment details, further complicating the data's representativeness. We hypothesize that enhanced data collection from first-responder agencies and

nonprofit organizations, including detailed records and metadata such as timestamps, would significantly improve model performance and the depth of analyses, such as temporal trends in rescue requests.

Despite the model's modest accuracy, its true utility lies in supporting early disaster response efforts, that are more beneficial to the long term success of the response. By providing data-driven insights at the onset of a disaster, the model enhances "time-zero" knowledge—critical for emergency crews and non-profits deploying from across the country without prior familiarity with the affected area. This enables more informed decision-making when time and resources are scarce.

Currently, most disaster response is reactive rather than proactive, and even the proactive measures taken often rely on a combination of local knowledge and fragmented incoming reports. In the crucial first 24 to 48 hours after a storm, reliable information is scarce, and the little that is available can quickly become outdated. As a result, early planning and resource allocation decisions are often made with incomplete or inaccurate data, leading to inefficiencies that can hinder response efforts. By leveraging a model such as this, decision makers such as the Incident Commander, Operations Section Chief, and Planning Section Chief working from an Incident Command Post can make faster, more informed decisions, ultimately improving outcomes for affected communities and reducing the overall impact of the disaster.

CHAPTER

6

CONCLUSIONS

6.1 Summary of Results

In this study, we developed a predictive model to help identify areas where individuals may need assistance after a disaster, focusing on factors that influence evacuation behavior and vulnerability. We trained machine learning models using a wide range of socio-economic, demographic, and geographic features, achieving a mean out-of-sample accuracy of 65% and a peak ROC-AUC score of 69%. Notably, our model performed better than many alternatives, especially when fine-tuned to minimize Type I errors.

The results show that socioeconomic factors like poverty, lack of vehicle access, and family dynamics are strong predictors of the likelihood of evacuation. Moreover, vulnerabilities specific to elderly individuals, people living alone, and residents in group quarters such as nursing homes or prisons further exacerbate the need for assistance. These insights highlight the importance of targeted interventions in disaster preparedness, acknowledging that certain groups face more significant challenges in evacuation and recovery. However, the model's performance also revealed some areas for improvement.

Despite its promising results, challenges in data quality, particularly the reliance on self-reported data and the limitations faced during Hurricane Harvey, impacted the accuracy and representativeness of the model. To address this, future work should explore integrating

additional data from first-responder agencies and non-profit organizations. Better data could improve the model's accuracy and its ability to predict locations of rescues more reliably.

Nonetheless, the model's ability to offer valuable insights in the early stages of a disaster response, when resources are limited, is crucial. By providing a data-driven approach, it aids in allocating resources efficiently and minimizing the response time, ultimately improving disaster preparedness and response strategies in critical situations.

6.2 Future Work

A key area for future work involves leveraging Large Language Models (LLMs) to enhance the extraction of relevant information from the "information" feature in the disaster dataset. This feature likely contains valuable text data from distress calls, shelter-in-place requests, and emergency reports, which can be difficult to analyze manually due to the volume and complexity of the unstructured text. By utilizing LLMs, we can automate the extraction of key keywords, phrases, and relevant entities from these reports, thereby improving the efficiency and accuracy of our analyses.

There is also a need to better understand the time data associated with the inputs in order to add a temporal component and gain more insight into how quickly after the disaster takes place, populations start reaching out asking for help. This could help differentiate between life-threatening and nonlife-threatening emergencies.

A third area for future work is by utilizing user input to add data points to the same location, for example a user may input that they are stranded with two other adults and three children, this would give the dataset a better sense of how many people in each census tract actually need assistance. The field of disaster management and mitigation is constantly expanding and with it better predictive models can assist in saving lives, time and resources.

Finally, future work includes leveraging this predictive model alongside an optimization framework for disaster preparation and response. Potential applications include optimizing mass transportation and decreasing traffic congestion to ensure the efficient evacuation of residents, enhancing disaster relief supply chain logistics to streamline the distribution of critical resources, and improving shelter placement strategies, police and EMS placement, and medical facilities to maximize accessibility in areas with higher areas of concentration of people who decided to shelter in place and need assistance after a disaster.

REFERENCES

Baker, E. J. (1991). Hurricane evacuation behavior. *International Journal of Mass Emergencies & Disasters*, 9(2):287–310.

Brodie, M., Weltzien, E., Altman, D., Blendon, R. J., and Benson, J. M. (2006). Experiences of Hurricane Katrina evacuees in Houston shelters: Implications for future planning. *American Journal of Public Health*, 96(8):1402–1408. https://ajph.aphapublications.org/doi/full/10.2105/AJPH.2005.084475.

Collins, J., Ersing, R., and Polen, A. (2017). Evacuation decision-making during Hurricane Matthew: An assessment of the effects of social connections. *Weather, Climate, and Society*, 9(4):769–776. https://journals.ametsoc.org/view/journals/wcas/9/4/wcas-d-17-0047_1.xml.

Cutter, S. L., Boruff, B. J., and Shirley, W. L. (2003). Social vulnerability to environmental hazards. *Social Science Quarterly*, 84(2):242–261.

FEMA (2017). Historic disaster response to Hurricane Harvey in Texas. https://www.fema.gov/press-release/20250121/historic-disaster-response-hurricane-harvey-texas.

Flanagan, B. E., Gregory, E. W., Hallisey, E. J., Heitgerd, J. L., and Lewis, B. (2011). A social vulnerability index for disaster management. *Journal of Homeland Security and Emergency Management*, 8(1). https://doi.org/10.2202/1547-7355.1792.

Friedman, J. H. (2001). Greedy function approximation: A gradient boosting machine. *Annals of Statistics*, 29(5):1189–1232.

Jonkman, S. N., Maaskant, B., Boyd, E., and Levitan, M. L. (2009). Loss of life caused by the flooding of New Orleans after Hurricane Katrina: Analysis of the relationship between flood characteristics and mortality. *Risk Analysis*, 29(5):676–698. https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1539-6924.2008.01190.x.

Logan, T. M. and Guikema, S. D. (2020). Reframing resilience: Equitable access to essential services. *Risk Analysis*, 40(8):1538–1553.

Maas, P., Almquist, Z., Giraudy, E., and Schneider, J. W. (2020). Using social media to measure demographic responses to natural disaster: Insights from a large-scale Facebook survey following the 2019 Australia bushfires. http://arxiv.org/abs/2008.03665.

Mesa-Arango, R., Hasan, S., Ukkusuri, S. V., and Murray-Tuite, P. (2012). Household-level model for hurricane evacuation destination type choice using Hurricane Ivan data. *Natural Hazards Review*, 14(1):11–20. https://ascelibrary.org/doi/10.1061/%28ASCE%29NH.1527-6996.0000083.

Metaxa-Kakavouli, D., Maas, P., and Aldrich, D. P. (2018). How social ties influence hurricane evacuation behavior. *Proceedings of the ACM on Human-Computer Interaction*, 2:1–16. https://dl.acm.org/doi/10.1145/3274391.

Morrow, B. H. (2002). Identifying and mapping community vulnerability. *Disasters*, 23(1):1–18. https://onlinelibrary.wiley.com/doi/abs/10.1111/1467-7717.00102.

Morss, R. E., Turner, L. R., Tate, E., Wilhelmi, O., and Dabbers, N. (2016). An examination of factors that influence hurricane evacuation decision-making and behavior: A review of the literature. *Natural Hazards Review*, 17(3):04016002.

Nakagawa, Y. and Shaw, R. (2004). Social capital: A missing link to disaster recovery. *International Journal of Mass Emergencies & Disasters*, 22(1):5–34. http://journals.sagepub.com/doi/10.1177/028072700402200101.

National Hurricane Center (2018). Tropical cyclone report: Hurricane Harvey. Available at: https://www.nhc.noaa.gov/data/tcr/AL092017_Harvey.pdf.

Peek-Asa, C., Ramirez, M., Seligson, H., and Shoaf, K. (2003). Seismic, structural, and individual factors associated with earthquake related injury. *Injury Prevention*, 9(1):62–66. https://injuryprevention.bmj.com/content/9/1/62.

Pucher, J. and Renne, J. L. (2004). URBAN-RURAL DIFFERENCES IN MOBILITY AND MODE CHOICE: EVIDENCE FROM THE 2001 NHTS. https://www.vtpi.org/pucher_ur.pdf.

Quarantelli, E. L. (1980). Evacuation behavior and problems: Findings and implications from the research literature. Disaster Research Center.

Rachka, S. and Liu, Y. (2022). Machine learning with pytorch and scikit-learn. *Birmingham: Packt Publishing*.

Rosenkoetter, M. M., Covan, E. K., Cobb, B. K., Bunting, S., and Weinrich, M. (2007). Perceptions of older adults regarding evacuation in the event of a natural disaster. *Public Health Nursing*, 24(2):160–168. https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1525-1446.2007.00620.x.

Sankar, A. R., Doshi, P., and Goodie, A. (2020). Evacuate or not? a POMDP model of the decision making of individuals in hurricane evacuation zones. In *Proceedings of The 35th Uncertainty in Artificial Intelligence Conference*, pages 669–678. PMLR. https://proceedings.mlr.press/v115/sankar20a.html.

Schmidlin, T. W. and King, P. S. (1995). Risk factors for death in the 27 March 1994 Georgia and Alabama tornadoes. *Disasters*, 19(2):170–177. https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1467-7717.1995.tb00367.x.

Smith, S. K. and McCarty, C. (2009). Fleeing the storm(s): an examination of evacuation behavior during Florida's 2004 hurricane season. *Demography*, 46(1):127–145. https://doi.org/10.1353/dem.0.0048.

Vieira, L. (2022). Limited english proficiency individuals: Assessing the inclusion of risk communication aimed at LEP populations in emergency and disaster management plans across the United States. Master's thesis, Georgetown University, Washington, D.C. https://repository.digital.georgetown.edu/handle/10822/1079094.

Vogt, B. M. (1990). Evacuation of institutionalized and specialized populations. Technical report, Oak Ridge National Lab.(ORNL), Oak Ridge, TN (United States).

Yabe, T. and Ukkusuri, S. V. (2020). Effects of income inequality on evacuation, reentry and segregation after disasters. *Transportation Research Part D: Transport and Environment*, 82:102260. https://www.sciencedirect.com/science/article/pii/S1361920919311101.